Contents

Chapter 9 — Online Content

Articles

- An Optimized Grounded Base Oscillator Design for VHF/UHF by Dr. Ulrich Rohde, N1UL, and Ajay Poddar, AC2KG
- Automatic Tracking Filter for DDS Generator by Riccardo Gionetti, IØFDH
- Calculation of FM and AM Noise Signals of Colpitts Oscillators in the Time Domain by Dr. Ulrich Rohde, N1UL
- Crystal Test Oscillators by Fred Brown, W6HPH
- Frequency Synthesis: Current and Future Projections by Alexander Chenakin
- Novel Grounded Base Oscillator Design for VHF-UHF by Dr. Ulrich Rohde, N1UL
- Optimized Oscillator Design by Dr. Ulrich Rohde, N1UL
- Oscillator Design Using LTSpice by David Stockton, GM4ZNX (also see LTSpice files in SwissRoll folder)
- Oscillator Phase Noise by Dr. Ulrich Rohde, N1UL
- Programming the DDS AD9850 Signal Generator Module by James Kretzschmar, AE7AX
- Simulation of the Low Noise Oscillator from Solid State Design for the Radio Amateur by Linley Gumm, K7HFD
- Some Thoughts on Crystal Oscillator Design by Dr. Ulrich Rohde, N1UL
- Some Thoughts on Designing Very High Performance VHF Oscillators by Dr. Ulrich Rohde, N1UL
- What You Always Wanted to Know About Colpitts Oscillators by Dr. Ulrich Rohde, N1UL, and Anisha M. Apte

Tools and Data

- Simulation files for Oscillator Design Using LTSpice (SwissRoll folder)

Chapter 9
Oscillators and Synthesizers

RF signals all begin with oscillators — circuits that generate a periodic output signal without any input signal. In general, we use sinusoidal waveforms for communication signals, and square-wave or pulse outputs for digital signals (clocks). Amateurs care greatly about oscillator design. In a crowded band, a stable and clean oscillator is necessary for the receiver to resolve each of the signals present. The same is required of oscillators in transmitters so that our signals use only the amount of spectrum necessary for communication. While this chapter includes a section on phase noise, the **RF Techniques** chapter discusses noise and noise metrics more generally.

This chapter has been updated by Bob Clarke, N1RC, including material from Earl McCune, WA-6SUH; Ulrich Rohde, N1UL; and David Stockton, GM4ZNX. These authors and others have contributed articles on phase noise and oscillator design that are included in the online content for the *Handbook*.

The sheer number of different oscillator circuits seen in the literature can be intimidating, but a manageably small number of oscillator types will cover all but very special requirements. Look at an oscillator circuit and "read" it: What form of filter — resonator — does it use? What form of amplifier? How have the amplifier's input and output been coupled into the filter? How is the filter tuned? These questions are easily answered, putting oscillator types into appropriate categories to make them understandable. The questions themselves make more sense when we understand the mechanics of oscillation, in which resonance plays a major role.

9.1 How Oscillators Work

9.1.1 Resonance

Figure 9.1 shows a simple LC circuit, also called a *tank circuit* or resonator. ("Tank" is used because the L and C components store energy.) The circuit is resonant (oscillates) at the frequency where the inductive reactance and capacitive reactance are equal, that is:

$Z = X_L + X_C = 0$. For this to occur, the magnitudes of the two reactances must be equal, or

$|X_L| = |X_C|$. Since $|X_L| = 2\pi fL$ and $|X_C| = 1/2\pi fC$, we can solve for the frequency, f, which yields

$$f^2 = \frac{1}{4\pi^2 LC}$$

Taking the square root of both sides give the familiar formula for the resonant frequency of an LC circuit,

$$f = \frac{1}{2\pi\sqrt{LC}}$$

Figure 9.2 illustrates the oscillating voltage and current in the circuit created by the regular exchange of energy between the components in Figure 9.1.

9.1.2 Maintained Resonance

The LC values of a tank circuit determine the resonant frequency but like all real-world circuits, the inductor and capacitor have real-world losses; the energy lost per cycle must be replaced for oscillation to occur. What we need is a way to make up for the resonator losses, but only just enough to maintain the oscillation indefinitely. By adding an amplifier to the tank circuit or *resonator*, we can take a small amount of the energy from the resonator, amplify it, and inject a part of the amplified tank signal back into the resonator. The resonator acts as a filter and the energy fed back to the amplifier input is *feedback*. The amount of feedback is represented as β. This process is shown in **Figure 9.3A**.

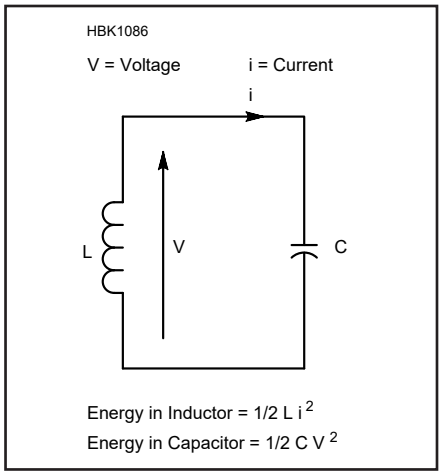

Figure 9.1 — The LC circuit stores energy in the inductor's magnetic field and capacitor's electric field.

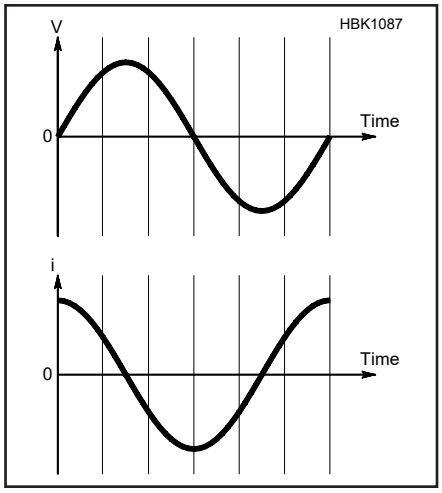

Figure 9.2 — The regular exchange of energy between the components in Figure 9.1 creates the oscillating voltage and current in the circuit.

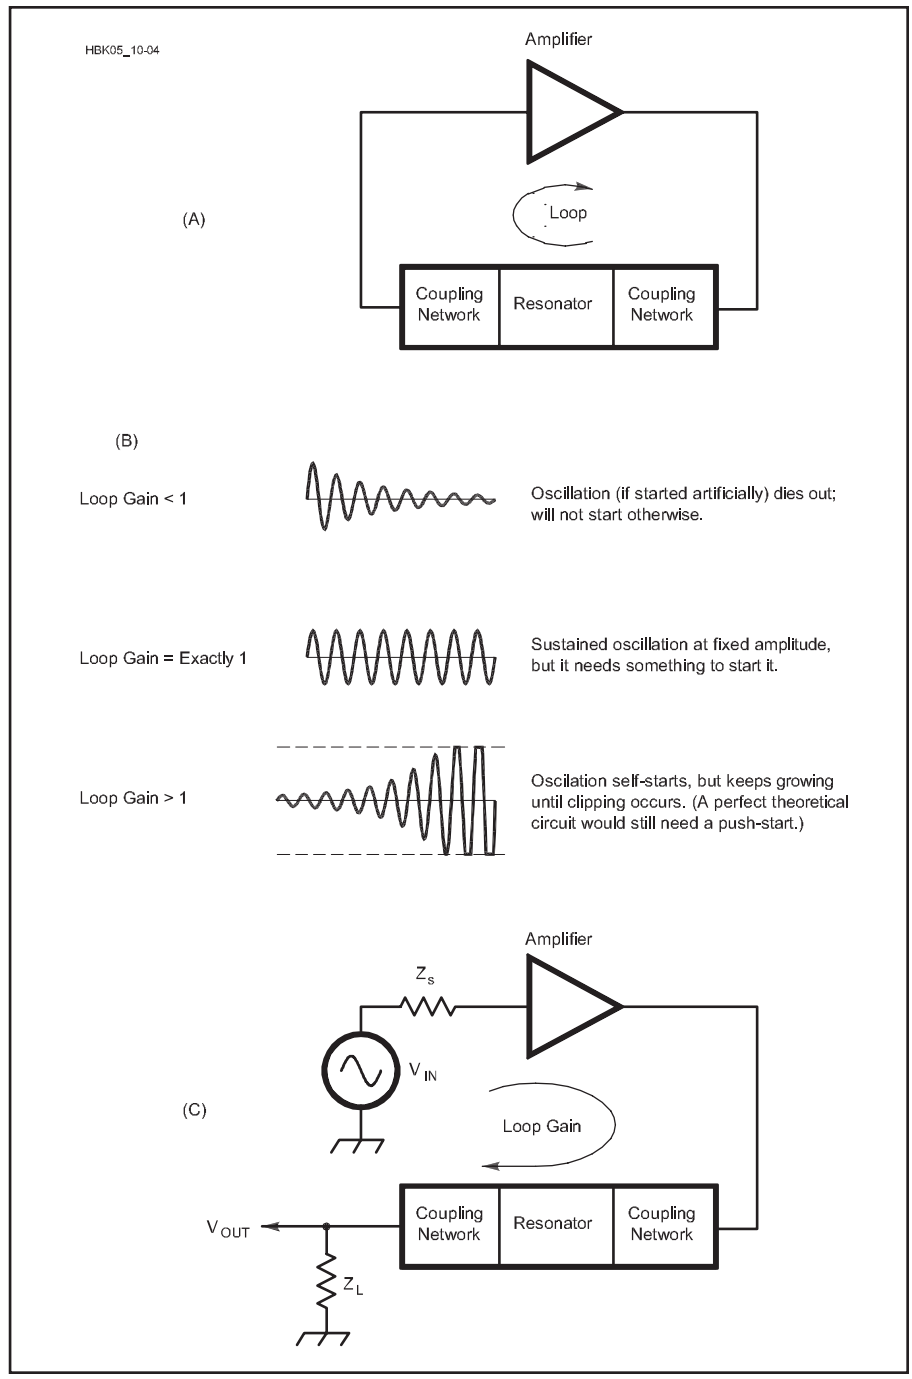

Figure 9.3 — A practical oscillator requires networks to couple power in and out of the resonator (A). At (B), we see the amount of loop gain determines whether the oscillations will die out (loop gain < 1), grow until the waveform is no longer a sine wave (>1), or sustain at a steady amplitude (=1). Breaking the loop, inserting a test signal and measuring the loop's overall gain allows us to determine whether the system can oscillate, sustain oscillation or clip (C).

Unlike in general amplifier design, here we are not interested in maximizing power transfer. Instead, we want to couple just enough energy into the resonator to overcome its losses, and to take only the minimum amount of energy out of the resonator needed to generate the restoring energy. Thus, Figure 9.3A shows coupling networks instead of matching networks. We end up with a loop: The amplifier input comes from the output of the resonator, and a portion of the amplifier output goes to the input of the resonator as feedback.

The trick to oscillator design is to get the coupling networks and the amplifier gain, *A*, working together just right. What this means is shown in Figure 9.3B. If we want to get a sine wave output, we need to get the loop gain — gain computed as the signal passes through the combination of amplifier, resonator, and both coupling networks — exactly equal to one, also called *unity*. If the loop gain is even slightly less than one, then not enough energy is supplied to overcome resonator losses and the damping still happens. If the loop gain is greater than one, then the magnitude of the sine wave will grow until the waveform is no longer sinusoidal.

There is more to making an oscillator work than getting the loop gain right. We also need to be sure that the energy from the amplifier output is applied to the resonator having the correct phase alignment with the oscillating signal in the resonator. Taken together these requirements are known as the *Barkhausen criteria*: to achieve oscillation, 1) the magni-

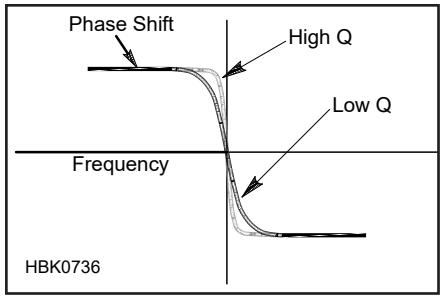

Figure 9.4 — Phase shift through the resonator depends on resonator Q. The higher resonator Q, the more abruptly phase changes through the resonator as frequency changes.

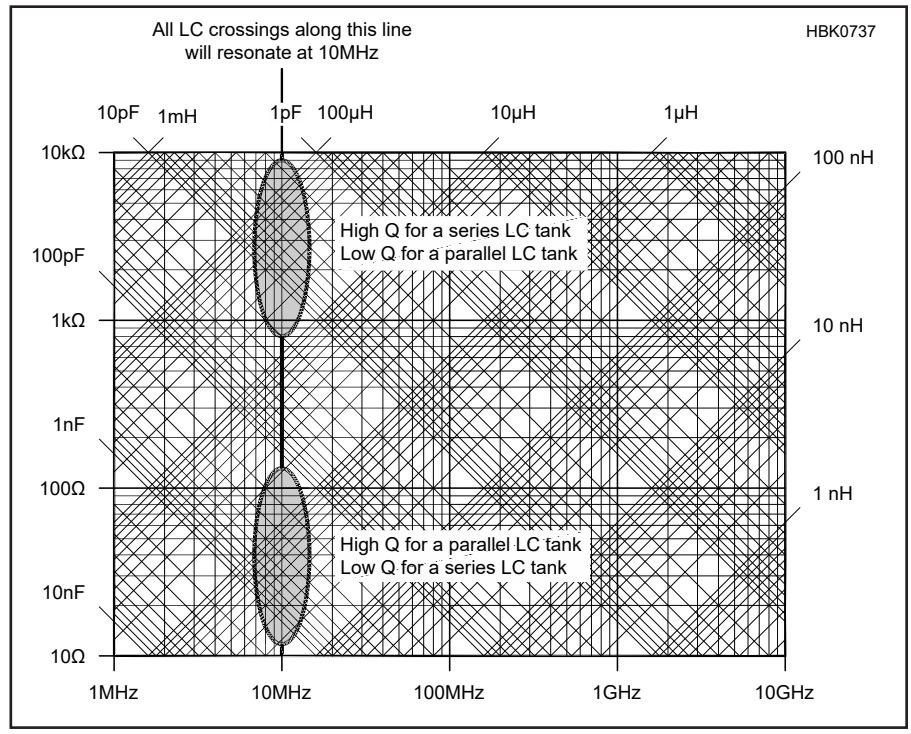

Figure 9.5 — There are an infinite number of combinations of L and C for a given resonant frequency. For a series-LC tank circuit, larger ratios of L to C result in higher values of Q and vice versa for parallel-LC tank circuits.

tude of the loop gain must be exactly equal to unity, $|\beta A| = 1$; and 2) the phase shift through the loop, $\angle \beta A$, must be zero or an integer multiple of 360°. Both of these are critically important. If either of these criteria is not met, the circuit will not oscillate. (See the following section on Oscillator Start-Up regarding loop gain greater than unity.)

How do we find out if we are meeting the Barkhausen criteria? We need to break the loop, usually at the input of the amplifier as shown in Figure 9.3C. We apply a signal slightly below the desired oscillation frequency, then sweep the frequency through and slightly past the desired oscillation frequency. During the sweep we monitor both the magnitude and phase response at the output of the resonator output coupling network. If everything is right, then at the desired frequency of oscillation the input and output signals will look exactly the same on an oscilloscope. If that doesn't happen, we have work to do.

Getting the phase response of the loop correct is usually the harder problem. Much of this difficulty comes from the many types of resonators that can be used, including:

- LC tank circuit (parallel or series)
- Quartz crystal (and other piezoelectric materials)
- Transmission line (stripline, microstrip, open-wire, coax, and so on)
- Microwave cavities, YIG spheres, dielectric resonators
- Surface-acoustic-wave (SAW) devices
- Ceramic Resonators

Each of these also can be called a filter. It is true that any filter can also be used as a resonator in an oscillator. The more complicated phase responses of filters make meeting both of the Barkhausen criteria more challenging, but certainly possible.

The phase response of the resonator and the energy lost per cycle are related to its Q (quality factor — see the **Radio Fundamentals** chapter). If we have a parallel LC tank circuit, then at frequencies well below resonance the tank circuit acts inductive ($X_L \ll X_C$) and the current lags the voltage. At frequencies well above resonance the tank acts capacitive ($X_C \ll X_L$) so the current leads the voltage. In between, the phase shift of the current relative to the voltage depends on the actual frequency. The rapidity of the phase shift with changing frequency is governed by resonator Q: the phase shift happens very rapidly when Q is high. This is shown in **Figure 9.4**. When the resonator Q is low then the phase changes much more slowly at different frequencies.

If we have a series resonator then the phase change profile is reversed (capacitive to inductive) but otherwise the shape is the same. There are an infinite number of combinations of L and C values that provide the same resonant frequency. If we are interested in high Q, for a series tank circuit we want to select large inductor values and small capacitor values. For a high-Q parallel tank circuit we want the opposite — large capacitor values and small inductor values. **Figure 9.5**, a modified version of the reactance vs frequency chart in the **Radio Fundamentals** chapter, shows these regions for a resonant frequency of 10 MHz.

9.1.3 Oscillator Start-Up

Looking at Figure 9.3B we see that the oscillation will build up in magnitude only when the loop gain is greater than unity. Thus, this is an important additional criterion for oscillator design. The loop gain must be slightly greater than unity for it to start oscillating. Otherwise, we have the undesired condition in which there is no signal in the resonator, so we sample nothing from it and put nothing back into it.

Still, we need one more thing — noise. Some kind of signal needs to be injected into the resonator for the oscillation to start building up. We are fortunate that all amplifiers have output noise when power is applied, even if there is no input signal. It is this noise that allows any oscillator to start. Amplifier noise initiates a resonant signal in the tank circuit, and if loop gain is slightly greater than unity, more signal is fed back into the resonator than the resonator loses. The output signal builds up until something stops the process.

Here we take advantage of an amplifier characteristic where gain is reduced as the output reaches some predefined value. This is referred to as *compression* or *limiting*: When the output signal gets large enough, the amplifier gain is reduced a little bit. When compression balances steady output amplitude with a loop gain of exactly unity, the

oscillator has reached steady state operation.

When the loop gain and phase shift have stabilized at the correct values, the oscillation will be sustained at a steady value. This is called a *continuous wave* (CW). Amateurs use the term to refer to Morse transmissions (see the section on On-Off Keying (OOK) in the **Modulation** chapter) but it really means a steady, unchanging signal with a constant amplitude and frequency.

Why don't we have noise at the oscillator output instead of a sine wave? Because of the filtering action of the resonator, the filtered noise waveform appears mostly sinusoidal. The narrower the bandwidth of the resonator (the higher its Q), the more sinusoidal the waveform, but it actually consists of very well-filtered noise. It is impossible to get only a pure sine wave. Noise is inevitable, and tremendous efforts are spent in reducing this noise. This is discussed in this chapter's section on Phase Noise. A demonstration by David Stockton, GM4ZNX, of oscillator design is based on a simulation of how oscillators start up by the free *LTSpice* simulation program, which is available in the book's online material.

9.2 LC Variable Frequency Oscillator (VFO) Circuits

In this section we introduce some important oscillator circuits and provide well-tested guidelines on how to successfully build one. The following section presents a design example.

There are thousands of oscillator circuits, but just a few principal designs. One of the principal oscillator circuits is shown in **Figure 9.6**. An LC tank circuit is shown in Figure 9.6A. The usual single capacitance is replaced with two series capacitors having the same equivalent capacitance, C_{EQUIV}. The two capacitors act as an ac voltage divider, so that the voltage V1 at the midpoint is less than the total ac voltage of the tank, V_{TANK}.

What happens when we try to force V1 to be the same as V_{TANK} by adding an amplifier with unity voltage gain as shown in Figure 9.6B? The voltage division action of the capacitive divider does not change. If V1 is forced to be equal to V_{TANK} by the amplifier, then V_{TANK} will become greater than before. But then V1 will take on the new value of V_{TANK} and so forth. This is positive feedback and we have created an oscillator. Connecting an amplifier with low voltage gain to a split tank capacitance leads to the *Colpitts* group of oscillator designs.

It is certainly possible to "split" the inductor instead of the capacitor in a tank circuit and apply the same amplifier trick as before. This is shown in Figure 9.6C and describes the *Hartley* group of oscillators.

9.2.1 Basic Oscillator Circuits

(See also the papers "What You Always Wanted to Know About Colpitts Oscillators" by Ulrich Rohde, N1UL, and Anisha M. Apte, and "Calculation of FM and AM Noise Signals of Colpitts Oscillators in the Time Domain" by Ulrich Rohde, N1UL, included with the online material accompanying this book.)

The LC oscillators used in radio equipment are usually designed to be variable frequency oscillators (VFOs). Tuning is achieved by either varying part of the capacitance of the resonator or, less commonly, by using a movable magnetic core to vary the inductance. Since the early days of radio, there has been a huge quest for the ideal, low-drift VFO. Amateurs and professionals have made immense efforts in this pursuit. A brief search of the literature reveals a large number of VFO designs, many accompanied by claims of high stability. The quest for stability has been solved by the development of low-cost frequency synthesizers, which give crystal-controlled stability. Synthesizers are generally noisier than a good VFO, so the VFO still has much to offer in terms of signal cleanliness, cost, and power consumption, making it attractive for homebrew construction. No single VFO circuit has any overwhelming advantage over any other — component quality, mechanical design and the care taken in assembly are much more important.

Let's take a look at three popular oscillator circuits stripped of any non-essential components so they can be easily compared. The original Colpitts circuit (**Figure 9.7**) is now often referred to as the parallel-tuned Colpitts because its series-tuned derivative (**Figure 9.8**) has become the most common. The Hartley version of the oscillator is shown in **Figure 9.9**. All three of these circuits use an amplifier with a voltage gain less than unity, but large current gain. The N-channel JFET source follower shown is the most popular current choice for the transistor amplifier circuit.

A rule of thumb is that C3 and C4 in Figures 9.7 and 9.8 should be equal and valued such that their $X_C = 45\ \Omega$ at the operating frequency; likewise, for C2 in Figure 9.7, $X_C = 100\ \Omega$. For best stability, use C0G or NP0 units for all capacitors associated with the FET gates and sources. Depending on the FET chosen, the 1-kΩ source-bias-resistor value shown may require adjustment for reliable starting.

PARALLEL-TUNED COLPITTS VFO

In the parallel-tuned Colpitts, C3 and C4 are large values, perhaps 10 times larger than typical values chosen for C1 and C2 to resonate L at the desired frequency. This means only a small fraction of the total tank voltage is applied to the FET gate, and the FET can be considered to be only lightly coupled into the tank circuit. Equally important, the values of C3 and C4 must be much larger than the FET device internal capacitances to provide good stability for the output frequency. This

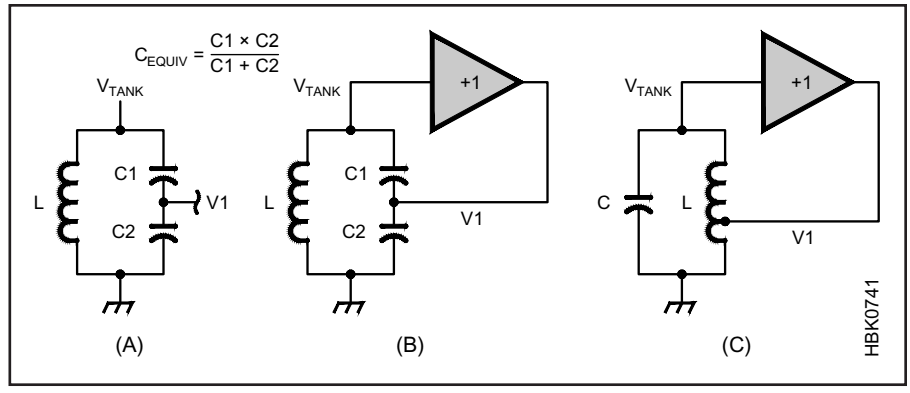

Figure 9.6 — The capacitor in an LC tank circuit (A) can be split into two capacitors with the same equivalent capacitance, C_{EQUIV}, so that the resonant frequency of the tank circuit is unchanged. (B) shows a unity-gain amplifier connected so that it forces V1 = V_{TANK}, creating positive feedback and a Colpitts oscillator results. (C) shows the same technique applied to the tank circuit inductor, creating a Hartley oscillator.

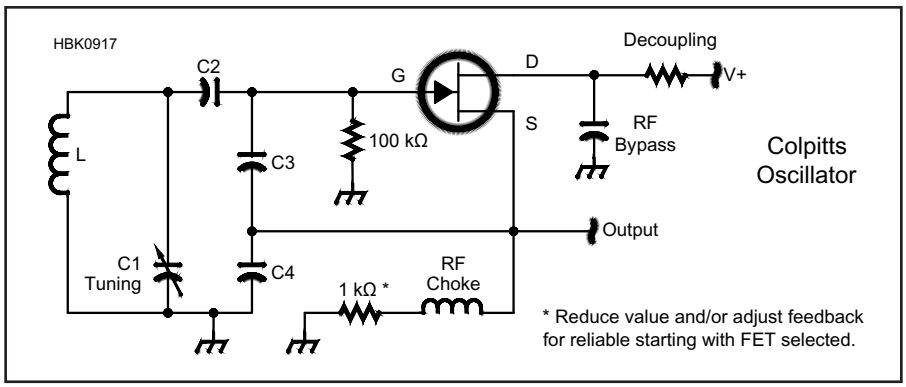

Figure 9.7 — The Colpitts oscillator circuit.

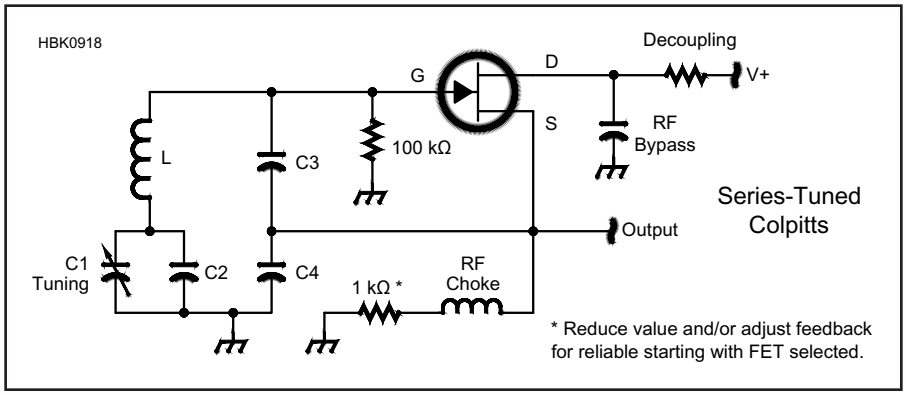

Figure 9.8 — The series-tuned Colpitts oscillator circuit.

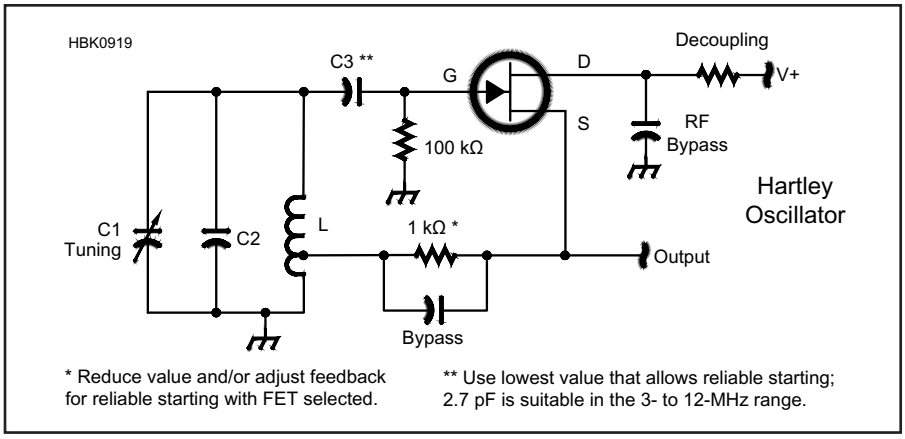

Figure 9.9 — The Hartley oscillator circuit.

keeps small variations in the FET capacitances from having a significant effect on oscillator operation.

The FET is driven by the sum of the voltages across C3 and C4, while it drives the voltage across C4 alone. This means that the tank circuit operates as a resonant, voltage-step-up transformer compensating for the less-than-unity-voltage-gain amplifier. The resonant circuit consists of L, C1, C2, C3, and C4. The resonant frequency can be calculated by using the formula derived at the beginning of this section.

The equation holds for all cases. Its accuracy is dependent on the designer's ability to account for all of the contributions to tank inductance and capacitance. There are several rules of thumb that help a VFO designer identify the various "stray" contributions to tank inductance and capacitance:

1. Wherever there is voltage, stray capacitance must be considered.

2. Wherever there is current, stray mutual inductance must be considered.

3. All currents form loops, creating inductance.

These "rules" will help point out where to look for contributions to inductance and capacitance. At frequencies below 10 MHz the actual components you see and touch tend to be all you need to consider. As the oscillator frequency increases, consideration of stray capacitance and stray inductance — reactance effects that are not associated directly with a visible component — becomes increasingly important. If the oscillator frequency is noticeably lower than what you predict, it is almost certain that there are important stray reactances that you need to account for.

Getting back to the circuit of Figure 9.7, for a wide tuning range C2 must be kept small to reduce the swamping effect of C3 and C4 on the variable capacitor C1. (For more information on component selection, the chapters on oscillators in *Experimental Methods for RF Design* and *Introduction to Radio Frequency Design* listed in the Reference section provide excellent material.) "Swamping" refers to a much larger value component reducing the effect of a small component connected to it.

A parallel-tuned Colpitts oscillator is the subject of the detailed paper "A Design Example for an Oscillator for Best Phase Noise and Good Output Power" by Ulrich Rohde, N1UL, available in the online information for this book. The paper shows the design process by which both noise and power can be optimized in a simple oscillator.

SERIES-TUNED COLPITTS VFO

The series-tuned Colpitts circuit (Figure 9.8) works in much the same way. The difference is that the variable capacitor, C1, is positioned so that it is well-protected from being swamped by the large values of C3 and C4. In fact, small values of C3 and C4 would act to limit the tuning range. Fixed capacitance, C2, is often added across C1 to allow the tuning range to be reduced to that required without interfering with C3 and C4, which set the amplifier coupling. The series-tuned Colpitts has a reputation for better stability than the parallel-tuned original. Note how C3 and C4 swamp the much smaller capacitances of the JFET amplifier in both versions.

HARTLEY VFO

The Hartley oscillator of Figure 9.9 is similar to the parallel-tuned Colpitts, but the JFET amplifier source is connected to a tap on the tank inductance instead of the tank capacitance. A typical tap placement is 10% to 20% of the total turns up from the cold end of the inductor. (It's common to refer to the lowest-signal-voltage end of an inductor as "cold" and the other, with the highest signal voltage as "hot.") C2 limits the tuning range as required.

C3 is reduced to the minimum value that allows reliable starting. This is necessary

because the Hartley's lack of the Colpitts' capacitive divider would otherwise couple the FET's internal capacitances to the tank more strongly than in the Colpitts, potentially affecting the circuit's frequency stability.

9.2.2 VFO Design Topics

FREQUENCY SCALING

Generally, VFOs can be adapted to work at other frequencies (within the limits of the active device). To do so, compute an adjustment factor: f_{old}/f_{new}. Multiply the value of each frequency determining, feedback L, or C by that factor. As frequency increases and amplifier gain drops, it may be required to increase feedback more than indicated by the factor.

BIAS STABILIZATION

In all three circuits, there is a 1 kΩ resistor in series with the source bias circuit. This resistor does a number of desirable things. It reduces the Q of the inevitable low-frequency resonance of the choke with the tank tap circuit. It decreases tuning drift due to choke impedance and winding capacitance variations. It also protects against spurious oscillation at undesired frequencies due to internal choke resonances. Less obviously, it acts to stabilize the loop gain of the built-in AGC action of this oscillator. Stable operating conditions act to reduce frequency drift.

Some variations of these circuits may be found with added resistors providing a dc bias to stabilize the system quiescent current. More elaborate still are variations characterized by a constant-current source providing bias. This can be driven from a separate AGC detector system to give very tight level control.

GATE CLAMPING DIODE

It is important to note that the gate-to-ground clamping diode (1N914 or similar) long used by radio amateurs as a means of avoiding gate-source conduction has been shown by Ulrich Rohde, N1UL, to degrade oscillator phase-noise performance and its use is virtually unknown in professional circles.

OTHER AMPLIFYING DEVICES

Figure 9.10 shows two more VFOs to illustrate the use of a vacuum tube triode and a BJT as amplifying devices. Compared to the more frequently used JFET, bipolar transistors like the one used in Figure 9.10B are relatively uncommon in oscillators because their relatively low input and output impedances are more difficult to couple into a high-Q tank without loading it excessively. Bipolar devices do tend to give better sample-to-sample amplitude uniformity for a given oscillator circuit, because many of the bipolar transistor's characteristics are due more to device physics and are less influenced by manufacturing variations. This is not true for JFETs of any given type; JFETs tend to vary more in their characteristics because manufacturing variations directly affect their threshold or pinch-off voltage.

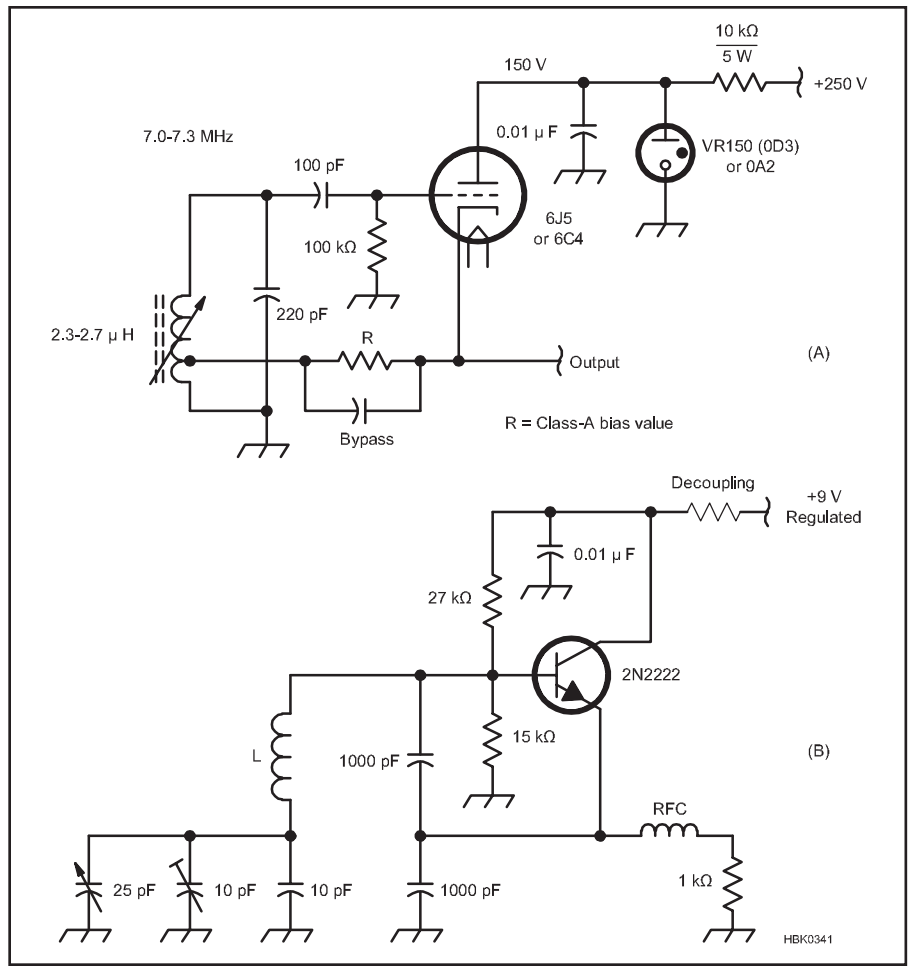

Figure 9.10 — Two additional oscillator examples: at A, a triode-tube Hartley; at B, a bipolar junction transistor in a series-tuned Colpitts.

SQUEGGING

The effect called *squegging* (or *squeeging*) can be loosely defined as oscillation on more than one frequency at a time, but may also manifest itself as RF oscillation interrupted at an audio frequency rate, as in a super-regenerative detector. One form of squegging occurs when an oscillator is fed from a power supply with high source impedance. The power supply charges up the decoupling capacitor until oscillation starts. The oscillator draws current and pulls down the capacitor voltage until it has starved itself of power and oscillation stops. The oscillator stops drawing current and the decoupling capacitor then recharges until oscillation restarts. The process, the low-frequency cycling of which is a form of relaxation oscillation, repeats indefinitely. Examination with an oscilloscope reveals the oscillator output is pulse modulated.

Squegging can be a consequence of poor design in battery-powered radios. As the battery becomes discharged, its internal resistance often rises and circuits that it powers can begin to misbehave. In audio stages, such misbehavior may manifest itself in the "putt-putt" sound of the slow relaxation oscillation called "motorboating."

9.2.3 Two Low-Noise VHF Oscillators

The following section is excerpted from the May/June 2018 QEX article "An Optimized Grounded Base Oscillator Design for VHF/UHF," by Ulrich Rohde, N1UL, and Ajay Poddar, AC2KG. The full article is available as a PDF file in this book's online content. Interested readers are encouraged to read the original article, which includes a comprehensive set of references and all design details.

The design of VHF/UHF oscillators has been described in many books and journals with most of the emphasis on frequency stability and, to a lesser extent, on output/efficiency.

Vacuum Tube Oscillators

Vacuum-tube radios are still maintained in operation and occasionally constructed for enjoyment but the semiconductor long ago achieved dominance in VFOs. Vacuum tubes cannot run as cool as competitive semiconductor circuits, so care is needed to keep the tank circuit away from the tube heat. LC VFOs are also no longer used in commercial transceivers today with the advent of DDS technology that is described elsewhere in this chapter.

In many amateur and commercial vacuum-tube oscillators, oscillation drives the tube into grid current at the positive peaks, causing rectification and producing a negative grid bias. The oscillator thus runs in Class C, in which the conduction angle reduces as the signal amplitude increases until the amplitude stabilizes. As in the JFET circuits of Figures 9.7 through 9.9, better stability and phase-noise performance can be achieved in a vacuum-tube oscillator by moving its operating point out of true Class C — that is, by including a bypassed cathode-bias resistor (the resistance appropriate for Class A operation is a good starting value).

The vacuum-tube triode Hartley shown in Figure 9.10A features permeability tuning, which has no sliding contact like that of a capacitor's rotor and can be made reasonably linear by artfully spacing the coil turns. The slow-motion drive can be done with a screw thread. The disadvantage is that special care is needed to avoid backlash and eccentric rotation of the core. If a non-rotating core is used, the slides have to be carefully designed to prevent motion in unwanted directions. The Collins Radio Company made extensive use of vacuum tube-based permeability tuners, and a semiconductor version was used in a number of Ten-Tec radios.

Since the introduction of reliable phase noise measurements and the ability to predict/simulate the phase noise, the improvement of this important parameter with the help of CAD and analytic equations has gained more attention. The novel design concept described in the article with the online content uses a time domain approach to provide both the best output power and the best phase noise performance. This very simple but powerfully scalable set of formulas for the oscillator synthesis provides extremely good results. In addition, the design principle for fixed or narrowband oscillators discussed here also applies to the broadband VCO design. This design methodology works over a multi-octave tuning range. The circuits here are shown as fixed-frequency designs. To use them as a VCO, variable capacitance diodes (varicaps) should be used in parallel with the output capacitors at P1 with very light coupling to the frequency-controlling LC circuit.

CIRCUIT DESIGN GUIDELINES

The results we have obtained so far were based on mathematical calculations. Some of these calculations are difficult to obtain (see the original *QEX* article). However, there are certain relationships between the values of the capacitance of the tuned circuit and the two feedback capacitors, the collector-emitter capacitor and the emitter-to-ground capacitor. The following shows the set of recommended steps for easy design of such oscillators.

A very popular circuit for VHF/UHF oscillators is the grounded-base configuration, which is shown in **Figure 9.11**. Its phase noise can be made very good, since the RF voltage swing at the collector can be close to the supply voltage. The circuit is simple and has been used for decades. This oscillator type works well from RF frequencies such as 10 MHz to above 1000 MHz.

The oscillator function is based on the principle that power from the output is taken and fed to the emitter. This feedback arrangement generates a negative resistance at the output, thus compensating the losses of the output-tuned circuit, starts oscillating, and then stabilizes the oscillation amplitude. The circuit has simple design rules where C_E and C_F are the feedback capacitors that generate the negative resistance to compensate for the loss resistance in the resonator network comprised of L_E and C^*_L.

Theoretically, the grounded-base configuration can be rotated to be the Colpitts circuit. If we look at the performance, in the case of the Colpitts oscillator the RF voltage swing is limited by the base-emitter and emitter-to-ground voltage and as a result there is less energy stored in the circuit and because of loading the operational Q can be less than in the grounded-base configuration. The Colpitts oscillator is popular because of its simplicity, and its perceived high isolation as the output power is taken at the collector. However, due to the strong Miller effect at very high frequencies, this is not a true statement.

SIMPLE DESIGN RULES

By setting the L/C ratio to a fixed value of 1,200 (for optimum energy storage, group delay, and energy transfer for a given cycle in the resonator network), the following should be used.

$$\frac{L}{C} = \left[\frac{L_E}{C^*_L}\right]_{\text{Grounded-Base}} = 1,200$$

$$\Rightarrow Z_0 = \sqrt{1,200} \cong 34.6\ \Omega$$

$$L = 1,200 \times C$$

$$\Rightarrow L_E = 1,200 \times C^*_L$$

$$f = \frac{1}{2\pi\sqrt{LC}} = \frac{1}{2\pi C\sqrt{1,200}}$$

$$C^*_L = C_L + \frac{C_E C_F}{C_E + C_F}$$

where C_E and C_F are feedback capacitors.

$$C_L = \frac{C_A C_B}{C_A + C_B}$$

where C_A and C_F are feedback capacitors.

$$C_B = 10\ C_A$$

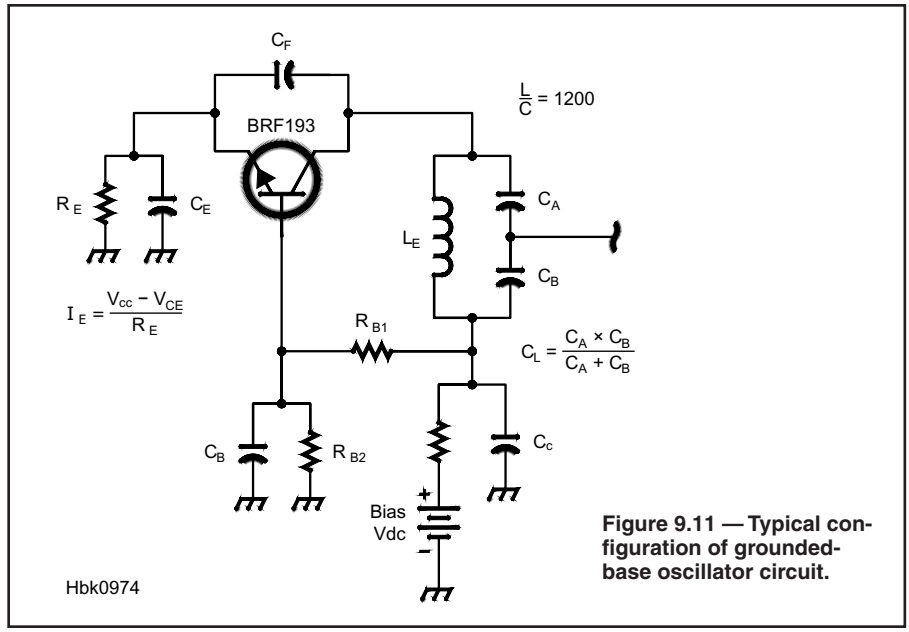

Figure 9.11 — Typical configuration of grounded-base oscillator circuit.

The reader is reminded that this "recipe" was developed based on the Y-parameters, transconductance, and package parasitics for the BRF193 transistor. Using a transistor greatly different than the BRF193 requires the full design process to be performed from the beginning using the new transistor as described in the original article.

144 MHz GROUNDED-BASE OSCILLATOR

The oscillator circuit in **Figure 9.12** is designed using the procedure of the preceding section.

$$C_L^* = \frac{1}{2\pi f\sqrt{1,200}} \Rightarrow C \cong 31\ \text{pF}$$

$$L_L = \frac{1}{(2\pi f)^2 C} \Rightarrow L \cong 39\ \text{nH}$$

$$C_F = 0.3 \times C_L^* \cong 11\ \text{pF}$$

$$C_L = 2 \times C_F \cong 22\ \text{pF}$$

$$C_E = 4 \times C_F \cong 44\ \text{pF}$$

$$C_A = 22\ \text{pF}$$

$$C_B = 220\ \text{pF}$$

These results are frequency scalable with minor corrections possibly necessary.

Figure 9.12A shows the schematic of the 144 MHz oscillator at I_C = 10 mA. The oscillator uses a lumped inductor of 39 nH and an unloaded Q of 200 at the operating frequency. Even at these frequencies the layout is quite critical. The example layout in Figure 9.12B shows an assembly of components where the lead inductances have been kept small. The inductor is a standard off the shelf component. (Note that the sample layout is given only as a guideline and not as a template.)

Figure 9.12C shows the CAD simulated phase noise plot. The output power is 11.5 dBm and the second and third harmonics are about –28 dBm and –34 dBm. Using phase noise simulation, the result is –134 dBc/Hz and –94 dBc/Hz at 10 kHz and 100 Hz offset. Calculated, simulated, and measured results closely agree within 1 dB.

If the same transistor is operated at 30 mA, the phase noise at 10 kHz offset will improve to be –144 dBc/Hz and the output power is increased to 20 dBm. This shows that for low phase noise design a more powerful transistor is a good choice. It is important to keep the dc dissipation of the device in mind, as the CAD process will probably not flag a misuse of the device.

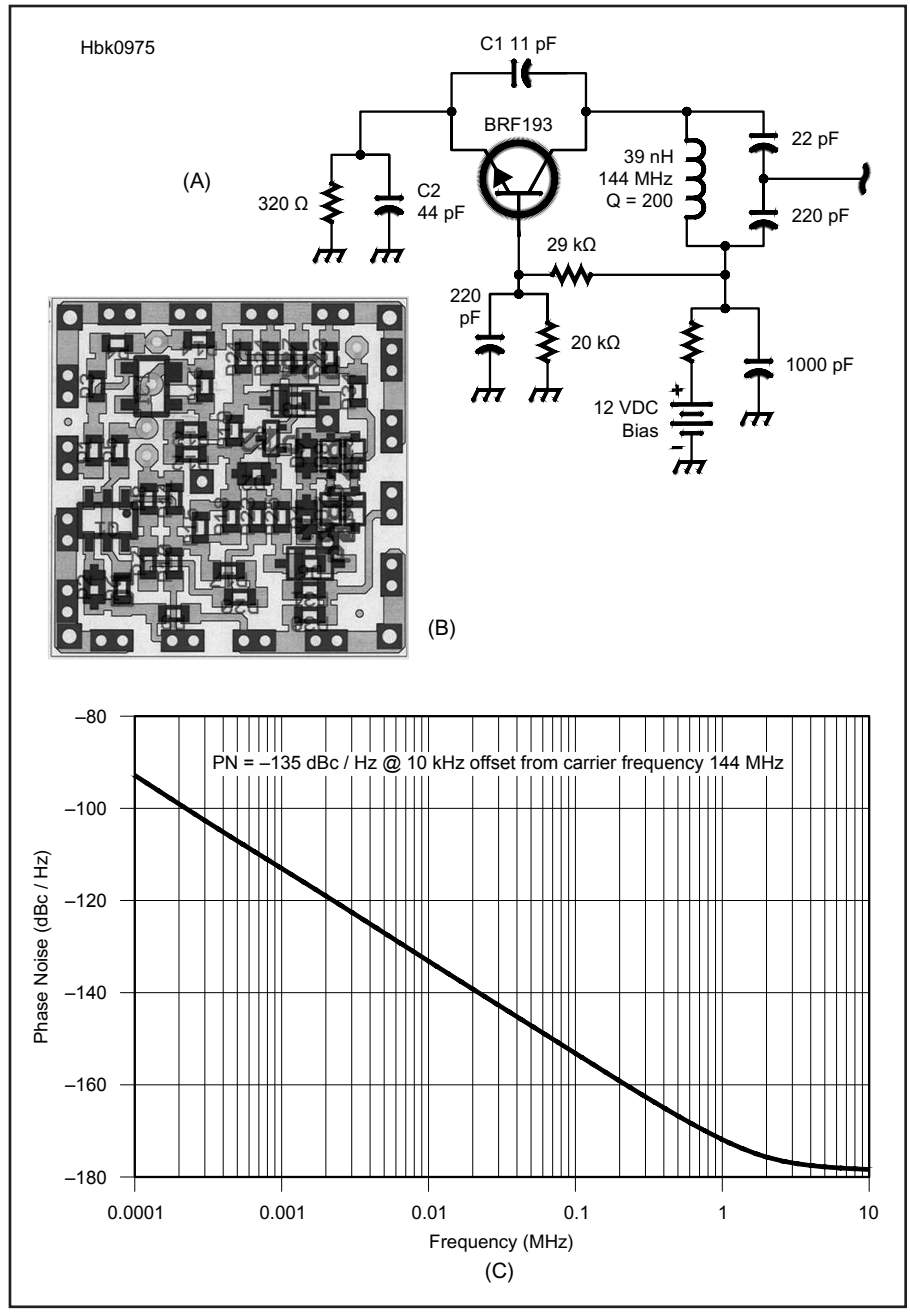

Figure 9.12 — 144 MHz oscillator circuit for I_E = 10 mA (A), an example of the circuit layout using LC lumped inductor-capacitor resonator network (B), and simulated phase noise plot (C).

MODIFIED CIRCUIT FOR UHF (432 MHz) AND HIGHER CURRENT

At 432 MHz and at I_E = 30 mA, the loading of the tank circuit decreases the operating Q significantly. The way around this to use a center-tapped inductor. As the coupling at these frequencies from winding to winding is not extremely high, two separate identical inductors can be used successfully.

Figure 9.13A shows the schematic of 432 MHz grounded base oscillator using the tapped inductor. This is a modification of the circuit we have used previously. In the case of a VCO, it would be advantageous to use a different output coupling scheme because in this configuration, the loading would vary with frequency. This can easily be achieved by adding some inductive coupling to the circuit. In case of a printed resonator this can be accomplished quite simply.

Figure 9.13B shows a sample guideline layout of the 432 MHz oscillator circuit using

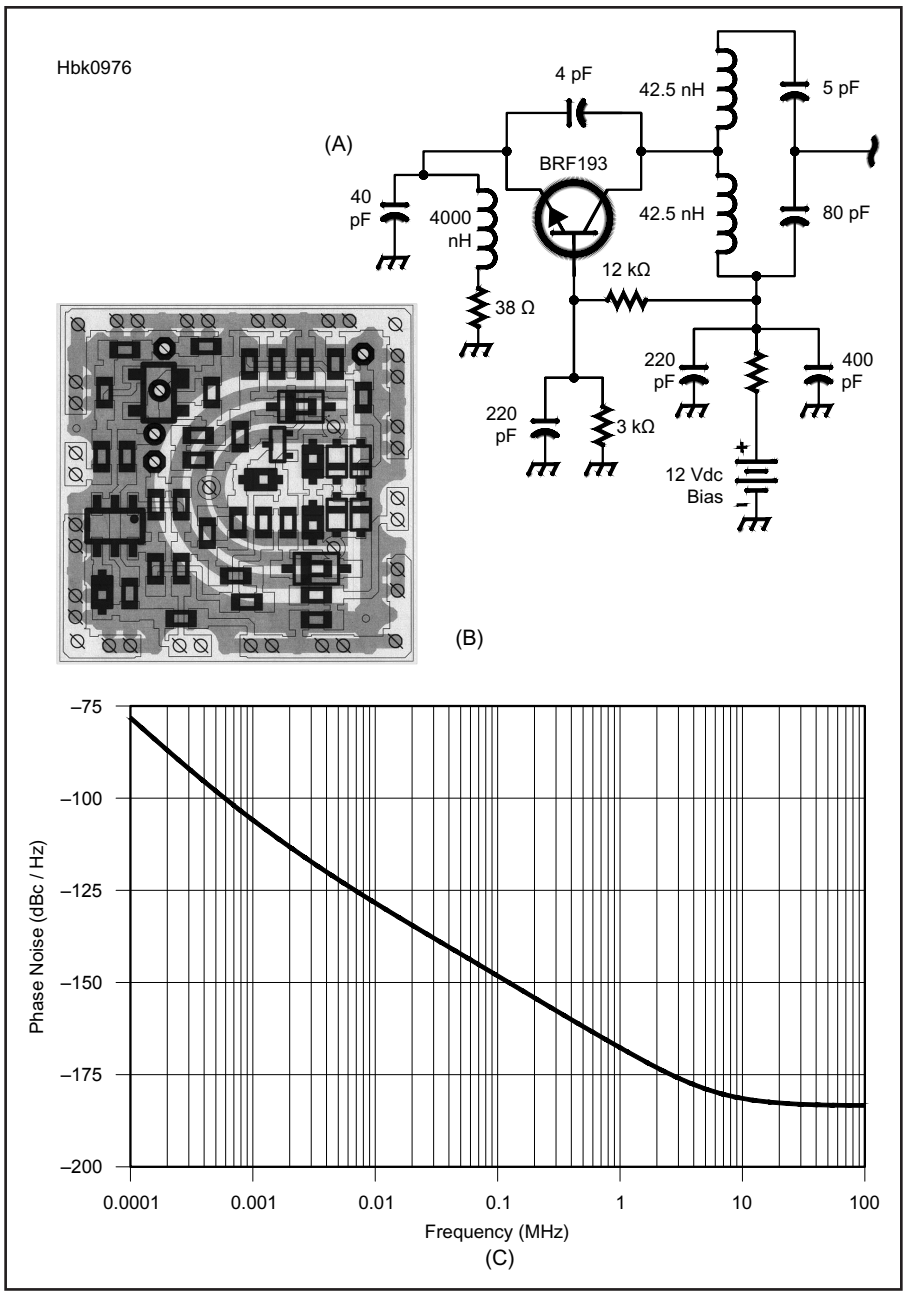

Figure 9.13 — 432 MHz oscillator circuit using tapped inductor and I_E = 30 mA (A), an example of the circuit layout using buried printed coupled line resonator (stripline resonator) in the middle layer of the board (B), and simulated phase noise plot (C).

a buried printed coupled line resonator network (stripline resonator in the middle layer). The actual resonator would not be visible if the oscillator is visually inspected.

Figure 9.13C shows the simulated phase noise plot with the expected noise degradation of 9 dB, as the frequency is approximately three times higher. The resulting simulated output power at 432 MHz is 16 dBm, compared to 18 dBm at 144 MHz. This is due to internal package parasitics, which could not be compensated externally. The second harmonic is suppressed by 38 dB; this is due to the higher operating Q.

9.2.4 Three High-Performance HF VFOs

THE N1UL MODIFIED VACKAR VFO

The oscillator circuit of **Figure 9.14A** is contributed by Ulrich Rohde, N1UL. It is a modified Vackar design (see the sidebar) in which a small coupling capacitor (8 pF) and voltage divider capacitor (18 pF) isolate the resonator circuit (10 µH and 50 pF tuning capacitor) from the oscillator transistor.

The oscillator transistor is followed by a buffer stage to isolate the oscillator from the load. Because the coupling between the transistor base and the resonator is fixed and light, stability of the oscillator is high across a wide tuning range from 5.5 to 6.6 MHz. Either the inductor or capacitor may be varied to tune the oscillator, but a variable capacitor is recommended as more practical and gives better performance.

Because of the oscillator transistor's large capacitors from base to ground (220 pF) and collector to ground (680 pF), the various parameters of the oscillator transistor have little practical influence on circuit performance. The widely available 2N3904 performs well for both the oscillator and buffer transistors. The BC546 is a suitable substitute.

Practical resonator coil and the tuning capacitors will have a positive temperature coefficient. The 8 pF and the 18 pF capacitors should have an N150 temperature coefficient to partially compensate for their drift. After 1 hour, the observed frequency drift for this circuit was less than 10 Hz/hour.

THE K7HFD LOW-NOISE-DIFFERENTIAL OSCILLATOR

The other high performance oscillator example, shown in **Figure 9.15**, is designed for low-noise performance by Linley Gumm, K7HFD, and appears on page 126 of the ARRL's *Solid State Design for the Radio Amateur* (out of print, but available used and through libraries). This circuit uses no unusual components and looks simple, yet it is a subtle and sophisticated circuit. (An analysis and simulation of this circuit by its designer is included in the online material for this book.)

The effects of limiting in reducing AM oscillator noise were covered previously. However, because AM noise sidebands can get translated into PM noise sidebands by imperfect limiting, there is an advantage to stripping off the AM as early as possible, in

The Vackar Oscillator

The original Vackar oscillator is named for Jirí Vackár, who invented the circuit in the late 1940s. The circuit's description — a refinement of the Clapp oscillator — can be found in older editions of the Radio Society of Great Britain's Radio Communication Handbook, with some further comments on the oscillator in RSGB's Amateur Radio Techniques. The circuit is also described in the Nov 1955 *QST* "Technical Correspondence" column by W9IK. The Vackar circuit optimized the Clapp oscillator for frequency stability: the oscillator transistor is isolated from the resonator, tuning does not affect the feedback coupling, and the transistor's collector output impedance is kept low so that gain is the minimum necessary to sustain oscillation.

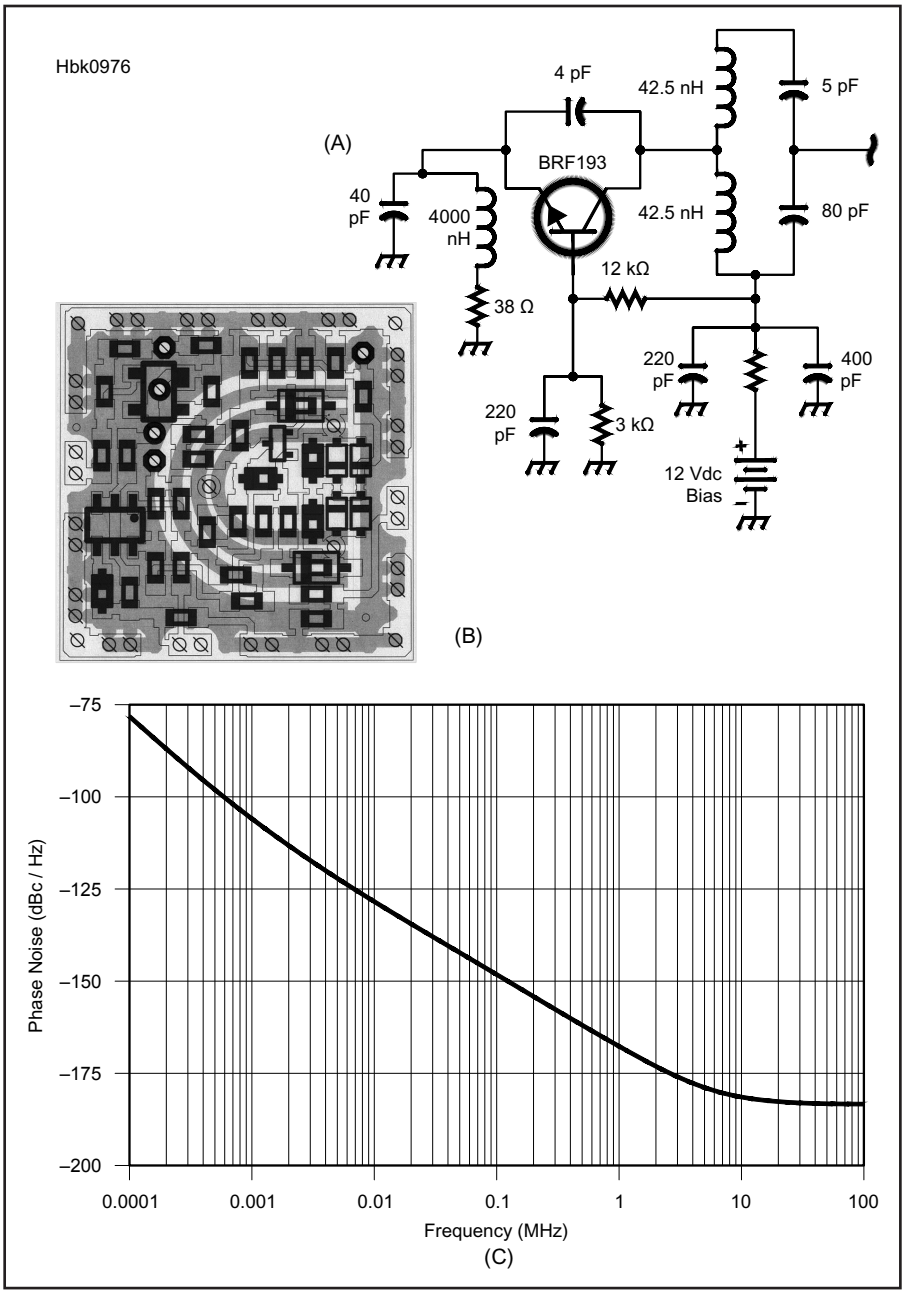

Figure 9.13 — 432 MHz oscillator circuit using tapped inductor and I_E = 30 mA (A), an example of the circuit layout using buried printed coupled line resonator (stripline resonator) in the middle layer of the board (B), and simulated phase noise plot (C).

the oscillator itself. An ALC system in the oscillator will counteract and cancel only the AM components within its bandwidth, but an oscillator based on a limiter will do this over a broad bandwidth. K7HFD's oscillator uses a differential pair of bipolar transistors as a limiting amplifier. The dc bias voltage at the bases and the resistor in the common emitter path to ground establishes a controlled dc bias current, here 25 to 27 mA. The ac voltage between the bases switches this current between the two collectors. This applies a rectangular pulse of current into link winding L2, which drives the series-resonant tank L1-C1.

The output impedance of the collector is high in both the current-on and current-off states. Along with the small number of turns of the link winding, this presents very high impedance to the tank circuit, which minimizes degradation of the tank Q. The input impedance of this limiter is also quite high and is applied across only a one-turn tap of L1, which similarly minimizes any impact on the tank Q. The input transistor base is driven into conduction only on one peak of the tank waveform. The output transformer has the inverse of the current pulse applied to it, so the output is not a low distortion sine wave, although the output harmonics will not be as extensive as simple theory would suggest because the circuit's high output impedance allows stray capacitances to attenuate high-frequency components. The low-frequency transistors used here also act to reduce the harmonic power.

With an output of +17 dBm, this is a power oscillator, running with nearly 300 mW of dc input power, so appreciable heating is present that can cause temperature-induced drift. The circuit's high-power operation is a deliberate ploy to create a high signal-to-noise ratio by having as high a signal power as possible. This also reduces the problem of the oscillator's broadband noise output. The limitation on the signal level in the tank is the transistors' base-emitter-junction breakdown voltage. The circuit runs with a few volts peak-to-peak across the one-turn tap, so the full tank is running at over 50 V_{P-P}.

Excessive voltage levels for the transistors can easily be generated by this circuit. The single easiest way to damage a bipolar transistor is to reverse bias the base-emitter junction until it avalanches. Most devices are rated to withstand only 5 V applied this way, the current needed to do damage is small, and very little power is needed. If the avalanche current is limited to less than that needed to perform immediate destruction of the transistor, it is likely that there will be some degradation of the device, a reduction in its bandwidth and gain along with an increase in its noise. These changes are irreversible and cumulative. Small, fast signal diodes have breakdown voltages of over 30 V and less capacitance than the transistor bases, so one possible experiment would be to try the effect of adding a fast signal diode in series with the base of each transistor and running the circuit at even higher levels.

The oscillation amplitude is controlled by the drive current limit. The voltage on L2 must never allow the collector of the transistor driving it to go into saturation, for if this happens the transistor presents very low impedance to L2 and badly loads the tank, wrecking the Q and the noise performance. The circuit can be checked to verify the margin from saturation by probing the hot end of L2 and the emitter with an oscilloscope. Another, less obvious, test is to vary the power-supply voltage and monitor the output power. While the circuit is under current control, there is very little change in output power, but if the supply is low enough to allow saturation, the output power will change significantly with varying supply voltage.

The use of the inexpensive, general-purpose 2N3904 is interesting, as it is not nor-

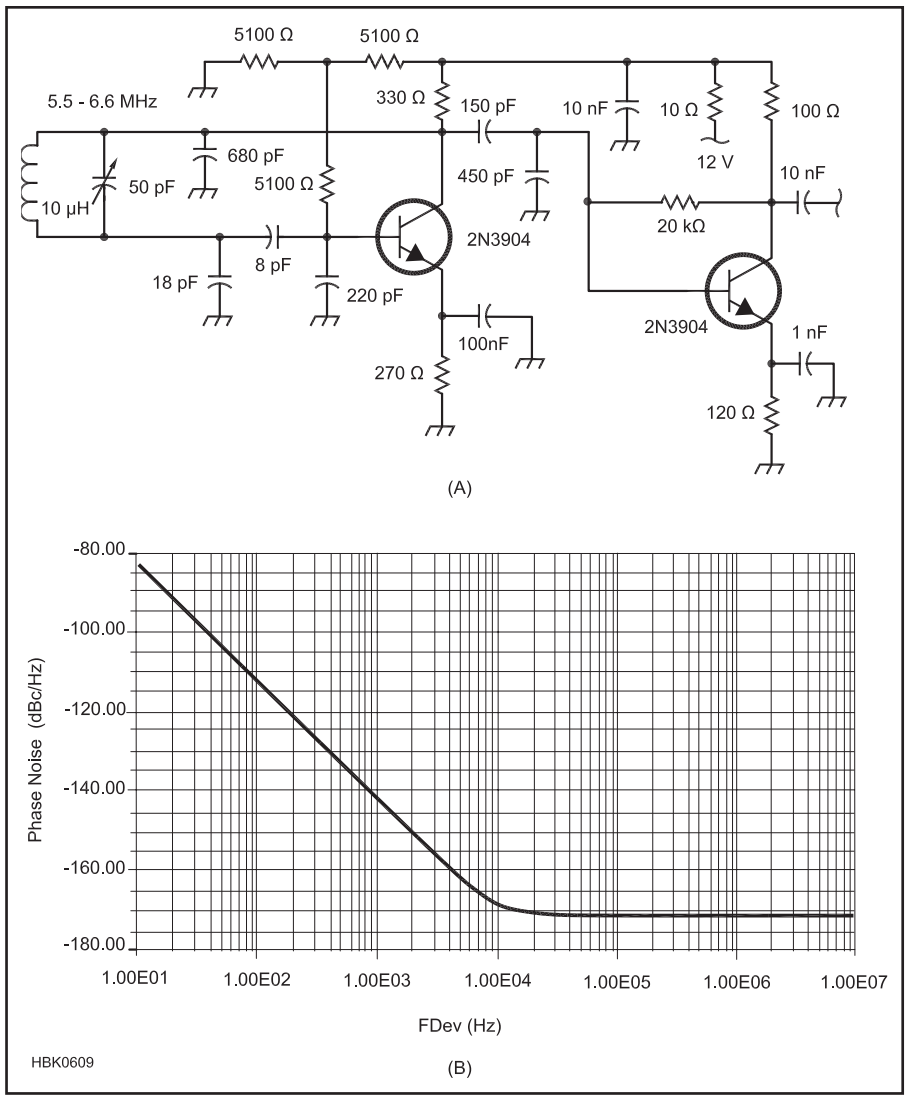

Figure 9.14 — At A, N1UL's Modified Vackar VFO is tuned from 5.5 to 6.6 MHz using the 50-pF capacitor. Tuning may be restricted to narrower ranges by placing a fixed capacitor in parallel with a smaller variable capacitor. The resonant frequency of the oscillator is determined by the 10 µH inductor and 50 pF tuning capacitor. B shows the excellent phase noise performance of the modified Vackar VFO in this *Harmonica* simulation. At 1 kHz from the carrier, noise is –144 dBc.

mally associated with RF oscillators, more often used at dc or audio frequencies. There is evidence that suggests some transistors that have good noise performance at RF have worse noise performance at low frequencies, and that the low-frequency noise they create can modulate an oscillator, creating noise sidebands. Experiments with low-noise audio transistors may be worthwhile, but many such devices have high junction capacitances.

In the description of this circuit in *Solid State Design for the Radio Amateur*, the results of a phase-noise test made using a spectrum analyzer with a crystal filter as a preselector are given. Ten kilohertz away from the carrier, in a 3 kHz measurement bandwidth, the noise was more than 120 dB below the carrier level. This translates into better than –120 – 10 log (3000), which equals –154.8 dBc/Hz, SSB, consistent with the modeled phase noise performance shown in Figure 9.15B. At this offset, –140 dBc is usually considered to be excellent. This VFO provides state-of-the art performance by today's standards — in a 1977 publication.

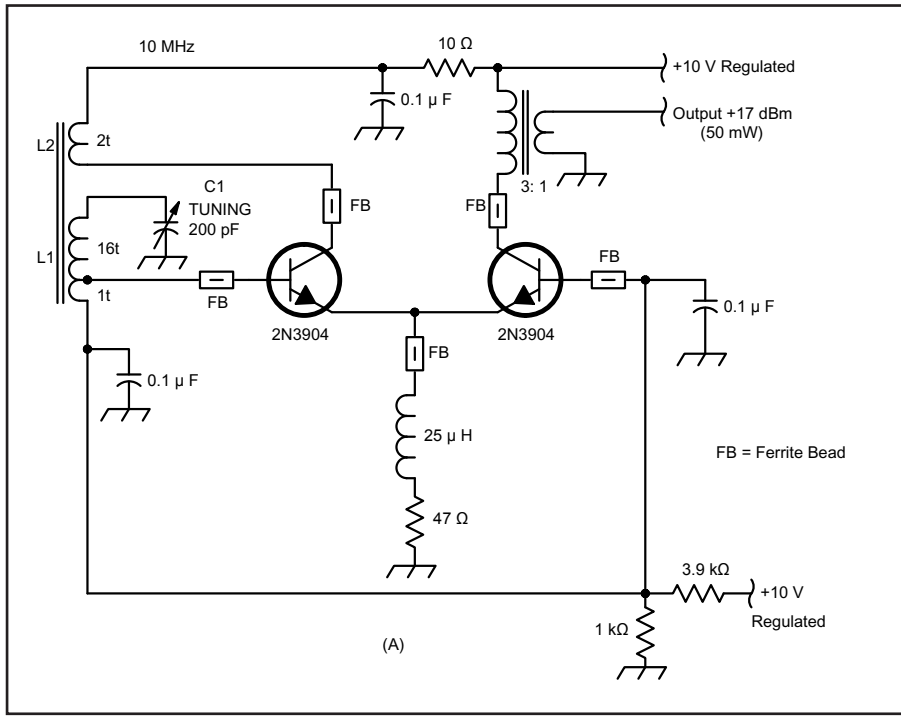

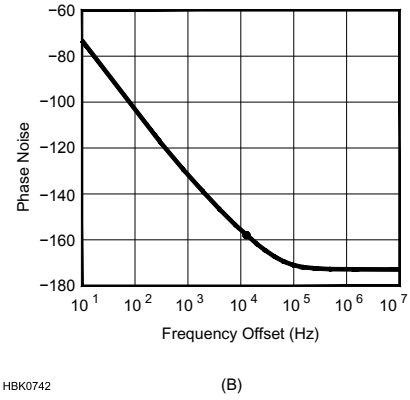

Figure 9.15 — At A, this low-noise oscillator design by K7HFD operates at an unusually high power level to achieve a high C/N (carrier-to-noise) ratio. L1 is 1.2 µH and uses 17 turns of wire on a T-68-6 toroid core. The tap is at 1 turn. Q at 10 MHz is 160. L2 is a 2-turn link over L1. At B, modeling of the differential oscillator by Ulrich Rohde, N1UL, shows its excellent phase-noise performance.

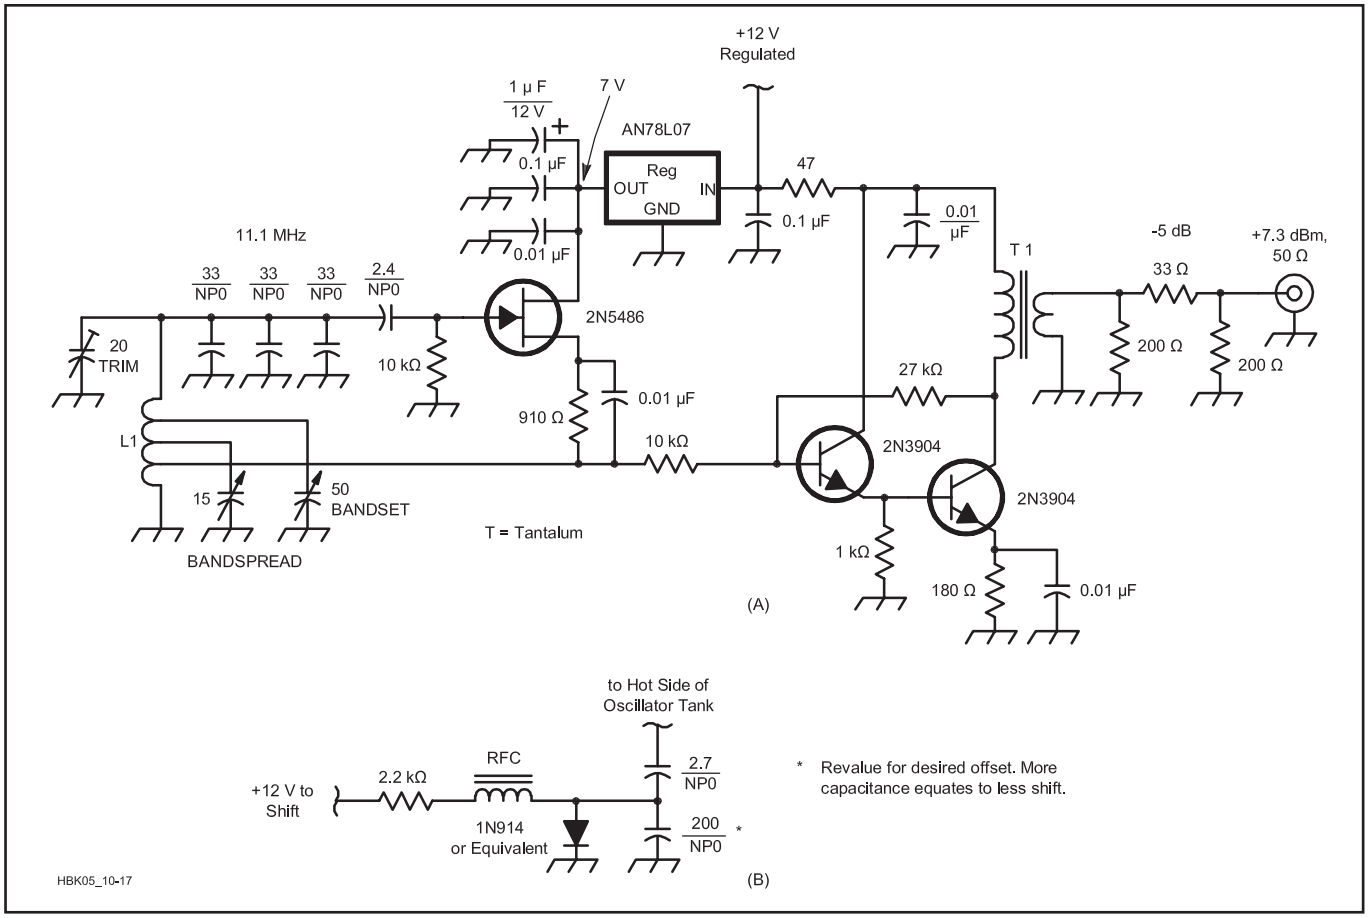

Figure 9.16 — Incorporating ideas from N1UL, KA7EXM, W7ZOI and W7EL, the oscillator at A achieves excellent stability and output at 11.1 MHz without the use of a gate-clamping diode, as well as end-running the shrinking availability of reduction drives through the use of bandset and bandspread capacitors. L1 consists of 10 turns of B & W #3041 Miniductor (#22 tinned wire, 5/8 inch in diameter, 24 turns per inch). The source tap is 2-1/2 turns above ground; the tuning-capacitor taps are positioned as necessary for bandset and bandspread ranges required. T1's primary consists of 15 turns of #28 enameled wire on an FT-37-72 ferrite core; its secondary, 3 turns over the primary. B shows a system for adding fixed TR offset that can be applied to any LC oscillator. The RF choke consists of 20 turns of #26 enameled wire on an FT-37-43 core.

A JFET HARTLEY VFO

Figure 9.16 shows an 11.1 MHz version of a VFO and buffer closely patterned after that used in 7 MHz transceiver designs published by Roger Hayward, KA7EXM, and Wes Hayward, W7ZOI ("The Ugly Weekender") and Roy Lewallen, W7EL ("The Optimized QRP Transceiver"). In it, a Hartley oscillator using a 2N5486 JFET drives the two-2N3904 buffer attributed to Lewallen. This version diverges from the originals in that its JFET uses source bias (the bypassed 910 Ω resistor) instead of a gate-clamping diode and is powered from a low-current 7 V regulator IC instead of a Zener diode and dropping resistor. The 5 dB pad sets the buffer's output to a level appropriate for "Level 7" (+7 dBm LO) diode ring mixers.

The circuit shown was originally built with a gate-clamping diode, no source bias and a 3 dB output pad. Adjusting the oscillator bias as shown increased its output by 2 dB without degrading its frequency stability (200 to 300 Hz drift at power up, stability within ±20 Hz thereafter at a constant room temperature).

In recognition that precision mechanical tuning components are hard to obtain, the resonator uses "Bandset" and "Bandspread" variable capacitors. These terms are from the early days of radio: *bandset* is for coarse-tuning and *bandspread* is for fine-tuning.

9.3 Building an Oscillator

We've covered a lot of ground about how oscillators work, their limitations and a number of interesting circuits, so the inevitable question arises of how to design one. Let's make an embarrassing confession right here: Very few oscillators you see in published circuits or commercial equipment were designed by the equipment's designer. Almost all have been adopted from other sources. While recycling in general is important for the environment, it means in this case that very few professional or amateur designers have ever designed an oscillator from scratch. We all have collections of circuits we've "harvested," and we adjust a few values or change a device type to produce something to suit a new project.

Oscillators aren't designed, they evolve. They seem to have a life of their own. The Clarke & Hess book listed in the references contains one of the few published classical design processes. The ARRL book *Experimental Methods in RF Design* contains extensive material on oscillator circuits that is well worth reading.

9.3.1 VFO Components and Construction

TUNING CAPACITORS AND REDUCTION DRIVES

As most commercially made radios now use DSP frequency synthesis, it has become increasingly difficult to find certain key components needed to construct a good VFO. Slow-motion drives and variable capacitors are available from *QST* advertiser National RF (**www.nationalrf.com**), Dan's Small Parts and Kits (**www.danssmallpartsandkits.net**), and Antique Electronic Supply (**www.tubesandmore.com**). Variable capacitors should be a high-quality component with double ball-bearings and silver-plated surfaces and contacts.

An alternate approach is also available: Scavenge suitable parts from old equipment; use tuning diodes instead of variable capacitors — an approach that, if uncorrected through phase locking, generally degrades stability and phase-noise performance; or use two tuning capacitors, one with a capacitance range ⅕ to ⅒ that of the other, in a bandset/bandspread approach.

Assembling a variable capacitor to a chassis and its reduction drive to a front panel can result in *backlash* — an annoying tuning effect in which rotating the capacitor shaft deforms the chassis and/or panel rather than tuning the capacitor. One way of minimizing this is to use the reduction drive to support the capacitor, and use the capacitor to support the oscillator circuit board, which is then attached to the chassis.

FIXED CAPACITORS

Use silvered-mica or other highly-stable capacitors for all fixed-value capacitors in the oscillator circuit. Power-supply decoupling capacitors may be any convenient type, such as ceramic or film.

Traditionally, silver-mica fixed capacitors have been used extensively in oscillators, but their temperature coefficient is not as low as can be achieved with other types, and some silver micas have been known to behave erratically. Polystyrene film has become a proven alternative. One warning is worth noting: polystyrene capacitors exhibit a permanent change in value should they ever be exposed to temperatures much above 70 °C; they do not return to their old value on cooling.

Particularly suitable for oscillator construction are the low-temperature-coefficient ceramic capacitors, often described as NP0 or C0G types. (NP0 and C0G are equivalent) These abbreviations are actually temperature-coefficient codes. **Figure 9.17** contains graphs showing the behavior of three common temperature coefficients. Some ceramic capacitors are available with deliberate, controlled temperature coefficients so that they can be used to compensate for other causes of frequency drift with temperature. For example, the code N750 denotes a part with a temperature coefficient of –750 parts per million per degree Celsius. These parts are now somewhat difficult to obtain, so other methods are needed. (Values for temperature coefficients and other attributes of capacitors are presented in the Component Information section of the **Construction Techniques** chapter.)

In a Colpitts circuit, the two large-value capacitors that form the voltage divider for the active device still need careful selection. It is tempting to use any available capacitor of the right value, because the effect of these components on the tank frequency is reduced by the proportions of the capacitance values in the circuit. This reduction is not as great as the difference between the temperature stability of an NP0 ceramic part and some of the low-cost, decoupling-quality X7R-dielectric ceramic capacitors. It's worth using low-temperature coefficient parts even in the seemingly less-critical parts of a VFO circuit — even for the bypass capacitors. Chasing the cause of temperature drift is more challenging than fun. Buy critical components like high-stability capacitors from trustworthy sources.

TUNING CAPACITOR NETWORKS

Often an available variable capacitor has greater capacitance than required for a desired frequency range. While plates can sometimes be removed, a better solution embeds the variable capacitor in a network of fixed capacitors. The evolution of this network is shown in the middle section of **Figure 9.18**. The variable capacitor, C_V, and C_2 are paralleled to form the equivalent C_{2V}. This is then placed in series with C_1 for the equivalent C_{12V}. This is, in turn, paralleled by C_3 to form the total capacitor, C_{NET}. The overall frequency is calculated from the usual resonance relationship. The equations are shown, with capacitance in farads, inductance in henrys, and frequency in Hz.

There is considerable flexibility available to the designer, afforded by picking C_1 and C_2 values. Some combinations with C_1 much smaller than the variable capacitor can produce highly nonlinear tuning.

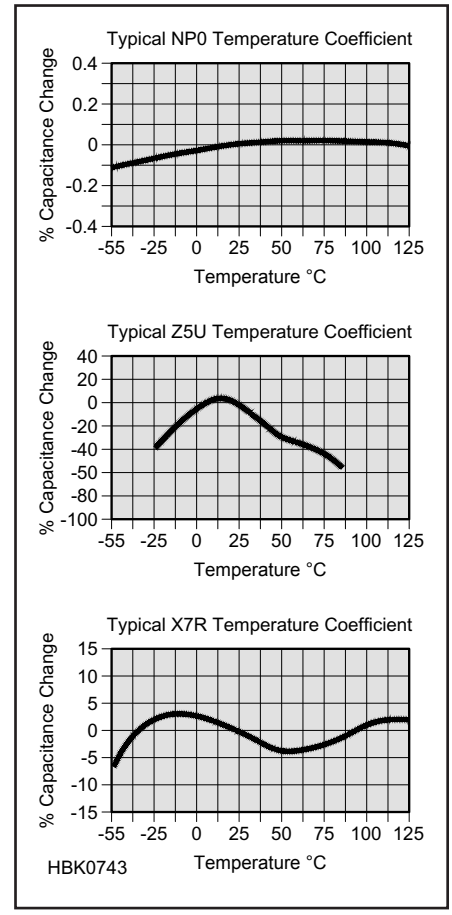

Figure 9.17 — EIA capacitor temperature coefficients specify change in capacitance with temperature. See the Component Data and References chapter for a complete table of temperature coefficient identifiers and characteristics.

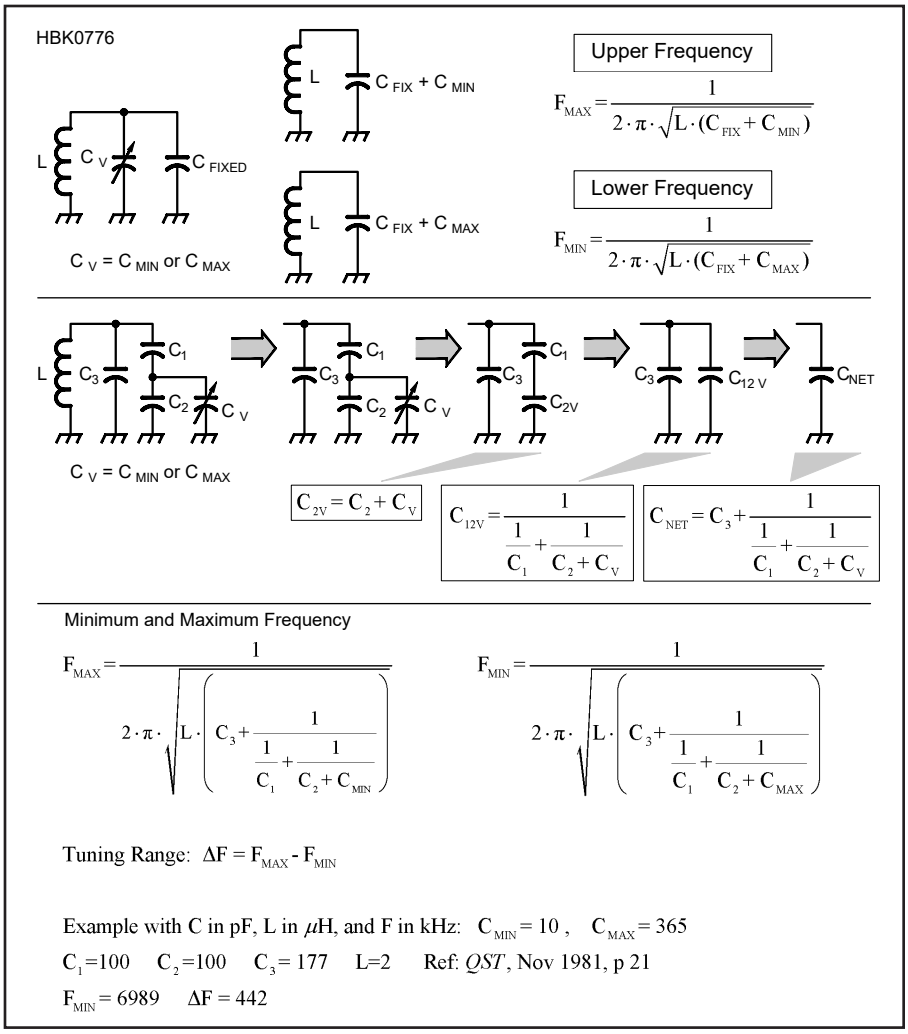

Figure 9.18 — A simple resonant circuit is tuned with parallel capacitors as shown in the top section. The tuning range is controlled by the ratio of the variable capacitance to the fixed capacitance.

INDUCTORS

Ceramic coil forms can give excellent results. If the required inductance is small enough, wind the coil on a ceramic form that is securely mounted. Self-supporting air-wound coils can also give good results if securely mounted and braced against vibration. B&W Miniductor (**www.bwantennas.com/coilcat.html** is the best-known commercially available product. It is not available to individuals from the manufacturer but there is a good selection available through **eBay.com**.

If you use a magnetic core, use powdered iron and support it securely. Do not use ferrite because of its temperature instability. Toroidal cores are preferred due to their self-shielding nature. Micrometals mix #6 has a low temperature coefficient and works well in conjunction with NP0 ceramic capacitors. Other materials have to be assessed on an individual basis.

A material's temperature stability will not be apparent until you try it in an oscillator, but you can apply a quick test to identify those nonmetallic materials that are lossy enough to spoil a coil's Q. Put a sample of the coil-form material into a microwave oven next to a glass of water and cook it about 10 seconds on low power. Do not include any metal fittings or ferromagnetic cores. Good materials will be completely unaffected; poor ones will heat and may even melt, smoke, or burst into flame. (This operation is a fire hazard if you try more than a tiny sample of an unknown material. Observe your experiment continuously and do not leave it unattended!)

W7ZOI suggests annealing toroidal VFO coils after winding. W7EL reports achieving success with this method by boiling his coils in water and letting them cool in air.

VOLTAGE REGULATORS

VFO circuits must be powered by well-regulated, low-noise supplies. Three-terminal regulators are inexpensive and have low output noise. It is easy to include them locally with the oscillator circuit. If the regulator is not part of the oscillator circuit, include a decoupling R-C combination to act as a noise filter.

Many oscillator circuits include Zener diode regulators. (See the **Power Sources** chapter.) Zener diodes have some idiosyncrasies that could spoil the oscillator. They are noisy, so decoupling is needed down to audio frequencies to filter this out. Zener diodes are often run at much less than their specified optimum bias current. Although this saves power, it results in a lower output voltage than intended and the diode's impedance is much greater, increasing its sensitivity to variations in input voltage, output current and temperature. Some common Zener types may be designed to run at as much as 20 mA; check the data sheet for your diode family to find the optimum current.

True Zener diodes are low-voltage devices; above a couple of volts, so-called Zener diodes are actually avalanche types. The temperature coefficient of these diodes depends on their breakdown voltage and crosses through zero for diodes rated at about 5 V. If you intend to use nothing fancier than a common-variety Zener, designing the oscillator to run from 5 V and using a 5.1 V Zener will give you a free advantage in voltage-versus-temperature stability.

There are some diodes available with especially low temperature coefficients, usually referred to as reference or temperature-compensated diodes. These usually consist of a series pair of diodes designed to cancel each other's temperature drift. The 1N821A diode has a temperature coefficient of ±100 ppm/°C. Running at 7.5 mA, the 1N829A provides 6.2 V ±5% and a temperature coefficient of just ±5 parts per million (ppm) maximum per degree Celsius. A change in bias current of 10% will shift the voltage less than 7.5 mV, but this increases rapidly for greater current variation. The curves in **Figure 9.19** show how the temperature coefficients of these diodes are dependent on bias current.

The LM399 is a complex IC that behaves like a superb Zener at 6.95 V, ±0.3 ppm/°C. Precision, low-power, three-terminal regulators are also available that are designed to be used as voltage references, some of which can provide enough current to run a VFO. There are comprehensive tables of all these devices on pages 614 and 615 of Horowitz and Hill, *The Art of Electronics, Third Edition*.

Over the last decade, specialized, low-noise, low dropout (LDO) regulators have been developed for oscillator circuits and are available from multiple manufacturers, including Analog Devices (including LTC and Maxim), Texas Instruments, and others. The noise performance of these parts is com-

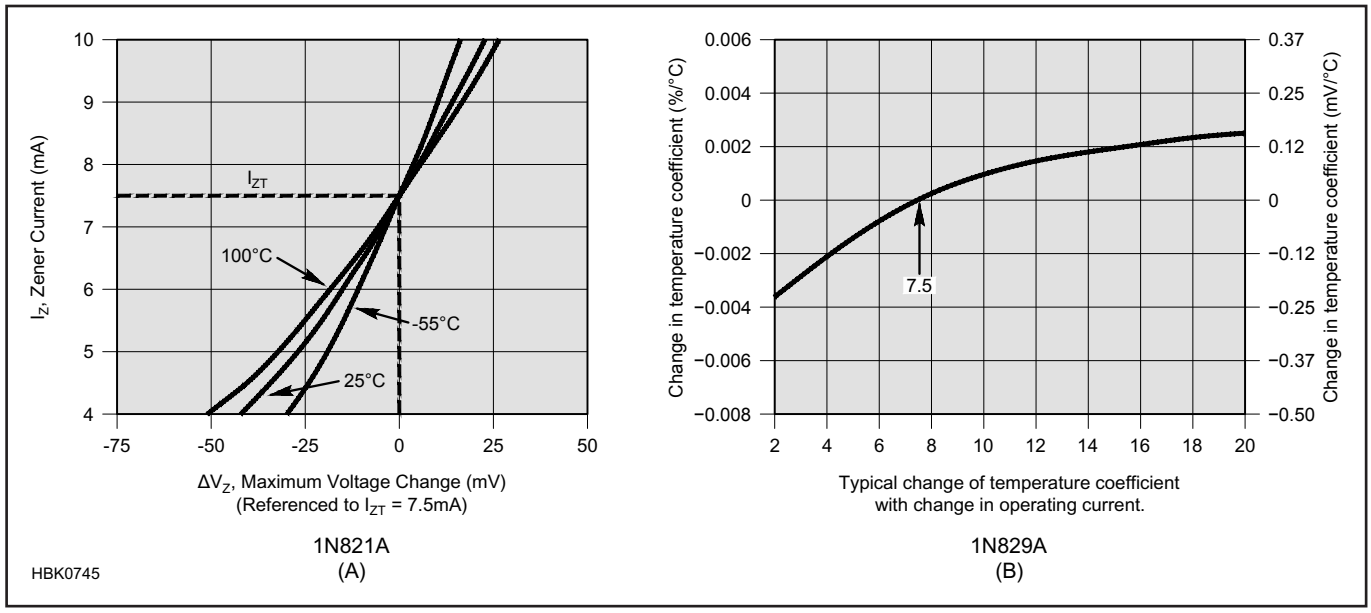

Figure 9.19 — The temperature coefficient of temperature-compensated diodes varies with bias current. To obtain the best temperature performance, use the specific bias current for the diode.

parable to that of batteries in bench tests comparing performance of a battery supply versus a supply using the low noise LDO. Texas Instruments' tutorial "LDO Noise Demystified – SLAA412B" (see **ti.com**) by Pithadia and Verma provides a good technical introduction to this product area. Note that you need a plot of noise versus bandwidth to make a useful decision on which part to use. Noise close to dc will modulate the VCO's output and be within the loop filter's bandwidth while noise at 100 kHz may be attenuated by the loop filter.

OSCILLATOR DEVICES

The 2N3819 FET, a classic from the 1960s, has proven to work well in VFOs but, like the MPF102, which is also long-popular with ham builders, it is manufactured to wide tolerances. Considering an oscillator's importance in receiver stability, you should not hesitate to spend a bit more on a better device. The 2N5484, 2N5485 and 2N5486 are worth considering; together, their transconductance ranges span that of the MPF102, but each is a better-controlled subset of that range. The 2N5245 is a more recent device with better-than-average noise performance that runs at low currents like the 2N3819. The 2N4416/A, also available as the plastic-cased PN4416, is a low-noise device, designed for VHF/UHF amplifier use, which has been featured in a number of good oscillators up to the VHF region. Its low internal capacitance contributes to low frequency drift. The J310 (plastic; the metal-cased U310 is similar) is another popular JFET for use in oscillators.

The 2N5179 (plastic, PN5179 or MPS5179) is a bipolar transistor capable of good performance in oscillators up to the top of the VHF region. Care is needed because its absolute-maximum collector-emitter voltage is only 12 V, and its collector current must not exceed 50 mA. Although these characteristics may seem to convey fragility, the 2N5179 is sufficient for circuits powered by stabilized 6 V power supplies.

Although discrete VCOs have been the mainstay of radio amateurs from the beginning of amateur radio; both the tubes and JFETs that were used to design oscillators are discontinued, or "no longer available" in industry language. Fortunately, Linear Systems (**www.linearsystems.com**) continues to make replacement JFETs for such VCO stalwarts as the MPF102 and U310 for fans of discrete oscillator designs from previous articles and books. Another supplier of JFETs, Interfet (**www.interfet.com**), still supplies the 2N3819.

VHF-UHF capable transistors are not really necessary in LC VFOs because such circuits are rarely used above 10 MHz. (High-bandwidth transistors also increase high-frequency harmonic content in the output signal.) Absolute frequency stability is progressively harder to achieve with increasing frequency, so free-running oscillators are used only rarely to generate VHF-UHF signals for radio communication. Instead, VHF-UHF radios usually use voltage-tuned, phase-locked oscillators in some form of synthesizer. Bipolar devices like the BFR90 and MRF901, with f_T in the 5 GHz region and mounted in stripline packages, are needed for successful oscillator design at UHF.

The popular SA/NE602 mixer IC has a built-in oscillator and can be found in many published circuits. This device has separate input and output pins to the tank and has proved to be quite tame. It may not have been "improved" yet (so far, it has progressed from the SA/NE602 to the SA/NE602A, the A version affording somewhat higher dynamic range than the original SA/NE602). It might be a good idea for anyone laying out a board using one to take a little extra care to keep PCB traces short in the oscillator section to build in some safety margin so that the board can be used reliably in the future. Professional designers know that their designs are going to be built for possibly more than 10 years and have learned to make allowances for the progressive improvement of semiconductor manufacture.

The number of sources for ICs has also shrunk due to consolidation in the semiconductor industry: As of early 2022, Analog Devices (**www.analog.com**) acquired Hittite, Linear Technology, and Maxim; Skyworks (**www.skyworksinc.com**) acquired Silicon Labs' clock products; and Texas Instruments (**www.ti.com**) acquired Burr-Brown and National Semiconductor; Hewlett Packard begat Agilent, which begat Avago (originally HP's semiconductor operation) and Keysight, and, finally, Avago acquired Broadcom and operates under that name, just to name a few. Fortunately, integrated solutions are mainstream products and new products are released on a regular basis.

MECHANICAL CONSTRUCTION

• All oscillator components should be clean and attached to a solid support to minimize thermal changes and mechanical vibration.

• The enclosure should be solid and isolated

from mechanical vibration.

• Keep component leads short and if point-to-point wiring is employed, use heavy wire (#16 to #18 AWG).

• Single-point grounding of the oscillator components is recommended to avoid stray inductance and to minimize noise introduced from other sources. If a PCB is used, include a ground plane.

It is often instructive to look at commercial or military equipment to see what techniques and materials are used for those demanding applications. The mechanical assemblies and parts can be removed from surplus equipment for home-built VFOs, as well.

9.3.2 Temperature Compensation

The general principle for creating a high-stability VFO is to use components with minimal temperature coefficients in circuits that are as insensitive as possible to changes in components' secondary characteristics. Even after careful minimization of the causes of temperature sensitivity, further improvement can still be desirable. The traditional method was to split one of the capacitors in the tank so that it could be apportioned between low-temperature-coefficient parts and parts with deliberate temperature dependency. Only a limited number of different, controlled temperature coefficients are available, so the proportioning between low coefficient and controlled coefficient parts was varied to "dilute" the temperature sensitivity of a part more sensitive than desired. This is a tedious process, involving much trial and error, an undertaking made more complicated by the difficulty of arranging means of heating and cooling the unit being compensated. (Hayward described such a means in December 1993 *QST*.) As commercial and military equipment have been based on frequency synthesizers for some time, supplies of capacitors with controlled temperature sensitivity are drying up. An alternative approach is needed.

A temperature-compensated crystal oscillator (TCXO) is an improved-stability version of a crystal oscillator that is used widely in industry. Instead of using controlled-temperature coefficient capacitors, most TCXOs use a network of thermistors and normal resistors to control the bias of a tuning diode. Manufacturers measure the temperature vs. frequency characteristic of sample oscillators, and then use a computer program to calculate the optimum normal resistor values for production. This can reliably achieve at least a tenfold improvement in stability. We here are not interested in mass manufacture, but the idea of a thermistor tuning a varactor is worth adopting. The parts involved are likely to be available for a long time.

Browsing through component suppliers' catalogs shows ready availability of 4.5 to 5 kΩ bead thermistors intended for temperature-compensation purposes, at less than a dollar each. **Figure 9.20** shows a circuit based on this form of temperature compensation. Commonly available thermistors have negative temperature coefficients, so as temperature rises, the voltage at the counterclockwise (CCW) end of R8 increases, while that at the clockwise (CW) end drops. Somewhere near the center there is no change. Increasing the voltage on the tuning diode decreases its capacitance, so settings toward R8's CCW end simulate a negative-temperature-coefficient capacitor; toward its clockwise end, a positive-temperature-coefficient part. Choose R1 to pass 8.5 mA from whatever supply voltage is available to the 6.2 V reference diode, D1. The 1N821A/1N829A-family diode used has a very low temperature coefficient and needs 7.5 mA bias for best performance; the bridge takes the other 1 mA. R7 and R8 should be good-quality multi-turn trimmers. D2 and C1 need to be chosen to suit the oscillator circuit. Choose the least capacitance that provides enough compensation range. This reduces the noise added to the oscillator. (It is possible, though tedious, to solve for the differential varactor voltage with respect to R2 and R5, via differential calculus and circuit theory. The equations in Hayward's 1993 article can then be modified to accommodate the additional capacitors formed by D2 and C1.) Use a single ground point near D2 to reduce the influence of ground currents from other circuits. Use good-quality metal-film components for the circuit's fixed resistors.

The novelty of this circuit is that it is designed to have an easy and direct adjustment process. The circuit requires two adjustments, one at each of two different temperatures, and achieving them requires a stable frequency counter that can be kept far enough from the radio so that the radio, not the counter, is subjected to the temperature extremes. (Using a receiver to listen to the oscillator under test can speed the adjustments.) After connecting the counter to the oscillator to be corrected, run the radio containing the oscillator and compensator in a room-temperature, draft-free environment until the oscillator's frequency reaches its stable operating temperature (rise over the ambient temperature). Lock its tuning, if possible. Adjust R7 to balance the bridge. This causes a drop of 0 V across R8, a condition you can reach by winding R8 back and forth across its range while slowly adjusting R7. When the bridge is balanced and 0 V appears across R8, adjusting R8 causes no frequency shift. When you've found this R7 setting, leave it there, set R8 to the exact center of its range and record the oscillator frequency.

Run the radio in a hot environment and allow its frequency to stabilize. Adjust R8 to restore the frequency to the recorded value. The sensitivity of the oscillator to temperature should now be significantly reduced between the temperatures at which you performed the adjustments. You will also have somewhat improved the oscillator's stability outside this range.

For best results with any temperature-compensation scheme, it's important to group all the oscillator and compensator components in the same enclosure, avoiding differ-

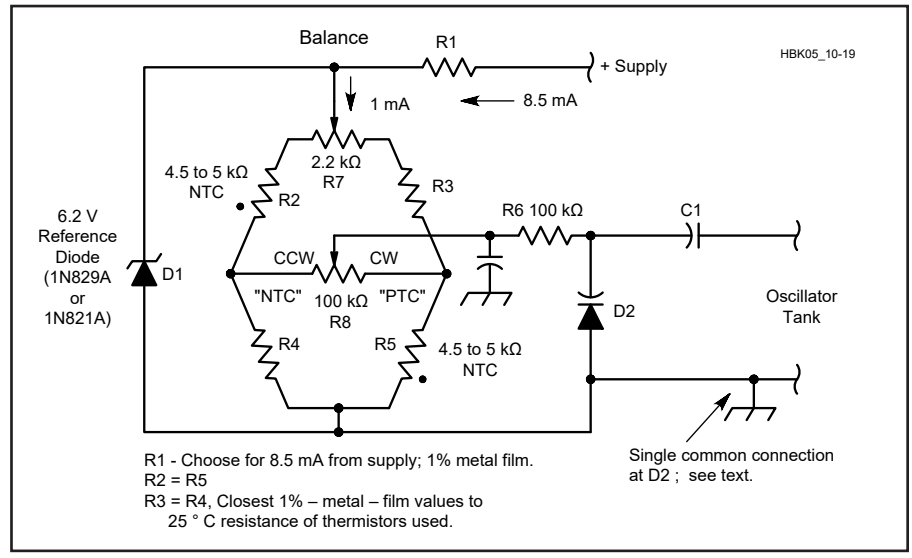

Figure 9.20 — Oscillator temperature compensation is difficult because of the scarcity of negative-temperature-coefficient capacitors. This circuit by GM4ZNX uses a bridge containing two identical thermistors to steer a tuning diode for drift correction. The 6.2 V Zener diode used (a 1N821A or 1N829A) must be a temperature-compensated part; just any 6.2 V Zener will not do.

ences in airflow over components. A good oscillator should not dissipate much power, so it's feasible, even advisable, to mount all of the oscillator components in an unventilated box. In the real world, temperatures change and if the components being compensated and the components doing the compensating have different thermal time constants, a change in temperature can cause a temporary change in frequency until the slower components have caught up. One cure for this is to build the oscillator in a thick-walled metal box that's slow to heat or cool, and so dominates and reduces the possible rate of change of temperature of the circuits inside. This is sometimes called a *cold oven*.

9.3.3 Shielding and Isolation

It is important to remember that any inductor acts as half of a transformer. Oscillators contain inductors running at moderate power levels and so can radiate strong enough signals to cause interference with other parts of a radio or with other radios. This is the tank (or other) inductor behaving as a transformer primary. Oscillators are also sensitive to radiated signals or other nearby varying magnetic fields. This is the tank (or other) inductor also behaving as a transformer secondary. Effective shielding is therefore vital.

Any oscillator is particularly sensitive to interference on the same or very nearby frequency. If this interference is strong enough an undesirable effect called *injection locking* will occur. The oscillator effectively stops oscillating and instead directly follows the interfering signal. For example, a VFO used to directly drive a power amplifier and antenna (to form a simple CW transmitter) can prove surprisingly difficult to shield well enough because of injection locking from any leakage of the power amplifier's high-level signal back into the oscillator. Even if injection locking does not fully kick in, the wrestling inside the VFO between its own oscillation and the interference can affect its frequency, resulting in an unstable transmitted signal. If the radio gear is in the antenna's near field, there are also strong fields that are coherent with the VFO oscillation, making sufficient shielding even more difficult.

The following rules of thumb continue to serve ham builders well:

• Use a complete metal box, with as few holes drilled in it as possible, with good contact around surface(s) where its lid(s) fit(s) on.

• Use feedthrough capacitors on power and control lines that pass in and out of the VFO enclosure and on the transmitter or transceiver enclosure as well.

• Use buffer amplifier circuitry that amplifies the signal by the desired amount and provides sufficient attenuation of signal energy flowing in the reverse direction. This is known as *reverse isolation* and is a frequently overlooked loophole in shielding. Figures 9.14 and 9.16 include buffer circuitry of proven performance. Another (and higher-cost) option is to consider using a high-speed buffer-amplifier IC (such as the LM6321N from Texas Instruments, a part that combines the high input impedance of an op amp with the ability to drive 50-Ω loads directly up into the VHF range).

• Use a mixing-based frequency-generation scheme instead of one that operates straight through or by means of multiplication. Such a system's oscillator stages can operate on frequencies with no direct frequency relationship to its output frequency. This essentially eliminates the possibility of injection locking the VFO.

• Use the time-tested technique of running your VFO at a sub-harmonic of the output signal desired — say, 3.5 MHz in a 7 MHz transmitter — and multiply its output frequency in a suitably nonlinear stage for further amplification at the desired frequency. This does reduce the tendency to injection lock.

9.4 Crystal Oscillators

Because crystals afford Q values and frequency stabilities that are orders of magnitude better than those achievable with LC circuits, fixed-frequency oscillators usually use quartz crystal resonators. Master references for frequency counters and synthesizers are always based on crystal oscillators.

So glowing is the crystal's reputation for stability that newcomers to radio experimentation naturally believe that the presence of a crystal in an oscillator will force oscillation at the frequency stamped on the can. This impression is usually revised after the first few experiences to the contrary! There is no sure-fire crystal oscillator circuit (although some are better than others); reading and experience soon provide a learner with plenty of anecdotes to the effect that:

• Some circuits have a reputation of being temperamental, even to the point of not always starting.

• Crystals sometimes mysteriously oscillate on unexpected frequencies.

Even crystal manufacturers have these problems, so don't be discouraged from building crystal oscillators. The occasional uncooperative oscillator is a nuisance, not a disaster, and it just needs a little individual attention. Knowing how a crystal behaves is the key to a cure.

Ulrich Rohde, N1UL, has generously contributed a pair of detailed papers that discuss the crystal oscillator along with several HF and VHF designs. Both papers, "Quartz Crystal Oscillator Design" and "A Novel Grounded Base Oscillator Design for VHF/UHF Frequencies" are included with the supplemental material accompanying this book.

9.4.1 Quartz and the Piezoelectric Effect

Quartz is a crystalline material with a regular atomic structure that can be distorted by the simple application of force. Remove the force, and the distorted structure springs back to its original form with very little energy loss. This property allows *acoustic waves* — sound — to propagate rapidly through quartz with very little attenuation, because the velocity of an acoustic wave depends on the elasticity and density (mass/volume) of the medium through which the wave travels.

If you heat a material, it expands. Heating may cause other characteristics of a material to change — such as elasticity, which affects the speed of sound in the material. In quartz, however, expansion and change in the speed of sound are very small and tend to cancel, which means that the transit time for sound to pass through a piece of quartz is very stable.

The third property of this wonder material is that it is *piezoelectric*. Apply an electric field to a piece of quartz and the crystal lattice distorts just as if a force had been applied. The electric field applies a force to electrical charges locked in the lattice structure. These charges are captive and cannot move around in the lattice as they can in a semiconductor, for quartz is an insulator. A capacitor's dielectric stores energy by creating physical distortion on an atomic or molecular scale. In a piezoelectric crystal's lattice, the distortion

An Oscillator for Crystal Testing

A test oscillator for determining the condition of a crystal and its actual frequency is quite useful. The circuit described in "Crystal Test Oscillators" by Fred Brown, W6HPH, will work with crystals from below 25 kHz to above 25 MHz in their fundamental mode. This article is available in the online material.

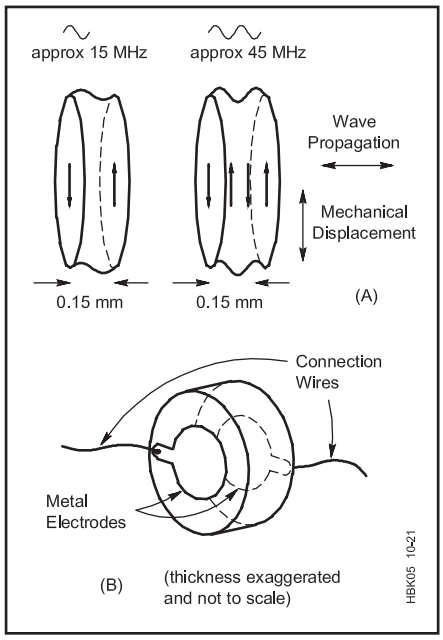

Figure 9.21 — Thickness-shear vibration at a crystal's fundamental and third overtone (A); B shows how the modern crystals commonly used by radio amateurs consist of etched quartz discs with electrodes deposited directly on the crystal surface.

affects the entire structure. In some piezoelectric materials, this effect is sufficiently pronounced that special shapes can be made that bend *visibly* when a field is applied.

Consider a rod made of quartz. Any sound wave propagating along it eventually hits an end, where there is a large and abrupt change in acoustic impedance. Just as when an RF wave hits the end of an unterminated transmission line, a strong reflection occurs. The rod's other end similarly reflects the wave. At some frequency, the phase shift of a round trip will be such that waves from successive round trips exactly coincide in phase and reinforce each other, dramatically increasing the wave's amplitude. This is *resonance*.

The passage of waves in opposite directions forms a standing wave with antinodes at the rod ends. Here we encounter a seeming ambiguity: not just one, but a family of different frequencies, causes standing waves — a family fitting the pattern of ½, ⅔, ⅝, ¾ and so on, wavelengths into the length of the rod. And this is the case: A quartz rod can resonate at any and all of these frequencies.

The lowest of these frequencies, where the crystal is ½ wavelength long, is called the *fundamental* mode. The others are named the third, fifth, seventh and so on, *overtones*. There is a small phase-shift error during reflection at the ends, which causes the frequencies of the overtone modes to differ slightly from odd integer multiplies of the fundamental. Thus, a crystal's third overtone is very close to, but not exactly, three times its fundamental frequency. Many people are confused by overtones and harmonics. Harmonics are additional signals at exact integer multiples of the fundamental frequency. Overtones are not signals at all; they are additional resonances that can be exploited if a circuit is configured to excite them.

The crystals we use most often resonate in the 1 to 30 MHz region and are of the *AT-cut, thickness shear* type, although these last two characteristics are rarely mentioned. A 15 MHz fundamental crystal of this type is about 0.15 mm thick. Because of the widespread use of pressure-mounted FT-243 crystals, you may think of crystals as small rectangles on the order of a half-inch in size. The crystals we commonly use today are discs, etched and/or doped to their final dimensions, with metal electrodes deposited directly on the quartz. A crystal's diameter does not directly affect its frequency; diameters of 8 to 15 mm are typical. (Quartz crystals are also discussed in the **Analog and Digital Filtering** chapter.)

AT-cut is one of a number of possible standard designations for the orientation at which a crystal disc is sawed from the original quartz crystal. The crystal lattice atomic structure is asymmetric, and the orientation of this with respect to the faces of the disc influences the crystal's performance. *Thickness shear* is one of a number of possible orientations of the crystal's mechanical vibration with respect to the disc. In this case, the crystal vibrates perpendicularly to its thickness. This is not easy to visualize, and diagrams don't help much, but **Figure 9.21** is an attempt at illustrating this. Place a moist bathroom sponge between the palms of your hands, move one hand up and down, and you'll see thickness shear in action.

There is a limit to how thin a disc can be made, given requirements of accuracy and price. Traditionally, fundamental-mode crystals have been made up to 20 MHz, although 30 MHz is now common at a moderately raised price. Using techniques pioneered in the semiconductor industry, crystals have been made with a central region etched down to a thin membrane, surrounded by a thick ring for robustness. This approach can push fundamental resonances to over 100 MHz, but these are more lab curiosities than parts for everyday use. The easy solution for higher frequencies is to use a manufacturably-thick crystal on an overtone mode. All crystals have multiple modes, so if you order a 28.060 MHz, third-overtone unit for a little QRP transmitter, you'll get a crystal with a fundamental resonance somewhere near 9.353333 MHz, but its manufacturer will have adjusted the thickness to plant the third overtone exactly on the ordered frequency. An accomplished manufacturer can do tricks with the flatness of the disc faces to make the wanted overtone mode a little more active and the other modes a little less active. (As some builders discover, however, this does not *guarantee* that the wanted mode is the most active!)

Quartz's piezoelectric property provides a simple way of driving the crystal electrically. Early crystals were placed between a pair of electrodes in a case. This gave amateurs the opportunity to buy surplus crystals, open them and grind them a little to reduce their thickness, thus moving them to higher frequencies. The frequency could be reduced very slightly by loading the face with extra mass, such as by blackening it with a soft pencil. Modern crystals have metal electrodes deposited directly onto their surfaces (Figure 9.21B), and such tricks no longer work.

The piezoelectric effect works both ways. Deformation of the crystal produces voltage across its electrodes, so the mechanical energy in the resonating crystal can also be extracted electrically by the same electrodes. Seen electrically, at the electrodes, the mechanical resonances look like electrical resonances. Their Q is very high. A Q of 10,000 would characterize a poor crystal today; 100,000 is often reached by high-quality parts. For comparison, a Q of over 200 for an LC tank is considered good.

9.4.2 Frequency Accuracy

A crystal's frequency accuracy is as outstanding as its Q. Several factors determine a crystal's frequency accuracy. First, the manufacturer makes parts with certain tolerances: ±200 ppm for a low-quality crystal for use as in a microprocessor clock oscillator, ±10 ppm for a good-quality part for professional radio use. Anything much better than this starts to get expensive! A crystal's resonant frequency is influenced by the impedance presented to its terminals, and manufacturers assume that once a crystal is brought within several parts per million of the nominal frequency, its user will perform fine adjustments electrically.

Second, a crystal ages after manufacture. Aging could give increasing or decreasing frequency; whichever, a given crystal usually keeps aging in the same direction. Aging is rapid at first and then slows down. Aging is influenced by the care in polishing the surface of the crystal (time and money) and by its holder style. The cheapest holder is a soldered-together, two-part metal can with glass bead insulation for the connection pins. Soldering debris lands on the crystal and affects its frequency. Alternatively, a two-part metal case can be made with flanges that are pressed together until they fuse, a process called *cold-welding*. This is much cleaner and improves aging rates roughly fivefold compared to soldered cans. An all-glass case can be made in two parts and fused together by heating in a

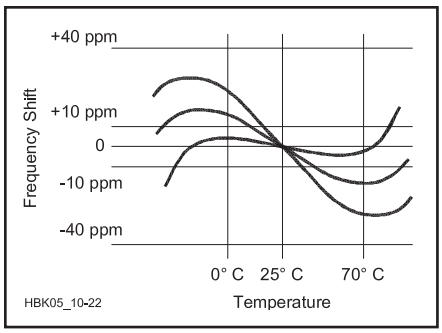

Figure 9.22 — Slight changes in a crystal cut's orientation shift its frequency-versus-temperature curve.

vacuum. The vacuum raises the Q, and the cleanliness results in aging that's roughly 10 times slower than that achievable with a soldered can. The best crystal holders borrow from vacuum-tube assembly processes and have a *getter*, a highly reactive chemical substance that traps remaining gas molecules, but such crystals are used only for special purposes.

Third, temperature influences a crystal. A reasonable, professional quality part might be specified to shift not more than ±10 ppm over 0 to 70 °C. An AT-cut crystal has an *S*-shaped frequency-versus-temperature characteristic, which can be varied by slightly changing the crystal cut's orientation. **Figure 9.22** shows the general shape and the effect of changing the cut angle by only a few seconds of arc. Notice how all the curves converge at 25 °C. This is because this temperature is normally chosen as the reference for specifying a crystal. The temperature stability specification sets how accurate the manufacturer must make the cut. Better stability may be needed for a crystal used as a receiver frequency standard, frequency counter clock and so on. A crystal's temperature characteristic shows a little hysteresis. In other words, there's a bit of offset to the curve depending on whether temperature is increasing or decreasing. This is usually of no consequence except in the highest-precision circuits.

It is the temperature of the quartz that is important, and as the usual holders for crystals all give effective thermal insulation, only a couple of milliwatts dissipation by the crystal itself can be tolerated before self-heating becomes troublesome. Because such heating occurs in the quartz itself and does not come from the surrounding environment, it defeats the effects of temperature compensators and ovens.

The techniques shown earlier for VFO temperature compensation can also be applied to crystal oscillators. An after-compensation drift of 1 ppm is routine and 0.5 ppm is good. The result is a *temperature-compensated crystal oscillator* (*TCXO*). Recently, oscillators have appeared with built-in digital thermometers, microprocessors and ROM look-up tables customized on a unit-by-unit basis to control a tuning diode via a digital-to-analog converter (DAC) for temperature compensation. These *digitally temperature-compensated oscillators* (*DTCXOs*) can reach 0.1 ppm over the temperature range. With automated production and adjustment, they promise to become the cheapest way to achieve this level of stability.

Oscillators have long been placed in temperature-controlled *ovens*, which are typically held at 80 °C. Stability of several parts per billion can be achieved over temperature, but this is a limited benefit as aging can easily dominate the accuracy. These are usually called *oven-controlled crystal oscillators* (*OCXOs*).

Fourth, the crystal is influenced by the impedance presented to it by the circuit in which it is used. This means that care is needed to make the rest of an oscillator circuit stable, in terms of impedance and phase shift.

Gravity can slightly affect crystal resonance. Turning an oscillator upside down usually produces a small frequency shift, usually much less than 1 ppm; turning the oscillator back over reverses this. This effect is quantified for the highest-quality reference oscillators.

9.4.3 The Equivalent Circuit of a Crystal

Because a crystal is a passive, two-terminal device, its electrical appearance is that of an impedance that varies with frequency. **Figure 9.23A** shows a very simplified sketch of the magnitude (phase is ignored) of the impedance of a quartz crystal. The general trend of dropping impedance with increasing frequency implies capacitance across the crystal. The sharp fall to a low value resembles a series-tuned tank, and the sharp peak resembles a parallel-tuned tank. These are referred to as series and parallel resonances. Figure 9.23B shows a simple circuit that will produce this impedance characteristic. The impedance looks purely resistive at the exact centers of both resonances, and the region between them has impedance increasing with frequency, which looks inductive.

C1 (sometimes called *motional capacitance*, C_m, to distinguish it from the lumped

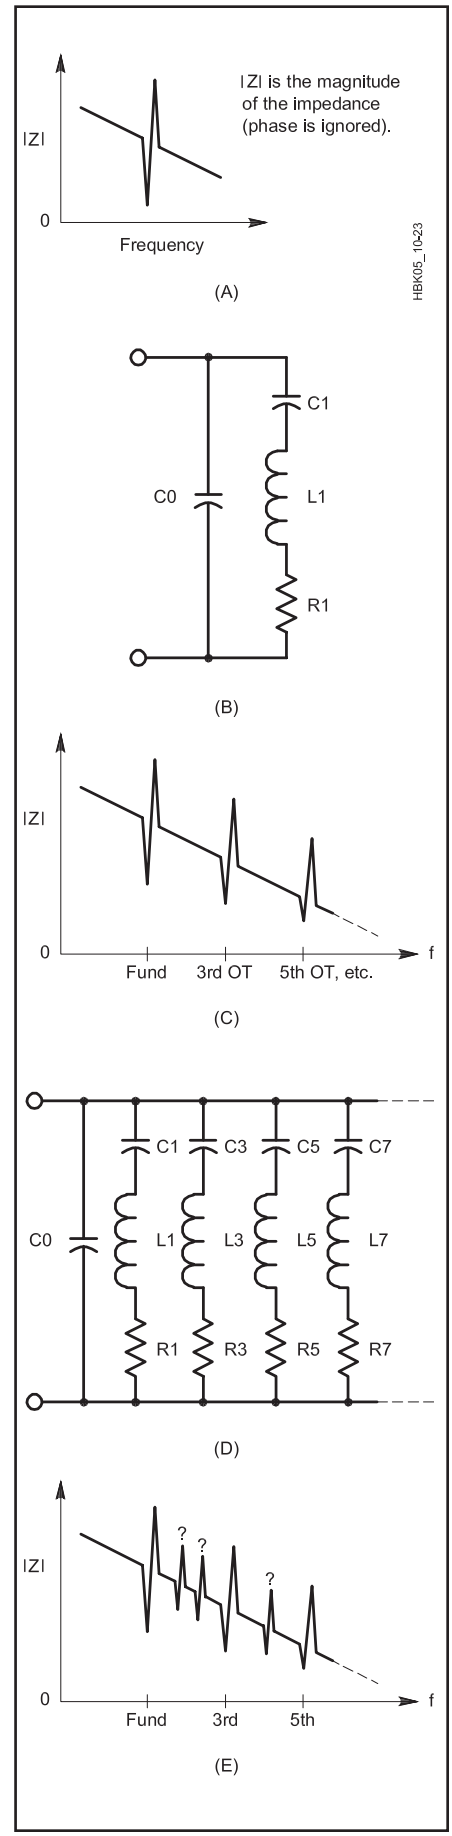

Figure 9.23 — Exploring a crystal's impedance (A) and equivalent circuit (B) through simplified diagrams. C and D extend the investigation to include overtones; E, to spurious responses not easily predictable by theory or controllable through manufacture. A crystal may oscillate on any of its resonances under the right conditions.

capacitance it approximates) and L1 (*motional inductance, L_m*) create the series resonance, and as C0 and R1 are both fairly small, the impedance at the bottom of the dip is very close to R1. At parallel resonance, L1 is resonating with C1 and C0 in series, hence the higher frequency. The impedance of the parallel tank is extremely high; the terminals are connected to a capacitive tap, which causes them to see only a small fraction of what is still a very large impedance. The overtones should not be neglected, so Figures 9.23C and 9.23D include them. Each overtone has series and parallel resonances and so appears as a series tank in the equivalent circuit. C0 again provides the shifted parallel resonance.

This is still simplified, because real-life crystals have a number of spurious, unwanted modes that add yet more resonances, as shown in Figure 9.23E. These are not well controlled and may vary a lot even between crystals made to the same specification. Crystal manufacturers work hard to suppress these spurs and have evolved a number of recipes for shaping crystals to minimize them. Just where they switch from one design to another varies from manufacturer to manufacturer.

Always remember that the equivalent circuit is just a representation of crystal behavior and does not represent circuit components actually present. Its only use is as an aid in designing and analyzing circuits using crystals. **Table 9.1** lists typical equivalent-circuit values for a variety of crystals. It is impossible to build a circuit with 0.026 to 0.0006 pF capacitors; such values would simply be swamped by strays. Similarly, the inductor must have a Q that is orders of magnitude better than is practically achievable, and impossibly low stray C in its winding.

The values given in Table 9.1 are nothing more than rough guides. A crystal's frequency is tightly specified, but this still allows inductance to be traded for capacitance. A good manufacturer could hold these characteristics within a ±25% band or could vary them over a 5:1 range by special design. Similarly marked parts from different sources vary widely in motional inductance and capacitance.

Quartz is not the only material that behaves in this way, but it is the best. Resonators can be made out of lithium tantalate and a group of similar materials that have lower Q, allowing them to be *pulled* over a larger frequency range in VXOs. Much more common, however, are ceramic resonators based on the technology of the well-known ceramic IF filters. These have much lower Q than quartz and much poorer frequency precision. They serve mainly as clock resonators for cheap microprocessor systems in which every last cent must be saved. A ceramic resonator could be used as the basis of a wide range, cheap VXO, but its frequency stability would not be as good as a good LC VFO.

Table 9.1

Typical Equivalent Circuit Values for a Variety of Crystals

Crystal Type	*Series L*	*Series C (pF)*	*Series R (Ω)*	*Shunt C (pF)*
1 MHz fundamental	3.5 H	0.007	340	3.0
10 MHz fundamental	9.8 mH	0.026	7	6.3
30 MHz third overtone	14.9 mH	0.0018	27	6.2
100 MHz fifth overtone	4.28 mH	0.0006	45	7.0

9.4.4 Crystal Oscillator Circuits

(See also the papers "What You Always Wanted to Know About Colpitts Oscillators," "Quartz Crystal Oscillator Design," "A Novel Grounded Base Oscillator Design for VHF/UHF Frequencies," and "Some Thoughts on Designing Very High Performance VHF Oscillators" by Ulrich Rohde, N1UL, included with the online material accompanying this book.)

Crystal oscillator circuits are usually categorized as series- or parallel-mode types, depending on whether the crystal's low- or high-impedance resonance comes into play at the operating frequency. The series mode is now the most common; parallel-mode operation was more often used with vacuum tubes. **Figure 9.24** shows a basic series-mode oscillator. Some people would say that it is an overtone circuit, used to run a crystal on one of its overtones, but this is not necessarily true. The tank (L-C1-C2) tunes the collector of the common-base amplifier. C1 is larger than C2, so the tank is tapped in a way that transforms to lower impedance, decreasing signal voltage, but increasing current. The current is fed back into the emitter via the crystal. The common-base stage provides a current gain of less than unity, so the transformer in the form of the tapped tank is essential to give loop gain. There are *two* tuned circuits, the obvious collector tank and the series-mode one "in" the crystal. The tank kills the amplifier's gain away from its tuned frequency, and the crystal will only pass current at the series resonant frequencies of its many modes. The tank resonance is much broader than any of the crystal's modes, so it can be thought of as the crystal setting the frequency, but the tank selecting which of the crystal's modes is active. The tank could be tuned to the crystal's fundamental, or one of its overtones.

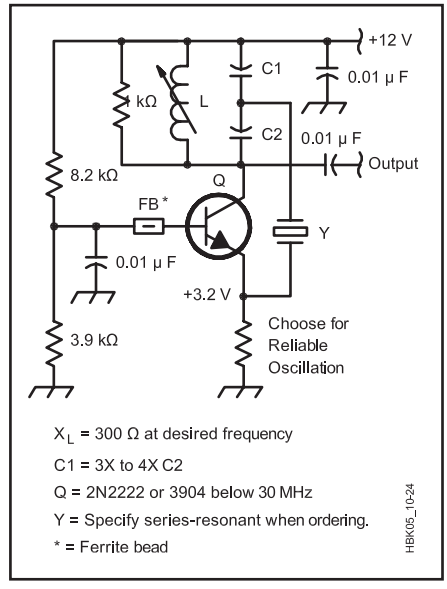

Figure 9.24 — A basic series-mode crystal oscillator. A 2N5179 can be used in this circuit if a lower supply voltage is used; see text.

Fundamental oscillators can be built without a tank quite successfully, but there is always the occasional one that starts up on an overtone or spurious mode. Some simple oscillators have been known to change modes while running (an effect triggered by changes in temperature or loading) or to not always start in the same mode! A series-mode oscillator should present low impedance to the crystal at the operating frequency. In Figure 9.24, the tapped collector tank presents a transformed fraction of the 1-kΩ collector load resistor to one end of the crystal, and the emitter presents a low impedance to the other. To build a practical oscillator from this circuit, choose an inductor with a reactance of about 300 Ω at the wanted frequency and calculate C1 in series with C2 to resonate with it. Choose C1 to be 3 to 4 times larger than C2. The amplifier's quiescent ("idling") current sets the gain and hence the operating level. This is not easily calculable, but can be found by experiment. Too little quiescent current and the oscillator will not start reliably; too much and the transistor can drive itself into saturation. If an oscilloscope is available, it can be used to check the collector waveform; otherwise, some form of RF voltmeter can be used to allow the collector voltage to be set to 2 to 3 V RMS. 3.3 kΩ would be a suitable starting point for the emitter bias resistor. The transistor type is not critical; 2N2222A or 2N3904 would be fine up to 30 MHz; a 2N5179 would allow operation as an overtone

oscillator to over 100 MHz (because of the low collector voltage rating of the 2N5179, a supply voltage lower than 12 V is required). The ferrite bead on the base gives some protection against parasitic oscillation at UHF.

If the crystal is shorted, this circuit should still oscillate. This gives an easy way of adjusting the tank; it is even better to temporarily replace the crystal with a small-value (tens of ohms) resistor to simulate its *equivalent series resistance* (ESR), and adjust L until the circuit oscillates close to the wanted frequency. Then restore the crystal and set the quiescent current. If a lot of these oscillators were built, it would sometimes be necessary to adjust the current individually due to the different equivalent series resistance of individual crystals. One variant of this circuit has the emitter connected directly to the C1/C2 junction, while the crystal is a decoupler for the transistor base (the existing capacitor and ferrite bead not being used). This works, but with a greater risk of parasitic oscillation.

We commonly want to trim a crystal oscillator's frequency. While off-tuning the tank a little will pull the frequency slightly, too much detuning spoils the mode control and can stop oscillation (or worse, make the circuit unreliable). The answer to this is to add a trimmer capacitor, which will act as part of the equivalent series tuned circuit, in series with the crystal. This will shift the frequency in one way only, so the crystal frequency must be re-specified to allow the frequency to be varying around the required value. It is common to specify a crystal's frequency with a standard load (30 pF is commonly specified), so that the manufacturer grinds the crystal such that the series resonance of the specified mode is accurate when measured with a capacitor of this value in series. A 15 to 50 pF trimmer can be used in series with the crystal to give fine frequency adjustment. Too little capacitance can stop oscillation or prevent reliable starting. The Q of crystals is so high that marginal oscillators can take several seconds to start!

This circuit can be improved by driving the crystal's lower series-resonant impedance with an emitter follower as in **Figure 9.25**. This is the *Butler* oscillator. Again, the tank controls the mode to either force the wanted overtone or protect the fundamental mode. The tank need not be tapped because Q2 provides current gain, although the circuit is sometimes seen with C split, driving Q2 from a tap. The position between the emitters offers a good, low-impedance environment to keep the crystal's in-circuit Q high. R, in the emitter of Q1, is again selected to give reliable oscillation. The circuit has been shown with a capacitive load for the crystal, to suit a unit specified for a 30 pF load. An alternative circuit to give electrical fine tuning is also shown. The diodes across the tank act as limiters to stabilize the operating amplitude and limit the power dissipated in the crystal by clipping the drive voltage to Q2. The tank should be adjusted to peak at the operating frequency, not used to trim the frequency. The capacitance in series with the crystal is the proper frequency trimmer.

The Butler circuit works well, and has been used in critical applications to 140 MHz (seventh-overtone crystal, 2N5179 transistor). Although the component count is high, the extra parts are cheap ones. Increasing the capacitance in series with the crystal reduces the oscillation frequency but has a progressively diminishing effect. Decreasing the capacitance pulls the frequency higher, to a point at which oscillation stops; before this point is reached, start-up will become unreliable. The possible amount of adjustment, called *pulling range*, depends on the crystal; it can range from less than 10 to several hundred parts per million. Overtone crystals have much less pulling range than fundamental crystals on the same frequency; the reduction in pulling is roughly proportional to the square of the overtone number.

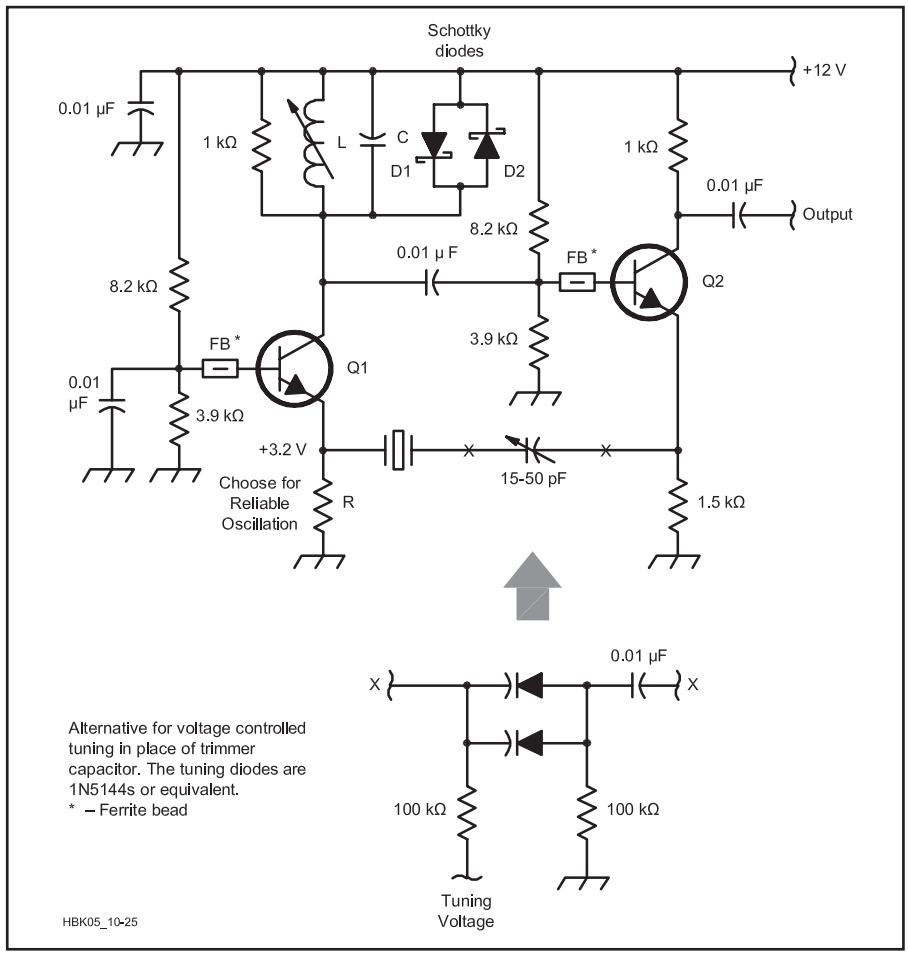

Figure 9.25 — A Butler crystal oscillator with Q2 connected as an emitter follower to drive the crystal's low series-resonant impedance.

LOW-NOISE CRYSTAL OSCILLATORS

Figure 9.26A shows a crystal operating in its series mode in a series-tuned Colpitts circuit. Because it does not include an LC tank to prevent operation on unwanted modes, this circuit is intended for fundamental mode operation only and relies on that mode being the most active. If the crystal is ordered for 30 pF loading, the frequency trimming capacitor can be adjusted to compensate for the loading of the capacitive divider of the Colpitts circuit. An unloaded crystal without a trimmer would operate slightly off the exact series resonant frequency in order to create an inductive impedance to resonate with the divider capacitors. Ulrich Rohde, N1UL, in Figure 4-47 of his book *Digital PLL Frequency Synthesizers — Theory and Design*, published an elegant alternative method of extracting an output signal from this type of circuit, shown in Figure 9.26B. This taps off a signal from the current in the crystal itself. This can be thought of as using the crystal as a band-pass filter for the oscillator output. The RF choke in the emitter keeps the emitter bias resistor from loading the tank and degrading the Q. In this case (3 MHz operation), it has

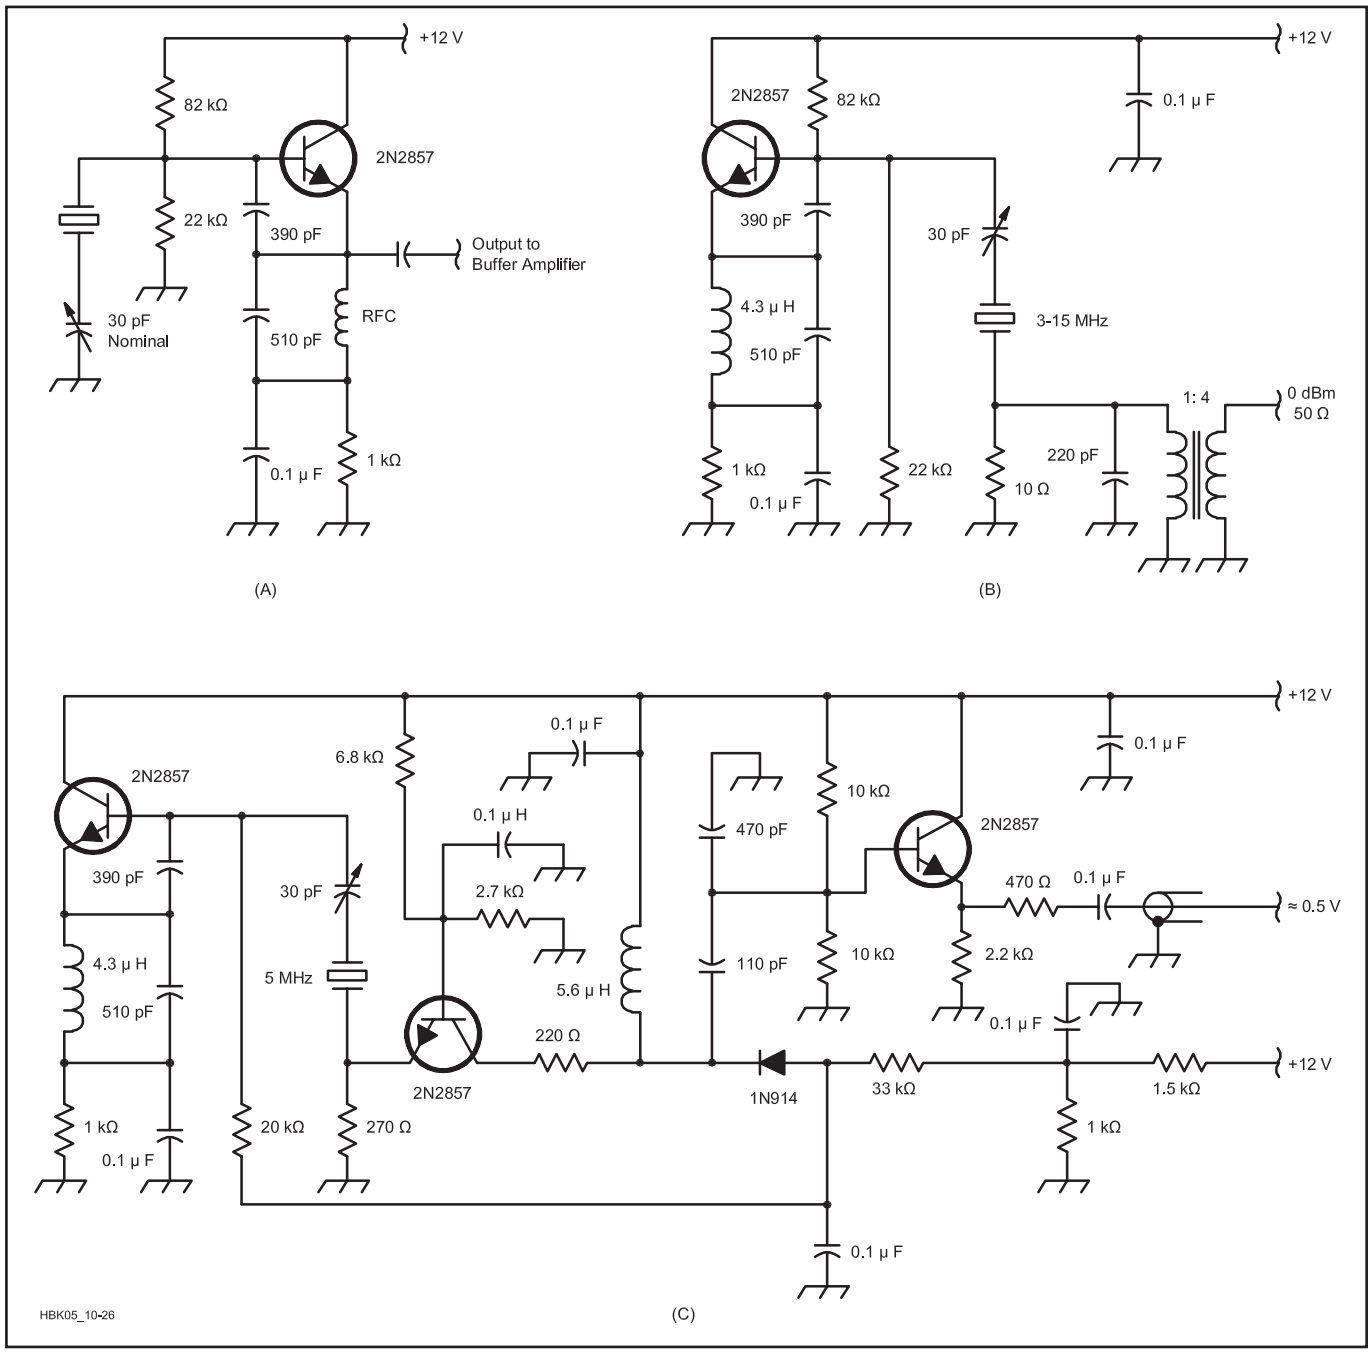

Figure 9.26 — The crystal in the series-tuned Colpitts oscillator at A operates in its series-resonant mode. B shows N1UL's low-noise version, which uses the crystal as a filter and features high harmonic suppression (from Rohde, *Microwave and Wireless Synthesizers Theory and Design*; see references). The circuit at C builds on the B version by adding a common-base output amplifier and ALC loop.

been chosen to resonate close to 3 MHz with the parallel capacitor (510 pF) as a means of forcing operation on the wanted mode. The 10-Ω resistor and the transformed load impedance will reduce the in-circuit Q of the crystal, so a further development substituted a common base amplifier for the resistor and transformer. This is shown in Figure 9.26C. The common-base amplifier is run at a large quiescent current to give a very low input impedance. Its collector is tuned to give an output with low harmonic content and an emitter follower is used to buffer this from the load. This oscillator sports a simple ALC system, in which the amplified and rectified signal is used to reduce the bias voltage on the oscillator transistor's base. This circuit is described as achieving a phase noise level of –168 dBc/Hz a few kilohertz out from the carrier. This may seem far beyond what may ever be needed, but frequency multiplication to high frequencies, whether by classic multipliers or by frequency synthesizers, multiplies the deviation of any FM/PM sidebands as well as the carrier frequency. This means that phase noise worsens by 20 dB for each tenfold multiplication of frequency. A clean crystal oscillator and a multiplier chain is still the best way of generating clean microwave signals for use with narrow-band modulation schemes.

It has already been mentioned that overtone crystals are much harder to pull than fundamental ones. This is another way of saying that overtone crystals are less influenced by their surrounding circuit, which is helpful in

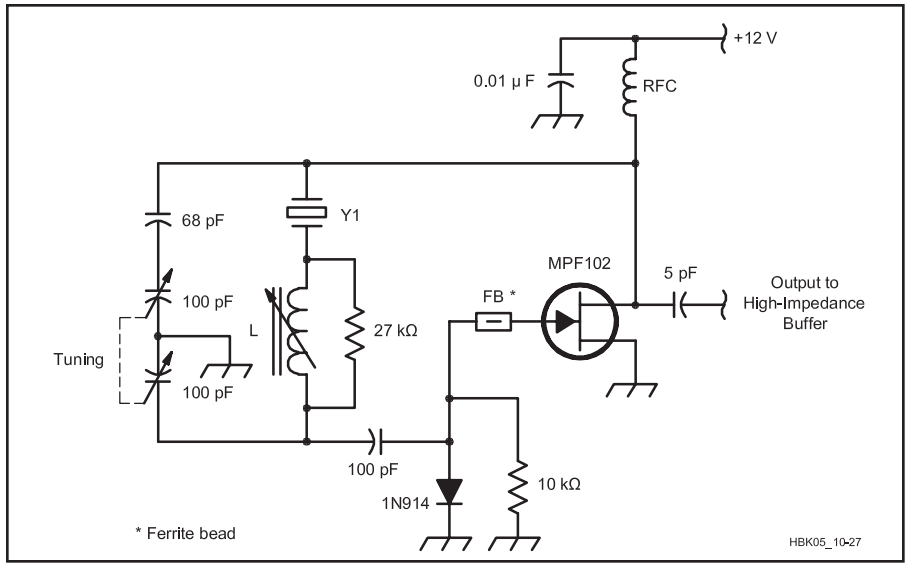

Figure 9.27 — A wide-range variable-crystal oscillator (VXO) by W7ZOI and W1FB. It was originally designed for use in low-power radios without the usual wide-range VFO.

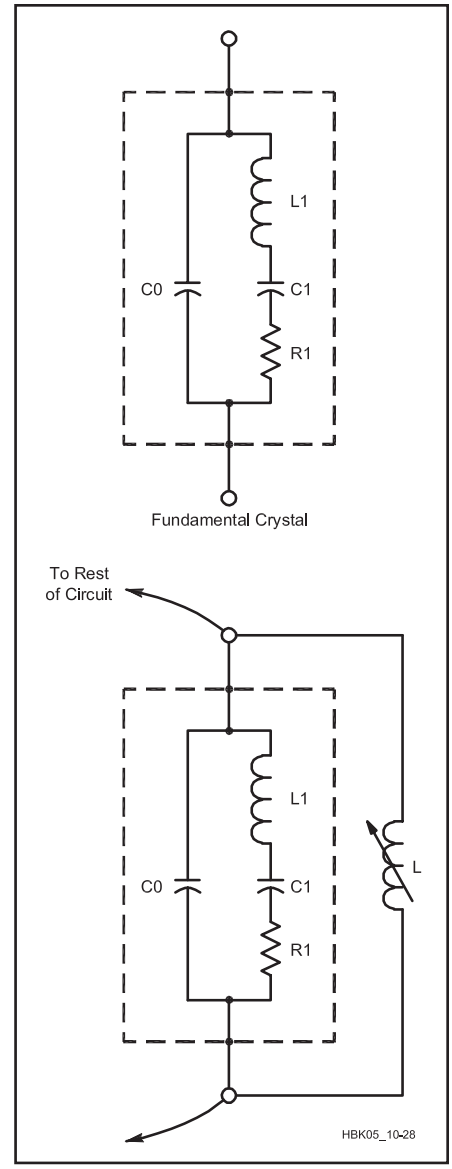

Figure 9.28 — Using an inductor to "tune out" C0 can increase a crystal oscillator's pulling range.

a frequency-standard oscillator like this one. Even though 5 MHz is in the main range of fundamental-mode crystals and this circuit will work well with them, an overtone crystal has been used. To further help stability, the power dissipated in the crystal is kept to about 50 µW. The common-base stage is effectively driven from a higher impedance than its own input impedance, under which conditions it gives a very low noise figure.

9.4.5 Variable-frequency Crystal Oscillators (VXOs)

Some crystal oscillators have frequency trimmers. If the trimmer is replaced by a variable capacitor as a front-panel control, we have a *variable crystal oscillator* (*VXO*): a crystal-based VFO with a narrow tuning range, but good stability and noise performance. VXOs are often used in small, simple QRP transmitters to tune a few kilohertz around common calling frequencies. Artful constructors, using optimized circuits and components, have achieved 1000-ppm tuning ranges. Poor-quality "soft" crystals are more pull-able than high-Q ones. Overtone crystals are not suited to VXOs. For frequencies beyond the usual limit for fundamental mode crystals, use a fundamental unit and frequency multipliers.

ICOM and Mizuho made some 2 meter SSB transceivers based on multiplied VXO local oscillators. This system is simple and can yield better performance than many expensive synthesized radios. SSB filters are available at 9 or 10.7 MHz, and yield sufficient image rejection with a single conversion. The choice of VXO frequency depends on whether the LO is to be above or below signal frequency and how much multiplication can be tolerated. Below 8 MHz, multiplier filtering is difficult. Above 15 MHz, the tuning range per crystal narrows. A 50-200 kHz range per crystal should work with a modern front-end design feeding a good 9 MHz IF, for a contest quality 2-meter SSB receiver.

The circuit in **Figure 9.27** is a JFET VXO from Wes Hayward, W7ZOI, and Doug DeMaw, W1FB, optimized for wide-range pulling. Published in *Solid State Design for the Radio Amateur*, many have been built and its ability to pull crystals as far as possible has been proven. Ulrich Rohde, N1UL, has shown that the diode arrangement as used here to make signal-dependent negative bias for the gate confers a phase-noise disadvantage, but oscillators like this that pull crystals as far as possible need any available means to stabilize their amplitude and aid start-up. In this case, the noise penalty is worth paying. This circuit can achieve a 2000-ppm tuning range with amenable crystals. If you have some overtone crystals in your junk box whose fundamental frequency is close to the wanted value, they are worth trying.

This sort of circuit doesn't necessarily stop pulling at the extremes of the possible tuning range; sometimes the range is set by the onset of such undesirable behavior as jumping mode or simply stopping oscillating. L was a 16 µH slug-tuned inductor for 10 MHz operation. It is important to minimize the stray and interwinding capacitance of L since this dilutes the range of impedance presented to the crystal.

One trick that can be used to aid the pulling range of oscillators is to tune out the C0 of the equivalent circuit with an added inductor. **Figure 9.28** shows how. L is chosen to resonate with C0 for the individual crystal, turning it into a high-impedance parallel-tuned circuit. The Q of this circuit is orders of magnitude less than the Q of the true series resonance of the crystal, so its tuning is much broader. The value of C0 is usually just a few picofarads, so L has to be a fairly large value considering the frequency at which it is resonated. This means that L has to have low stray capacitance or else it will self-resonate at a lower frequency. The tolerance on C0 and the variations of the stray C of the inductor means that individual adjustment is needed. This technique can also work wonders in crystal ladder filters.

9.4.6 Logic-Gate Crystal Oscillators

The frequency-determining network of an RC oscillator has a Q of less than one, which means that phase shift changes very slowly with different frequencies (see the prior section on RC Oscillators). The Pierce crystal oscillator discussed previously is a converted phase-shift oscillator, with the crystal taking the place of one series resistor. The crystal provides a much higher phase shift than an RC stage, so the crystal "controls" the frequency of oscillation.

The actual frequency of oscillation is the frequency at which the Barkhausen criteria are met. The crystal must operate near its series resonance in order for the loop gain to be unity. Oscillation then settles at the frequency where the crystal phase shift, added to the RC phase shifts, equals 180 degrees. The crystal usually provides between 45 and 60 degrees of phase shift, so the oscillation frequency is above the series resonance and below the parallel resonance of the crystal where the crystal behaves as a large inductor. (See the figure showing crystal response in the section Quartz Crystal Filters in the **Analog and Digital Filtering** chapter.)

The Pierce circuit is rarely seen in this full form. Instead, a cut-down version is the most common circuit in many microprocessors and other digital ICs that need a crystal-controlled clock. **Figure 9.29** shows this minimalist Pierce, using a CMOS logic inverter as the amplifier. R_{bias} provides dc negative feedback to bias the gate into its linear region. This value is not critical, anything between 100 kΩ and 10 MΩ works fine. The RC phase shifts needed to make this work come from R_{out}-C1 and $R_{crystal}$-C2. This circuit has a reputation of being temperamental, mainly because neither R_{out} nor $R_{crystal}$ are well documented or controlled in manufacturing. There is also a general belief that this circuit requires C1 = C2, which is not true.

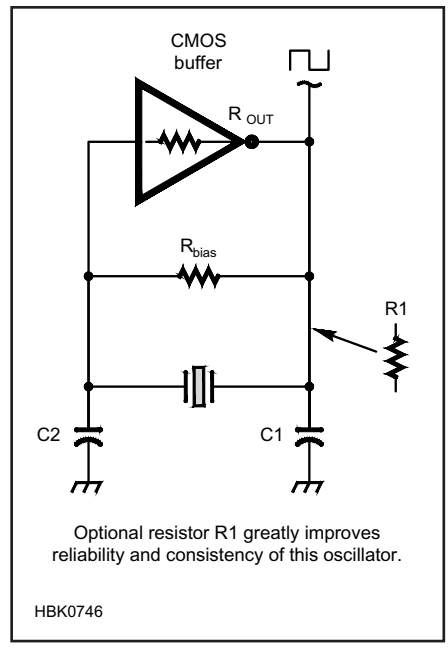

Figure 9.29 — The simplified Pierce oscillator using a logic-gate for the inverting amplifier. Adding R1 improves oscillator design and reliability.

A great improvement in performance is achieved by adding one resistor, R1, as shown in Figure 9.29. R_{bias} remains connected directly across the CMOS inverter. R1 is inserted in series from the gate output to the feedback network. The benefits are multiple:

- The edge speed of modern CMOS gates is extremely fast, so R1C1 eliminates the possibility of this fast edge exciting overtone operation.
- Phase shift into the crystal can be intentionally designed by R1 and C1 value selection.
- Drive from the inverter output is reduced into the crystal, possibly saving it from being damaged.
- Loop gain can now be controlled for best waveform and startup characteristics.

Design values are strongly dependent on the actual crystal frequency needed. C2 is selected first, to provide around 60 degrees of phase shift working against the crystal equivalent series resistance (usually a few tens of ohms, but the value needs to be verified!). The time constant R1 × C1 is usefully chosen to be the reciprocal of the crystal radian frequency ($\frac{1}{2}pf_{XTAL}$). Higher values of R1 reduce the loop gain and provide better protection for the crystal, until the loop gain gets too small and oscillation stops.

9.5 Oscillators at UHF and Above

The traditional way to make signals at higher frequencies is to make a signal at a lower frequency (where oscillators are easier) and multiply it up to the wanted range. Frequency multiplication is still one of the easiest and best ways of making a clean UHF/microwave signal. The design of a multiplier depends on whether the multiplication factor is an odd or even number. For odd multiplication, a Class-C biased amplifier can be used to create a series of harmonics; a filter selects the one wanted. For even multiplication factors, a full-wave-rectifier arrangement of diodes or other non-linear devices can be used to create a series of harmonics with strong even-order components, with a filter selecting the wanted component. At higher frequencies, diode-based passive circuits are commonly used. Oscillators using some of the LC circuits already described can, with care in construction, be used in the VHF range. At UHF, different approaches become necessary.

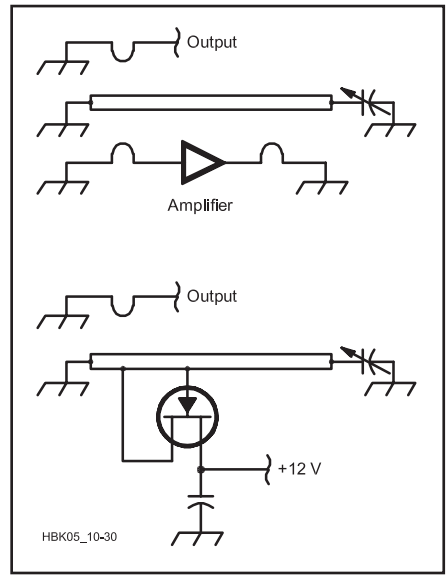

Figure 9.30 — Oscillators that use transmission-line segments as resonators. Such oscillators are more common than many of us may think, as Figure 9.31 reveals.

9.5.1 UHF Oscillators: Intentional and Accidental

The biggest change when working at UHF and microwave frequencies is that stray reactances and resistance dominate everything, making lumped circuit elements impractical. Success at these high frequencies requires making peace with the situation and developing a very good sense of where stray circuit elements reside in your layout and their approximate magnitude.

Figure 9.30 shows a pair of oscillators based on a resonant length of line, which is a distributed circuit element. The first one is a return to basics: a resonator, an amplifier and a pair of coupling loops. The amplifier can be a single bipolar or FET device or a monolithic microwave integrated circuit (MMIC) amplifier. The second circuit is really a Hartley oscillator; this one is a test oscillator for the 70 cm band made from a 10 cm length of wire suspended 10 mm over an unetched PC

board as a ground plane, bent down, and soldered at one end, with a trimmer at the other end. The FET is a BF981 dual-gate device used as a source follower.

No free-running oscillator is really stable enough to be practical on these bands. Practical oscillators in this range are almost invariably tuned with tuning diodes controlled by phase-locked-loop synthesizers, which are themselves controlled by a crystal oscillator. This transfers the same stability of the crystal oscillator to the UHF oscillator.

There is one extremely common UHF oscillator that is almost always an undesired accident, being a very common form of spurious VHF/UHF oscillation in circuitry intended to process lower-frequency signals. **Figure 9.31A** shows the circuit in its simplest form. Analyzing this circuit using a comprehensive model of the UHF transistor reveals that the emitter presents an impedance that is small, resistive, and negative to the outside world. When this negative resistance is large enough to cancel the effective series resistance of a tank placed on the emitter, the circuit will oscillate.

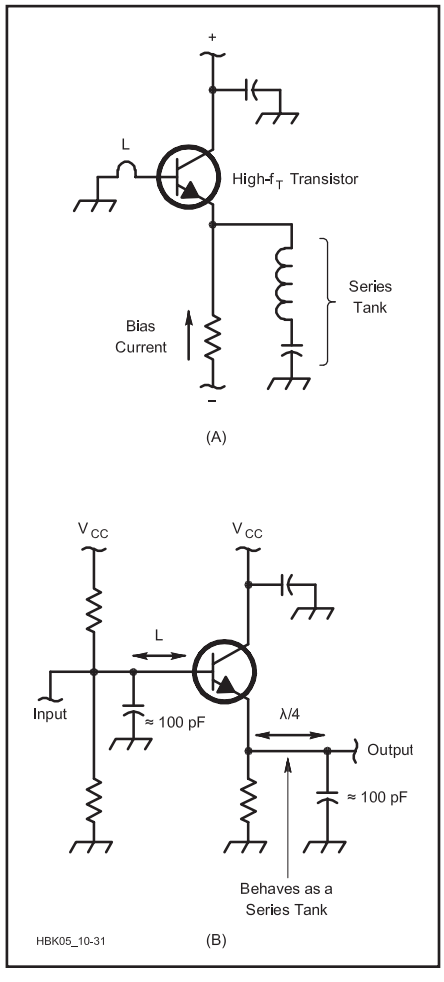

Figure 9.31 — High device gain at UHF and resonances in circuit board traces can result in spurious oscillations even in non-RF equipment.

At UHF, your schematic diagram will probably not show a tank circuit. But the high frequency transistor will know it is there, since it "sees" all of the stray reactances present in the layout. Figure 9.31B shows a very basic emitter-follower circuit with some capacitance to ground on both the input and output. If the capacitor shunting the input is a distance away from the transistor, the trace to the transistor's base looks like an inductor. The trace at the emitter of the transistor also looks like an inductor, and any nearby conductor will look like a capacitor to that trace (two conductors separated by a dielectric, which here is air and the PC board material). Any intentional capacitors add to this. If the transistor has gain at any frequency where the phasing from all of these stray reactances is right for oscillation (see the Barkhausen criteria in previous sections) then the circuit oscillates. This circuit is effectively the same as that in Figure 9.31A. This is a good reason to use the lowest-frequency transistor that you can for any application.

Stray reactances do not always have to cause headaches. Indeed, you can use them, knowing that they are there. The circuit of Figure 9.31A can be deliberately built as a useful wide-tuning oscillator covering say, 500 MHz to 1 GHz. This circuit is well-suited to construction with printed-circuit inductors. Common FR4 glass-epoxy board is lossy at these frequencies; better performance is achieved by using (the much more expensive) glass-Teflon board. If you can get surplus pieces of this type of material, it has many uses at UHF and microwave, but it is difficult to use as the adhesion between the copper and the substrate can be weak. A high-UHF transistor with a 5 GHz f_T such as the BFR90 is suitable; the base inductor can be 30 mm of 1 mm trace folded into a hairpin shape (inductance: less than 10 nH).

The upshot of this is that there is no longer any branch of electronics where RF design and layout techniques can be ignored safely. A circuit must not just be designed to do what it should do; it must also be designed so that it cannot do what it should not do.

If you have an accidental oscillator, there are three ways of taming such a circuit: adding a small resistor of perhaps 50 to 100 Ω in the collector lead close to the transistor, adding a similar resistor in the base lead, or by fitting a ferrite bead over the base lead under the transistor. Extra resistors can disturb dc conditions, depending on the circuit and its operating currents. Ferrite beads have the advantage that they can be easily added to existing equipment and have no effect at dc and low frequencies. Beware, however, that there are some electrically-conductive ferrite materials that can short transistor leads. If an HF oscillator uses beads to prevent any risk of spurs (such as shown in Figure 9.15), the beads should be anchored with a spot of adhesive to prevent movement which can cause small frequency shifts. Ferrite beads of Fair-Rite #43 material are especially suitable for this purpose; they are specified in terms of impedance, not inductance. Ferrites at frequencies above their inductive range become very lossy and can make a lead present a few tens of ohms of resistance.

9.5.2 Microwave Oscillators

Using conventional PC-board techniques with surface-mount components and extraordinarily careful layout allows the construction of circuits up to 4 GHz or so. Above this, commercial techniques and fancier materials become necessary unless we take the step to the ultimate distributed circuit — the cavity resonator.

Older than stripline techniques and far more amenable to home construction, cavity-based oscillators can give the highest possible performance at microwave frequencies. Air is a very low-loss dielectric with a dielectric constant of 1, so it gives high Q and does not force excessive miniaturization. **Figure 9.32** shows a series of structures used by G. R. Jessop, G6JP, to illustrate the evolution of a cavity from a tank made of lumped components. All cavities have a number of different modes of resonance, the orientation of the currents and fields are shown in **Figure 9.33**. The cavity can take different shapes, but the one shown here has proven to suppress

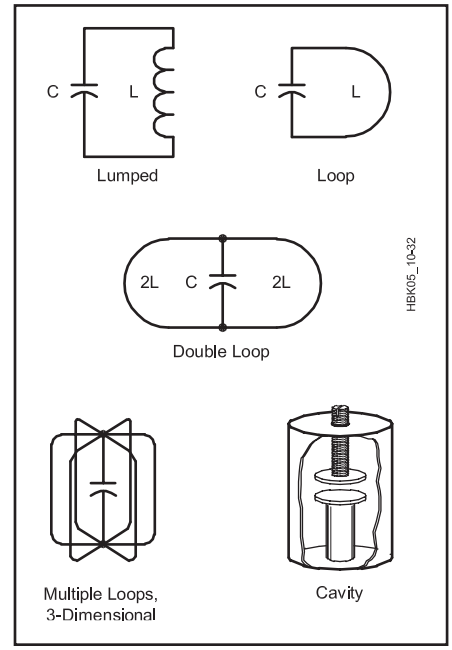

Figure 9.32 — Evolution of the cavity resonator.

unwanted modes well. The gap need not be central and is often right at the top. A screw can be fitted through the top, protruding into the gap, to adjust the frequency.

SEMICONDUCTOR CAVITY OSCILLATORS

To make an oscillator using a cavity, an amplifier is needed. Gunn and tunnel diodes have regions in their characteristics where their current falls with increasing bias voltage. This is negative resistance. If such a device is mounted in a loop in a cavity and bias is applied, the negative resistance can more than cancel the effective loss resistance of the cavity, causing oscillation. These diodes are capable of operating at extremely high frequencies and were discovered long before transistors were developed that had any gain at microwave frequencies.

A Gunn-diode cavity oscillator is the basis of many of the Doppler radar modules used to detect traffic or intruders and of the Gunnplexer 10 GHz transceiver modules used by amateurs. **Figure 9.34** shows a common configuration. The coupling loop and coax output connector could be replaced with a simple aperture to couple into waveguide or a mixer cavity. **Figure 9.35** shows a transistor cavity oscillator version using a modern microwave transistor, which can be either a FET or bipolar device. The two coupling loops are electrically completed by the capacitance of the feedthrough capacitors.

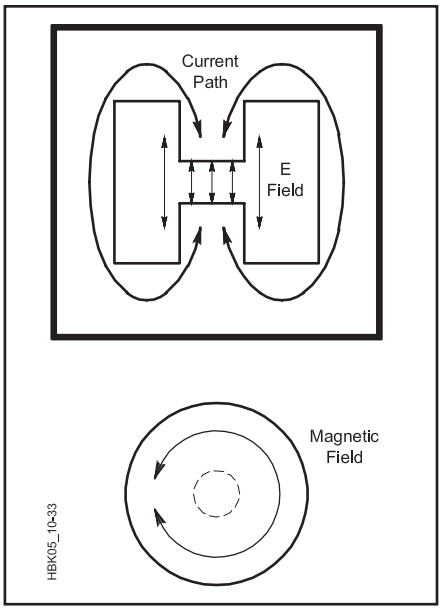

Figure 9.33 — Currents and fields in a cavity.

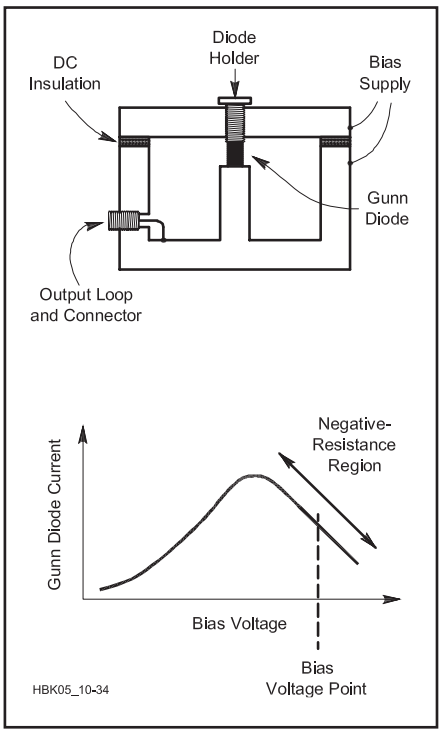

Figure 9.34 — A Gunn diode oscillator uses negative resistance and a cavity resonator to produce radio energy.

DIELECTRIC-RESONATOR OSCILLATORS (DRO)

The dielectric-resonator oscillator (DRO) is a very common microwave oscillator, as it is used in the downconverter of satellite TV receivers. The dielectric resonator itself is a ceramic cylinder, like a miniature hockey puck, several millimeters in diameter. The ceramic has a very high dielectric constant, so the surface (where ceramic meets air in an abrupt dielectric mismatch) reflects electromagnetic waves and makes the ceramic body act as a resonant cavity. It is mounted on a substrate and coupled to the active device of the oscillator by a stripline that runs past it. At 10 GHz, a FET made of gallium arsenide (GaAsFET), rather than silicon, is normally used. The dielectric resonator elements are made at frequencies appropriate to mass applications like satellite TV. The setup charge to manufacture small quantities at special frequencies is likely to be prohibitive for the foreseeable future.

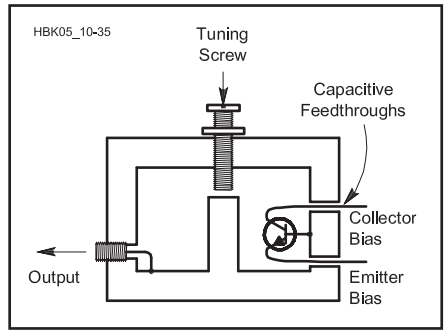

Figure 9.35 — A transistor can also directly excite a cavity resonator.

The challenge with these devices is to devise new ways of using oscillators on industry standard frequencies. Their chief attraction is their low cost in large quantities and compatibility with microwave stripline (microstrip) techniques. Frequency stability and Q are competitive with good cavities, but are inferior to that achievable with a crystal oscillator and chain of frequency multipliers.

YIG TUNED OSCILLATORS (YTO)

The yttrium-iron garnet (YIG) oscillator is a fundamental microwave source with a very wide tuning range and a linear tuning characteristic. Many YIG oscillators can be tuned over more than an octave and some tune more than five octaves. They are complete units that appear as heavy blocks of metal with low-frequency connections for power supplies and tuning, and an SMA connector for the RF output. The manufacturer's label usually states the tuning range and often the power supply voltages. This is very helpful because, with new units being very expensive, it is important to be able to identify the characteristics of surplus units. The majority of YTOs operate within the 2 to 18 GHz region, although units down to 500 MHz and up to 40 GHz are occasionally found. At microwave frequencies there is no octave-tunable device that can equal their signal cleanliness and stability. A typical stability for a 3 GHz YTO is less than 1 kHz per second drift. This may seem poor — until we realize that this is 0.33 ppm per second. Still, any YIG application involving narrow-band modulation will require some form of frequency stabilizer.

Nearly all surplus RF spectrum analyzers use YTOs as their first LO. For example, a 0 to 1500 MHz analyzer usually uses a 2 to 3.5 GHz YTO with a 2.0 GHz first IF. To understand the YTO tuning circuits, should there be need for troubleshooting and repair, a basic understanding of YIG oscillators definitely helps.

Figure 9.36 shows the construction of a YIG oscillator. A YTO is based on a YIG sphere that is carefully oriented within a coupling loop. This resonator is connected to a negative-resistance device and the whole assembly is placed between the poles of an electromagnet. Negative-resistance (Gunn) diodes were originally used in these oscillators, but transistor circuits are now essentially universal and use much less power. The support for the YIG sphere often contains a controlled heater to reduce temperature variation. YIG spheres are resonant at a frequency controlled not only by their physical dimensions, but also by any magnetic field around them. Hence the electromagnet: by varying the current through the electromagnet's windings, a controlled variable magnetic field is applied across the YIG sphere. This tunes the oscil-

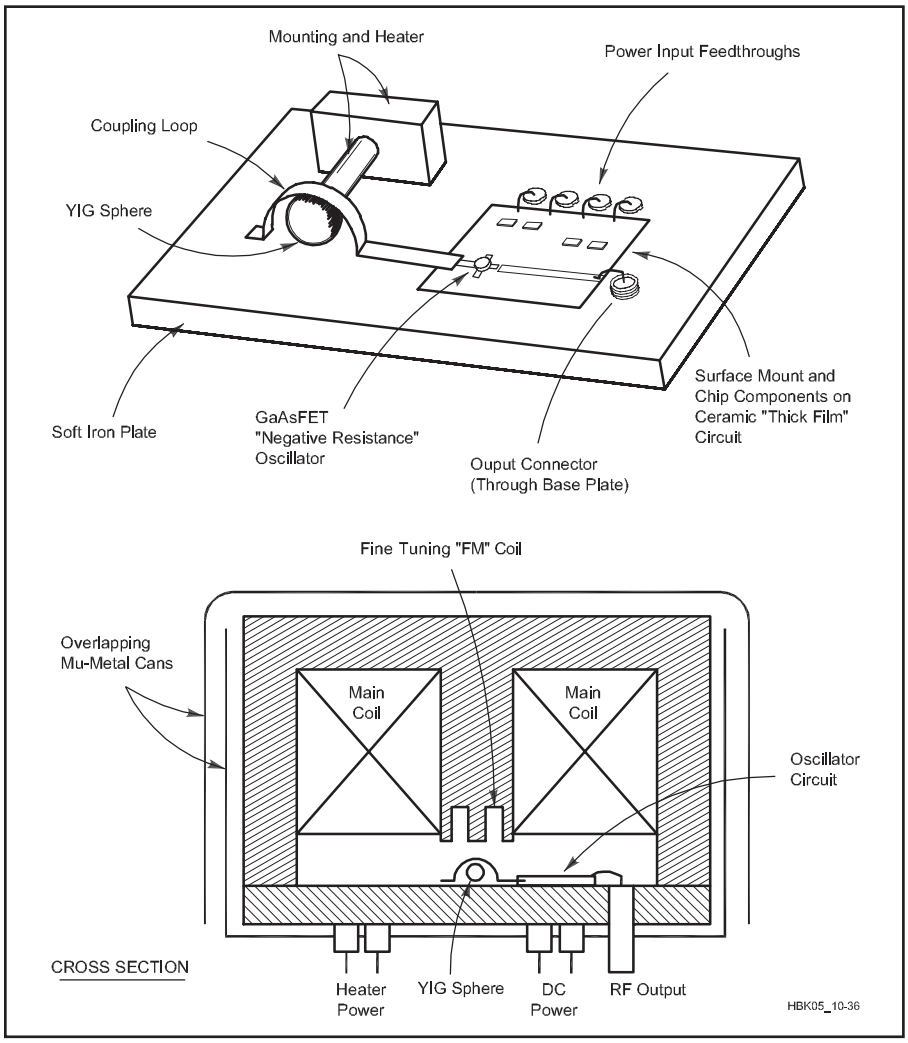

Figure 9.36 — A yttrium-iron-garnet (YIG) sphere serves as the resonator in the sweep oscillators used in many spectrum analyzers.

lator across a very wide range. The frequency/current relationship can have excellent linearity, typically around 20 MHz/mA. A YIG oscillator can be controlled by a PLL with the appropriate interface circuitry.

Magnetically-tuned oscillators bring some unique problems with them. The first problem is that magnetic fields, especially at low frequencies, are extremely difficult to shield. Therefore, YTO tuning is influenced by any local magnetic fields, causing frequency modulation. The YTO's magnetic core must be carefully designed to be all-enclosing in an attempt at self-shielding, and then one or more nested mu-metal cans are fitted around everything. It is still important to place the YTO away from obvious sources of magnetic fields, like power transformers. Cooling fans are also sources of fluctuating magnetic fields, with some fans generating fields 20 dB higher than from a well-designed 200 W 50/60-Hz transformer.

The second new problem is that the YTO's internal tuning coil needs significant current from the power supply to create the strong tuning field. This can be eased by adding a permanent magnet as a fixed "bias" field, but the bias shifts as the magnet ages. The main tuning coil still has many turns, and therefore high inductance (often more than 0.1 henrys). Large inductances require a high supply voltage for rapid tuning, with correspondingly high power consumption. The usual compromise is to have dual coils: One with many turns for slow tuning over a wide range, and a second coil with far fewer turns for fast tuning or FM over a limited range. This "FM coil" has a sensitivity around 1% to 2% of the main coil, perhaps 500 kHz/mA.

9.5.3 Klystrons, Magnetrons, and Traveling Wave Tubes

There are a number of thermionic (vacuum-tube) devices that are widely used as amplifiers or fundamental oscillators at microwave frequencies. Standard vacuum tubes (see the **RF Power Amplifiers** chapter for an introduction to vacuum tubes) work well for frequencies up to hundreds of megahertz. At frequencies higher than this, the amount of time that the electrons take to move between the cathode and the plate becomes a limiting factor. There are several special tubes designed to work at microwave frequencies, usually providing more power than can be obtained from solid-state devices.

Two of the following tubes (klystrons and traveling wave tubes) use the principle of *velocity-modulation* to extract RF energy from an electron beam. The general principles of velocity modulation and basic properties of devices using it are presented in the on-line tutorial **www.radartutorial.eu/druck/Book5.pdf**. Additional resources are described below.

THE KLYSTRON

The klystron tube uses the principle of velocity modulation of the electrons to avoid transit time limitations. The beam of electrons travels down a metal drift tube that has interaction gaps along its sides. RF voltages are applied to the gaps and the electric fields that they generate accelerate or decelerate the passing electrons. The relative positions of the electrons shift due to their changing velocities, causing the electron density of the beam to vary. This variation of electron beam density is used to perform amplification or oscillation.

Klystron tubes can be relatively large, and they can easily provide hundreds of watts to hundreds of kilowatts of microwave power. These power levels are useful for UHF broadcasting and particle accelerators, for example. Unfortunately, klystrons have relatively narrow bandwidths, and may not be re-tunable for operation on amateur frequencies. A video titled "How a Klystron Tube Works" can be found online on YouTube.com and a detailed history and tutorial is available at **www.slac.stanford.edu/cgi-bin/getdoc/slac-pub-7731.pdf**.

THE MAGNETRON

The magnetron tube is an efficient oscillator for microwave frequencies. Magnetrons are most commonly found in microwave ovens and high-powered radar equipment. The anode of a magnetron is made up of a number of coupled resonant cavities that surround the cathode. The applied magnetic field causes the electrons to rotate around the cathode and the energy that they give off as they approach the anode adds to the RF electric field. The RF power is obtained from the anode through a vacuum window.

Magnetrons are self-oscillating with the frequency determined by the construction of

their anodes; however, they can be tuned by coupling either inductance or capacitance to the resonant anode. The range of frequencies depends on how fast the tuning must be accomplished. The tube may be tuned slowly over a range of approximately 10% of the center frequency. If faster tuning is necessary, such as is required for frequency modulation, the range decreases to about 5%.

Excellent drawings showing how magnetrons work are available at **hyperphysics.phy-astr.gsu.edu/hbase/waves/magnetron.html** and a thorough introduction for the interested reader is can be downloaded from **www.cpii.com/docs/related/2/Mag%20tech%20art.pdf**.

THE TRAVELING WAVE TUBE

A third type of tube operating in the microwave range is the traveling wave tube (TWT). For wide-band amplifiers in the microwave range this is the tube of choice. Either permanent magnets or electromagnets are used to focus the beam of electrons that travels through the TWT. This electron beam passes through a helical slow-wave structure, in which electrons are accelerated or decelerated, providing density modulation due to the applied RF signal, similar to that in the klystron. The modulated electron beam induces voltages in the helix that provides an amplified output signal whose gain is proportional to the length of the slow-wave structure. After the RF energy is extracted from the electron beam by the helix, the electrons are collected and recycled to the cathode. Traveling wave tubes can often be operated outside their designed frequencies by carefully optimizing the beam voltage.

9.6 Frequency Synthesizers

Like many of our modern technologies, the origins of frequency synthesis can be traced back to WWII. The driving force was the desire for stable, rapidly switchable, and accurate frequency control technology to meet the demands of narrow-band, frequency-agile HF communications systems without resorting to large banks of switched crystals. Early synthesizers were cumbersome and expensive, and therefore their use was limited to only the most sophisticated communications systems. With the help of the same technologies that have taken computers from room-sized to now fitting into the palms of our hands, frequency synthesis techniques have become one of the most enabling technologies in modern communications equipment.

Every communications device manufactured today, be it a handheld transceiver, cell phone, pager, AM/FM entertainment radio, scanner, television, HF communications equipment, or test equipment, contains a frequency synthesizer. Synthesis is the technology that allows an easy interface with both computers and microprocessor controllers. It provides amateurs with many desirable features, such as the feel of an analog knob with precision 10-Hz frequency increments, wide-band accuracy and stability determined by a single precision crystal oscillator, frequency memories, and continuously variable precision frequency splits. Now reduced in size to small integrated circuits, frequency synthesizers have long replaced the cumbersome chains of frequency multipliers and filters in VHF, UHF, and microwave equipment, giving rise to the highly portable communications devices we use today. Frequency synthesis has also had a major impact in lowering the cost of modern equipment, particularly by reducing manufacturing complexity.

Frequency synthesizers are categorized in two general types: *direct synthesizers*, where the output signal is the result of operations directly performed on the input signal; and *indirect synthesizers*, where selected characteristics of the input signal are transferred onto a separately generated output signal. One defining feature of any direct synthesizer is that no feedback is used, so there is never any dynamic stability problem. Indirect synthesizers always include feedback control loops, so dynamic stability is a major design concern.

Direct synthesizers include the major techniques of *direct analog synthesis* (DAS) and *direct digital synthesis* (DDS). Direct analog synthesizers consist primarily of frequency multipliers, frequency dividers, mixers, and filters. DAS is very useful for generating small numbers of signals, with widely spaced frequencies, from a reference oscillator. DAS is particularly useful when more than one output signal is required at the same time.

Direct digital synthesizers are essentially dedicated microprocessors that have one function, to create an output signal waveform given a desired frequency (in digital form) using the applied reference signal. Unlike DAS, the output frequencies from a DDS can easily be separated by millihertz (0.001 Hz) while keeping all of the stability of the reference oscillator. Both being direct techniques, neither DAS nor DDS use feedback control loops, so switching from one frequency to another happens in nanoseconds.

Indirect synthesizers include the major techniques of frequency-locked loops (FLL) and phase-locked loops (PLL). The whole idea behind any indirect synthesizer is to transfer characteristics of one signal onto another separate signal. An FLL transfers only frequency characteristics of the input signal onto the output signal. Major FLL applications used by radio amateurs include *automatic frequency control* (AFC) loops and tone decoders. PLLs are more precise because not only frequency characteristics of the reference oscillator, but also its phase characteristics, are transferred to the oscillator generating the output signal.

There are two main application classes in which PLLs find wide use. The first is as a *frequency generator*, where we call it a *frequency synthesizer*. In transmitters this synthesizer may also include modulation, particularly for FM and FSK signals. The second PLL application class is as an angle (frequency or phase) demodulator, where we call it a *tracking demodulator* or *synchronizer*. Tracking demodulators used in deep space communication and clock recovery used in digital communication are major applications today.

It is curious to note that using modern digital circuitry, it actually is easier to make a PLL than an FLL. This turns out to be fortuitous. This section will focus on the PLL used as a frequency synthesizer. For readers interested in PLL use as a demodulator or as a synchronizer, there are many textbooks available that discuss these applications in great detail.

9.6.1 Digital Frequency Synthesis

A digital process is defined to be quantized in both value and time. (See the **DSP and SDR Fundamentals** chapter for background on digital signals and systems.) This is important to remember because there are sampled signal processes that are not digital. One is the sample-and-hold, which is discrete in time but analog in value. The opposite is the phase-frequency detector used in phase-locked loops, which is discrete in value but analog in time. Digital frequency synthesis is a true digital process.

The origins of digital frequency synthesis go back to 1972, even before the widespread use of microprocessors. In essence, these techniques directly calculate the desired out-

put waveform, including any modulations. When no feedback is used these are called Direct Digital Synthesis (DDS), and sometimes you also see the name Direct Digital Frequency Synthesis (DDFS) used. The DDS name is older, though DDFS is more specific. Both names survive, but here we will use the original, DDS.

DDS has many attractive properties and two major drawbacks. Its attractions include very fine frequency resolution (microhertz is practical), essentially instant switching speed, no frequency drift, and absolute frequency set-ability. If modulation is included, this also happens with digital precision. Direct control of frequency, phase shift, and amplitude are available and independent of each other. For frequency and phase shift, this is new. Analog oscillators do not have the ability to independently address frequency and phase.

DDS drawbacks are also significant. Being a digital process, DDS is a sampled data system and so must adhere to the Nyquist criterion for waveform reconstruction as explained in the **DSP and SDR Fundamentals** chapter. This effectively means that the output frequency from a DDS is restricted to being less than half of the DDS clock frequency. The practical upper limit is around 35 – 40% of the clock frequency. Interestingly, there is no effective lower limit. Any DDS can tune all the way down to zero hertz (dc) since 0 is a legitimate tuning command.

The second drawback is more operationally serious. We demand a high-quality sine wave from our oscillators and DDS is no exception. But as a sampled data system, the DDS output is a sampled sine wave. Sampling a sine wave without error requires infinite resolution due to the properties of the sinusoid. We do not have infinite resolution available to our designs, so some errors creep in as the signal is constructed. These errors manifest themselves as non-harmonic spurious output signals, a.k.a. "spurs," which can be a huge problem. Much DDS design spends a lot of time dealing with spurs.

DDS PRINCIPLES

DDS is based on the simple counter structure of an accumulator, as shown in **Figure 9.37**. We choose to interpret the digital outputs from this counter as signal phase, and commonly represent the states as points around a circle, which means that one cycle through the accumulator states represents one output cycle from this synthesizer. The figure shows an example of the 16 digital states from a 4-bit accumulator. Action of this accumulator when its input number M is set to 3 is shown in (C), showing how the "wraparound" from one cycle to the next continues when the adder overflows. It is common to have the state progression from one output cycle to the next be a different state sequence. What is vitally important is that the jump size from one state to the next be exactly the same. This sets the phase change to be:

$$\frac{\Delta\theta}{\Lambda t} = \frac{\Delta \text{ PHASE}}{1 \text{ CLOCK CYCLE}} = M$$

Frequency is simply how fast phase changes with time, so this accumulator input M is our direct tuning control. For an N-bit accumulator then the DDS frequency control is given by:

$$f_{DDS} = \frac{f_{CLOCK}}{2^N} \times M$$

Note that this relationship is perfectly linear. Also note that f_{DDS} is quantized, because the tuning number M must be an integer. There are some frequencies that a DDS cannot generate. But 0 Hz at M = 0 is legitimate.

One important property of DDS, which makes it particularly useful, is that N can be a reasonably large number. As examples, for N = 24 the denominator 2^N = 16,777,216. N = 32 is common, and then 2^N = 4,294,967,296. This means that the DDS frequency resolution, if a 100 MHz clock frequency is used, is 5.96 Hz when N = 24, and 0.023 Hz when N = 32. This is extremely fine frequency resolution and it is as accurate as the clock frequency. The only way that the DDS output frequency can drift is if the clock frequency drifts. Because everything else is digital, that cannot drift.

However, we need to have a sine wave output from our synthesizer. The full DDS block diagram is shown in **Figure 9.38**, where a new block called the waveform map is included. This is a digital process, usually implemented as a non-volatile memory, that converts the phase information from the accumulator into sinusoid waveform values which are then converted to a sampled waveform in a DAC. The output low-pass filter is critical to getting the desired sine wave.

DDS NOISE

DDS is a frequency divider and by the rules of frequency division the output phase noise is lower than the input phase noise, which in

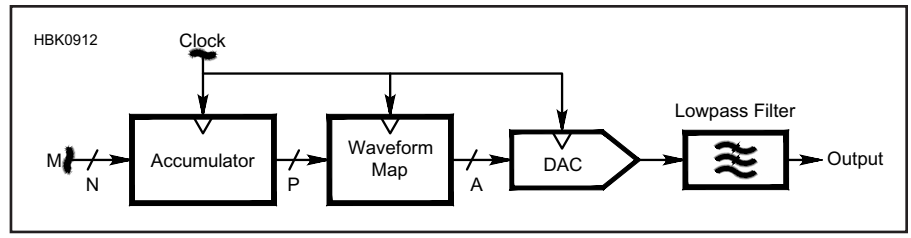

Figure 9.38 — Complete block diagram of a DDS.

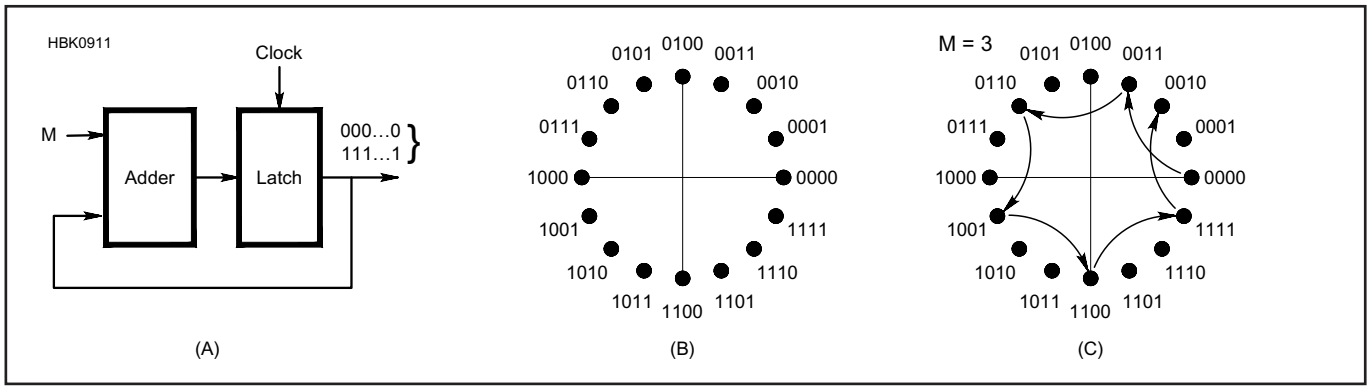

Figure 9.37 — Direct digital synthesis is based on the accumulator circuit (A). For our purposes, we assign all of the possible states that the accumulator can have around a circle and call them "phase" (B). The accumulator continuously adds its input number, once each clock cycle, to the present count. When the adder overflows, the overflow is dropped and counting continues around the circle as in (C).

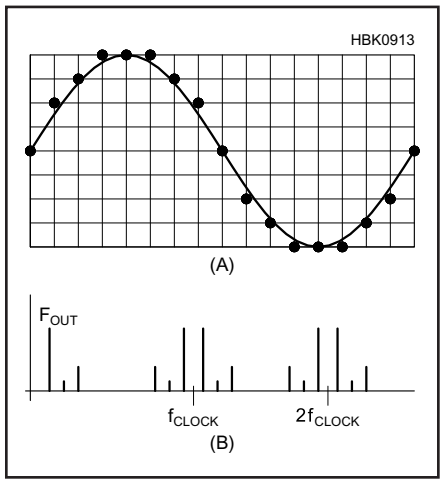

Figure 9.39 — Waveform map quantization forces output waveform errors because the output signal values must remain at the intersection of the grid lines (A). This results in output non-harmonic spurious signals from the DDS (B).

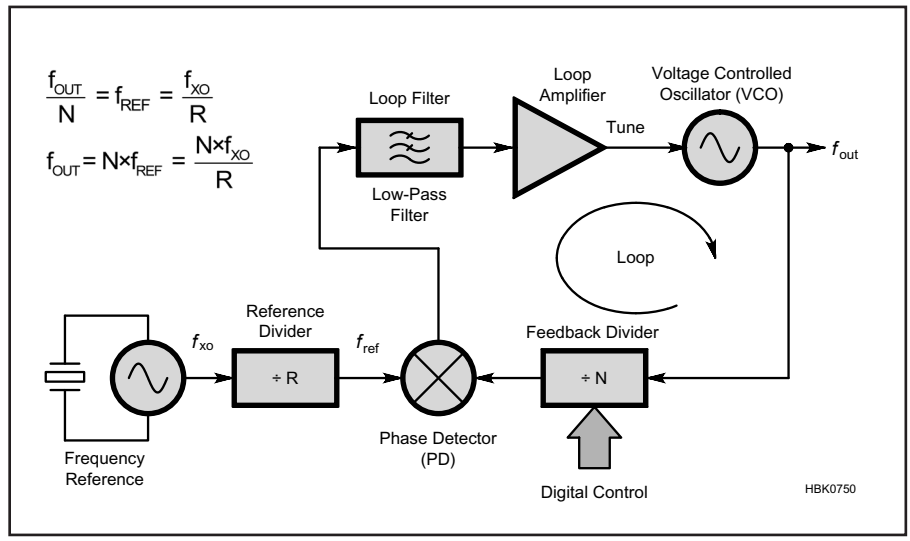

Figure 9.40 — A basic phase-locked-loop (PLL) synthesizer acts to keep the divided-down signal from its voltage-controlled oscillator (VCO) phase-locked to the divided-down signal from its reference oscillator. Fine tuning steps are therefore possible without the complication of direct synthesis.

any DDS comes from the clock. There is no VCO to add noise to the output signal, like we have in PLLs.

JITTER AND WANDER

Clock signals also exhibit *jitter* which is an irregular deviation from the exact clock period that creates noise in the DDS output. *Jitter period* is the interval between the maximum and minimum deviation from the nominal clock period. *Jitter frequency* is the inverse of jitter period. Jitter frequencies below 10 Hz are classified as *wander* and frequencies at or above 10 Hz as *jitter*.

ADDITIONAL DDS PRINCIPLES

In addition to very fine frequency resolution, extremely good tuning linearity, very low output phase noise, and tuning bandwidth that goes all the way down to dc, an additional important benefit of DDS is that it tunes very quickly. Indeed, as soon as the tuning number M is changed, the DDS immediately shifts to calculating the new output waveform. The tuning dynamics of PLLs are not present here. Likewise, the stability issues of a PLL do not happen in DDS designs.

The presence of non-harmonic spurious signals remains the biggest issue in successful application of DDS to amateur equipment. In principle, a 32-bit DDS can cover all HF amateur bands with 0.23 Hz resolution when clocked at 100 MHz. No VCOs, no band-switching, no tuning — it sounds like a dream! — but we need to apply our RF design skills to get rid of those pesky spurs. The quantization level of modern waveform maps is sufficient to keep the sinusoid quantization noise spurs below –100 dBc. The largest remaining problem is the DAC.

The signal sampling theorem does not predict spurs other than those generated by quantization. That means any additional spurs are a result of the actual output wave form not being exactly what the signal digital quantization requires. (See **Figure 9.39**.) Normal DAC specifications of *integral nonlinearity* (INL) and *differential nonlinearity* (DNL) turn out to not be very important in the DDS application. More important are the symmetry of output transition times for up and down steps, and having no overshoot on either transition direction. DACs with the latter two properties generally work very well for precision DDS use.

9.6.2 Phase-Locked Loops (PLL)

As mentioned in the introduction, the PLL is a key component of any indirect synthesizer. This section presents a detailed discussion of design and performance topics of PLL synthesizers and the circuits used to construct them.

ARCHITECTURE

The principle of the PLL synthesizer is straightforward. A tunable oscillator is first built to cover the required output frequency range; then the PLL system is constructed around this oscillator to keep it tuned to the desired frequency. This is done by continuously comparing the phase of the oscillator to a stable reference, such as a crystal oscillator, and using the result to steer tuning of the oscillator. If the oscillator frequency is too high, the phase of its output will start to lead that of the reference by an increasing amount. This is detected and is used to decrease the frequency of the oscillator. Too low a frequency causes an increasing phase lag, which is detected and used to increase the frequency of the oscillator. In this way, the oscillator is locked into a fixed-phase relationship with the reference, which means that their frequencies are held exactly equal.

A representative block diagram of a PLL is shown in **Figure 9.40**. Measurement of any error in the output signal phase is made by the *phase detector* (PD). The measured error is supplied to the *loop filter* and *loop amplifier* which work together to tune the *voltage-controlled oscillator* (VCO) just enough to make the error go to zero. The phase detector determines the actual frequency and phase of the VCO through the *feedback divider* (N). The phase detector operates at a lower frequency which is a sub-multiple of both the crystal oscillator frequency reference (f_{XO}) and the output frequency (f_{OUT}). This lower frequency is correctly called the PLL *reference frequency* (f_{REF}) because it sets the operating conditions of the PLL.

There is unfortunate history about the term reference frequency. It is often used to refer to both the operating frequency of the phase detector and to the frequency reference applied to the PLL. If the *reference divider* (R) is not present then of course these two frequencies are the same. But in general, the reference divider is present and these frequencies are different. Ambiguity here is widespread and a big problem, even today, many decades into

the widespread use of PLLs. We must be very careful in the words we choose to use.

FREQUENCY RESOLUTION (STEP SIZE)

Frequency of the output signal can be easily changed by changing the divide ratio of the feedback divider. For example, if the reference frequency (at the phase detector) is 100 kHz and the output frequency is 147.5 MHz, the feedback divider number must be N = 1475. If we changed the value of N to 1474, the frequency at the output of the feedback divider becomes 100.068 kHz. The phase detector notices that the frequency of the VCO is too high, and "tells" the loop filter to reduce the frequency of the VCO until the frequency out of the feedback divider becomes exactly 100.000 kHz. This will happen when the VCO is retuned to a frequency of 147.4 MHz.

By changing the value of N by 1, we have just tuned the frequency synthesizer by its smallest available step. This is called the *frequency resolution* of the synthesizer, or equivalently the synthesizer *step size*. Here we note that this step size is (147.5 – 147.4) = 0.1 MHz. It is no accident that this frequency is exactly equal to the PLL reference frequency f_{REF}. This result is an important reason to avoid all ambiguity in the term "reference frequency." f_{REF} is equal to the synthesizer output frequency step size.

There certainly are applications where we want the step size to be a small number, say 10 Hz. You might guess — correctly — that there are problems in building a PLL synthesizer for a very low f_{REF}. The usual solution is to combine multiple PLLs with some DAS techniques to get around these problems.

DYNAMIC STABILITY AND LOOP BANDWIDTH

Because feedback control is required to transfer the frequency and phase characteristics from the crystal oscillator onto the VCO, all of the stability problems inherent in feedback control apply to PLL design. If the gain and phase responses are not well designed, the PLL can oscillate. In this case, instead of getting the stability transfer we want, the output is essentially a frequency modulated signal spread across the entire tuning bandwidth of the synthesizer — clearly a very bad thing. How to perform this design is discussed below.

For the moment let us assume that our PLL dynamics are properly designed. One other characteristic of feedback control is now important: how fast can the output frequency be changed from one value to another? This is controlled by the *loop bandwidth* of the PLL. For practical reasons the loop bandwidth should be less than 5% of f_{REF}. Using the example above in which f_{REF} = 100 kHz, the loop bandwidth we design for cannot exceed 5 kHz. To answer our question, we can use the rule of thumb that in a well-designed loop, the settling time should be between 1 and 2 times the reciprocal of the loop bandwidth. For a 5 kHz loop bandwidth the settling time should not exceed 2/5,000 = 400 µs. We can begin to see one problem with PLL design if f_{REF} is very small: To get a 1 kHz step size from a PLL synthesizer the maximum loop bandwidth available is 50 Hz. This is impractically small.

This loop bandwidth also influences how we can modulate our PLL synthesizer in an FM transmitter. Imagine now that we apply a very small amount of FM to the reference oscillator f_{XO}. The amount of deviation will be amplified by N/R — but this is true only for low modulating frequencies. If the modulating frequency is increased, the VCO has to change frequency at a higher rate. As the modulating frequency continues to increase, eventually the VCO has to move faster than the available *settling time*. The PLL now acts as a low-pass filter, reducing the available deviation from higher modulating frequencies. The cut-off frequency of this PLL low-pass action is equal to this same loop bandwidth.

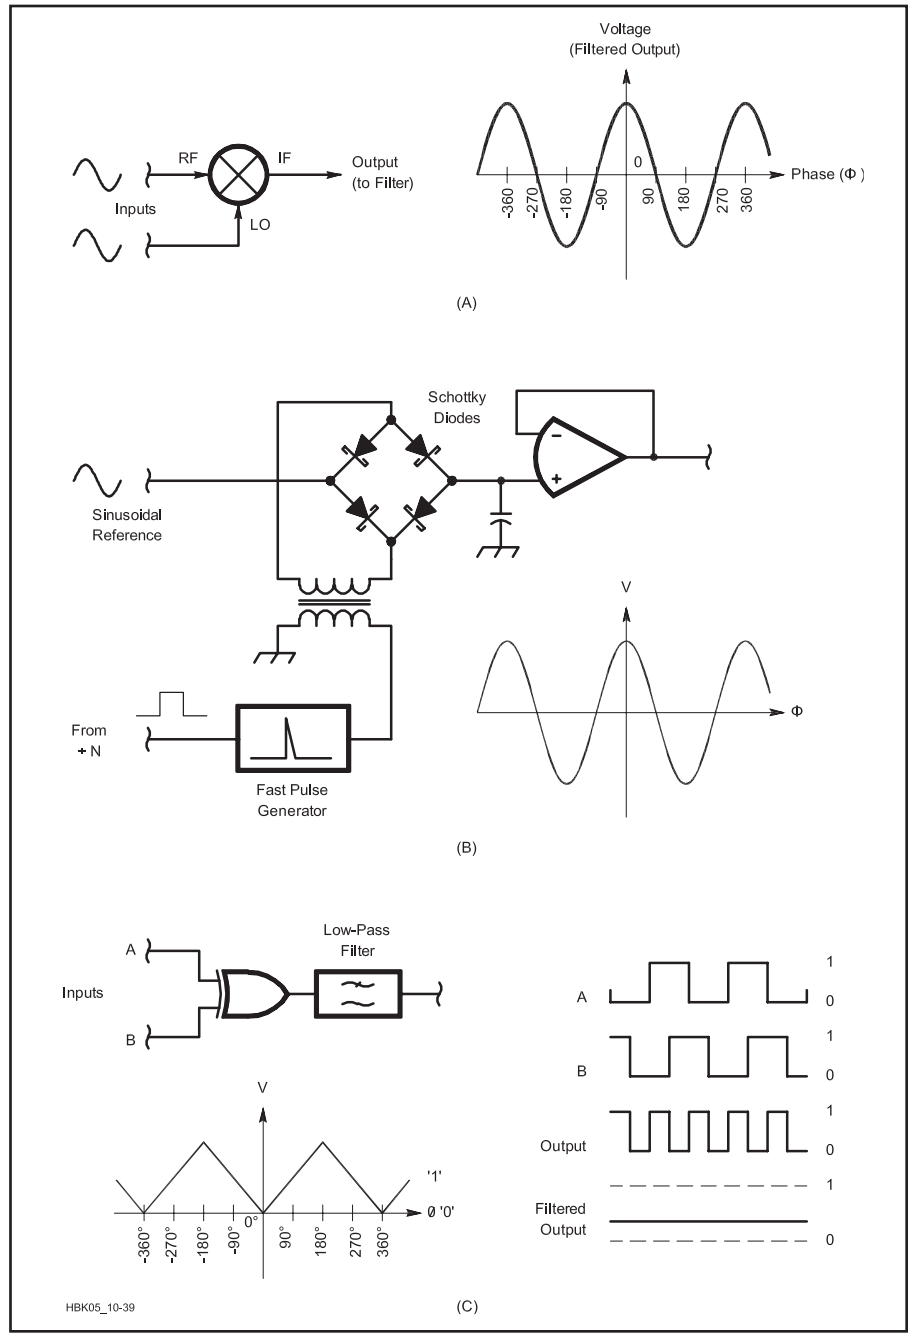

Figure 9.41 — Simple phase detectors: a mixer (A), a sampler (B) and an exclusive-OR gate (C).

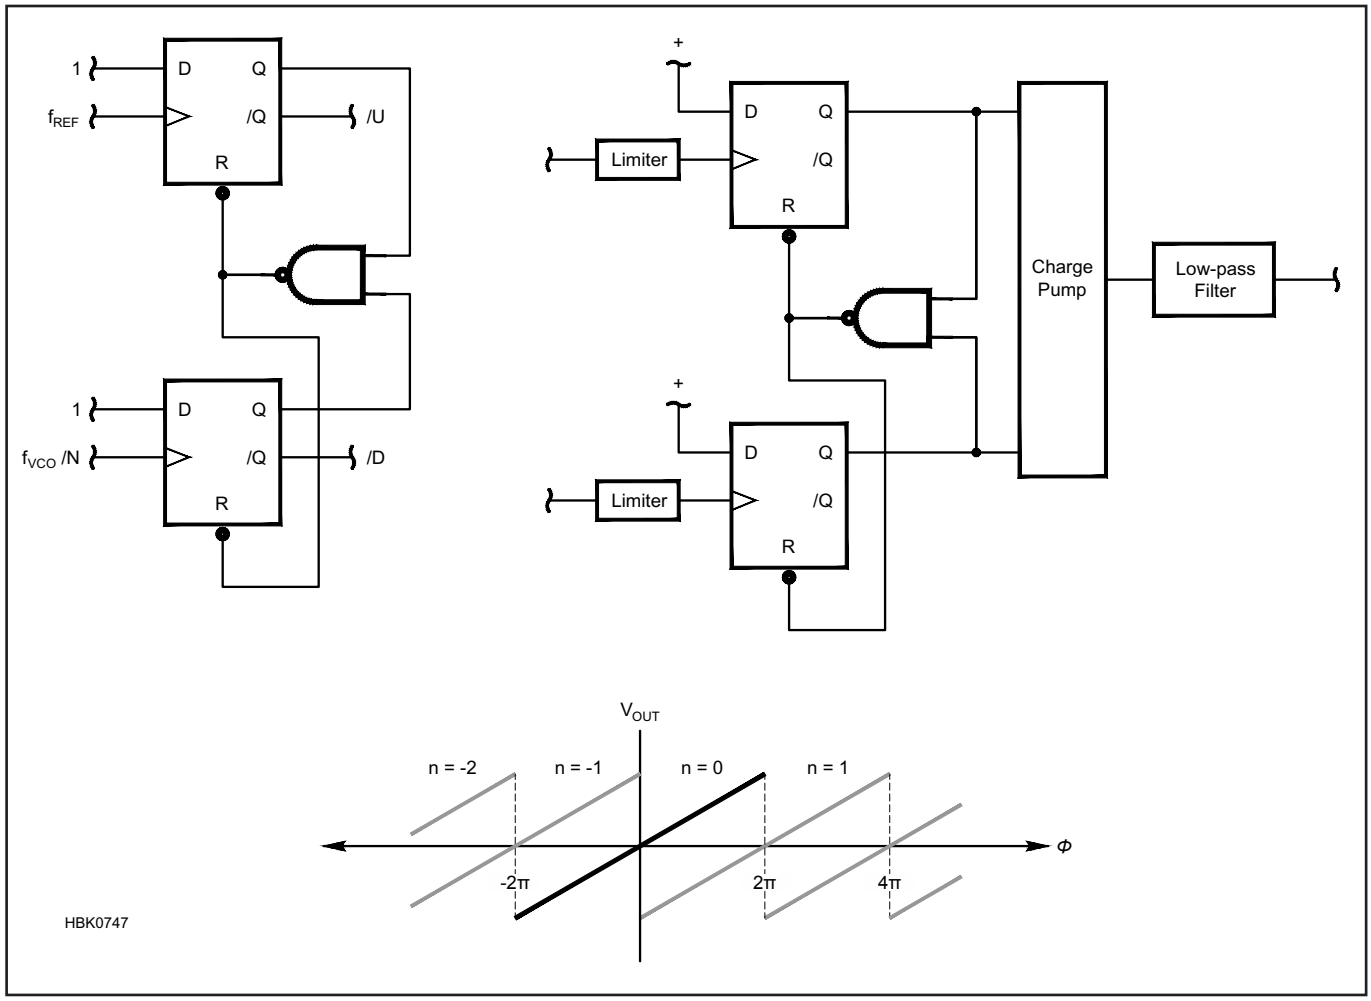

Figure 9.42 — The phase frequency detector (PFD) makes a measurement of the time difference between two rising edges of the input signals.

PLL COMPONENTS

Let us continue our discussion of phase-locked loop synthesizers by examining the role of each of the component pieces of the system. They are the phase detector, the VCO, the dividers (with possible prescalers), the loop filter, and of course the reference oscillator.

Phase Detectors

Phase detection is the key operation in any PLL. Remembering that we are making a phase locked loop, what we really need is not *absolute* phase measurement but a *relative* phase measurement. We need to know only how the phase of the VCO output signal f_{OUT} is changing with respect to the phase of the crystal reference f_{XO}.

As usual, there are both analog and digital circuit techniques available to perform phase detection. For frequency synthesizer design the analog techniques are seldom used. Analog phase detectors are used only in receiver and demodulator applications. There are two that sometimes still appear in ham equipment, the multiplier/mixer of **Figure 9.41A** and the sampling phase detector shown in Figure 9.41B.

The mixer is a very low noise device when used as a phase detector, which explains why it is not yet completely gone from synthesizer design. It remains only in PLL designs where f_{REF} is very high, greater than 10 – 100 MHz. The main reason for this is that a PLL using a mixer for the PD is subject to a phenomenon called *false lock*, where the PLL may act like it is locked even though it really isn't. Detecting and eliminating a false lock condition is much easier to do at higher frequencies. The sampling phase detector is now used only in synthesizers with output frequencies in the high microwave region, typically well above 10 GHz. The gain of this sampling circuit is very low, so it is used only if there is no viable alternative.

The most widely used digital phase detector is the exclusive-OR (XOR) gate shown in Figure 9.41C. If inputs A and B are almost in phase, the output will be low most of the time, and its average filtered value will be close to the logic 0 level. If inputs A and B are almost in phase opposition, the output will be high most of the time, and its average voltage will be close to logic 1. The average voltage is near its mid-value when the inputs are shifted in phase by 90 degrees. In this respect the XOR gate is much like the analog mixer. To achieve this circuit's full output-voltage range, it's important that the reference and VCO signals at its input have a 50% duty cycle.

Phase-Frequency Detector

One problem common to *all* phase detectors is the possibility of false lock. This happens because the output of any phase detector can have an average value of zero even when the frequencies at its inputs are not equal. This is a very serious problem that requires any PLL using them to have additional circuitry to serve as an acquisition aid. If the PLL tries to settle when the PD inputs are not at the same frequency, this acquisition aid must prod the PLL to keep moving toward the true lock frequency. Such acquisition aids are not

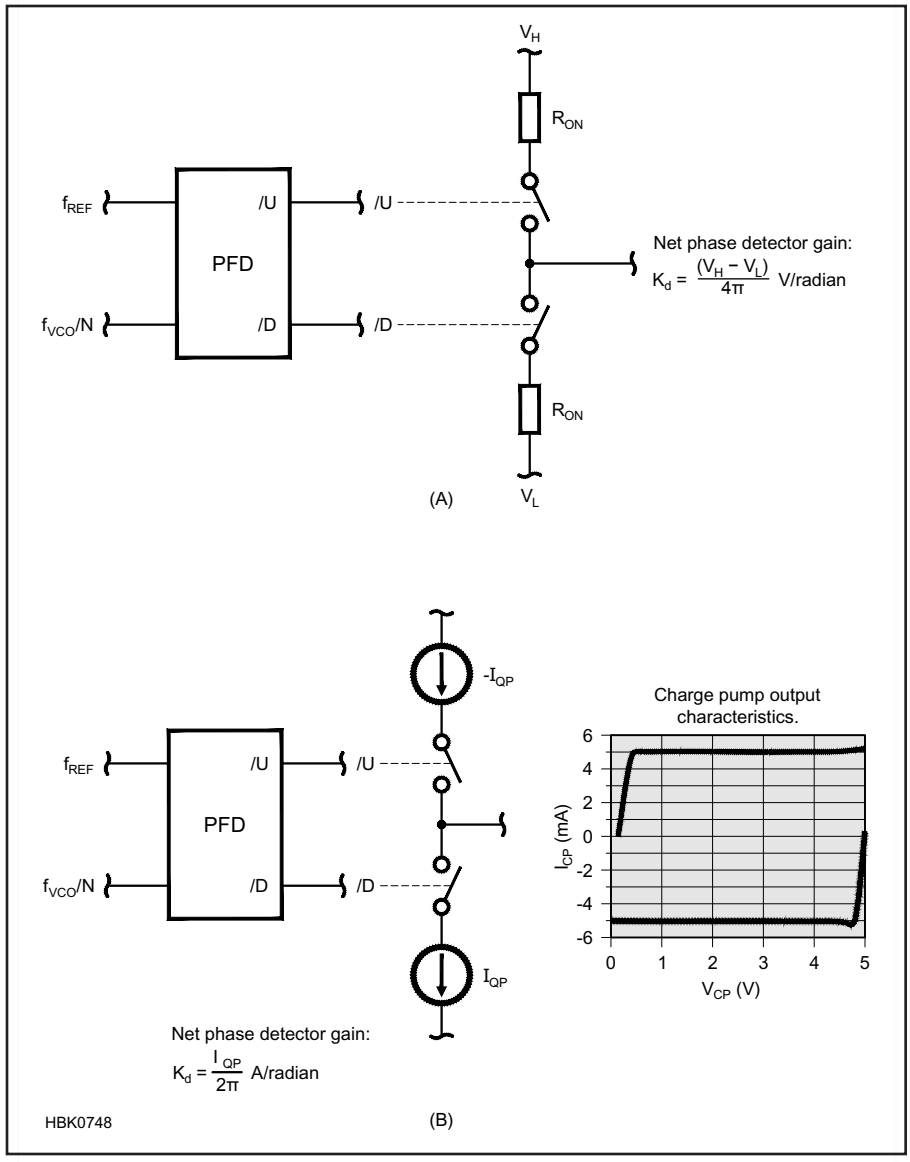

Figure 9.43 — The charge pump delivers short bursts of current to the loop filter under the control of the phase frequency detector (PFD) outputs. A voltage mode charge pump is shown at A and a current mode charge pump at B.

needed when a *phase-frequency detector* (PFD) is used.

The benefit of the PFD is that if the input signals are at different frequencies, even if they are at very different frequencies, the output is never zero. The PFD output always produces a dc shift in the direction the PLL needs to move until the lock condition is achieved. No acquisition aids are ever needed. This alone is enough to explain its nearly universal adoption today.

As seen in **Figure 9.42A**, the PFD is a simple digital circuit. This also helps make it attractive to use in this era of CMOS integrated circuits. Other designs for the PFD exist, but they all follow the logic shown in Figure 9.42A. Of particular importance is that the PFD has two outputs.

The PFD operates by measuring the time difference between the rising edges of the reference signal and the divided VCO. The first signal to arrive sets its flip-flop. The second edge to arrive also sets its flip-flop, which immediately causes both flip-flops to be reset, ready for the next set of signal edges. By operating at high speed, the PFD is very sensitive to phase differences between the input signals. The PFD does require an equal number of edges for each signal, so it is not useful when the signals have varying numbers of edges, such in a receiver's demodulator or a synchronizer.

If the VCO frequency is too low only the /U output is active, causing the loop filter to raise the VCO frequency. On the other hand, when the VCO frequency is high only the /D output is active, causing the loop filter to lower the VCO frequency. It is interesting to note that unlike the mixer and XOR phase detectors, which require the inputs to be in phase quadrature, the PFD locks when the input signals are in exact phase alignment.

Being an edge triggered circuit, PFD operation is essentially independent of input signal duty cycle. This eliminates the need for input waveform processing. If the input frequencies are different, the time differences measured by the PFD will be either constantly increasing or decreasing. As a result, only one of the PFD output signals will be active, informing the PLL unambiguously which way to correct the controlled source to achieve phase lock.

Charge Pumps

Because the PFD has two outputs and most loop filters have only one input, something is needed to bridge the PFD and the loop filter. The usual technique is the *charge pump*, shown in Figure 9.42B. The charge pump gets its name by allowing current to flow only for the brief interval, here when the /D or /U outputs are active. Current flow in short bursts (pulses) is equivalent to a finite charge transfer.

Charge pumps come in two configurations, voltage mode and current mode, as shown in **Figure 9.43**. Voltage mode (Figure 9.43A) was the original version, though now current mode (Figure 9.43B) is most widely used with integrated circuit PLLs.

The voltage mode charge pump directly connects voltage sources to the loop filter when the PFD output signals are active. When the PFD output signals are at rest (both high) both voltage sources are disconnected from the loop filter. These voltage spikes transfer charge to the integrator in accordance with the PLL loop filter time constant. One voltage source inputs current to the loop filter and the other removes it.

The current mode charge pump is architecturally similar to the voltage mode charge pump, with voltage sources replaced by current sources. Like the voltage mode charge pump, this circuit disconnects both sources when the PFD output signals are at rest. In this case, current spikes transfer charge to and from the loop filter integrator. This transfer is independent of the loop filter series resistances.

The current mode charge pump has some advantages to the frequency synthesizer designer. Three are particularly significant:

• Pump current flows at a constant value independent of the loop filter voltage. Unless V_H and V_L are both very large compared with any possible loop filter voltage, the amount of charge transfer will vary with different loop filter voltage values.

• During the PFD reset time, both output sources are active. The voltage mode charge

pump will attempt to connect both voltage sources to the loop filter during this brief period, which can generate some undesirably large power supply current spikes. During PFD reset, the current mode charge pump loads the power supply with a single current value of I_{QP}. If both current sources are well matched, their currents completely cancel at the loop filter input, effectively disconnecting them that much earlier.

• At low offset frequencies, the loop gain flattens out with a voltage mode charge pump. Loop gain continues to increase with a current mode charge pump, effectively matching active filter performance.

VCO

The design and characteristics of oscillators, including tunable oscillators, is the major topic of this chapter so it will not be repeated here. Of particular importance to PLL design are the tuning characteristic of the VCO and the VCO tuning bandwidth.

The tuning characteristic of an oscillator shows the output frequency versus the tuning voltage, such as the example of **Figure 9.44A**. The slope of this curve is called the VCO's *modulation gain* (K_0) and it is this modulation gain that is most important to PLL design. For wideband VCOs the variation of K_0 can exceed 5:1 from minimum to maximum frequency. The design method below shows how a wide gain variation such as this is handled with a single loop filter design.

It is important to realize that the VCO has no idea that it is controlled by a PLL. While it is oscillating it continues to drift and jitter with time and temperature like any other oscillator. You must characterize these variations to be assured that the tuning capability has sufficient range for the PLL to "find" a tuning voltage value that will retune the VCO to the required frequency, under all circumstances.

The primary VCO characteristic of interest to the synthesizer designer is the tuning gain "constant," K_0. Almost never a constant, K_0 is a local measure of how much the VCO output frequency changes with a change in tuning voltage.

VCO designers go to great lengths to make K_0 an actual constant. This additional effort often increases the VCO design complexity and usually results in a more expensive design. If no linearization effort is undertaken, K_0 can easily vary from 3:1 to 10:1 over the tuning band. This also complicates the frequency synthesizer design. Specific applications will dictate where effort needs to be expended: in the VCO design or the PLL design.

All PLL design algorithms assume that there is no delay between when the loop filter presents a tuning change command to the VCO and when the VCO frequency actually changes. This situation is true when the VCO tuning bandwidth is much greater than the PLL loop bandwidth. If the VCO tuning bandwidth is not much greater than the PLL loop bandwidth, there are additional phase shifts that will cause problems with PLL stability. It is usually much easier to narrow the loop bandwidth than to redesign the VCO. If redesigning the VCO is a possibility, reduce the capacitance seen at the tuning input to increase the tuning bandwidth.

The VCO tuning bandwidth is important mainly for PLL designs that have wide loop bandwidths. Measuring the frequency deviation as the modulating frequency varies is best done by looking at the FM sidebands on a spectrum analyzer, as shown in Figure 9.46B. For constant deviation, FM sidebands follow the top profile. When the modulating frequency doubles, the sidebands drop below the profile by 6 dB, and if it triples, by 10 dB. The first FM sideband must exceed all other sidebands by 10 dB for proper VCO *characterization* — the process of determining the VCO's characteristic parameters such as the tuning sensitivity, noise spectrum, dynamic response to changes in the tuning input signal, and so on. (See the Keysight application note "Boosting PLL Design Efficiency" in the References section.) VCO tuning bandwidth is determined by the point at which increasing the modulating frequency deviates from

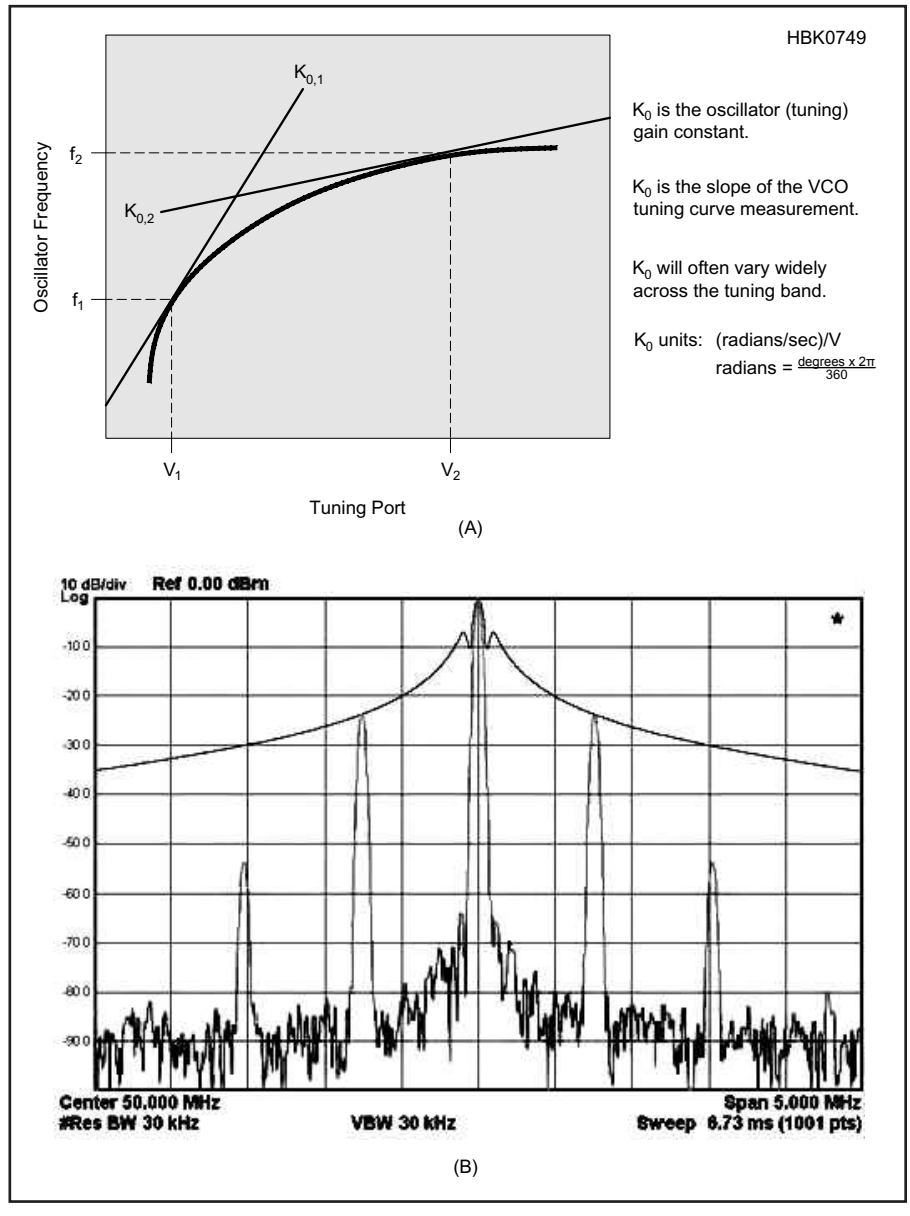

Figure 9.44 — (A) Measurement of the VCO tuning characteristic provides a curve describing frequency versus tuning voltage. The VCO tuning gain at a particular frequency (or tuning voltage) is the slope of this curve at that point. (B) A spectrum analyzer is used to measure frequency deviation of the VCO and the modulating frequency varies. FM sideband peak amplitudes follow the top profile when deviation is constant.

(drops below) the profile shown in Figure 9.44B.

Reference Divider

The reference divider is one of the easiest parts of a PLL to design, though technically it is outside the PLL. The input frequency is well known — it is the crystal reference. The output frequency is also well known, being the phase detector frequency f_{REF}. The ratio of these two frequencies is the required counter divide-by ratio or *modulus*.

When designing PLL synthesizers, f_{REF} is usually specified but f_{XO} is not. This provides some flexibility in design to use easy counter implementations if the resulting f_{XO} is also easy to get. The easiest digital counters to use for any divider are those operating with binary numbers (divide by 2, 4, 8, 16, …, 128, …). For example, if f_{REF} = 20 kHz and you choose R = 512 then f_{XO} = 10.24 MHz, which is a readily available crystal frequency.

Feedback Divider and Prescalers

The feedback divider is more of a challenge. First and foremost, this divider must work properly with whatever possible frequency it may ever see at its input. Whatever the highest and lowest frequencies the VCO may be, the divider must handle them all.

The feedback divider is almost always programmable so that the output frequency can be changed. Programmable counters are always slower than non-programmable designs, so some speed limitation occurs. Fortunately, in this era of BiCMOS integrated circuits the problems of past years in getting programmable counters to operate at hundreds of MHz are over. Older equipment still has feedback divider designs that are carefully crafted to handle the frequencies at which they must operate.

Fixed Prescaler

When the programmable counter does not have sufficient speed to handle the frequencies required, then the VCO frequency must be divided down ahead of the programmable counter to assure that everything works reliably. This additional divider or prescaler usually has a fixed binary value so it will be fast enough without using too much power. As long as the output frequency from the prescaler is always within the operating range of the programmable divider the PLL will be reliable.

Of course, there are consequences: when a prescaler is used, the PLL step size is multiplied by the prescaler divider value. For our f_{REF} = 100 kHz, if we adopt a divide-by-16 prescaler then the step size increases by a factor of 16 to 1.6 MHz. The direct way to correct for this is to reduce f_{REF} by the same factor of 16, which unfortunately means that the PLL loop bandwidth is also reduced by a factor of 16. Is there a solution with fewer compromises?

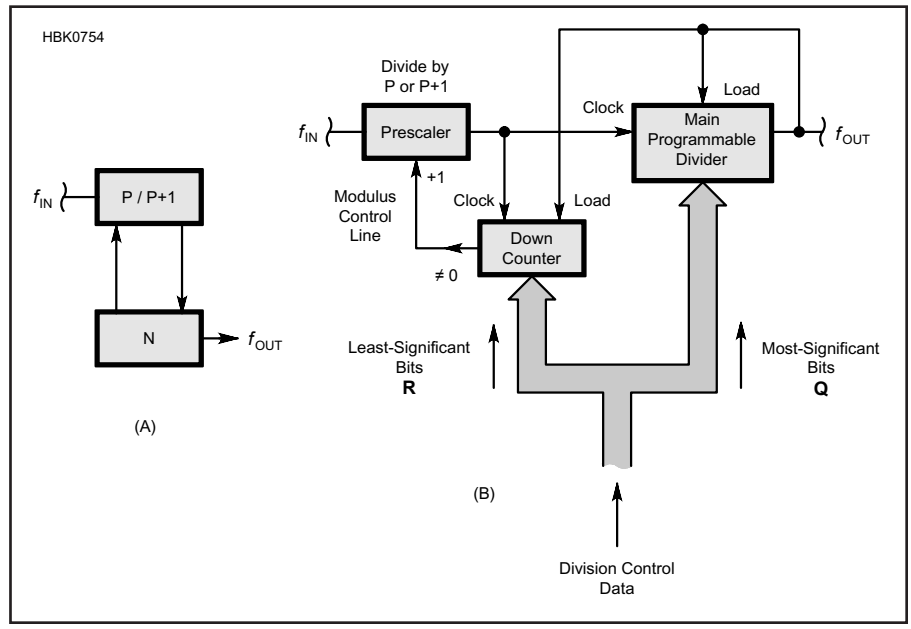

Figure 9.45 — A typical implementation of a dual-modulus prescaler.

Dual-Modulus Prescalers

Yes, there is such a solution. If the prescaler is designed to divide by two values separated by 1, say 16 and 17, or 8 and 9, then we call this a *dual-modulus P/P+1 prescaler*. Using a dual-modulus prescaler, we can leave f_{REF} unchanged, which not only keeps our step size and loop bandwidth but also gains a much higher operating frequency for the feedback divider.

A dual-modulus prescaler is a cooperative effort between a high speed prescaler and a much lower speed programmable counter. The result of this cooperation is a programmable counter that operates at a very high input frequency while maintaining unit division resolution. A dual-modulus prescaler works by allocating quotient and remainder values from the division N/P as shown in **Figure 9.45A**.

The improvement comes from viewing the operation of division in a slightly different way. The division ratio (N) of any two integers will always provide a quotient (Q) and a remainder (R). If Q counts of f_{in}/P are followed by R counts of f_{in} directly, then N = QP + R. One way to build this counter would be to first count a quotient's worth of prescaled input, and then count a remainder's amount of the input frequency directly. This last step is hard because of the high frequency.

A second approach is to add the remainder counts to the prescaled output as follows: If Q – R counts of prescaled f_{in}/P are followed by R counts of prescaled f_{in}/(P + 1), then N = (Q – R) P + R (P + 1). This is the same as N = QP – RP + RP + R = QP + R.

If the remainder counts are distributed among the quotient counts by increasing the prescaler modulus by one from P to P + 1 for the remainder's number of cycles, then returning the prescaler to its nominal division ratio while the rest of the quotient counts are made, the same result is achieved. This design always operates the programmable counter at a much lower frequency than the input frequency, which is a much more robust design.

The catch, however, is that there have to be enough quotient counts over which to distribute any amount of remainder counts. If you run out of quotient counts before remainder counts the technique falls apart and that particular loop divisor will not be *realizable*. For real dividers, all of the terms — Q, P, and R — must be positive integers, so Q – R must be greater than zero for N = QP+R to be realizable.

The maximum remainder, R_{max} = P – 1. The minimum quotient is equal to R_{max}. So the minimum remainder, R_{min} = 0. That means:

$$N_{min} = Q_{min}\,P + R_{min} = (P - 1)\,P + 0 = P2 - P$$

If the PLL design always requires N to be greater than P2 – P then there are no problems. If N does need to go below P2 – P, then the only solution is to choose a dual modulus prescaler with a smaller value of P. Various values of P and N_{min} are:

P	N_{min}
8	56
16	240
64	4032
128	16256

Dual-modulus prescalers therefore implies a minimum divisor value before continuous divider value coverage is realized. The value of this minimum divisor depends on the base modulus of the prescaler, P, and increases quadratically. Proper choice of prescaler modulus is one of the important decisions the frequency synthesizer designer has to make.

Figure 9.45B shows how a typical dual-modulus prescaler is implemented. Each cycle of the system begins with the last output pulse having loaded the frequency control word, shown as "Division Control Data," into both the main divider and the prescaler controller. If the division control data's least significant bits (which make up R) loaded into the prescaler controller are not zero, the prescaler is set to divide by P + 1, 1 greater than its normal ratio, P. The main divider counts Q pulses from the prescaler.

Each cycle out of the prescaler then clocks the down counter. Eventually, the down counter reaches zero, and two things happen: The counter is designed to hold (stop counting) at zero (and it will remain held until it is next reloaded) and the prescaler is switched back to its normal ratio, P, until the next reload.

Because the technique is widely used, dual-modulus prescaler ICs are widely used and widely available. Devices for use to a few hundred megahertz are cheap, and devices in the 2.5 GHz region are commonly available. Common prescaler IC division ratio pairs are: 8-9, 10-11, 16-17, 32-33, 64-65 and so on. Many ICs containing programmable dividers are available in versions with and without built-in prescaler controllers.

Loop Filter

When designing PLL synthesizers we usually just have to measure and accept the VCO characteristics of K_0 and the PFD output characteristics. Output frequency range and step size are also not flexible. We pull all of these together into a working and stable synthesizer by proper design of the loop filter.

Loop filter design is usually shrouded in mystery, which has given PLL synthesizer design the aura of a "black art." This is not at all warranted because very reliable design algorithms have been around for decades. Unfortunately, these have usually been published in obscure places so they are not well known. The best algorithms are presented below. These methods have served extremely well for over 30 years and should provide you with fast design times and very stable PLL synthesizers. No iteration should be needed.

Loop filters come in active and passive structures. In general, the easiest designs use passive (RC) loop filters if the output voltage range from the phase detector (including the PD or PFD and its charge pump) is sufficient to tune the VCO over its required frequency range. If the VCO needs a larger tuning voltage range then an active loop filter structure is necessary.

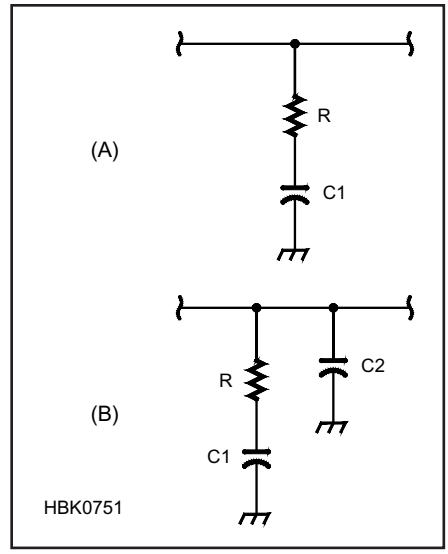

Figure 9.46 — Passive loop filters for current mode charge pump PFD outputs.

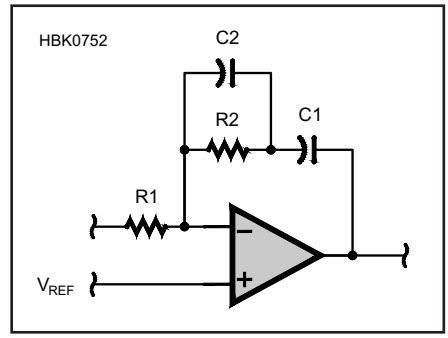

Figure 9.47 — Third-order active loop filter isolates the loop filter output voltage from the charge pump outputs.

Passive loop filter circuits for current mode charge pump outputs are shown in **Figure 9.46**. The loop filter time constant is set by the charge pump current and C1, independent of R and of the loop filter output voltage. R provides a phase leading zero for PLL stability. The first-order filter in Figure 9.46A has a minimum parts count and results in loop dynamics that are consistent with the output frequency.

A second-order filter has the undesirable characteristic of allowing step and impulse changes in the input to appear at the output. This results in significant reference sidebands — a bad thing. Figure 9.46B shows a third-order filter in which C2 changes step at the input to slower ramps by rolling off the high frequency response of the loop filter. This results in a significant reduction in reference sidebands.

Active loop filters offer more structures that work well and are easier to design, though more complicated to build. When a voltage mode charge pump is being used, an active loop filter is strongly desired since it removes the PLL response variations with different loop filter voltage values.

In **Figure 9.47**, the gain of the op-amp isolates the actual loop filter output voltage from the charge pump output voltage, controlling the loop so that the charge pump output voltage is V_{REF}. Each passive component has the same basic function here as in the passive filter.

The third-order active filter in Figure 9.47 adds an additional pole within the feedback loop and reference frequency sidebands are significantly reduced. Charge pump current impulses now cause ramps instead of steps in the output voltage, but they still cause problems at the amplifier inputs. Good design algorithms include this extra pole as part of the fundamental dynamics — not as a disturbance of conventional second-order dynamics.

Though the design in **Figure 9.48** adds another capacitor by splitting the input resistor, the added complexity is usually well worth it. Charge pumps put out very narrow, spiky signals that are usually of sufficient amplitude and duration to temporarily drive the op-amp input into saturation. If this happens the PLL response becomes non-linear and problems often arise.

The improvement of this design is to low-pass filter the charge pump pulses before the output is applied to the op amp input. The op amp far prefers the resulting ramps, remaining linear and therefore predictable in behavior. This results in further reduction in reference frequency sidebands and is easily adapted to difference-output digital PDs, including the PFD.

Design equations for these loop filters will be presented after the remaining PLL component blocks are discussed.

Reference Oscillator

Any PLL is simply a stability transfer mechanism, so the behavior of the reference oscillator will not be improved on. The fractional frequency error of the overall frequency synthesizer will match that of the reference oscillator. For example, if the output of the PLL synthesizer changes 1 kHz at an output frequency of 1 GHz [1000/1,000,000,000 = 1:1 million, or 1 part per million (ppm)] then the reference has changed that same fractional amount. If the crystal reference is 10 MHz (a very common frequency for crystal references) its frequency drift was (1 ppm) × (10 MHz) = 10 Hz. This is not much frequency drift, but it is significant to the ultimate output frequency.

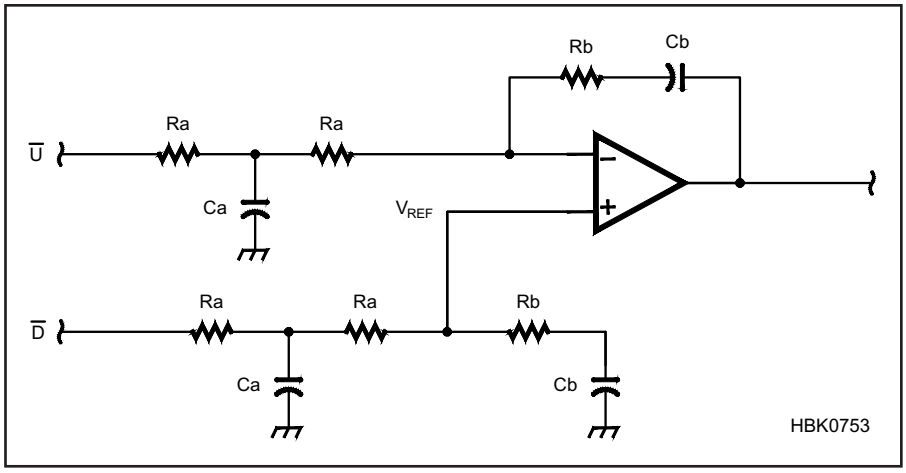

Figure 9.48 — Adding additional low-pass filtering by splitting the input resistor and adding C_a improves loop stability and reduces reference frequency sidebands. This design is well-suited for differential-output PFDs.

9.6.3 PLL Loop Filter Design

When it is time to design a PLL synthesizer, it is best to do it in steps. Begin with knowing how the hardware you are using behaves:

- VCO tuning characteristics, K_0
- PFD characteristic (with charge pump), K_s
- Crystal reference frequency, f_{XO}

Next, list the requirements of your application

- Output frequency range, f_{OUT}
- Output frequency step size, f_{STEP}

Now you are ready to choose the two main design parameters:

- Loop bandwidth — this should not exceed 5% of $f_{REF} = f_{STEP}$.
- Phase Margin — this sets the overall stability of the PLL and can be any value between 45 and 60 degrees. A phase margin of 45 degrees optimizes settling time for fast channel changes, as might be required in a frequency-hopping application. This level of speed is typically NOT needed in amateur radio applications. A phase margin of 60 degrees removes peaking in the transfer function and provides a small reduction in output phase noise, which is useful in amateur radio applications.

With all of this information in place it is time to calculate the values for the two frequency dividers and the loop filter components. Begin with the divider values:

$R = f_{XO}/f_{REF}$

$N = f_{OUT}/f_{REF}$

$N_{MIN} = f_{OUT,MIN}/f_{REF}$

$N_{MAX} = f_{OUT,MAX}/f_{REF}$

The best design algorithm for designing a stable PLL calls for using the geometric mean of the feedback divider value range in the loop filter design. This same idea is used to manage the range of VCO gain values for K_0.

$$N_{DESIGN} = \sqrt{N_{MAX} \times N_{MIN}}$$

$$K_0 \text{ design target value} = \sqrt{K_{0,MAX} \times K_{0,MIN}}$$

PFD gain depends on whether the charge pump is voltage mode or current mode:

(voltage mode) $K_d = (V_H - V_L)/4p$

(current mode) $K_d = I_{QP}/2\pi$

For all of the active loop filter structures, use the voltage mode value.

Everything is now ready. Place these design values into the appropriate design equations for the selected loop filter structure. For a passive loop filter use **Figure 9.49**. If you are using an active loop filter the appropriate equations are in **Figure 9.50** or **Figure 9.51**. The resistor and capacitor values from these calculations are not critical. Choose values within 30% of what you calculate, and the design will be fine.

Following any loop filter design, it is always a good idea to check the results by directly calculating the filter time constants from the components chosen. Compare these time constants with those derived theoretically. If they match within 10 – 20%, there should be no problem with loop stability.

The passive third order loop filter design equations of Figure 9.49 look very similar to those of its active filter counterpart. The major difference is that this procedure yields two time-constants instead of three, and a term equaling the sum of both filter capacitances.

Unlike the active filter form, there are no arbitrary component choices with the passive filter. From the time constants and the total capacitance, all three loop filter components are uniquely determined.

Since the concepts of natural frequency and damping no longer apply (by definition) in a third-order filter, the more general stability concepts of gain crossover and phase margin are used. A value of 50 – 60 degrees is usually used for the design phase margin.

The equations for Figure 9.50 and Figure 9.51 include the rolloff pole within the loop dynamics, rather than treating it as a loop perturbation. This allows a greater amount of high frequency rolloff to be achieved as it "kicks in" at a much lower frequency than perturbation techniques would allow.

$$\tau_2 = \frac{\sec\phi_0 - \tan\phi_0}{\omega_0} = \frac{1 - \sin\phi_0}{\omega_0 \cdot \cos\phi_0}$$

$$\tau_1 = \frac{1}{\omega_0^{\,2} \cdot \tau_2}$$

$$C_T = \frac{K_0 \cdot K_d}{N \cdot \omega_0^{\,2}} \cdot \sqrt{\frac{1+(\omega_0 \cdot \tau_1)^2}{1+(\omega_0 \cdot \tau_2)^2}} = \frac{K_0 \cdot K_d}{N} \cdot \frac{\tau_1}{\omega_0}$$

ϕ_0 = phase margin (°)
ω_0 = open loop gain crossover frequency (radians/sec)

$$\tau_1 = RC_1 \qquad C_2 = \frac{\tau_2}{\tau_1} \times C_T$$

$$\tau_2 = R\left(\frac{C_1 \times C_2}{C_1 + C_2}\right) \qquad C_1 = C_T - C_2$$

$$C_T = C_1 + C_2 \qquad R = \frac{\tau_1}{C_1}$$

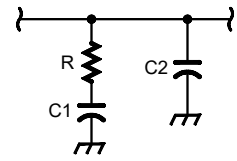

Figure 9.49 — Design equations for passive loop filters.

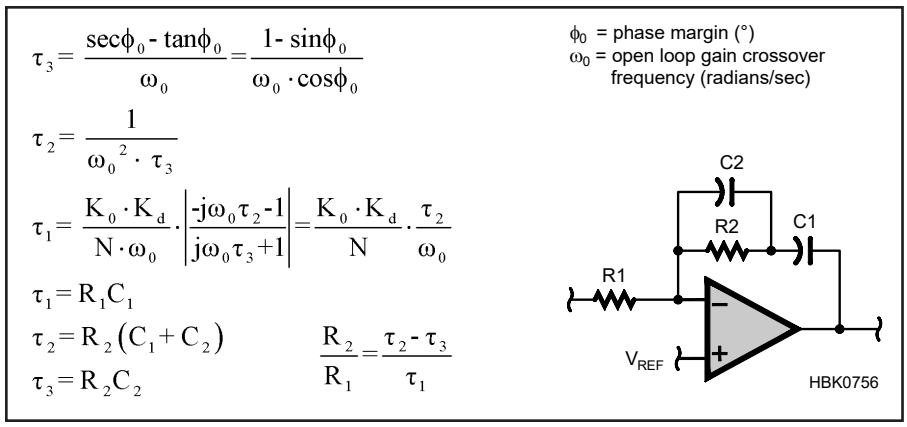

Figure 9.50 — Design equations for third-order active loop filters.

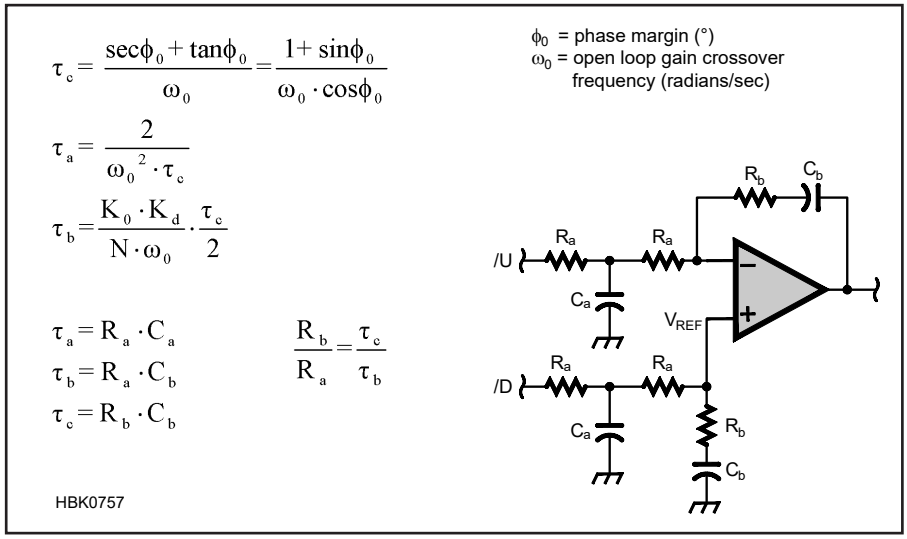

Figure 9.51 — Design equations for differential-input active loop filters.

A usual design strategy is to choose the resistors first to achieve the design ratio. This may provide several possible resistor pairs. Calculate the resulting capacitor values for each resistor pair possibility, and choose the set that provides the closest realizable values.

9.6.4 Fractional-N Synthesizers

The simplest frequency synthesizer is the single PLL with a programmable counter as a frequency divider in the feedback path. It can now all be integrated onto a single semiconductor device, though better performance may be available with some sections done discretely. These synthesizers are well suited to channelized radios where the channel spacing is fairly wide. They are common in VHF/UHF radios, cell phones, TV tuners and so on. However, apart from VHF/UHF FM transceivers, most amateur radio operation requires fine-resolution synthesizers to simulate the look and feel of a free-tuning VFO.

To this point in the chapter, the PLL synthesizer designs have assumed that the feedback divider operates as a single programmed integer value. This design is referred to as an *integer-N PLL*. These PLLs have an unavoidable trade-off of the step size versus every other performance parameter. Making a fine resolution loop by using a very high value of N forces a very low phase detector frequency and low open loop gain. This leads to extremely slow settling, bad phase noise, and poor suppression of spurs.

Multiple PLL loops can be used to give fine resolution without invoking other performance limitations other than cost, size, and power consumption. PLL/DDS hybrid designs can do the same job and save some cost, size, and power consumption. The holy grail is a fine resolution synthesizer with a clean output, small size, low cost and low power consumption but that has not yet been attained. There is another thread of development particularly suited to high levels of integration, although it, too, has limitations. It is called *fractional-N* or *frac-N synthesis.*

THE ORIGIN

A single PLL would have fine resolution, if the divide-by-N stage in its feedback path weren't constrained to divide only by integers. The dividers are implemented as digital counters, counting cycles of the VCO frequency. Counting in anything smaller than integers would need something running faster than the VCO to act as an interpolator. To get steps much finer than the increments given by integers would need an interpolator running at a very large multiple of the VCO frequency, which is impractical.

One solution is to vary the value of N so that the average value of N, taken over a long enough time period, gives the required resolution. This can be done, and it doesn't need impossibly fast logic. For example, consider the 2-meter synthesizer discussed in the previous section on Frequency Resolution: if N is varied in a repeating pattern of 8 divisions by 1474 followed by 2 divisions by 1475, the average is $(8 \times 1474 + 2 \times 1475)/10 = 1474.2$. A PLL operating with this type of variable division in the feedback loop is called a *fractional-N PLL.*

The resulting problem with such a PLL is that it is never exactly on the desired frequency. It switches between too-high and too-low, even though its average might be just right. We could view this as wide deviation FM, where the modulation is a rectangular pulse of controlled mark-space ratio. Such a spectrum will have high noise levels and huge sidebands which are quite undesirable.

Nevertheless, designers of such systems have tried to slow the loops so that they could not follow the switching and sat on the average frequency. The problem with obtaining sufficient filtering of the sidebands was that the loop became too slow and all the problems mentioned above appeared, making the solution worse than single-N synthesis.

ANALOG PHASE INTERPOLATION

Variable-N synthesizers were not used in amateur radio transceivers, but it was common in test equipment from the 1980s and 1990s without the tuning and noise requirements of radio equipment. This type of synthesizer is described in the context of using or repairing test equipment. (For example, see the HP 3335A Synthesizer service manual's Theory of Operation section.)

PLLs can be phase modulated, quite easily, within their bandwidths by summing in a modulating voltage after the phase detector. The phase effect of the switching of the N number can easily be calculated. A digital system can be made to do this, but an extra

analog system must be added to interpolate the results of the digital system. This hybrid circuit can compute a phase modulation waveform which cancels all the effects of the frequency switching, but leaves the average intact. Cancellation is never perfect but manufactured synthesizers were developed that suppressed the unwanted sidebands to around –80 dBc (decibels with respect to the carrier level).

This is a complex system and can be hard to understand. Sometimes a second way of viewing it may be easier: If a ramp waveform is added after a phase detector in a PLL, the loop will be phase modulated with a ramping phase, or in other words, a frequency offset. If the ramp slope is controlled, any amount of resolution is possible. The problem is that ramps cannot go on forever, and the phase detector will soon hit the end of its range. The solution is to increase the N number of the divider just once at the same time that the ramp is reset by a cycle's worth of phase. A sampling system is used to disguise the disturbance of the transient phase shift and ramping resumes.

This is a large system needing careful, individual, trimming to get the analog circuitry to mesh seamlessly with the digital circuitry. The digital parts are integrated into a custom IC, but the discrete analog parts require a lot of board space. The performance was acceptable for mid-range equipment at the time. State of the art equipment used it as part of a hybrid structure as the least significant of three PLLs. The resulting fractional-N loop gave the system much finer resolution than the previous generation of 5-loop hybrids.

NOISE SHAPING

These early schemes switched the PLL divider between N and N+1 in simple, repetitive patterns, so the inevitable sideband energy was concentrated into large, obvious components. One variant of the filtering scheme used "rate multiplier" logic devices to control N to N+1 switching and reduced the amplitude of the unwanted components by spreading them out. While still not clean enough to be generally useful, other technologies were maturing: oversampling 1-bit DACs for audio, pseudo-random binary sequences for simulating telephone traffic for error rate testing, and the addition of dither to ADCs.

An integer-N PLL is a sort of DAC. You put data in as modulation, and the result is a change of frequency rather than voltage. The principles of an extreme oversampling DAC can be applied to it just as much as to the output buffer of a 1-bit DAC in a CD player. These systems scramble the sideband energy and make it much less conspicuous — a good thing. The sideband energy is also spread over a wider frequency range than the bandwidth of the loop filter. This is even better because it means more energy can be filtered out. Further, the scrambling process can be engineered to control the spectrum of the noise-like sidebands, called *noise shaping*, pushing most of the noise energy high enough in frequency that filtering can do a good job. The logic systems which shape the noise are higher-order sigma-delta modulators (See the **DSP and SDR Fundamentals** chapter). They take up a lot of logic elements on an IC, but can be highly integrated and are therefore small, cheap, and have low power consumption. As a result, the noise-shaped fractional-N synthesizer looks like a simple PLL with a large amount of logic processing the data fed to its programmable divider.

Delta-modulators above the second order are inherently unstable. Audio converters have various proprietary schemes to stabilize them, and these little subtleties are carefully guarded secrets. The MASH (Multi-stAge noise SHaping) DAC arrangement adds together a series of low-order modulators, rather than trying to make one very high-order modulator. Fixing the stability problem requires that the output doesn't just jump between two states, N and N+1 in PLL terms, but it dances over a limited range.

LIMITATIONS

With the division factor varying in a pseudo-random way, the operating point of the phase detector is also varying over a wide range and the phase detector needs to have a similarly wide operating range. Since filtering only happens *after* the phase detector and large amounts of noise at high offset frequencies are present, the phase detector must be very linear.

The noise-shaping scrambler keeps the close-in frequency range clean, but any nonlinearity in the phase detector allows the strong high frequency noise components to intermodulate. Intermodulation creates products at close-in frequencies, spoiling the noise performance of the close-in range.

This limits how clean the noise-shaping synthesizer can be made. The best examples are useful for many purposes but still aren't good enough for a high-grade HF receiver or any system requiring superior adjacent channel rejection. For such uses, they are combined with a high-performance integer-N PLL in a hybrid structure similar to how the original fractional-N loop was used.

AVAILABILITY

Fractional-N synthesizers have been available in chip form for several years. Analog Devices, Texas Instruments, and other manufacturers have offered families of PLL ICs for a long time, and their ranges include devices with all the noise shaping logic on-board. Many even have GHz VCOs on-board as well, with dividers down to more mundane frequency ranges.

These synthesizers have become an RF design building block. The RFMD RFFC2071, for example has a pair of medium-level active mixers with an entire synthesizer all on one die. Handheld transceivers that also receive 0.5 to 999.999999 MHz are likely to be made with these parts. If you are choosing one of these devices, look carefully at the noise sideband performance: the higher frequency version of the RFFC2071 IC has noise sidebands a few dB lower than the lower frequency version.

WHAT NEXT?

The RF semiconductor industry is currently focused on mass-market applications such as mobile phones, WiFi and so on. In particular, the requirements for high-speed data links (e.g., 5G) and the RF Instrumentation needed to test them continue to drive advances in integrated synthesizer, VCO, and DDS development with increases in operating frequency as well as decreasing phase noise in PLL/VCOs and the release of integrated vector signal generators such as the AD9164 with a 48-bit DDS operating as high as 12 GHz.

Development of fractional-N systems is within the range of the interested amateur, either by using the general-purpose parts on the market, or by developing their own using programmable logic array devices. A moderately-sized FPGA contains enough logic for a number of fractional-N synthesizers. For example, see the July/August 1998 *QEX* article by Ulrich L. Rohde, N1UL "A High-Performance Fractional-N Synthesizer" which describes a design patented in 2001.

A figure of merit for a fractional-N system is the maximum frequency at which the phase detector can operate. Higher frequencies allow the scrambled noise to be spread over a greater frequency range, making it easier to filter. The obvious approach is to develop faster and faster logic, but there is an alternative: Several fractional-N dividers with several separate phase detectors may have their outputs combined. The phases of the outputs from the separate dividers may be offset from each other by programming registers differently in each scrambler, and the reference signals to the phase detectors may be offset to match.

Several phase detectors can thus act sequentially over each detector's cycling period, making a system with an effective phase detector frequency several times that of a single-phase detector. The effective phase detector frequency can even be higher than the output frequency. A further elaboration would be to seed the pseudo-random noise generators in each divider differently. In this way the noise components from each divider

will not correlate, though the tuning components correcting the VCO frequency will correlate. This means that the tuning components add more strongly, as voltages, while the noise components add more weakly as power. Each doubling of the number of dividers has the potential for a 3 dB improvement in noise performance of the system. This multiple divider arrangement is currently covered by a patent, but individuals are free to read them, to experiment and to perhaps find a better way. Patents eventually expire and their technology enters the public domain. Each increase in the density of programmable logic makes it easier.

9.6.5 A PLL Design Example

As our design example, let us consider a synthesized local oscillator chain for a 10 GHz transverter. **Figure 9.52** is a simplified block diagram of this 10 GHz converter. This example is chosen because it is a departure from the traditional multi-stage multiply-and-filter approach. It permits realization of the oscillator system with two simple loops and minimal RF hardware. It is also representative of what is achievable with current hardware, and can fit in a space of 2 to 3 square inches. This example is intended to be a vehicle to explore the loop design aspects and is not offered as a "construction project." The multiple design details required are beyond the scope of this chapter.

Two synthesized frequencies, 10 GHz and 340 MHz, are required. Since 10.368 GHz is one of the popular X-band traffic frequencies, we initially mix this with the 10 GHz LO to produce an IF of 368 MHz. The 368 MHz IF signal is subsequently mixed with the 340 MHz LO to produce a 28 MHz final IF, which can be fed into the 10-meter input of any amateur transceiver. We focus our attention on the design of the 10 GHz synthesizer only. Once this is done, the same principles are applied to the 340 MHz section. Our goal is to design a low-noise LO system (by adopting minimum division ratios) with a loop reference oscillator that is an integer multiple of 10 MHz. Using this technique allows the entire system to be locked at a later time to a 10 MHz standard for precise frequency control.

For the microwave synthesizer we consider using a line of microwave integrated circuits made by Analog Devices in gallium-arsenide (GaAs) material. These devices include a selection of prescalers operating to 12 GHz, a 5-bit counter that operates to 2.2 GHz and a phase/frequency detector that operates up to 1.3 GHz. If we use a crystal reference frequency f_{XO} of 100 MHz, and also apply that as f_{REF} into the PFD (therefore R = 1), then the feedback divider number needs to be N = 10,000/100 = 100. One way to realize this in hardware is to represent 100 = 4 × 25, starting with a divide-by-4 prescaler and finishing with the 5-bit counter programmed to divide by 25. The divide-by-4 prescaler output at lock will be 2500 MHz, which is too high for the 2200 MHz limited programmable counter. We need a design change.

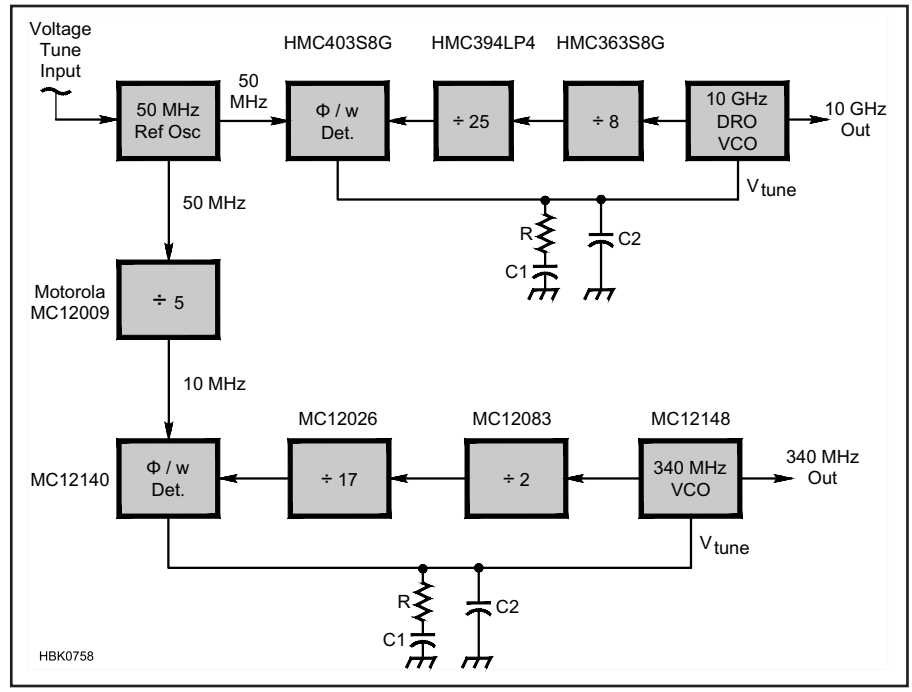

Figure 9.52 — A simplified block diagram of a PLL local oscillator for a 10 GHz converter.

The obvious solution is to select a divide-by-8 prescaler, providing a 1250 MHz output from the 10 GHz input. We cannot program the counter to half of 25, so we need to leave it programmed at 25. This makes the output from the feedback divider at 10,000/(8×25) = 50 MHz. We need to set R = 2.

Before even thinking about designing the loop filter, we need to know the VCO gain, VCO noise performance, divider noise performance, phase detector gain, phase detector noise performance and finally reference noise performance. For the 10 GHz VCO, we are always looking for parts that are easily available, useable, and economical, so salvaging a dielectric resonator from a Ku band LNB is promising (see "SHF Super Regenerative Reception" by Andre Jamet, F9HX, in Jan-Feb 2002 *QEX*). These high-Q oscillators can be fitted with a varactor and tuned over a limited range with good results. The tuning sensitivity of our dielectric resonator VCO is about 10 MHz per volt and the phase noise at 10 kHz offset is –87 dBc/Hz, and –107 dBc/Hz at 100 kHz.

The phase detector and divider information is available from the device data sheets. The HMC363 divide-by-8 operates up to 12 GHz and has a programmable charge pump that sets the phase detector gain K_d. It is usually good to keep this gain high, so we choose to use a 2-mA charge pump current, giving K_d = 0.32 mA/radian. The data sheet also says that this PFD has an output noise floor of –153 dBc/Hz measured at 100 kHz offset. The HMC394 programmable divider operated up to 2.2 GHz with the same output noise floor. The HMC984 PFD operated up to 350 MHz with a noise figure of merit (FOM) of –231 dBc/Hz/Hz-f_{REF}, so with f_{REF} = 50 MHz the PFD output noise floor is –231 + 10 log (50,000,000) = –154 dBc/Hz. Assuming the R divider is implemented in CMOS, the output noise floor from it is –163 + 10 log (50,000,000/125,000) = – 137 dBc/Hz. This noise is much higher than the noise from the GaAs dividers, so it is actually better here to use a GaAs divide-by-2 for the R counter instead of using CMOS.

The logic noise floor for the entire PLL is the sum of the individual noises from each divider and the PFD. The result is –147 dBc/Hz. This is determined at the PFD. To calculate how this noise will measure at the PLL output we need to add 20 log N = 46 dB. Thus, the output noise floor due to PLL logic devices is –101 dBc/Hz. This is close to the VCO phase noise at an offset of 100 kHz, so minimum noise design suggests that we select 100 kHz as our loop bandwidth.

An alternative is to eliminate the R counter and directly use a 50 MHz reference. An excellent choice for a low noise reference is the one described by John Stephensen, KD6OZH in Nov/Dec 1999 *QEX*. The noise

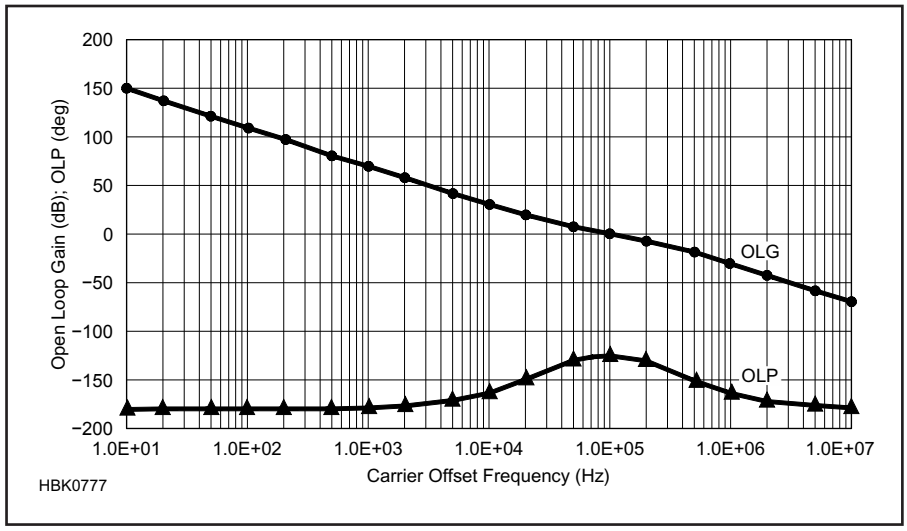

Figure 9.53 — Comparisons of the theoretical and actual PLL dynamics show extremely good correspondence.

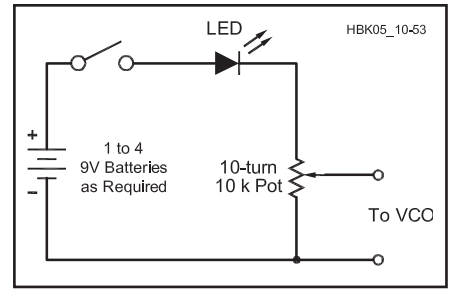

Figure 9.54 — Clean, variable dc voltage source used to measure VCO gain in a PLL design.

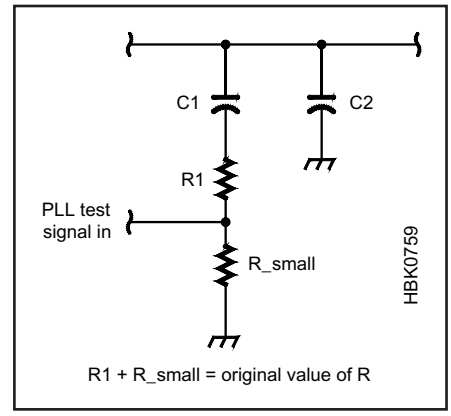

Figure 9.55 — A simple testing structure for PLL dynamics.

performance of this VCXO is in the order of –160 dBc/Hz at 10 kHz offset at the fundamental frequency. Translating this to the output frequency means adding 20 log N, for a result of –114 dBc/Hz. Eliminating the R counter also lowers the logic noise floor to –104 dBc/Hz at the output. The translated reference oscillator noise is 10 dB below the noise floor from the logic devices, so the logic noise floor dominates.

Using this information, we can now design the loop filter. Assuming that the VCO tuning voltage is between 0.5 and 4.5V we can directly use the charge pump included in this PFD and therefore select the passive third order loop filter from Figure 9.46. We use the equations in Figure 9.44 and calculate loop filter component values of R = 6,977 ohms, C1 = 724 pF, and C2 = 80 pF. We therefore select 6.8k, 820 pF, and 82 pF for these component values respectively. The Bode plot of the theoretical and actual PLL dynamics is shown in **Figure 9.53**. It is nearly impossible to tell them apart.

If the VCO requires higher tuning voltages, you can insert a non-inverting op-amp gain stage between the output of this loop filter and the VCO. Choose a low input current op-amp to minimize leakage from the loop filter capacitors. Be careful — the gain of this amplifier stage changes the effective VCO gain that the loop experiences. When designing the loop filter components, the "design" VCO gain must be replaced by the product of the actual VCO gain and the amplifier gain. It is also essential to assure that the 1 dB bandwidth of the voltage amplifier greatly exceeds the loop bandwidth so that it does not degrade the PLL phase response and hence loop stability.

9.6.6 PLL Measurements and Troubleshooting

VCO GAIN

One of the first things we need to measure when designing a PLL is the VCO gain. The tools needed include a voltmeter, some kind of frequency measuring device like a receiver or frequency counter and a clean source of variable dc voltage. The circuit in **Figure 9.54** containing one or more 9 V batteries and a 10-turn, 10 kΩ pot does nicely. One simply varies the voltage some amount and then records the associated frequency of the VCO. The gain of the VCO is then the change in f divided by the change in voltage (Hz/V).

LOOP BANDWIDTH

Measuring the actual bandwidth of the PLL usually means measuring the entire gain and phase responses of the PLL. There is a much easier way that needs only a square-wave generator and an oscilloscope. This method uses the PLL testing structure of **Figure 9.55** developed by Glenn Ewart that injects a low-value square wave into the loop at a low impedance point so that impulse response can be observed directly. The test signal is also ground-referenced and is independent of the loop locking voltage.

Start by swapping the series RC components in the loop filter so the resistor is connected to ground and the capacitor attaches to the VCO tuning line. Then split the resistor into two resistances that together add up to the total resistance needed by the filter. The bottom resistor connected to ground is very small — 50 Ω is a nice choice when the total resistor value is much greater than 50 Ω. The top resistor is temporarily replaced by a potentiometer that can make up the desired remaining resistance. Across the bottom resistor we connect the square-wave generator and set it to a small amplitude (100 or 200 mV is common) at a frequency a few percent of the loop bandwidth, and operate the PLL. Looking at the VCO tuning line with an oscilloscope we see the impulse response of the PLL.

Now adjust the potentiometer way off-value to the low side. The impulse response will now show ringing. The frequency of this ringing is a good approximation of the PLL loop bandwidth.

SETTLING TIME

Having measured the loop bandwidth, readjust the potentiometer back to its nominal value. The impulse response should look *much* cleaner! It is tempting to look at the settling time of the impulse response and say that this is the settling time of the PLL. This unfortunately is not quite true.

The PLL is not fully settled until capacitor C1 is fully charged. This usually takes slightly longer than what the impulse response itself shows. Move the scope probe to the "ground" side of C1 and measure when this voltage reaches zero. This means there is no more current flowing through the filter resistors and the capacitor is fully charged.

TROUBLESHOOTING PLLS

Here are some frequently encountered problems in PLL designs:

• The outputs of the phase detector are inverted. This results in the loop slewing to one or the other power supply rails. The loop cannot possibly lock in this condition. Solution: Swap the phase detector outputs.

• The loop cannot comply with the tuning voltage requirements of the VCO. If the loop runs out of tuning voltage before the required voltage for a lock is reached, the locked condition is not possible. Solution: Re-center the VCO at a lower tuning voltage or increase the rail voltages on the op amp.

The loop is very noisy and the tuning voltage is very low. The tuning voltage on the varactor diodes in the VCO should not drop below the RF voltage swing in the oscillator tank circuit. Solution: Adjust the VCO so that the loop locks with a higher tuning voltage.

9.6.7 Commercial Synthesizer ICs

In this section, we explore using commercially available synthesizer chips and the role of DDS in more recent hybrid architectures. The following is not intended to be project oriented, but rather is designed to expose the reader to additional concepts that cannot be fully explored here. The reader, being made aware of these ideas, may wish to examine them in detail using references at the end of this chapter.

Many synthesizer chips have been introduced in recent years for cellular and Wi-Fi applications and have some interesting properties and capabilities that can be exploited with additional design. Some properties that can be exploited are programmable charge pump current, a rich set of division options for creating multiple-loop synthesizers and typically low power consumption as will be shown in the next section.

For instance, an integrated PLL/VCO could be used as a 704 to 640 MHz synthesizer and its external divider set to 128. This division would reduce the phase noise profile by 40 dB, making a reasonably quiet synthesizer for the 5 to 5.5 MHz range. The step size would also be reduced by a factor of 100, making the spacing required in the 500 to 550 MHz range equal to 1 kHz for a step size of 10 Hz at 5 MHz. This would make a good local oscillator for a simple traditional radio that covers 80 and 20 meters using both mixer products against a 9 MHz SSB generator.

There are also other techniques that can prove helpful. Decoupling the VCO from the chip will permit one to avoid much of the "on chip" noise that VCOs are susceptible to at the expense of some more complexity. If the VCO is implemented externally, there is an opportunity to design it with increased operating Q, thereby improving the overall phase noise performance. An external VCO also opens other opportunities.

As mentioned earlier, these chips are designed for operation at a V_{CC} of 3 to 5 V, which limits the output tuning voltage swing to V_{CC}. One must be able to fit the entire tuning range voltage and the ac voltage excursion in the VCO tank circuit within the V_{CC} range. As mentioned at the end of the earlier section Improving VCO Noise Performance, this problem can be overcome with the addition of voltage gain in an external operational amplifier, running at a higher voltage. This allows the signal-to-noise ratio of the oscillator to be improved by increasing the tank voltage swing and still having adequate voltage range to perform the tuning function. There are certainly more examples of how the performance of these devices could be improved for amateur applications.

9.6.8 Synthesizer Architectures Using Commercial ICs

INTEGER-N VERSUS FRACTIONAL-N PLLS VERSUS DDS

One of the disadvantages of a single loop PLL frequency synthesizer is that the PFD frequency determines the step size of the output frequency, that is,

$$f_{OUT} = N \times f_{PFD}$$

where

$$f_{PFD} = \frac{f_{REF}}{R}$$

For two-way radio and other systems with fixed channel spacing this is generally not a problem. However, for CW and SSB operation, tuning steps of 10 Hz or smaller are almost a necessity. Both integer and fractional-N synthesizers can tune 10 Hz steps, but in a fractional-N PLL the cutoff frequency of the loop filter is still set by the channel spacing even if the PFD frequency is much higher. In the days of analog FM systems that used 25 kHz channel spacing, this wasn't an issue but a 10 Hz step is the same as 10 Hz channel spacing, which mandates an extremely long settling time in the loop filter. In an integer-N PLL, the PFD frequency is the same as the step size, so in this case both choices have the same settling time, set by the step size.

On the other hand, a DDS can easily tune in 10 Hz steps as fast as the steps can be programmed. However, it can have both phase truncation spurs (typically not an issue but worth mentioning) and (worth mentioning) spurs from the nonlinearity of its output DAC. Also, the resolution of the output DAC sets the noise floor, being approximately 6 dB per bit. Thus, the theoretical noise floors for 12-bit and 14-bit DACs are roughly 72 and 84 dB, respectively.

But the drawback of a DDS chip is its frequency coverage; the DDS chips used in amateur radio designs typically accept a 400 MHz clock input and provide a maximum output frequency that is 40% of that, or 160 MHz, using the industry rule-of-thumb of 2½ points per cycle. With none of these solutions having an obvious advantage, this brings us back to the original question of Integer-N versus fractional-N PLLs versus DDS: Which should we use?

Let's look at some modern synthesizer topologies, starting with the simple PLL/VCO combination shown in **Figure 9.56**, which is typically used for fixed-frequency clock or LO generation. Here, the VCO can be separate or integrated with the PLL and the PLL can be either an integer-N or fractional-N. Figure 9.56's output frequency is given by

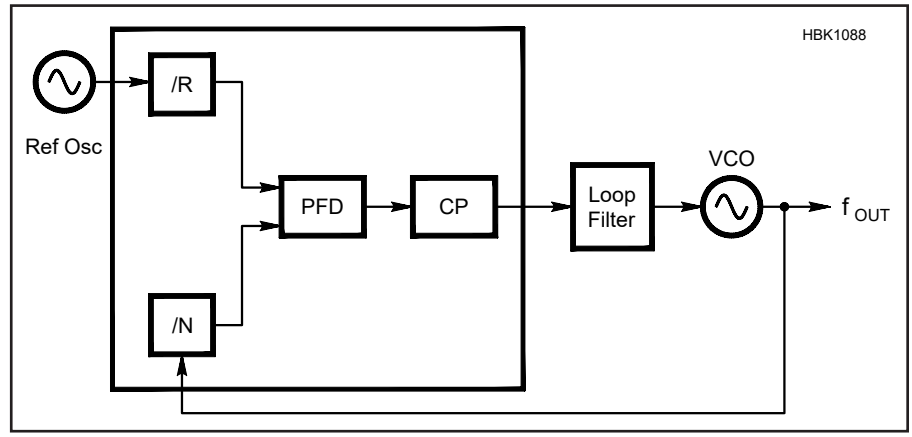

Figure 9.56 — An Integer-N PLL and VCO are the basic building blocks of frequency synthesizers.

$$f_{OUT} = \frac{N}{R} f_{REF}$$

where
f_{out} is the output frequency of the VCO,
N and R are the values of the N and R dividers, respectively, and
f_{ref} is the comparison frequency of the phase-frequency detector (PFD).

For the sake of example, let's say we are designing a 2 meter (144-148 MHz) all-mode receiver with a first IF of 45 MHz and high-side injection ($f_{LO} > f_{RF}$), giving us an LO tuning range of

144+45 = 189 MHz as the lowest LO frequency to

148+45=193 MHz as the highest LO frequency.

We will choose 25 MHz as f_{REF} and R=1 for several reasons. First, 25 MHz crystals are common. Second, no harmonics of 25 MHz fall either on the first IF (45 MHz), in the 2 meter band (144-148 MHz), or in the LO tuning range (189-193 MHz). And third, 25 MHz may also be used to clock the system's microprocessor, which minimizes both the number of oscillators to reduce cost and the number of "birdies" (low-level, in-band mixing products) heard in the receiver.

With R=1 and using the middle of the VCO's tuning range, 191 MHz, as f_{OUT},

$$N = \frac{f_{OUT}}{f_{REF}} = \frac{191}{24} = 7.64$$

where N_{INT}= 7 and N_{FRAC} = 0.64. Here, because N turns out to be a fractional value,

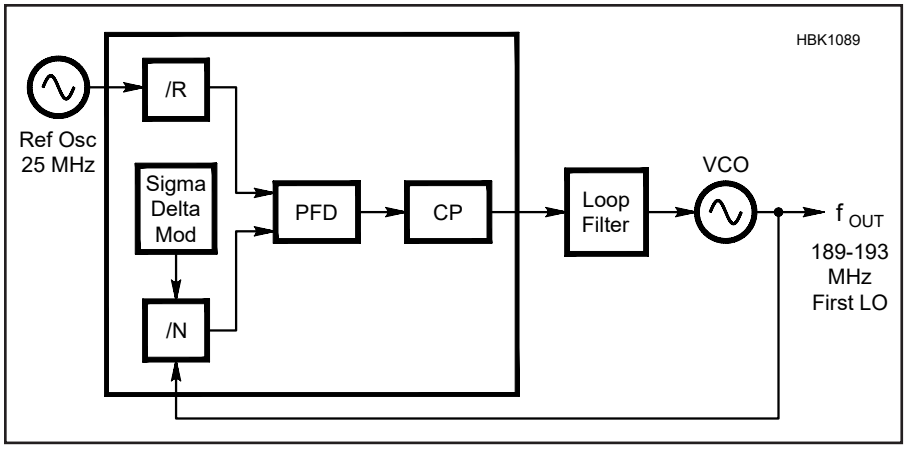

Figure 9.57 — A proposed fractional-N Synthesizer for the 2-meter band.

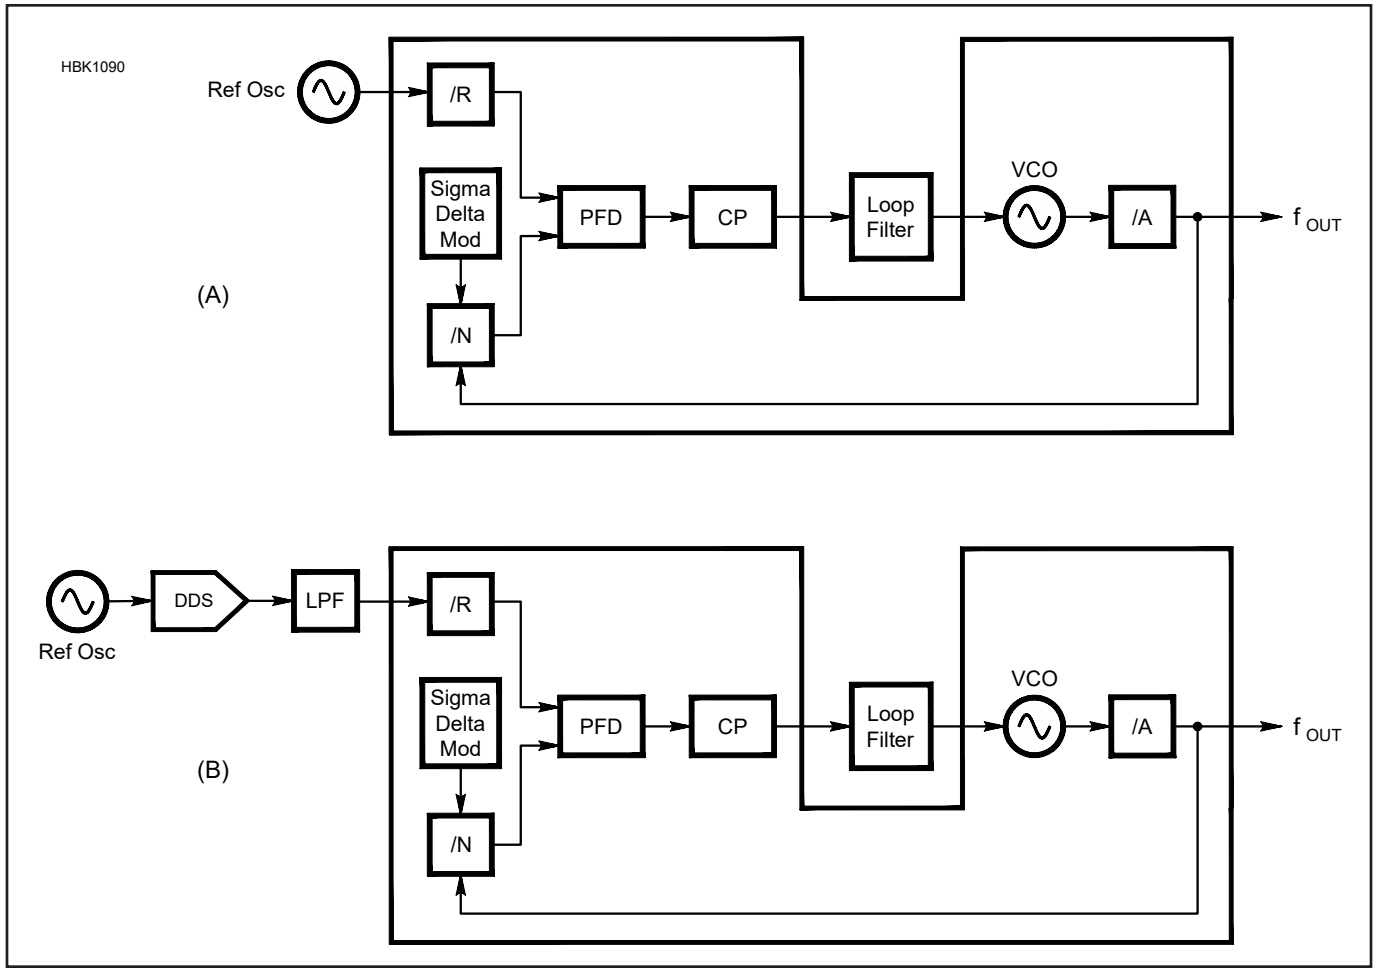

Figure 9.58 — A proposed synthesizer for the 2-meter band using a Fractional-N PLL with an integrated VCO and an auxiliary output divider. Suitable candidates include the Analog Devices ADF4351, which covers 35 – 4400 MHz; the STmicroelectronics STW81200, which covers 46.857 to 6000 MHz; and the Texas Instruments LMX2581, which covers 50 to 3760 MHz. Using a DDS as the reference allows 1 Hz resolution or better, depending on the resolution of the DDS.

we'll use a fractional-N PLL, which has a sigma-delta modulator.

A possible difficulty here is finding a fractional-N PLL with an N value as low as 7; many PLLs have prescalers that require a minimum N value in the 100s or even higher. In this case we would choose a PLL without a prescaler. **Figure 9.57** shows our completed design.

INTEGRATED PLL/VCOs

Fortunately, our baseline design and can also be implemented by using a fractional-N PLL with an integrated VCO and auxiliary divider A, as shown in **Figure 9.58**. The purpose of the auxiliary divider is twofold: it both improves the phase noise by 3 dB for each divide-by-two in the A divider and extends the frequency coverage of the PLL/VCO combination; it does this by dividing the integrated VCO's output frequency by 1, 2, 4, t etc. to turn a narrowband synthesizer into a wideband one. Figure 2's output frequency *at the output of the auxiliary divider A* is given by

$$f_{OUT} = \frac{f_{VCO}}{A} = \frac{N}{AR} f_{REF}$$

where

f_{OUT} is the output frequency of the PLL/VCO,

N and R are the values of the N and R dividers, respectively,

A is the value of the auxiliary divider, and

f_{ref} is the comparison frequency of the phase-frequency detector (PFD).

With this level of integration, integrated PLL/VCOs are specified by not only their available R, N, and A values but also their output range from the minimum possible output frequency to the maximum available output frequency.

Continuing our 2-meter design example, we need an integrated PLL/VCO with an output range that spans 189 to 193 MHz to provide the LO signal and a 25 MHz PFD frequency. Suitable candidates include the Analog Devices ADF4351, which covers 35-4400 MHz; the STmicroelectronics STW81200, which covers 46.857 to 6000 MHz; and the Texas Instruments LMX2581, which covers 50 to 3760 MHz. All three suppliers provide free PLL design software to use with these products.

The data sheet for the PLL/VCO may only have phase noise information (sometimes a single frequency measurement but more often a graph) with A=1 for the auxiliary output divider. In cases like this, you can use the following approximation:

$$PN_{dBc/Hz} = PN_{PLL/VCO\,(A=1)} - 3\frac{\log A}{\log 2}$$

where

$PN_{dBc/Hz}$ is the total output phase noise at the desired offset from the VCO frequency,

PN_{PLL} is the phase noise of the PLL/VCO at the same offset from the output frequency, usually listed as a typical specification on the PLL/VCO's data sheet, and

A is the value of the output divider.

The last piece of the equation (log A/log2) may need further explanation: it calculates the power of two used by the output divider. For example, if A=16 (2^4), then log A/log2 = log 16/log 2 = 4. Given that each divide-by-2 *decreases* the phase noise by 3 dB, the reduction in phase noise here is 4 × 3 dB = 12 dB.

So far, so good. But what if we want really small frequency steps, say 100 Hz or less. How do we get there? Let's use the circuit in Figure 9.58A and make tunable by using a DDS IC as the reference. Figure 9.58B's output frequency *at the output of the auxiliary divider* is now given by

$$f_{OUT} = \frac{N}{AR} f_{DDS}$$

where f_{DDS} is the output frequency of the DDS IC.

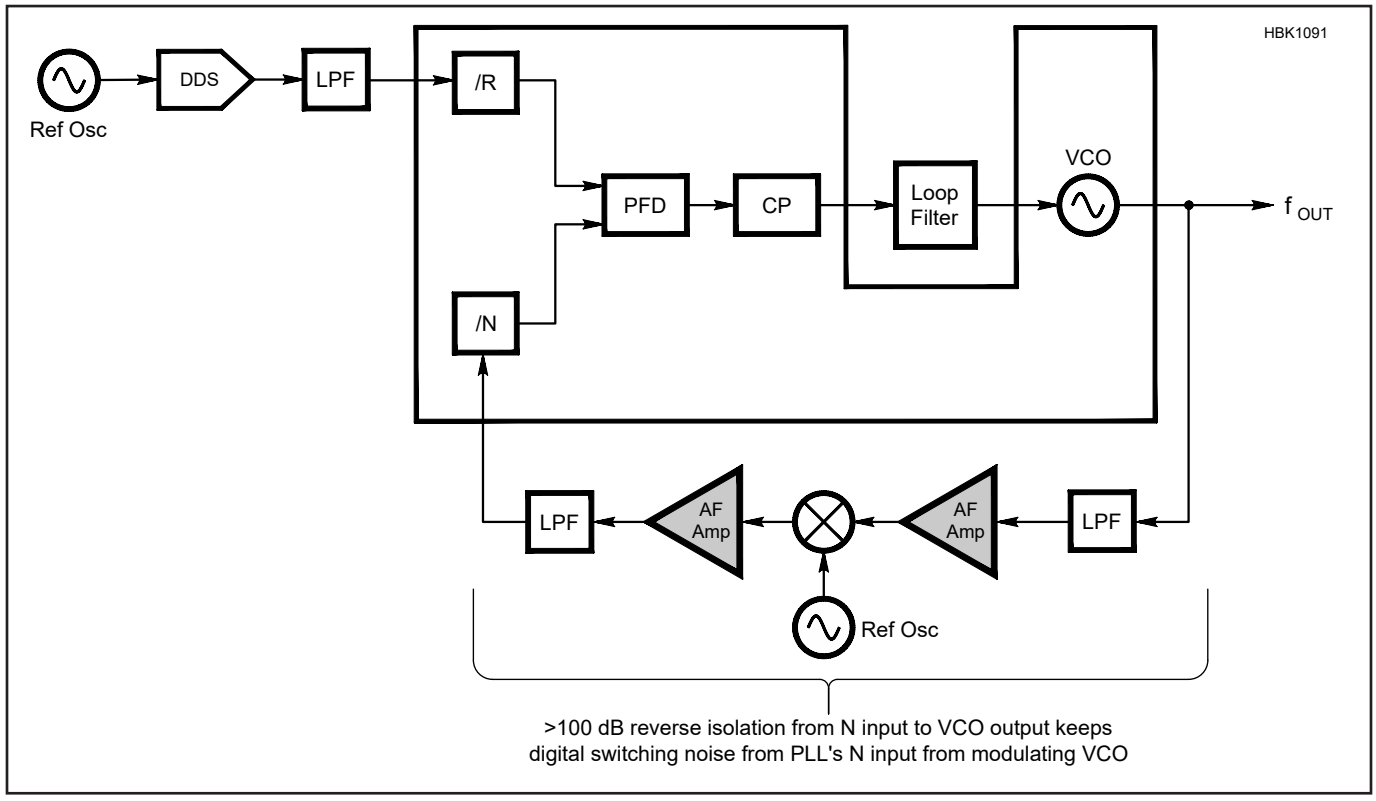

Figure 9.59 — A Frequency Translation Loop eliminates the noise introduced by the N divider in a frequency synthesizer at the expense of higher complexity.

For best performance, the DDS must be clocked by a low-phase-noise source such as a low-phase-noise crystal oscillator or a temperature-compensated crystal oscillator (TCXO) with a PLL/VCO and *not* a low-cost CMOS clock module used for microcontrollers.

Finally, note that the N divider multiplies the PFD frequency, which in turn comes from output of the R divider and before that the reference oscillator; any phase noise or drift from on the PFD's reference input signal is multiplied by N as part of the normal phase-locking process. Is there a way to reduce or eliminate this? Yes, by making N =1.

Figure 9.59 shows such a circuit, called a *frequency translation loop*. In this circuit, the N feedback divider is replaced by a mixer and the N divider is set to N=1. PLLs with built-in prescalers have minimum N values much greater than 1. If this is the case, then the R input can be used instead of the N input and the N input can be driven by the reference.

Assuming a PLL with N=R=1, Figure 9.59's output frequency can be shown to be

$$f_{OUT} = f_{PFD} + f_{EXT}$$

where

f_{OUT} = the synthesizer's output frequency in Hz,

f_{PFD} = the phase-frequency detector's comparison frequency, and

f_{EXT} = the external frequency driving the mixer in the feedback circuit

When f_{REF} is supplied by a DDS as shown here, then the output resolution of the synthesizer is that of the DDS for R =1 ($f_{PFD} = f_{DDS}$) and even smaller for R>1 (that is, $f_{PFD} = f_{DDS}/R$). Because the DDS determines the frequency resolution of this synthesizer, an integer-N PLL rated to only a few hundred MHz may be used with a high loop bandwidth for lower cost and faster frequency settling time. This architecture is used in such RF instruments as spectrum and network analyzers.

Typically, such circuits use either a PLL that allows N = R = 1 or a standalone PFD and charge pump. But what if N has a minimum value of, say, 100, due to a built-in prescaler? In this case, you could swap the N and R inputs: the R input is now the feedback input and the N input is now driven by the reference oscillator. Next use a reference oscillator based on a 2 GHz source, and set N at the lowest value that makes sense for the PFD frequency.

In a practical implementation of a translation loop synthesizer, the feedback path must have sufficient (e.g., 100 dB or more) isolation between the feedback input and the VCO to prevent digital switching noise from the PLL's digital dividers from feeding back to the VCO.

The choice of active or passive mixer can make a difference to the overall isolation in the feedback path. Let's look at the port-to-port isolation of a solid-state mixer compared to that of a diode ring mixer. To simplify the comparison, assume both mixers have identical 30 dB port-to-port isolation between all ports.

At first, all things seem equal with the same port-to port-isolation. However, the difference comes out when the LO drive levels are compared: an active mixer typically requires –10 to 0 dBm LO drive; a passive mixer typically requires +7 dBm or higher LO drive (**Figure 9.60**). Assuming identical port-to-port isolation specifications, the passive mixer can easily have 17 dB or more worse leakage between ports in this application. In practice, an active mixer may have 5 to 10 dB better port-to-port isolation than a passive mixer. Read the data sheets carefully and include the LO drive levels in your comparison.

A different approach is taken in **Figure 9.61**, showing a synthesizer with fractional division in the feedback loop but no fractional spurs. How does this work? The circuit uses a DDS, which functionally acts as a fractional divider. The key is that a DDS IC uses a DAC

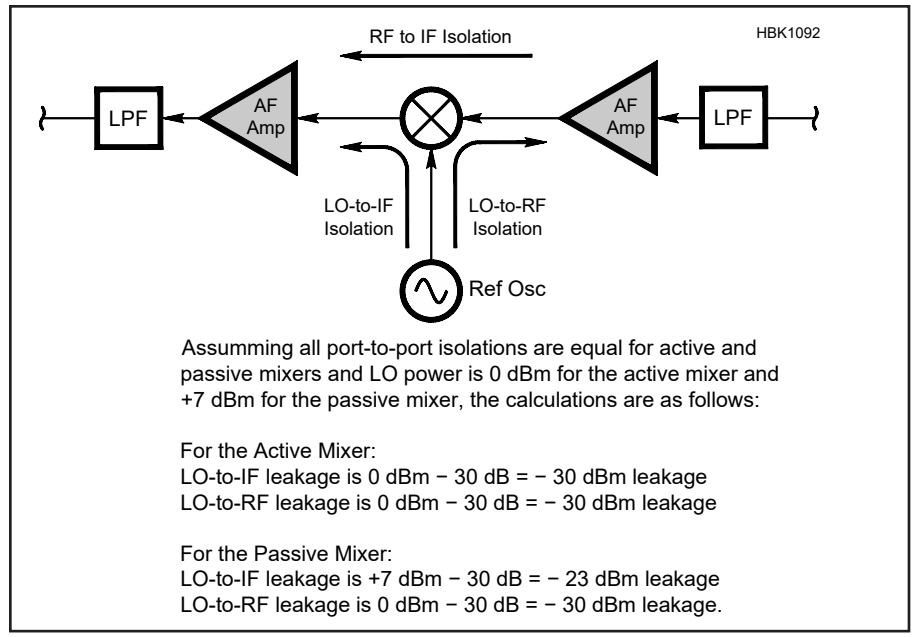

Figure 9.60 — Comparing the port-to-port isolation of active and passive mixers requires calculating the effect of the LO drive level on port-to-port leakage paths in a mixer.

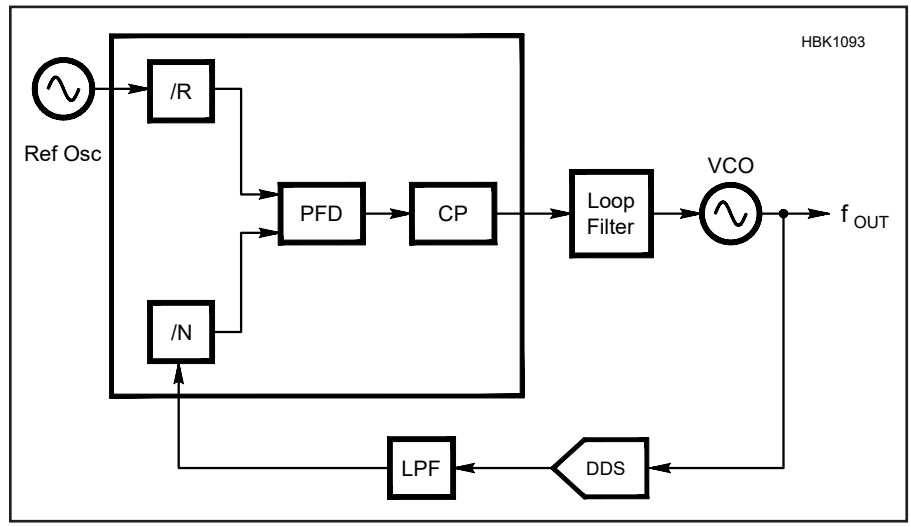

Figure 9.61 — Using a DDS as the "fractional divider" in the feedback path provides the functionality of a fractional-N PLL without the fractional spurs, resulting in a cleaner output spectrum: the DDS's output into the PFD is a single frequency synthesized by a DAC rather than the multiple frequencies created by a fractional-N divider.

and clock to generate a sinewave point-by-point; at all times the output is a single frequency. In contrast, a fractional-N PLL creates numerous output frequencies whose average is the desired frequency. Even with the averaging effect of the charge pump and loop filter, some vestige of the fractional spurs can modulate the VCO and show up in the output. Using the DDS allows a higher cutoff frequency in the loop filter and faster settling time.

A fractional-N divider, in contrast, uses the average of multiple divide-by-N cycles to create its fractional output frequency. The loop filter has to smooth out the variations in the VCO's tuning voltage caused by the multiple N divider values. What doesn't get filtered out becomes a fractional spur, hence the use of a DDS.

The VCO's output frequency also serves as the clock input, f_{clk}, to the DDS, that is, $f_{out} = f_{clk}$. The output of the DDS IC (f_{DDS}), and is one of the inputs to the PFD. Assuming R=1 so $f_{REF} = f_{PFD}$ and working our way back from the VCO to the PFD, these relationships give us

$$f_{REF} = f_{DDS} = \left\{\frac{FTW}{2^M}\right\} f_{CLK} = \left\{\frac{FTW}{2^M}\right\} f_{OUT}$$

Substituting for f_{DDS} and solving for f_{OUT} yields

$$f_{OUT} = \left\{\frac{2^M}{FTW}\right\} f_{REF}$$

Finally, solving for the *frequency tuning word* (FTW),

$$FTW = 2^M \frac{f_{REF}}{f_{OUT}}$$

where

f_{DDS} is the DDS's output frequency, in Hz,
f_{OUT} is the VCO's output frequency, in Hz,
f_{CLK} is the DDS's clock input frequency, in Hz, and
2^M is the fixed resolution of the DDS, here, m=32.

Using the same frequency plan as in our previous examples, that is, a 2 meter receiver covering 144-148 MHz with a 45 MHz first IF, high-side injection from a 189-193 MHz LO and a 25 MHz reference frequency, let's calculate the FTW. We'll use the midpoint of the VCO's 189-193 MHz tuning range in our calculations:

$$FTW = 2^M \frac{f_{REF}}{f_{OUT}} = 2^{32} \frac{25\,MHz}{191MHz} =$$

$$4,294,967,296 \times \frac{25}{191} = 562,168,494.241$$

Rounding off the FTW to the nearest integer (562,168,494)and checking our work we get

$$f_{OUT} = \left\{\frac{2^M}{FTW}\right\} f_{REF} =$$

$$\left\{\frac{2^{32}}{562,168,494}\right\} 25\,MHz = 191.000000082\,MHz$$

Note that the frequency error is 0.082 Hz, which, besides being difficult to measure, is more than adequate for this application!

It's worth noting that a standalone DDS IC can synthesize frequencies from VLF to HF and higher. Let's look at an example. To generate a unique frequency, the DDS must generate at least two points per cycle. Industry practice is to generate at least 2½ points per cycle, or 40% of the clock frequency. In the case of DDS used above, we had to generate a 25 MHz clock frequency, which required a minimum of 2.5 points/Hz × 25 MHz, which yields 62.5 MHz as the minimum clock frequency. A clock input of 189–191 MHz easily meets this requirement. What frequency ranges can a DDS cover? **Table 9.2** shows the required minimum DDS clock frequencies for a few amateur frequency allocations.

Using the information in Table 9.2, we see that a DDS clocked at 125 MHz or higher will cover VLF through the 6 meter band. Once again, we'll use a 45 MHz IF and high-side injection. Searching through online DDS selection guides, we settle on the AD9951, a 14-bit DDS that accepts a 400 MHz input

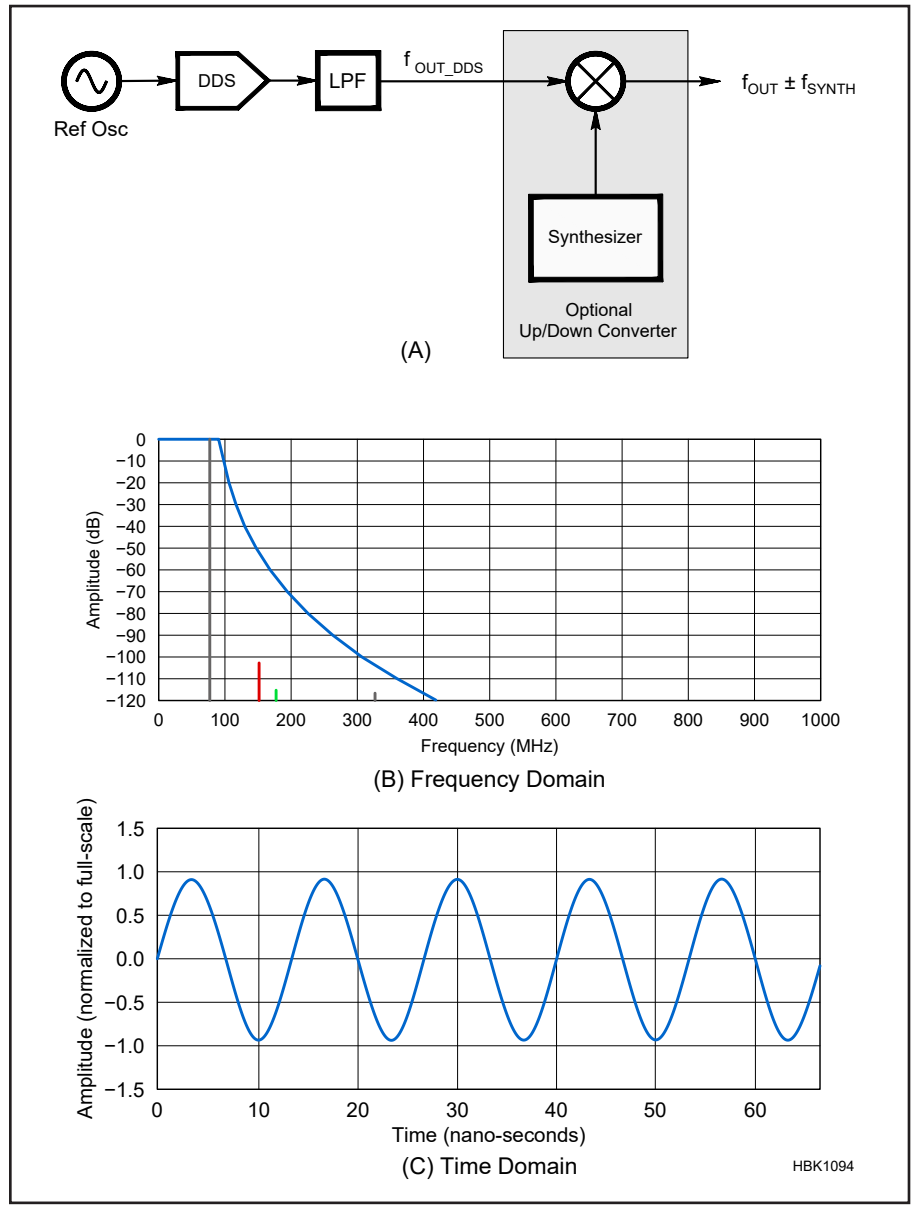

Figure 9.62 — Block diagram of a stand-alone DDS design using the AD9951 (A) with simulation results (B) from *ADIsimDDS* in the frequency and time domain. The filter is a 7th-order Chebyshev low-pass response.

Table 9.2

Frequency Band(s) and Minimum Required Clock for DDS using 2½ points per cycle.

Frequency Band(s)	*Minimum DDS Clock (MHz)*
VLF-30 MHz	75
50 MHz	125
70 MHz (UK)	175
144 MHz	360
225 MHz	562.5
432 MHz	1080

Phase Noise versus Jitter

The Si5351A was designed to replace the crystals used to generate the many clock frequencies required in digital systems: microprocessors, USB, video, data communications, etc. These clocks are specified by time jitter requirements rather than phase noise; that is, the system requirement is that the output signal has to fit within a time-based window for its leading and trailing edges. The phase noise in such a clock is higher than in its RF counterpart because there is no filtering of the output frequency in its fractional divider: the average output frequency is correct but all frequencies that were averaged will appear in the output waveform unless the output divider is programmed for integer mode. In addition, the integrated VCO is typically based on low-Q digital delay techniques (delay-locked loops, ring oscillators, etc.) rather an LC tank circuits, which again contributes to higher phase noise.

clock. We choose a DDS with a 14-bit DAC for better S/N ratio and lower spurious output than one with a 10- or 12-bit DAC. It also has an internal clock multiplier so we have the option of once again using a 25 MHz crystal (or slightly higher to keep its harmonics out of the 6 meter band). However, we'll assume a separate 400 MHz reference oscillator is available, which will have a lower noise floor than the crystal multiplier circuit. **Figure 9.62A** shows the block diagram for this synthesizer with an optional frequency converter.

The DDS manufacturer, Analog Devices, provides a web-based tool, *ADIsimDDS* for this and similar products. (**tools.analog.com/en/simdds**) Using this tool, we can generate an output spectrum and simulate several loop-pass filter topologies to choose the most effective. The tool also calculates the FTW value for a given combination of clock and output frequency in binary, decimal, and hex.

Figure 9.62A shows the simulator's output for the synthesizer in the frequency and time domain. In the frequency domain at the top, there are several spurious outputs shown as individual signals along with the low-pass filter's response. The output sine wave is also shown.

Figure 9.62's synthesizer architecture can also serve as a replacement VCO for vintage transceivers (for example, the Drake C-Line) that had optional external VCOs. Many of these transceivers typically used a tunable LO of 5-5.5 MHz based on a permeability-tuned oscillator (PTO, described previously) to tune a 10% frequency range with a linear frequency readout.

AFFORDABLE DESIGN SOLUTIONS

So far, we've looked at performance-driven synthesizer architectures. However, there are many designs where size and cost are the overriding concerns. For example, the Skyworks Si5153A is a low-cost, multiple-output digital clock synthesizer for crystal replacement in computers. It has three programable outputs that span 9 kHz to 200 MHz, which covers the VLF through 2m amateur bands.

The tradeoff in using a digital clock synthesizer is higher phase noise than a traditional PLL/VCO combination but satisfactory jitter for numerous digital applications. **Figure 9.63** shows a detailed block diagram based on the data sheet and includes references on using this particular IC. Note that the programmable output dividers are fractional-N dividers; this means that the average output frequency will have zero error but the instantaneous output frequency will vary. This and the noise from the integrated VCO accounts for the Si5153A's phase noise's being higher than that of a traditional PLL/VCO that uses an LC-based oscillator.

Not surprisingly, this combination of small size, high functionality, and low cost is used in numerous QRP designs to generate the LO

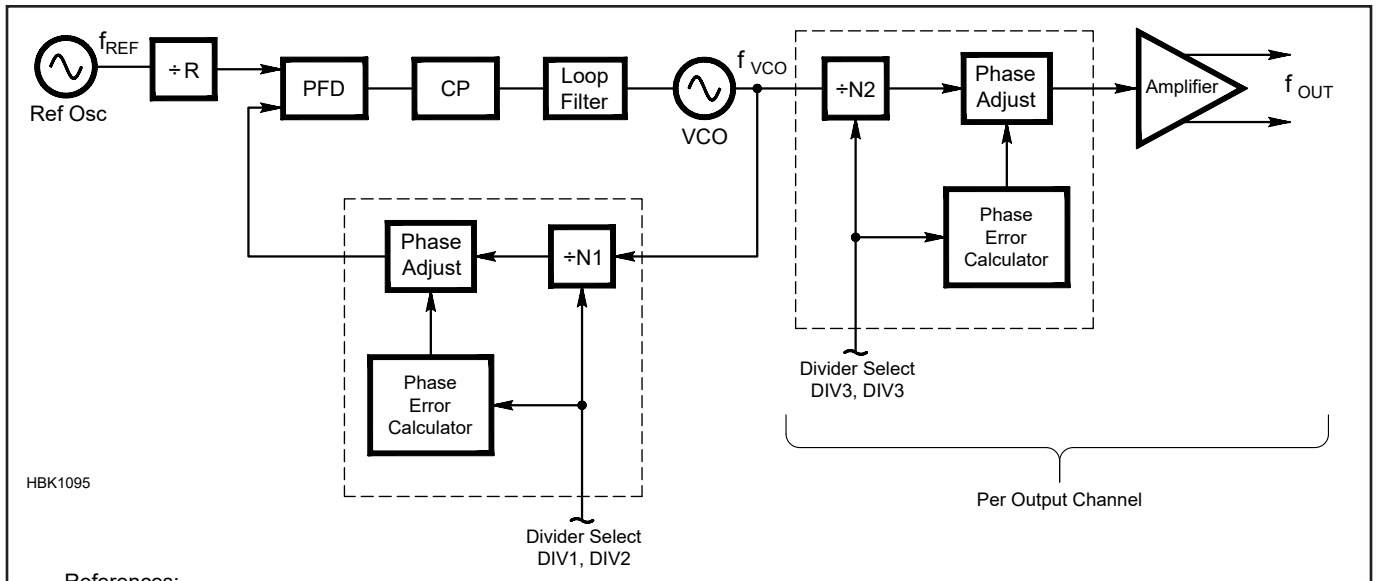

References:

1. Simplifying Timing Architectures with Flexible Clocks, Silicon Laboratories, Inc., Austin TX, now available from Skyworks, Inc.
2. Si5351A/B/C-BI2C-ProgrammableAny Frequency CMOS Clock Generator + VXCO Data Sheet.

Figure 9.63 — A detailed block diagram of the SI5351A's internal architecture.

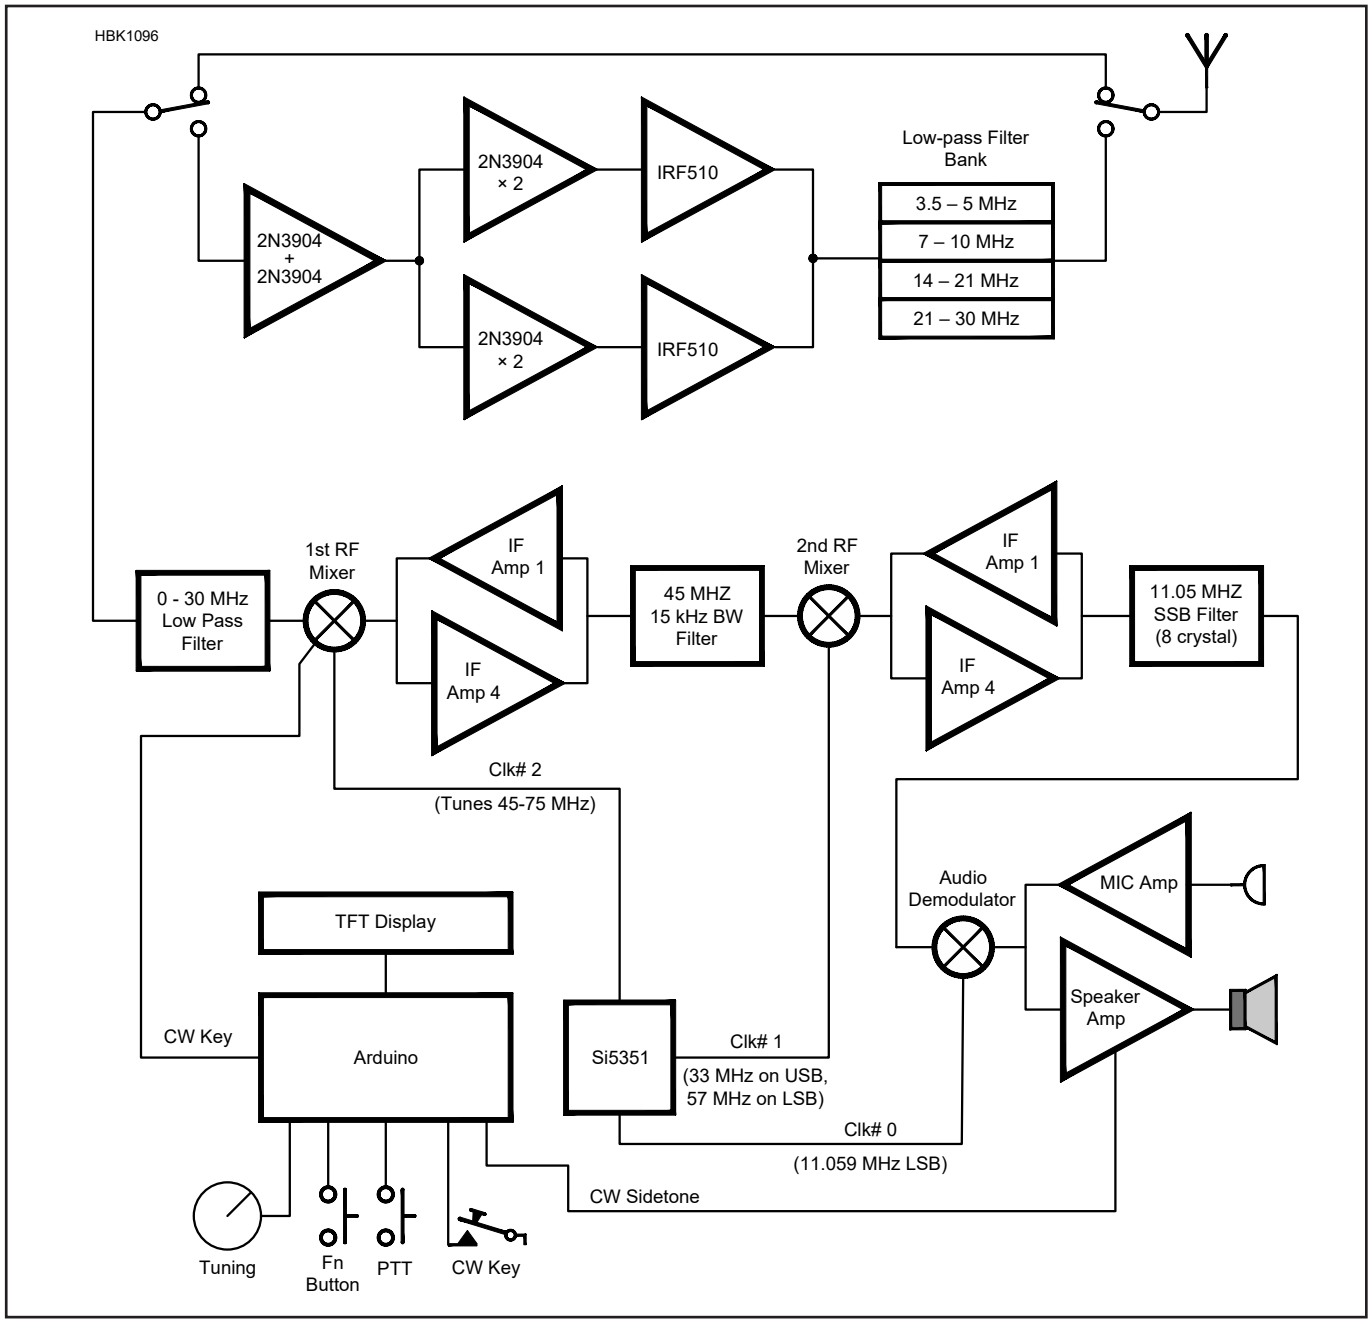

Figure 9.64 — A simplified block diagram of the uBitX6 and the Si5351a-based synthesizer, which is used in numerous QRP designs.

and BFO signals in single and dual conversion transceivers. One such example is the open-source uBitx family of SSB QRP transceivers designed by Ashhar Farhan, VU2ESE, and described in detail at **www.hfsignals.com**. In the synthesizer implementation used in the uBitx6 Five-Band Transceiver (**Figure 9.64**), the Si5351 provides three output frequencies: a 45-to-75 MHz first LO, a 57 MHz/33MHz LSB/USB second LO, and an 11.059 MHz BFO. Figure 9.64 shows a simplified version of the implementation of the receiver, including the synthesizer.

Another interesting synthesizer design is that of the AR7070 LF/HF Receiver, described in July 2013 *QEX* by Colin Horrabin, G3SBI. **Figure 9.65** shows the functional block diagram of the receiver section, including the synthesizer. Like the previous examples, it also uses a 45 MHz first IF. This synthesizer uses a mix of DDS, VCO, discrete PLL, frequency doublers to achieve its high performance.

The VCO is also worth mentioning; it uses a double-tuned tank circuit to achieve better performance than typical single-tuned tanks. Quoting from the article:

"The local oscillator in the HF7070 is a DDS/PLL design using a double-tank VCO. This was designed by John Thorpe and it is similar to the one used in the AR7030 receiver, but with improved performance. The basic double-tank oscillator is shown in **Figure 9.66**. The principle involved is that this circuit can only oscillate if the cold end of the active tank (the one with the J310) is a low impedance to ground. This can only occur if the dummy tank is series resonant at the same frequency as the active tank. This means that as you move away from the carrier two resonators are active, which increases the rate of phase noise fall-off with offset frequency. This circuit has never been analyzed from a phase noise point of view by a mathematician, but measurements show phase noise falls off at

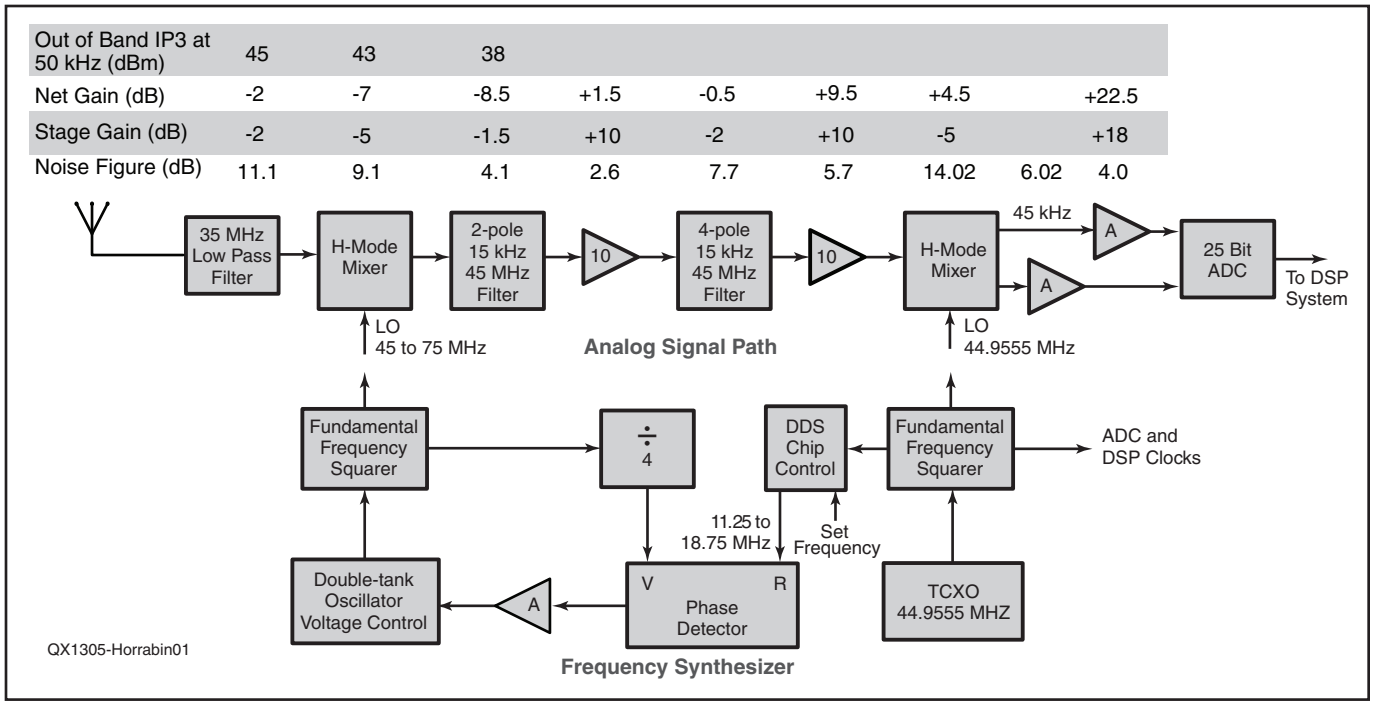

Figure 9.65 — Functional Block Diagram of the HF7070 LF/HF Receiver from *QEX* July 2013.

30 dB/decade compared to 20 dB/ decade in a single resonator oscillator."

Finally, let's look at a non-synthesizer design using a PLL as a programmable prescaler. Standalone prescalers are uncommon these days, yet a PLL can have a built-in programmable prescaler with several GHz of bandwidth. Fortunately, many PLLs have a MUXOUT pin, which lets the user access several internal signals during debug, including the output of the R and N programmable dividers. In a frequency counter, the R and N inputs can be used as user-programmable prescalers and the MUXOUT pin used to drive a microcontroller with a counter input to create a low-cost, wide bandwidth frequency counter. **Figure 9.67** shows an example.

INTEGRATING THE SYNTHESIZER COMPONENTS

Notwithstanding the previous block diagrams and text, there are still a few things missing from a typical design. These are the ubiquitous building blocks that are necessary to connect the elements in a synthesizer.

A PLL's feedback input (usually but not always driving the N counter input) requires both a sample of the VCO's output signal and the correct drive level at the PLL's input. A minimum-loss splitter (**Figure 9.68**) addresses the first issue. Many PLLs allow you to program the output power level to address the second issue.

There are other issues. For example, some PLLs require a minimum slew rate or a rail-to-rail input for proper operation of the R or N dividers when used at lower frequency inputs. For example, this type of requirement is stated as "For RFIN < 5 MHz, ensure slew rate (SR)> 4 V/us" and "For REFIN < 20 MHz, ensure slew rate (SR) > 50V/us". (These are minimum slew rate requirements from the ADF4002 data sheet.) In this case a high-speed Rail-to-Rail Input/Output comparator used to drive the input does the trick. The Texas Instruments TLV3601/TLV3603 or the Analog Devices LTC6752 are suitable choices for this application.

Many PLL/VCOs and DDS ICs have differential outputs; the output of the PLL is typically the differential output of a pair of NPN transistors and requires either inductors or resistors to the positive rail. Depending on the application, the output may connect directly to a mixer with differential inputs, in which case only a differential lowpass filter is needed (**Figure 9.69**).

In many cases, however, the differential output has to be converted to single-ended. Even though connecting only one output cheaply solves the problem, it would be a shame to lose 3 dB of the S/N ratio by discarding half the output signal. There are several ways to do this. **Figures 9.70** and **9.71** show how transformers and LC circuits may be used at higher frequencies.

For continuous coverage of the VLF through VHF bands (depending on the circuit), active circuits may be used, including a differential pair, active feedback circuits using op amps or ADC drivers, or even high-speed comparators when a square wave output is acceptable.

Figure 9.72A – D shows two circuit topologies for differential-to single ended conversion. A word of warning, a simulation is only as good as its designer and the models for the components; it can identify obvious design and wiring errors; and, finally, it lets you know that the circuit *may* work. There is no substitute for a physical prototype, such as breadboard version. Having said that, if you can simulate a circuit before construction, do so. It will give you an idea of how the circuit *should* behave (gain and bandwidth) and estimates for the voltages and currents you may end up measuring during debugging.

Figure 9.72A shows a simple, 4-transistor circuit comprising a differential pair (known as a long-tailed pair in the UK) of 2N3904s with an emitter follower output. Negative feedback (R15 and R16) reduces harmonics (simulation, Figure 9.72B) to about 30 dBc and distortion even more, to more than –45 dBc at frequencies at frequencies above 100 MHz. The circuit should be followed with a low pass filter. Figure 9.72C shows output (top) and input signals (bottom) simulated at 10 MHz. Bandwidth for the simulated circuit is about 30 MHz.

Figure 9.72D shows an op-amp-based circuit using ideal op amps for the simulation. The input here is really a discrete version of a differential ADC driver for high-speed

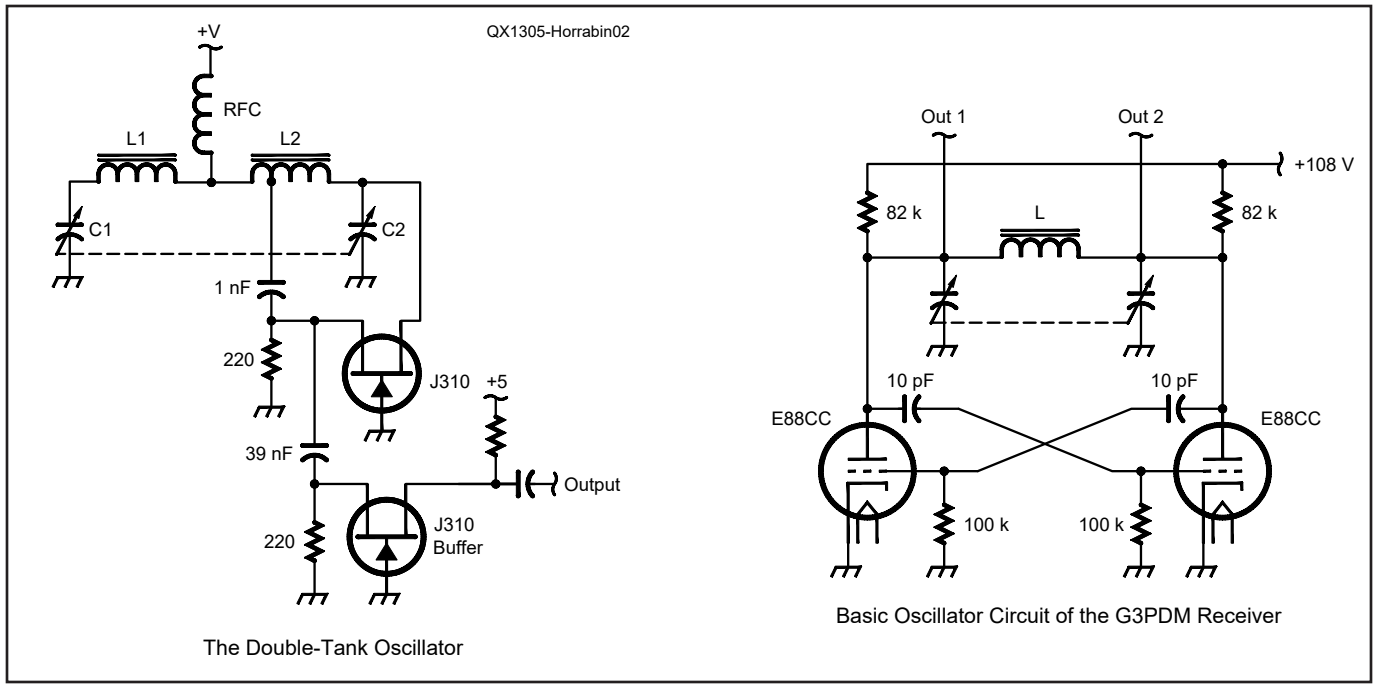

Figure 9.66 — The dual-tank circuit used in the HF7070's VCO achieves better phase noise than typical single-tank VCO designs.

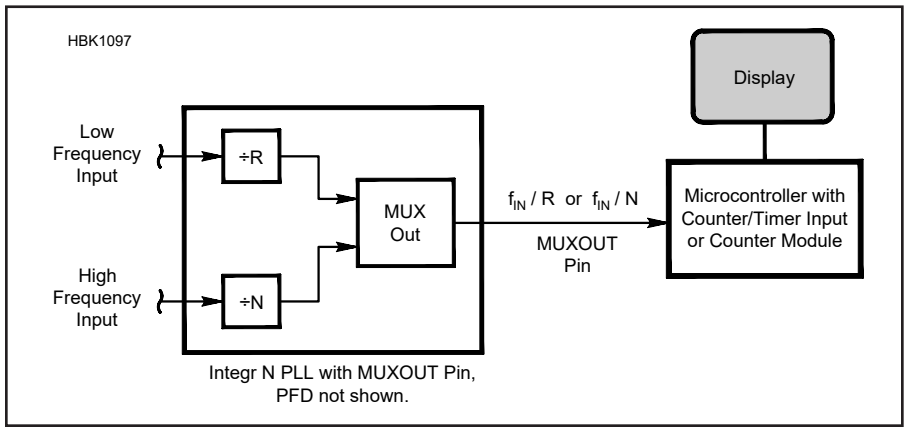

Figure 9.67 — A simple frequency counter uses a PLL as a programmable prescaler. Note that the MUXOUT pin has a limited frequency range (e.g., 10 MHz maximum). Check the data sheet for details.

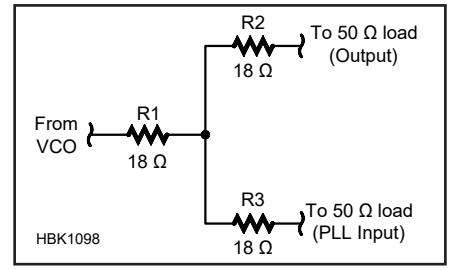

Figure 9.68 — A minimum-loss 3 dB signal splitter provides single-ended feedback from the PLL's output to its input and uses just three resistors.

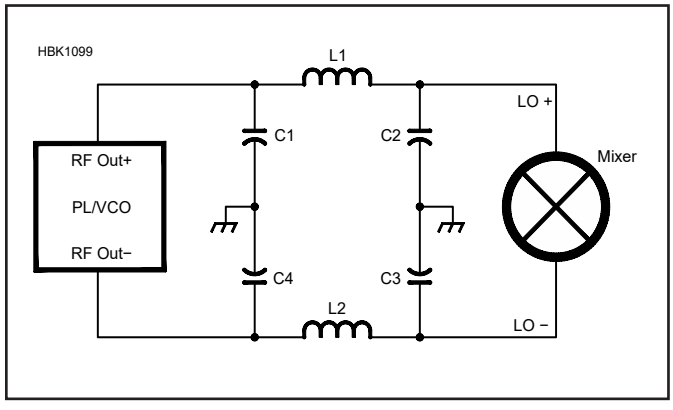

Figure 9.69 — The differential outputs of a PLL/VCO may connect directly to a mixer with a differential LO input. Pull-up inductors/resistors at the RFOUT pins are omitted for clarity.

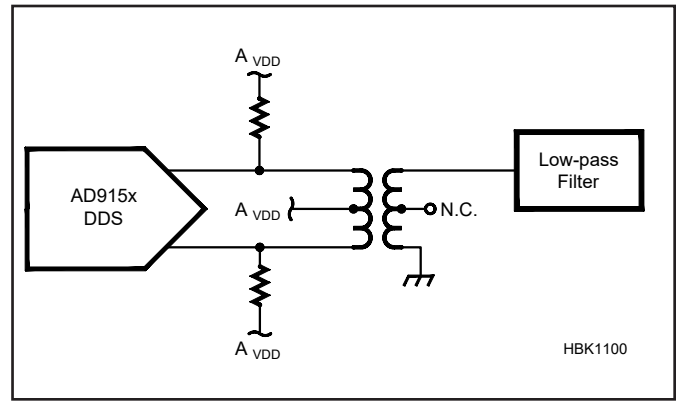

Figure 9.70 — A typical transformer circuit that converts the differential DAC outputs from a DDS IC to single-ended.

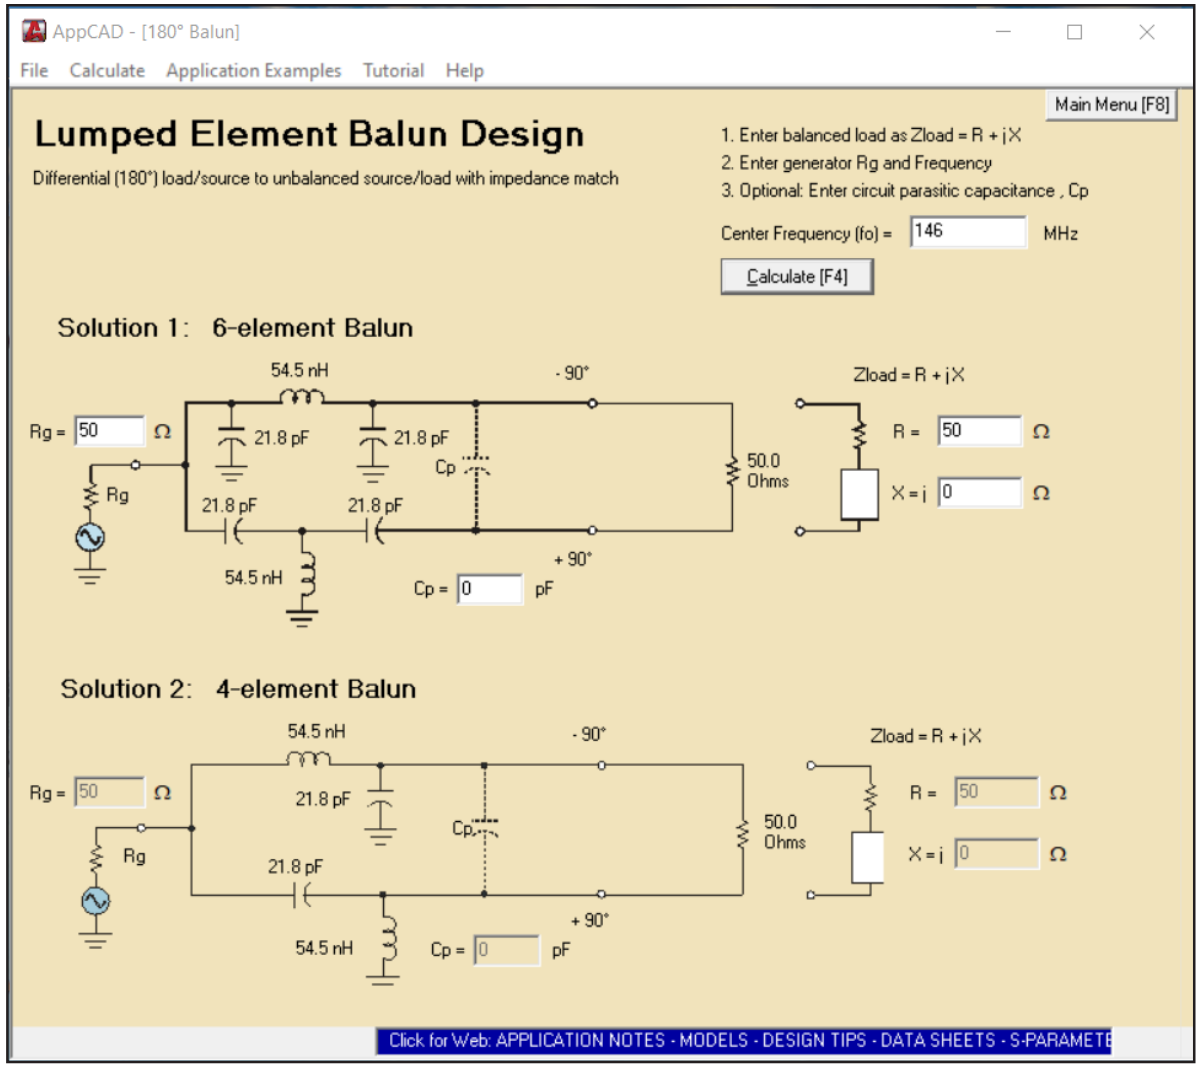

Figure 9.71 — Lumped-element baluns designed using *AppCAD* (now available free from Broadcom) convert differential outputs to single-ended signals. These circuits have about 5% bandwidth. Here the input for a differential output is to the right and the single-ended output is to the left. Suitable inductors are available from such vendors as Coilcraft, Mini-Circuits, Murata, and Toko.

Table 9.3
Popular Capacitor Types from Analog Devices' Tutorial MT-101

TECHNOLOGY	ADVANTAGES	DISADVANTAGES	APPLICATIONS
Aluminum Electrolytic, Switching Type. Avoid general purpose types	•High CV product/cost •Large energy storage •Best for 100V - 400V	•Temperature related wearout •High ESR/size •High ESR @ low temp	•Consumer products •Large bulk storage
Solid Tantalum	•High CV product/size •Stable @ cold temp •No wearout	•Fire hazard with reverse voltage •Expensive •Only rated up to 50V	•Popular in military •Concern for tantalum raw material supply
Aluminum-Polymer, Special-Polymer, Poscap, Os-Con	•Low ESR •Z stable over temp •Relatively small case	•Rapid degradation above 105°C •Relatively high cost	•Newest technology •CPU core regulators
Ceramic	•Lowest ESR, ESL •High ripple current •X7R good over wide temp	•CV product limited •Microphonics •C decreases with increasing voltage	•Excellent for HF decoupling •Good to 1GHz
Film (Polyester, Teflon, polypropylene, polystyrene, etc.	•Hi Q in large sizes •No wearout •High voltage	•CV product limited •Not popular in SMT •High cost	•High voltage, current •AC •Audio

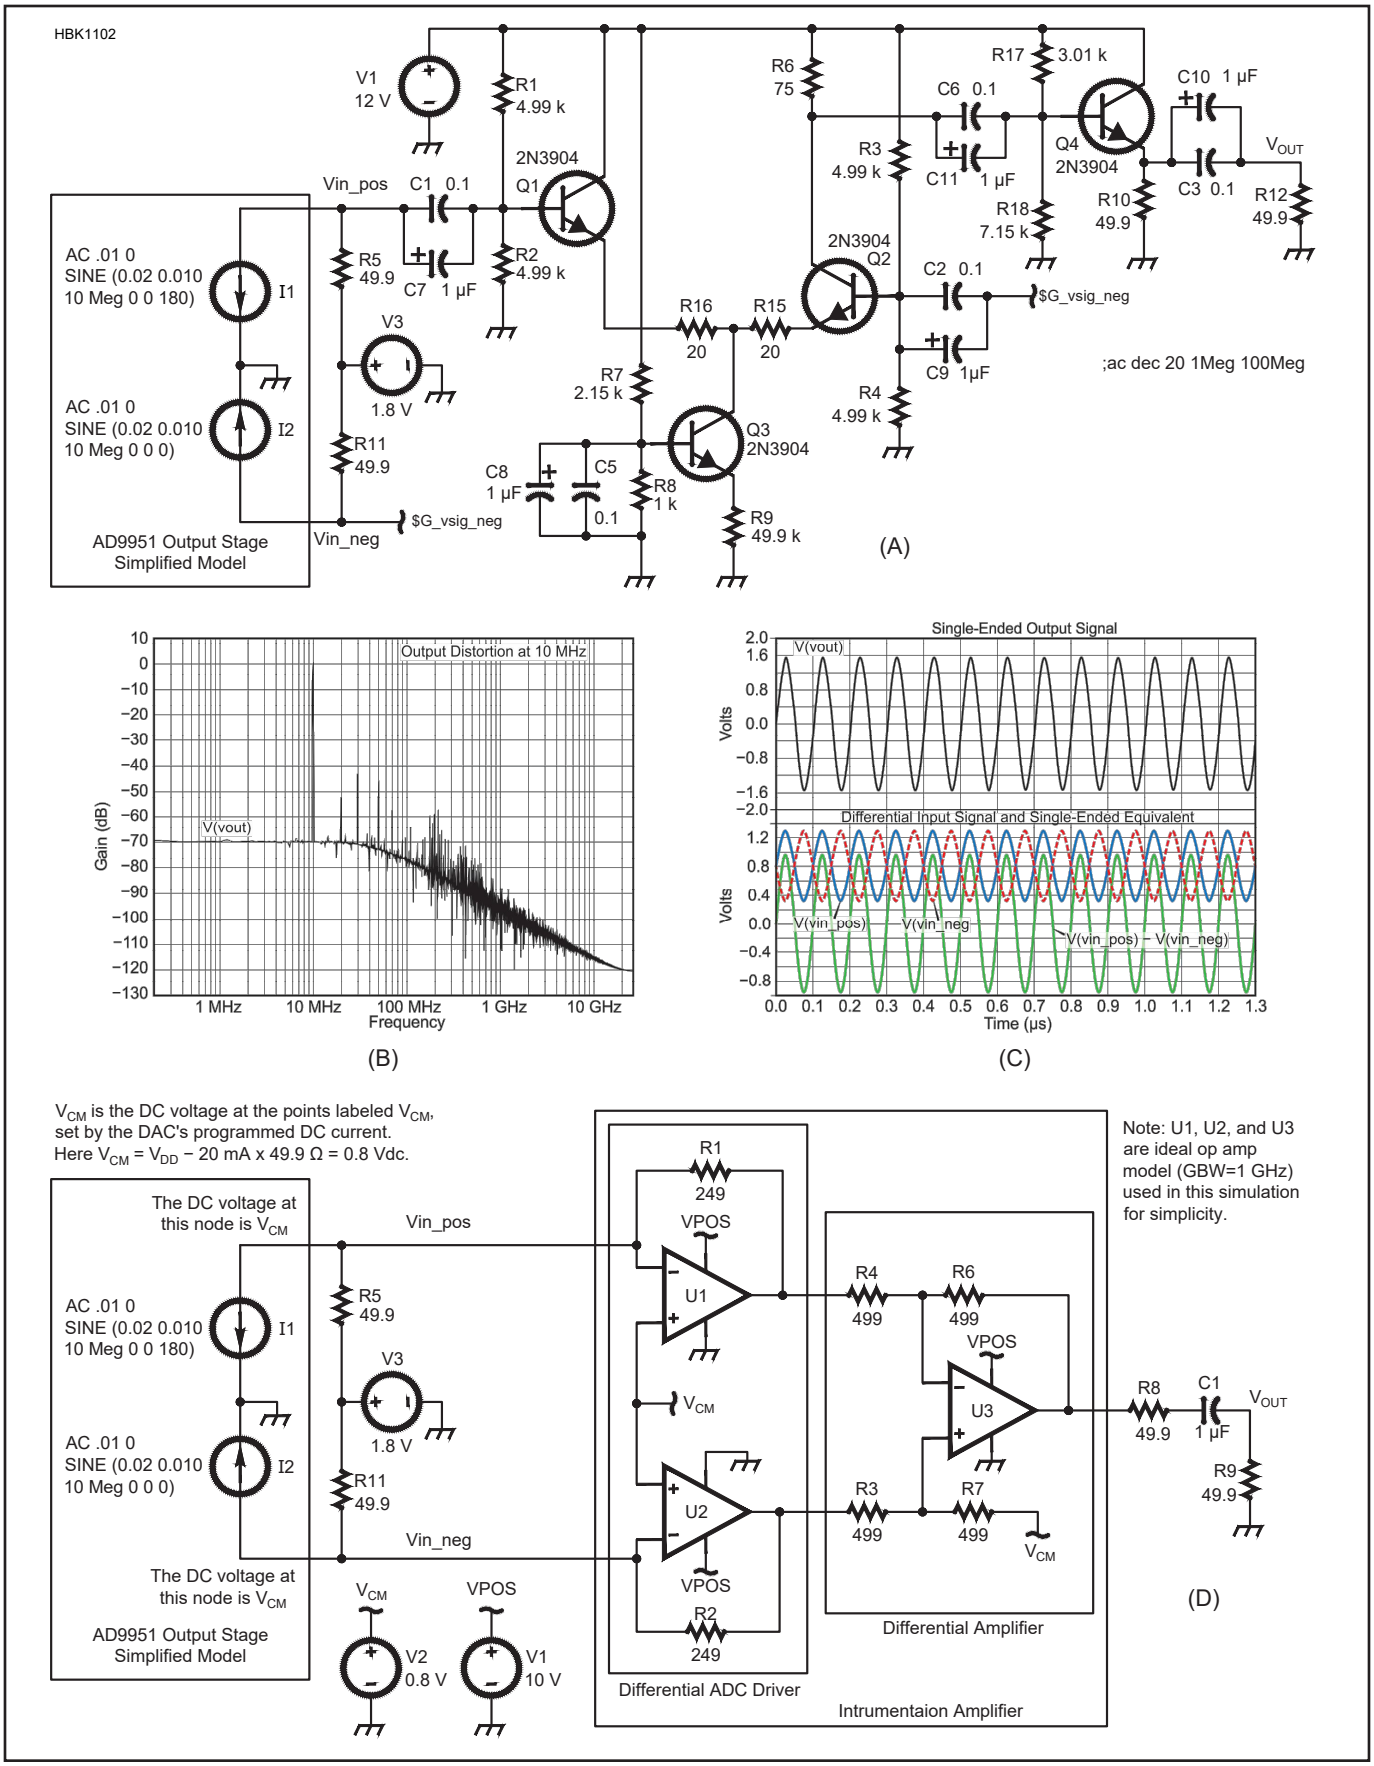

Figure 9.72 — Active differential-to-single-ended output topologies for audio through VHF operation of a DDS. See text for discussion of individual figures.

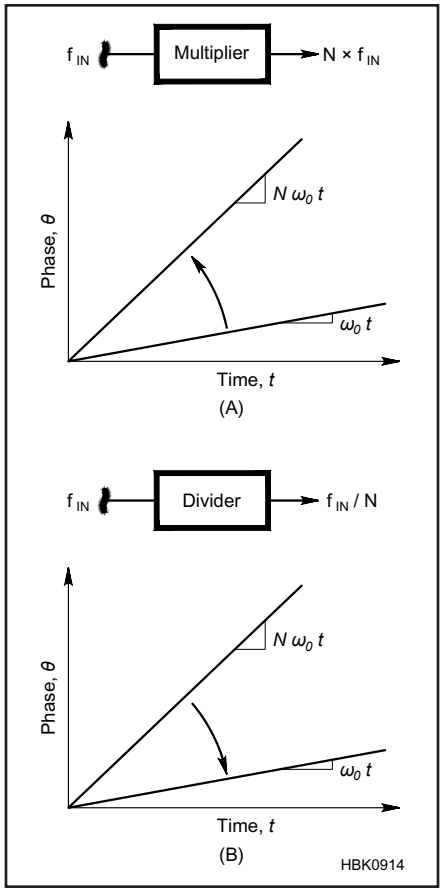

Figure 9.73 — Frequency multiplication and division: A — multiplication increases the rate that signal phase changes with time. This is often an integer (harmonic generation), but it can be non-integer in some applications (PLL); B — frequency division decreases the rate that signal phase changes with time. This is nearly always an integer.

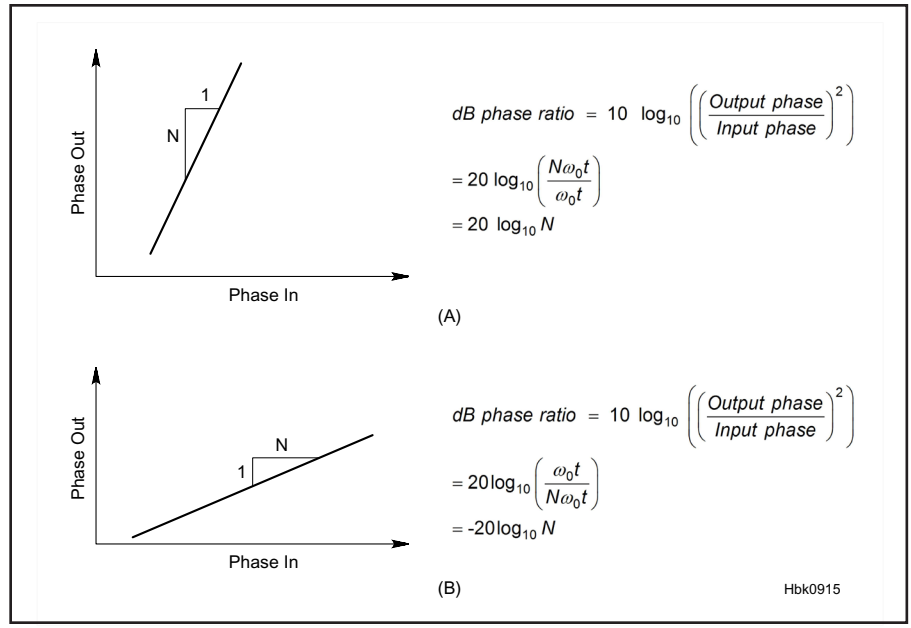

Figure 9.74 — Noise behavior of frequency multipliers and dividers: A — frequency multiplication increases output phase noise; B — frequency multiplication decreases output phase noise.

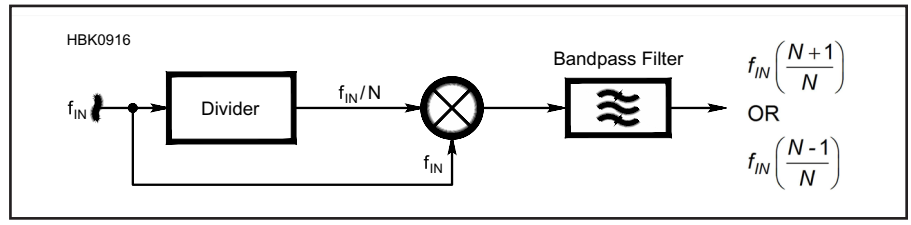

Figure 9.75 — One example of frequency mixing, where a signal is mixed with a divided version of itself to get a frequency ratio.

ADCs. Amplifiers A1 and A2 work as a pair of I/V converters, that is, current in, voltage out. The feedback resistors hold the amplifier's input voltages almost constant, which in turn stabilizes the voltage output of the DAC. A constant output voltage (see Figure 9.72E) minimizes the effect of voltage modulation on the DAC's output devices, which reduces distortion. It also maximizes output bandwidth: no change in voltage means any output capacitors have to be charged and discharged, minimizing their effect on distortion.

Figure 9.72D's circuit actually uses two standard circuits: a differential I/O ADC driver and a single op amp difference amplifier. Together, the two circuits form a third standard circuit: the classic three-op-amp instrumentation amplifier. This instrumentation amplifier, however, has a slightly unusual input configuration: the usual gain resistor is missing and in its place the instrumentation amplifier is connected to a pair of current inputs. In a real version of this circuit, op amps A1 and A2 would be current-feedback (transimpedance) op amps or a wideband-ADC driver and A3 would be a voltage-feedback op amp.

Although passive components have been covered elsewhere in this book, it doesn't hurt to list a summary of the advantages and disadvantages of different capacitor types (**Table 9.3**) from Analog Device's tutorial MT-101. Both ceramic (type X7R) and polystyrene capacitors have been popular in VFO tank circuits for HF. Many capacitor vendors (Murata and Toko, for example) have a line of ceramic capacitors designed specifically for RF applications. Murata, Toko, and Coilcraft, among others, also have inductors specifically for RF applications.

Ceramic capacitors are also used for grounding and bypassing. When drawing up schematics, consult the pin list on the data sheet, and determine a pin's function (input, output, analog or digital, power or ground) and its impedance (high impedance logic, or low impedance analog, etc.) and for each pin determine what should be connected to it and if bypassing is required.

At this point, review the schematic and layout of the evaluation board and see what the manufacturer does. Copy the relevant portions of the circuit, such as grounding, bypassing, routing the RF circuits, etc. Use the layout of the known-good evaluation board as a starting point in your board design. Do not be tempted to "improve" the layout by changing the routing other than removing what are obviously extra parts used only for evaluation.

The last component in a frequency synthesizer is the VCO. Although the text thus far has addressed both discrete (DIY) and integrated VCOs, many commercial products use VCO ICs or modules. **Table 9.4** lists some vendors.

Table 9.4

Manufacturers of VCO ICs and Modules

Manufacturer	*Website*
Analog Devices, Inc.	**www.analog.com**
Mini Circuits	**www.minicircuits.com**
Crystek Microwave, Inc.	**www.crystek.com**
Macom Technology Solutions	**www.macom.com**
Raltron	**www.raltron.com**
Synergy Microwave	**www.synergymicrowave.com**

Last but not least, there are many, many free tools for the designer and hobbyist provided by individuals, manufacturers, and universities. **Table 9.5** lists several, including some that have been mentioned previously.

Many of the preceding circuits can be used to generate a clock for an ADC. To conclude this section, it's worth looking at how the clock requirements are calculated. Oscillator signals for RF circuits are typically specified by their phase noise, shown as a curve in dBc/Hz below the LO across the spectrum of interest in the frequency domain. Jitter is how the total phase noise affects the same oscillator in the time domain and is measured in seconds.

For example, given an RMS jitter t_j in seconds and a sampling rate f_s in Hz, the SNR of an ADC is given by

$$SNR = 20log_{10}\left(\frac{1}{2\pi f_s t_j}\right)$$

Here, the time jitter has been calculated for the Nyquist frequency, or half the sampling rate (that is, $f_s/2$).

With a little bit of manipulation of the above equation, the jitter required for a given SNR and sampling frequency can be shown to be

$$t_j = \frac{1}{2\pi f_s 10^{\frac{SNR}{20}}}$$

A typical design starts by choosing an ADC that meets the system's requirements. Using the ADC's SNR specification, the designer next calculates the RMS jitter required to achieve that performance. For instance, when using an audio ADC specified for a 111 dB SNR at a sampling rate (f_s) of 192 kilosamples per second (ksps), the required RMS jitter,

$$t_j = \frac{1}{2\pi f_s 10^{\frac{SNR}{20}}} = \frac{1}{2\pi \times 192,000 \times 10^{\frac{111}{20}}} = 2.34 \ \mu s$$

These numbers would be typical of a SDR using an audio ADC for sampling signals at baseband.

For GHz sampling rates for a direct-sampling radio, the jitter requirements are in femtoseconds, depending on the number of bits in the ADC and the sampling rate. In this case, we'll use a hypothetical ADC with an 80 dB SNR at a sampling rate of 1 Gsps. Using these numbers, the required clock jitter is

$$t_j = \frac{1}{2\pi f_s 10^{\frac{SNR}{20}}} = \frac{1}{2\pi \times 1,000,000,000 \times 10^{\frac{80}{20}}} =$$

$$1.6 \times 10^{-14} s = 16 \, fs$$

The process of calculating time jitter from a phase noise curve is described in Analog Devices tutorials MT-007, "Aperture Time, Aperture Jitter, Aperture Delay Time—Removing the Confusion" and MT-008, "Converting Oscillator Phase Noise to Time Jitter" by Walt Kester, (**www.analog.com** — search for the tutorial IDs).

9.6.9 Analog Frequency Synthesis

Even though the term 'analog' seems old-fashioned in the present age, these techniques still have wide use in modern radio designs. They are briefly covered here.

FREQUENCY MULTIPLICATION

It is useful to view frequency multiplication as an increased slope of phase with time, as shown in **Figure 9.73A**. For harmonic generation, this multiplication factor is an integer. More complicated frequency multipliers, such as phase-locked loops, can readily provide frequency multiplication with non-integer values.

Table 9.5

Free RF Design Tools

Software	*Description*	*Download or Web Page*
ADIsimPLL	PLL design software	**www.analog.com**
ADIsimDAC	DDS design software	Web-based at preceding URL
AppCAD	Many useful RF design tools and calculators: lumped-element baluns, mixer spurs, etc.	**www.broadcom.com**
CPPSIM	System-level simulator for advanced users. Includes NGSpice analog simulator	**www.cppsim.com**
ELSIE (32-bit Windows)	Design and analyze lumped-element filters in the audio through microwave range. The free Student Version allows 7 components	**tonnesoftware.com/elsie.html**
Filter Design Tool	Audio frequency analog LP, BP, HP, Band Stop, and All-Pass filters using op amps	**www.TI.com**
GNU OCTAVE	Solves linear and nonlinear problems numerically using a language that is mostly compatible with Matlab An RF Toolbox is available.	**www.GNU.org**
KiCAD	Open-source schematic capture and layout software	**www.kicad.org**
LTspice	Analog Circuit Simulator	**www.analog.com**
PLLATINUMSIM-SW	PLL design software; TI PLLs only	**www.TI.com**
PSPICE for TI	Analog Circuit Simulator	**www.TI.com**
QUCS (Linux)	Quite Universal Circuit Simulator (QUCS) aims to support all kinds of circuit simulation types, e.g. DC, AC, S-parameter, Harmonic Balance analysis, noise analysis, etc.	**qucs.sourceforge.net**
RF Toolbox	Numerous RF Design Tools and Calculators	**RFtoolbox.dtu.dk** (Danish Technical University)
TinyCAD	Open-source schematic capture and drawing software	**www.tinycad.net**

FREQUENCY DIVISION

In the same way, frequency division is usefully viewed as a reduction in the slope of output signal phase with time. This is shown in Figure 9.73B. Frequency multiplication and division are inverse process of each other.

Noise behavior of frequency multipliers and dividers also has an inverse relationship. Phase noise encounters the same processes that the signal phase encounters. **Figure 9.74** shows that the phase noise out of any frequency multiplier is at least $+20\log_{10}(N)$ higher than the phase noise on the input signal. This is unavoidable. It is also a big problem when the multiplication factor N is large. We are motivated to have input signals into a phase multiplier with very low phase noise.

Fortunately, the output phase noise of a frequency divider is lower than that on its input signal. This effect is also symmetrical. This effect is often used in very low noise synthesizer design.

FREQUENCY MIXING

Mixers implement frequency addition and subtraction. The problem is that mixers do both at the same time, which means that some filtering is required to intentionally select whether the sum frequency, or the difference frequency, proceeds along to later stages of the radio. In **Figure 9.75** a particular example of frequency mixing is presented. Here the output is a close ratio of the input frequency, such a 2/3 or 3/2 when N = 2, or when N = 8 the outputs can be 7/8 or 9/8.

9.7 Phase Noise

(This section deals specifically with phase noise generated by oscillators. A more general discussion of noise can be found in the **RF Techniques** chapter.) No oscillator output signal is perfect. Viewing an oscillator as a filtered-noise generator is relatively modern. The older approach is to think of an oscillator making a pure sine wave with an added, unwanted noise signal. These are just different ways of visualizing the same thing. They are equally valid views which are used interchangeably, depending on which best makes some point clear.

For example, it is instructive to use the pure-sine-wave-plus-noise view to see relationships between AM (amplitude) noise and PM (phase) noise processes, shown in **Figure 9.76**. Adding Gaussian noise to a pure sine wave is usefully modeled using phasors as shown in Figure 9.76B. This generates both AM due to noise, and PM due to noise as shown in Figure 9.76C.

Passing the oscillator output signal through a limiter process, such as an amplifier in compression or through a switching mixer, leaves only the phase noise depicted in Figure 9.76D. Because it is so easy to remove AM noise but not phase noise, there is essentially no discussion about oscillator AM noise. Phase noise is the most important problem.

Phase modulation (PM) and frequency modulation (FM) are closely related. (See the **Modulation** chapter for a detailed description of each.) Phase is the integral of frequency, so phase modulation resembles frequency modulation, where the frequency deviation increases with increasing modulating frequency. Thus, there is no need to talk of "frequency noise" because phase noise already covers it.

A thorough analysis of oscillator phase noise is beyond the scope of this section. However, a detailed, state-of-the-art treatment by Ulrich Rohde, N1UL, of free-running oscillators using nonlinear harmonic-balance techniques and the sources of noise in the Colpitts oscillator circuit is presented in the online information accompanying this book. Also see the Reference entry for Dacus regarding low-noise synthesizer design.

Because of the dynamic range required to measure phase noise, it is one of the most difficult measurements. The section below on ARRL Lab Measurement of Transmitter Phase Noise illustrates the lengths to which one must go to obtain repeatable, reliable measurements. An additional article on measuring receiver phase noise is included with this book's online material.

9.7.1 Effects of Phase Noise

Phase noise becomes a problem when it is more noticeable than other limitations. It degrades all signals, but whether it is important or not depends on the application. For voice signals it sounds like background "hiss" in headphones or speakers. It also limits the dynamic range of receivers with closely separated signals, or receiving signals with widely different input powers.

Phase noise became a significant problem for amateurs when the use of frequency synthesizers supplanted conventional LC VFOs in amateur equipment. For reasons discussed in the Frequency Synthesizers section of this chapter, it is a major task to develop a synthesizer that tunes in steps fine enough for use with SSB and CW operation while competing with the phase-noise performance of a reasonable-quality LC VFO. Many synthesizers fall far short of this target. Along with the problems with frequency synthesizers, phase noise always gets worse at higher frequencies.

General-coverage, up-converting receiver architectures require local oscillators to operate at higher and higher frequencies, aggravating phase noise performance. As SDR techniques expand, however, phase noise performance is improving on both receive and transmit.

9.7.2 Reciprocal Mixing

All mixers are symmetrical, meaning that the output IF signal depends on the characteristics of both input signals: the local oscillator (LO) and the desired signal. A change to either signal shows up at the IF, where there is no way of knowing if the signal characteristics seen are from the input signal itself, or from the LO. We usually assume that the LO is very pure so only input signal modulations show up at the IF.

When there is phase noise on the LO signal, the situation changes. Noise on the LO transfers to *all* input signals at the mixer output as if it was originally present on each input signal and the LO was perfectly clean — the IF circuits can't tell the difference. This process is called *reciprocal mixing*, where noise on the LO appears as noise on the desired output signal. This is a serious limitation on a receiver's weak-signal ability.

How reciprocal mixing of LO phase noise can limit receiver dynamic range is shown in **Figure 9.77**. One possible scenario of band activity is shown as the set of input signals in Figure 9.77A. Figure 9.77B shows the phase noise profile for the LO of this receiver. Reciprocal mixing is illustrated in Figure 9.77C, showing that the LO phase noise is added to each of the signals in the band in proportion to its input power. The total noise seen by the receiver demodulator is the sum of all transferred phase noise profiles at the received frequency. This raises the apparent noise floor of the receiver.

Another receiver problem due to LO phase noise occurs if a very large signal appears in the receiver passband but at a frequency well removed from the desired signal. One example of this is shown in **Figure 9.78**. If the LO phase noise is excessive, then reciprocal mixing with this large signal can completely block the ability to demodulate the desired signals. Indeed, it is possible to make a wide swath of spectrum relatively useless.

Reciprocal mixing in a receiver does not affect the operating ability of other stations. The solution is to reduce the phase noise on the receiver's LO.

9.7.3 A Phase Noise Demonstration

Healthy curiosity demands some form of demonstration so the scale of a problem can

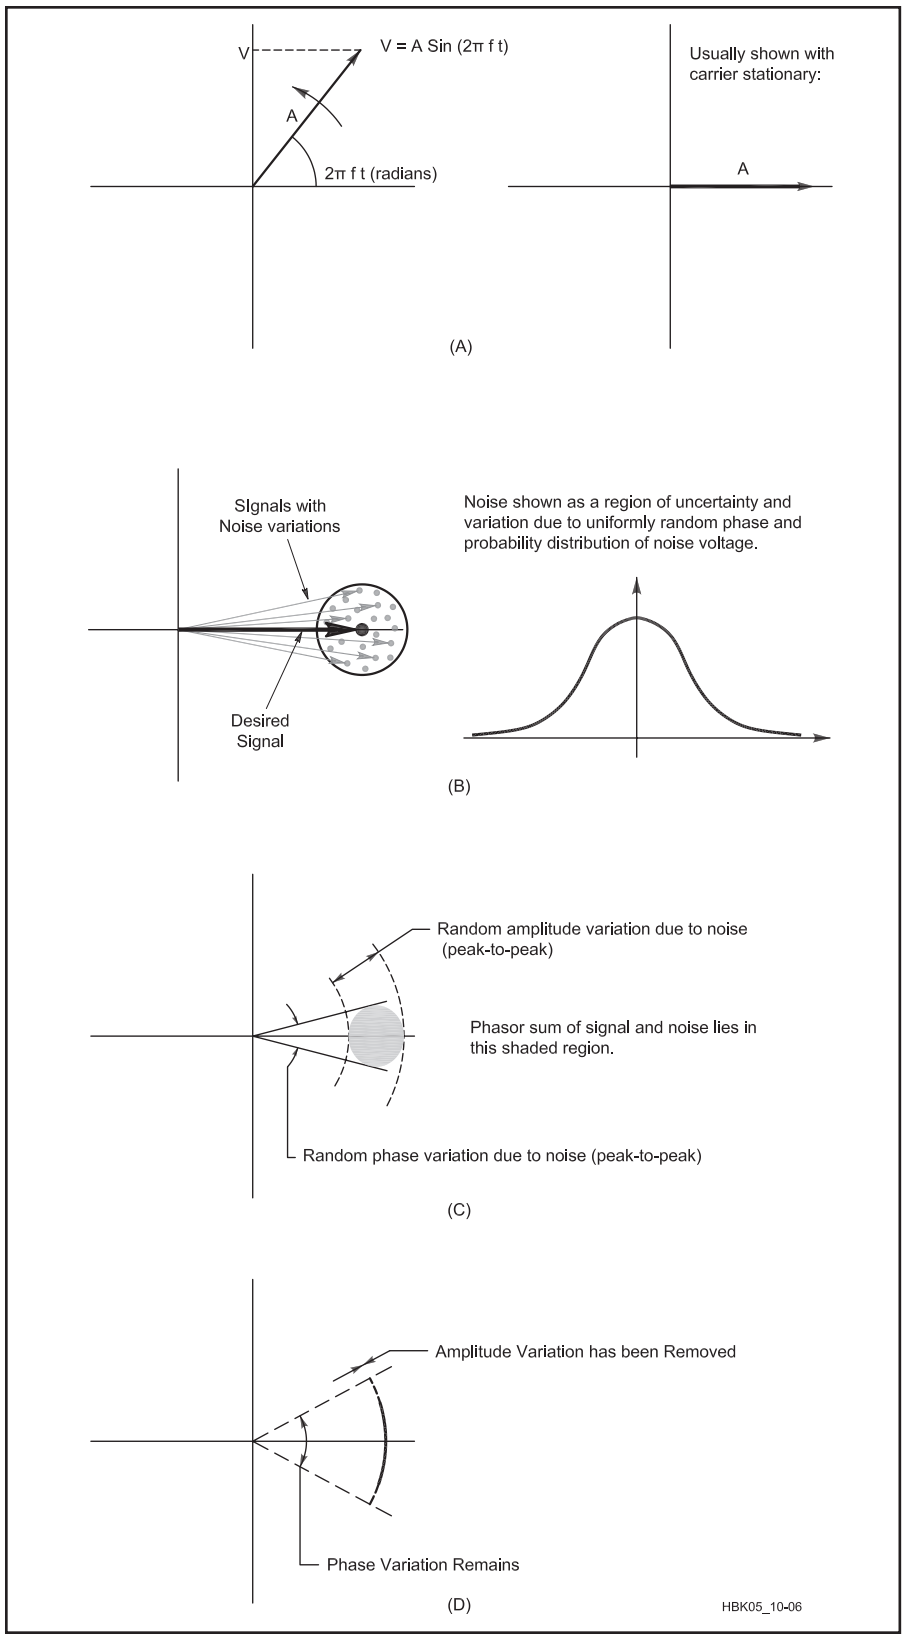

Figure 9.76 — At A, a vector (left) and phasor (right) diagram of an ideal oscillator with no noise. Added noise creates a region of uncertainty in the phasor's length and position (B). AM noise varies the phasor's length; PM noise varies the phasor's relative angular position (C). Limiting a signal that contains both AM and PM noise strips off the AM and leaves the PM (D).

be judged "by ear" before measurements are attempted. We need to be able to measure the noise of an oscillator alone (to aid in the development of quieter ones) and we also need to be able to measure the phase noise of the oscillators in a receiver (a transmitter can be treated as an oscillator). Conveniently, a receiver contains most of the functions needed to demonstrate its own phase noise.

Because reciprocal mixing adds the LO's sidebands to clean incoming signals, in the same proportion to the incoming carrier as they exist with respect to the LO carrier, all we need do is to apply a strong, clean signal wherever we want within the receiver's tuning range. This signal's generator must have lower phase noise than the radio being evaluated. A general-purpose signal generator is unlikely to be good enough; a crystal oscillator is needed.

It's appropriate to set the oscillator's signal level into the receiver to about that of a strong broadcast carrier, say S9 + 40 dB. Set the receiver's mode to SSB or CW and tune around the test signal, looking for an increasing noise floor (higher hiss level) as you tune closer toward the signal, as shown in Figure 9.77D. Switching in a narrow CW filter allows you to hear noise closer to the carrier than is possible with an SSB filter. This is also the technique used to measure a receiver's effective selectivity, and some equipment reviewers kindly publish their plots in this format. *QST* Product Reviews, done by the ARRL Lab, often include the results of specific phase-noise measurements.

9.7.4 Transmitted Phase Noise

Phase noise on an LO used to generate a transmitted signal will also be amplified and transmitted along with the desired output signal. This obscures reception by raising the noise floor even for receivers with low phase noise because they also receive the noise from the transmitter. In this case, the phase noise is generated externally to the receiver and must be removed at the transmitter.

If the transmitter is operating linearly, the strength of the transmitted noise is of the same proportion to transmitter output power as the phase noise is to the oscillator signal power. The noise may even extend well beyond the band in which the desired signal is transmitted unless the signal passes through narrow-band filtering that limits its bandwidth.

This transmitted noise is wasted power that is not useful for communication and unfortunately makes for a noisier band. If you are working a weak station and a nearby transmitter with a noisy oscillator comes on the air with a high-power signal on a different frequency, it is possible that the output noise from this off-frequency transmitter may completely block your ability to continue working that weak station.

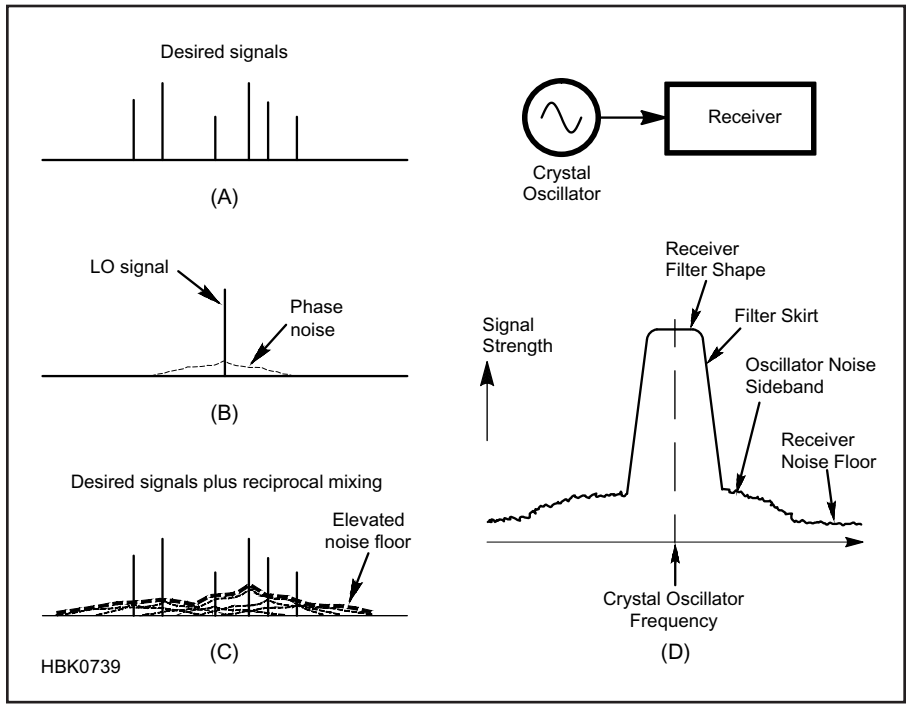

Figure 9.77 — A typical set of input signals is shown at A and an LO signal with phase noise at B. When the LO signal at B is mixed with the input signals at A, the result is a set of mixing products each having phase noise added, raising the noise floor across the band as shown at C. Phase noise in your receiver can be heard by tuning to a strong, clean crystal-oscillator signal as shown in D.

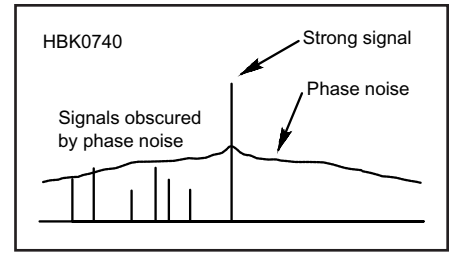

Figure 9.78 — A strong received signal can also cause such severe reciprocal mixing that desired signals are completely obscured by noise. If the signal is transmitted with phase noise from the transmitter LO, the effect is the same as noise covers the desired signals, sometimes across a very wide range of frequencies.

In bad cases, reception of nearby stations can be blocked over many tens of kilohertz above and below the frequency of the offending station. This is seen in Figure 9.80, where the offending signal and its associated noise are from a nearby transmitter. Frequencies close to that of the nearby transmitter are suddenly useless.

Transmitted phase noise can present serious problems for multi-station operation such as at Field Day or during emergency communications where several transmitters and receivers are in close proximity. This is a particular problem with transceivers that use early PLL-synthesized VFOs. If you are planning such an operation, be sure the level of transmitted phase noise is acceptable for the transceivers you plan on using.

At your receiver there is nothing you can do about transmitted phase noise. When there is noise power present at the same frequency as a weak (far) signal you are receiving, your receiver cannot separate them. The only solution is for the problem station to use a transmitter with a "cleaner" LO (less phase noise). This is a more serious problem than reciprocal mixing since this transmitted noise affects the ability of many other stations to use that band.

9.7.5 PLL Synthesizer Phase Noise

Differences in resonator Q usually make the phase-noise sidebands of a loop's reference oscillator much smaller than those of the VCO. Within its loop bandwidth, a PLL acts toward matching the phase-noise components of its VCO to those of the reference. There are several processes which get in the way of this actually happening, which are outlined here.

Dividing the reference oscillator signal f_{XO} to produce the signal f_{REF} which is applied to the phase detector also divides the deviation of the reference oscillator's phase-noise sidebands. (See the previous section on PLL.) This division results in a 20 dB reduction in phase noise per decade of division. This models as a factor of –20 log (R) dB, where R is the reference divisor value, since f_{REF} is less than f_{XO}. We also know that within its loop bandwidth the PLL acts as a frequency multiplier and this multiplies the deviation of the phase noise sidebands present at the PFD, again by 20 dB per decade or equivalently by a factor of 20 log (N) dB, where N is the loop divider's value. Between the reference oscillator and the PLL output, the sidebands on the reference signal are increased by 20 log (N/R) dB.

COUNTER AND PFD NOISE

Logic circuits have a noise characteristic that is extremely important to PLL design. Any circuit, when you really get down to it, is always an analog circuit. In this case the switching threshold of the logic circuits has a tiny amount of noise on it. This small amount of noise manifests itself in a very slight variation in when the logic circuit actually switches, even if the input clock is perfectly clean. We can think of this small timing jitter as a small shift in the output signal "zero crossing," which is directly equivalent to a phase modulation.

The amount of this timing jitter is constant, no matter the operating frequency. Thus, the amount of phase shift this represents depends on the operating frequency: 1 nanosecond of jitter at 1 kHz is very small, but represents 36 degrees at 100 MHz. As a result, we get different measures of the noise floor from digital circuits that depend on the frequency at which they are operating. This digital noise floor interferes with the PLL action to match the VCO noise to the divided reference oscillator noise as measured at the PFD inputs.

VCO NOISE

The PLL can track only slowly moving noise, well within its bandwidth. Therefore, VCO noise components that change slowly enough for the PLL to measure and track will be increasingly canceled successfully. Any VCO noise component that changes at a rate faster than the PLL bandwidth will not be changed at all. At offset frequencies beyond the PLL bandwidth the VCO noise is still present as if the PLL were not even there.

Does this imply that to get a low-noise synthesizer we are encouraged to make the PLL bandwidth very wide? Well, yes, to a certain degree. When we build PLL synthesizers we soon find that the logic noise from the digital section begins to get in the way if the loop bandwidth gets too wide. When this happens the total PLL output noise actually begins to increase as the bandwidth gets wider. There is a range of loop bandwidth values where the total PLL output noise has a minimum. This range is centered at the offset frequency where the logic noise floor, multiplied by 20 log (f_{OUT}/f_{REF}), has the same value as the VCO intrinsic phase noise.

Clearly, having a VCO design with minimum phase noise is very important to a good, low-noise PLL synthesizer. Much effort has been expended on this area in industry. A summary of important results from these noise

reduction efforts is presented in the section Improving VCO Noise Performance below.

FRACTIONAL-DIVISION NOISE

From the brief discussion on fractional-N principles, we note that this whole idea is based on changing the feedback divider value in a semi-randomized way. This inherently jitters the output from the feedback divider, which is a source of phase noise into the PFD. This noise source is algorithmic instead of physical, but the noise is a big problem nonetheless.

OTHER SOURCES OF NOISE

Phase noise can be introduced into a PLL by other means. Any amplifier stages between the VCO and the circuits that follow it (such as the loop divider) will contribute some noise, as will microphonic effects in loop- and reference-filter components (such as those due to the piezoelectric properties of ceramic capacitors and the crystal filters sometimes used for reference-oscillator filtering). Noise on the power supply to the system's active components can modulate the loop. The fundamental and harmonics of the system's ac line supply can be coupled into the VCO directly or by means of ground loops.

9.7.6 Improving VCO Noise Performance

It is tempting to regard VCO design as a matter of coming up with a suitable oscillator topology with a variable capacitor, simply replacing the variable capacitor with a suitable varactor diode and applying a tuning voltage to the diode. Unfortunately, things are not this simple (as if even this is simple!). There is the matter of applying the tuning voltage to the diode without significantly disturbing the oscillator performance. There is also an issue of not introducing parasitic oscillation with the varactor circuit.

Next, as mentioned earlier, one would not like the tuning voltage to drop below the voltage swing in the oscillator tank. If this is allowed to occur, the tuning diodes go into conduction and the oscillator noise gets worse. A good first approach is to use varactor diodes back to back as shown in **Figure 9.79**. This allows the tank voltage swing to be developed across two diodes instead of one, as well as allow for a more balanced loading of the oscillator's tank. The semiconductor industry has realized this and there are a number of varactors available prepackaged in this configuration today.

IMPROVING INTEGRATED VCO OPERATING Q

The IC designer can improve operating Q with a variety of techniques. First, the VCO's tuning range can be divided into multiple, overlapping frequency ranges, with an optimized VCO per range and automatic switching between ranges. For example, improving the IC's metalization by lowering its resistance yields lower-loss inductors with improved Q. Moving the inductor away from lossy materials in the IC by placing the inductor on top of the glass passivation that coats the die also increases the inductor's Q.

Using Metal-Intrinsic-Metal capacitors improves the capacitor's Q. Finally, by multiplying the number of VCO cores and dividing the output frequency, the designer can trade die area and power for an improvement in phase noise. For example, by injection-locking two adjacent VCO cores, they will both oscillate in phase, which gives a 3 dB improvement in their phase noise. Four injection-locked VCOs gives a 6 dB improvement, 8 cores gives a 9 dB improvement, and so forth.

Each division by 2 of the output frequency also improves the phase noise by 3 dB per division. Thus, a division by 2^4 provides a 12 dB improvement. The combination of 8 VCO cores and dividing the output frequency by 8 yields a 24 dB improvement in phase noise at the expense of complexity, die area, and power consumption.

IMPROVING DISCRETE VCO OPERATING Q

Often the Q of passive capacitors available for use in an oscillator tank circuit exceeds the Q of available varactors. One way of reducing the influence of the varactor is to use only the amount of varactor capacitance required to tune the oscillator over its desired range. The balance of the capacitance is supplied to the tank in the form of a higher-Q fixed capacitor. This has the advantage of not requiring the varactor to supply all the capacitance needed to make the circuit function, and often allows for the use of a lower capacitance varactor.

Lower capacitance varactors typically exhibit higher Q values than their larger capacitance counterparts. This concept can be extended by splitting the oscillator tuning range, say 70 MHz to 98 MHz (for a typical lower-side up-converter receiver design popular in recent analog architectures), into multiple bands. For this example, let us consider four oscillators, each with 7 MHz of tuning range. We desire the varactor Q effects to be swamped by high-Q fixed capacitors. This can be further improved through the use of a "segment-tuned VCO" discussed below. A secondary but important benefit of this is to reduce the effective tuning gain (MHz/volt, K_0) of the oscillators, making them less susceptible to other noise voltage sources in the synthesizer loop. These noise sources can come from a variety of places including, but not limited to, varactor leakage current, varactor tuning drive impedance and output noise of the driving operational amplifier or charge pump.

Segment-tuned VCOs provide the designer with additional benefits, but also with additional challenges. By segmented we mean that circuit elements, which could be both inductance and capacitance, are selected for each range that the VCO is expected to tune. **Figure 9.80** shows the frequency range-switching section of a typical VCO in which several diode switches are used to alter the total tank circuit inductance in small steps. These segments create a type of "coarse tuning" for the VCO, and the output of the PLL loop filter performs "fine tuning" with the varactor diode. Usually, the component values are arranged in some sort of binary tree. In some integrated designs, 64 or even 128 sub-bands are available from the VCO.

Many times, the segmented-VCO concept is applied to an oscillator design that is required to cover several octaves in frequency. For example, this would be the case with a single-IF radio with the IF at about 8 MHz. The VCO

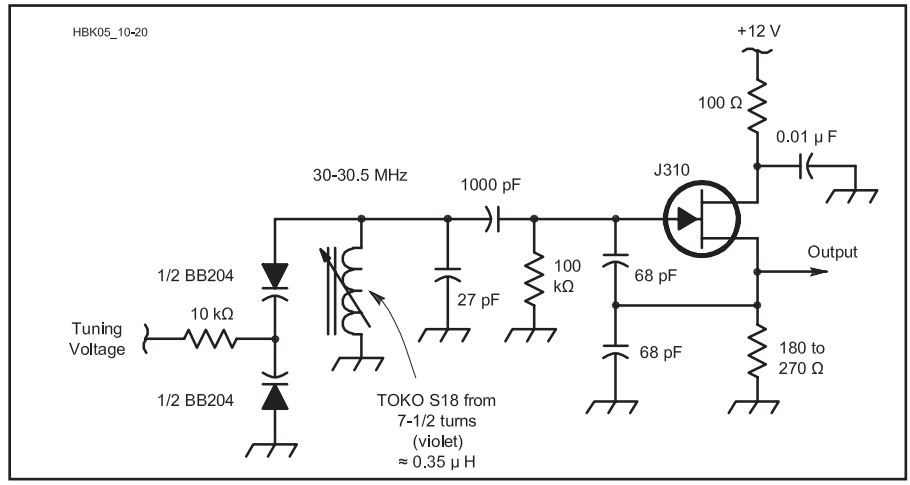

Figure 9.79 — A practical VCO. The tuning diodes are halves of a BB204 dual, common-cathode tuning diode (capacitance per section at 3 V, 39 pF) or equivalent. The ECG617, NTE617 and MV104 are suitable dual-diode substitutes, or use pairs of 1N5451s (39 pF at 4 V) or MV2109s (33 pF at 4 V).

would be required to cover 8 MHz to 48 MHz for 100 kHz to 50 MHz coverage of a typical transceiver.

While segmentation allows one to build a multi-octave tunable oscillator, segmentation also imposes some additional constraints. Recalling the earlier section in this chapter on oscillators, the condition for oscillation is an oscillator loop gain of 1 at a phase angle of 0 degrees. Over several octaves, the gain of the oscillating device (typically a transistor) varies, decreasing as operating frequency increases. While conventional limiting in the oscillator circuit will deal with some of this, it is not good practice to let the normal limiting process handle the entire gain variation. This becomes the role of automatic gain control (AGC) in multi-octave oscillator designs. Application of AGC allows for the maintenance of the oscillation criteria, a uniform output and good starting characteristics, and lower noise across the operating bandwidth.

Finally, a properly designed segmented VCO can improve synthesizer switching speed. The segmentation allows the designer to "pre-steer" the PLL system and reduce the time required for the loop filter to slew the VCO to the desired frequency for lock. Frequency-locked loops (FLL) are often employed for this steering operation to be sure that it happens properly.

9.7.7 ARRL Lab Transmitter Phase Noise Measurement

(This section was written by the ARRL Lab Staff.) Receiver performance has improved dramatically over the past generation or two of transceivers. With the widespread use of better filtering and improvements in digital signal processing, the effects of intermodulation distortion (IMD) at close signal spacing have been reduced greatly. Many software defined receivers (SDRs) now experience little to no reciprocal mixing or reduction of audio level on the received frequency from strong adjacent signals. Signals generated within the receiver during reception of one or more strong nearby signals are not the problem they once were, and it's possible to operate effectively in a crowded band.

Although today's receivers can hear weak signals very close to adjacent strong signals, even the best receivers cannot eliminate the effects of a wide signal from an adjacent transmitter. Transmitted signal issues include excessive keying bandwidth on CW, poor suppression of transmitted IMD products (which causes splatter on SSB), or transmitted phase noise (which raises a receiver's noise floor) in the speaker. ARRL Lab Product Review test reports include transmitter IMD products (both typical and worst case), a plot showing keying sidebands, and a plot showing transmitted phase noise.

PHASE NOISE

Phase noise is essentially the noise generated above and below an oscillator's frequency, also called sideband noise. All oscillators generate some level of phase noise. This is most evident when a close adjacent signal is very strong and the receiver's background noise increases. This is due to the mixing of the first local oscillator's phase noise with the incoming signal at the first IF (reciprocal mixing). The Lab reports the effects of reciprocal mixing in the receiver as "reciprocal mixing dynamic range" or RMDR. In many receivers we've tested, RMDR is the limiting dynamic range. In other words, third order IMD dynamic range and blocking dynamic range are better than RMDR. (See the **Receiving** chapter.)

A transmit oscillator, the heart of the transmitter, has phase noise too. A transmitter's phase noise is a fixed characteristic and, at times, can be a nuisance to other operators. A good example of observed phase noise can happen at Field Day or other environments where several transmitters are operating in close proximity. When two stations are operating on one band, the CW station blasts the phone operator's ears with bursts of noise, and vice versa. Using a transmitter exhibiting high phase noise with an RF amplifier magnifies the problem.

The solution for the reduction of both transmitted phase noise and receiver reciprocal mixing is the employment of high-quality oscillators by the manufacturer. Generally, the better the oscillator (the lower the phase noise), the better the RMDR and the lower the transmitted phase noise.

ARRL LAB TESTING

In the ARRL Laboratory, we use a Rohde & Schwarz FSUP 26 Signal Source Analyzer (**Figure 9.81**), which allows us to measure phase noise on any frequency up to 26.5 GHz. A crystal oscillator with very low phase noise is used to calibrate the system. We use an attenuator after the transmitter to bring the signal down to a suitable level for the phase noise test set.

A sample phase-noise plot for an amateur transceiver is shown in **Figure 9.82**. It was produced with the test setup shown in Figure 9.83. In this case, measurements for 14 MHz and 50 MHz are shown. These plots do not necessarily reflect the phase-noise characteristics of all units of a particular model.

The reference level (the top horizontal line on the scale in the plot) represents 0 dBc/Hz. Because each vertical division represents 20 dB, the plot shows the noise level between 0 dBc/Hz (the top horizontal line) and –180 dBc/Hz (the bottom horizontal line). The horizontal scale is logarithmic, with one decade per division (the first division shows noise from 100 Hz to 1 kHz Hz offset, whereas the last division shows noise from 100 kHz through 1 MHz offset).

WHAT DO THE PHASE-NOISE PLOTS MEAN?

Although they are useful for comparing different radios, plots can also be used to calculate the amount of interference you may receive from a nearby transmitter with known phase-noise characteristics. An approximation is given by

$$A_{QRM} = NL + 10 \times \log (BW)$$

where

A_{QRM} = Interfering signal level, dBc
NL = noise level on the receive frequency, dBc
BW = receiver IF bandwidth, in Hz

For instance, if the noise level is –90 dBc/Hz and you are using a 2.5 kHz SSB filter, the approximate interfering signal will be –56 dBc. In other words, if the transmitted signal is 20 dB over S-9, and each S unit is 6 dB, the interfering signal will be as strong as an S-3 signal.

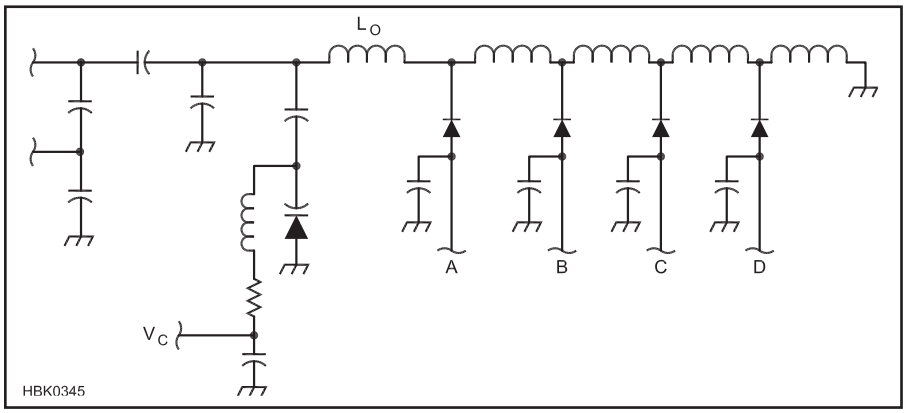

Figure 9.80 — The resonator portion of an inductor-switched segment-tuned VCO. A varactor diode provides tuning over a small range. Diode switches will alter the total inductance, changing the frequency in larger steps. Inductor-switched, capacitor-switched, and combinations are all used in VCO designs. (From Hayward, *Introduction to Radio Frequency Design*, Chapter 7. See references.)

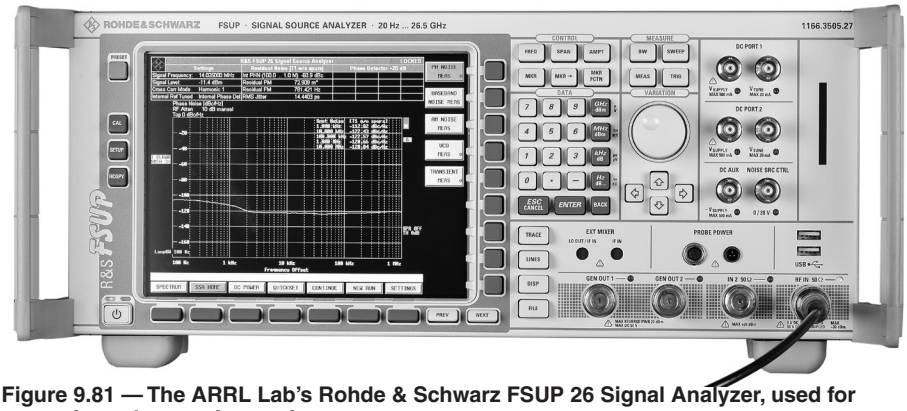

Figure 9.81 — The ARRL Lab's Rohde & Schwarz FSUP 26 Signal Analyzer, used for transmitter phase noise testing.

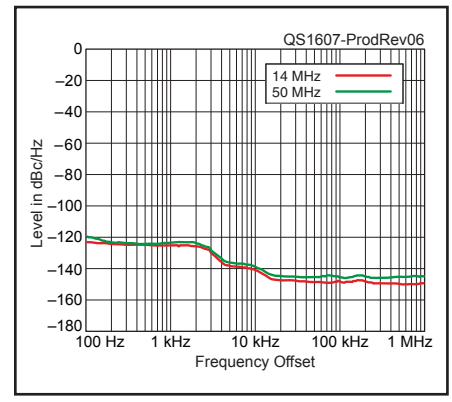

Figure 9.82 — Sample phase noise plot for an amateur HF transceiver as published in Product Review in *QST*. This is the Icom IC-7851.

The measurements made in the ARRL Lab apply only to transmitted signals. It is reasonable to assume that the phase-noise characteristics of most transceivers are similar on transmit and receive because the same oscillators are generally used in the local-oscillator (LO) chain or for generating DDS signals.

In some cases, the receiver may have better phase noise characteristics than the transmitter. Why the possible difference? The most obvious reason is that circuits often perform less than optimally in strong RF fields, as anyone who has experienced RFI problems can tell you. A less obvious reason results from the way that many high-dynamic-range superheterodyne receivers work.

To get good dynamic range in a superhet or in a hybrid receiver with mixers ahead of the DSP stages, a sharp crystal filter called a *roofing filter* is often placed immediately after the first mixer in the receive line. This filter removes all but a small slice of spectrum for further signal processing. If the desired filtered signal is a product of mixing an incoming signal with a noisy oscillator, signals far away from the desired one can end up in this slice. Once this slice of spectrum is obtained, however, unwanted signals cannot be reintroduced, no matter how noisy the oscillators used in further signal processing. As a result, some oscillators in receivers don't affect phase noise.

The difference between this situation and that in transmitters is that crystal filters are seldom used for reduction of phase noise in transmitting because of the high cost involved. Equipment designers have enough trouble getting smooth, click-free break-in operation in transceivers without having to worry about switching crystal filters in and out of circuits at 40 WPM keying speeds!

9.8 References and Bibliography

Analog Devices (**www.analog.com**), "ADI Converter Seminar."

Banerjee, D., *PLL Performance, Simulation, and Design, 6th Edition*, **deanbanerjeepll.com**

Brown, F., W6HPH, "Crystal Test Oscillators," *QEX*, Sep./Oct. 2016, pp. 11 – 13.

Clarke and Hess, *Communications Circuits, Analysis and Design* (Addison-Wesley, 1971). Wide coverage of transistor circuit design, including techniques suited to the design of integrated circuits. Its age shows, but it is especially valuable for its good mathematical treatment of oscillator circuits, covering both frequency- and amplitude-determining mechanisms. Look for a copy at a university library or initiate an interlibrary loan.

Gardner, F., *Phaselock Techniques*, (John Wiley and Sons, 1966).

Grebenkemper, J., KI6WX, "Phase Noise and Its Effects on Amateur Communications," *QST*, Part 1, Mar. 1988, pp. 14 – 20; *Part 2,* Apr. 1988 pp. 22 – 25. Also see Feedback, *QST,* May 1988, p. 44.

Hayward, W., W7ZOI, and DeMaw, D., W1CER, *Solid State Design for the Radio Amateur* (ARRL, 1986). Out of print, this is dated but a good source of RF design ideas.

Hayward, W., W7ZOI, Campbell, R., KK7B, and Larkin, B., W7PUA, *Experimental Methods in RF Design* (ARRL, 2003). A good source of RF design ideas, with good explanation of the reasoning behind design decisions.

Hayward, W., W7ZOI, *Introduction to Radio Frequency Design* (ARRL, 1994). Out of print, good in-depth treatments of circuits and techniques used at RF.

Hewlett-Packard Application Note 150-4, "Spectrum Analysis…Random Noise Measurements." **literature.cdn.keysight.com/litweb/pdf/5952-1147.pdf?id=800596**. Source of correction factors for noise measurements using spectrum analyzers.

Keyser, I., "An Easy to Set Up Amateur Band Synthesizer," *RADCOM* (RSGB), Dec. 1993, pp. 33 – 36.

"Boosting PLL Design Efficiency," Keysight application note 5989-9848EN **www.keysight.com/us/en/assets/7018-01964/application-notes/5989-9848.pdf**. Describes the characterization of VCOs using the E5202B Signal Source Analyzer.

Manassewitsch, V., *Frequency Synthesizers Theory and Design* (Wiley-Interscience, 1987).

Matthys, R., *Crystal Oscillator Circuits* (Wiley, 1983). Mini-Circuits collection of Application Notes on oscillators, synthesizers, and associated components and techniques. **blog.minicircuits.com/engineering-resources/**

Motorola Application Note AN-551, "Tuning Diode Design Techniques" No longer supplied by Freescale but available through various on-line sources. A concise explanation of varactor diodes, their characteristics and use.

Pappenfus. E., Bruene, W., Schoenike, E., *Single Sideband Circuits and Systems*, (McGraw-Hill, 1964). A book on HF SSB transmitters, receivers and accoutrements by Rockwell-Collins staff. Contains chapters on synthesizers and frequency standards. The frequency synthesizer chapter predates the rise of the DDS and other recent techniques, but good information about the effects of synthesizer performance on communi-cations is spread throughout the book.

Parzen, B., and Ballato. A., *Design of Crystal and Other Harmonic Oscillators* (Wiley-

Interscience, 1983). This book shows a thorough treatment of modern crystal oscillators including the optimum design. It is also the only book that looks at the base and collector limiting effect. The second author is affiliated with the US Army Electronics Technology & Devices Laboratory, Fort Monmouth, New Jersey.

Pontius, B., "Measurement of Signal Source Phase Noise with Low-Cost Equipment," *QEX,* May/Jun. 1998, pp. 38 – 49.

Rohde, U., N1UL, *Microwave and Wireless Synthesizers Theory and Design* (John Wiley and Sons, 1997). This book contains the textbook-standard mathematical analyses of frequency synthesizers combined with unusually good insight into what makes a better synthesizer and *a lot* of practical circuits to entertain serious constructors. A good place to look for low-noise circuits and techniques.

Rohde, U., N1UL, Newkirk, D., *RF/Microwave Circuit Design for Wireless Applications* (John Wiley and Sons, 2000). While giving a deep insight into circuit design for wireless applications, Chapter 5 (RF/Wireless Oscillators) shows a complete treatment of both discrete and integrated circuit-based oscillators. Chapter 6 (Wireless Synthesizers) gives more insight in synthesizers and large list of useful references.

Rohde, U., N1UL, Whitaker, J., *Communications Receivers DSP, Software Radios, and Design*, 3rd Edition (McGraw-Hill, 2001). While mostly dedicated to communication receivers, this book has a useful chapter on Frequency Control and Local Oscillators, specifically on Fractional Division Synthesizers.

Rohde, U., N1UL, Poddar, A., Boeck, G., *The Design of Modern Microwave Oscillators for Wireless Applications Theory and Optimization* (Wiley-Interscience, 2005). This is the latest and most advanced textbook on oscillator design for high frequency/microwave application. It shows in detail how to design and optimize high frequency and microwave VCOs. It also contains a thorough analysis of the phase noise as generated in oscillators. While highly mathematical, the appendix gives numerical solutions that can be easily applied to various designs. It covers both bipolar transistors and FETs.

Rohde, U., N1UL, "All About Phase Noise in Oscillators," *QEX*, Part 1, Dec. 1993; Part 2, Jan. 1994; Part 3, Feb. 1994.

Rohde, U., N1UL, "Key Components of Modern Receiver Design," *QST,* Part 1, May 1994; Part 2, Jun. 1994; Part 3 Jul. 1994. Includes discussion of phase-noise reduction techniques in synthesizers and oscillators.

Rohde, U., N1UL, "A High-Performance Hybrid Frequency Synthesizer," *QST*, Mar. 1995, pp. 30 – 38.

Stockton, D., "Polyphase noise-shaping fractional-N frequency synthesizer", US Patent 6,509,800.

Telewski, F., and Drucker, E., "Noise Reduction Method and Apparatus for Phase-locked Loops," US Patent 5,216,387.

Vendelin, G., Pavio, A., Rohde, U., *Microwave Circuit Design Using Linear and Nonlinear Techniques*, 2nd Edition, (Wiley-Interscience, 2005). This is a general handbook on microwave circuit design. It covers (in 200 pages) a large variety of oscillator design problems including microwave applications. It provides a step-by-step design guide for both linear and nonlinear oscillator designs.

Williams, J., *The Art and Science of Analog Circuit Design* (Butterworth-Heinemann, 1995).

Williams, J., *Analog Circuit Design: Art, Science and Personalities* (Butterworth-Heinemann, 1991).

INTERESTING MODERN JOURNAL PUBLICATIONS

Oscillator References

Analog Devices, "Fundamentals of Direct Digital Synthesis," Analog Devices (**www.analog.com**), MT-085.

Berny, A., Niknejad, A., and Meyer, R., "A 1.8 GHz LC VCO with 1.3 GHz Tuning Range and Digital Amplitude Calibration," *IEEE Journal of SSC*, vol. 40.
no. 4, 2005, pp. 909 – 917.

Chenakin, A., "Phase Noise Reduction in Microwave Oscillators," *Microwave Journal*, Oct. 2009, pp. 124 – 140.

Kester, W., "Aperture Time, Aperture Jitter, Aperture Delay Time—Removing the Confusion," Analog Devices, MT-007.

Khanna, A., "Microwave Oscillators: The State of The Technology," *Microwave Journal*, Apr. 2006, pp. 22 – 42.

Nomura, N., Itagaki, M., and Aoyagi, Y., "Small packaged VCSO for 10 Gbit Ethernet Application," IEEE IUFFCS, 2004, pp. 418 – 421.

Poddar, A., "A Novel Approach for Designing Integrated Ultra Low Noise Microwave Wideband Voltage-Controlled Oscillators," Dr.-Ing. Dissertation, TU-Berlin, Germany, 14 Dec. 2004.

Ravi, A., Carlton, B., Banerjee, G., Soumyanath, K., "A 1.4V, 2.4/5.2 GHz, 90 nm CMOS system in a Package Transreceiver for Next Generation WLAN," *IEEE Symp. on VLSI Tech*., 2005.

Rohde, U., N1UL, "A New Efficient Method of Designing Low Noise Microwave Oscillators," Dr.-Ing. Dissertation, TU-Berlin, Germany, 12 Feb 2004.

Rohde, U., N1UL, "Calculation of FM and AM Noise Signals of Colpitts Oscillators in the Time Domain," *QEX*, Mar./Apr. 2016, pp. 22 – 39.

Rohde, U., N1UL, "Some Thoughts on Designing Very High Performance VHF Oscillators," *QEX*, Nov./Dec. 2015, pp. 32 – 40.

Sheng Sun and Lei Zhu, "Guided-Wave Characteristics of Periodically Nonuniform Coupled Microstrip Lines-Even and Odd Modes," *IEEE Transaction on MTT*, Vol. 53, No. 4, Apr. 2005, pp. 1221 – 1227.

Wei, C., Chiu, H., and Feng, W., "An Ultra-Wideband CMOS VCO with 3-5 GHz Tuning Range," IEEE, *International Workshop on Radio-Frequency Integration Technology*, Nov. 2005, pp. 87 – 90.

Yim, M., "Switched Resonators and Their Applications in a Dual Band Monolithic CMOS LC Tuned VCO," *IEEE MTT*, Vol. 54, Jan. 2006, pp. 74 – 81.

Synthesizer References

H. Arora, N. Klemmer, J. C. Morizio, and P. D. Wolf, "Enhanced Phase Noise Modeling of Fractional-N Frequency Synthesizer," *IEEE Trans. Circuits and Systems-I: Regular Papers*, Vol. 52, Feb. 2005, pp. 379 – 395.

van de Beek, R., Leenaerts, D., van der Weide, G., "A Fast-Hopping Single-PLL 3-band UWB Synthesizer in 0.25um SiGe BiCMOS," *Proc. European Solid-State Circuit Conf.* 2005, pp. 173 – 176.

Best, R., *Phase-Locked Loops: Design, Simulation, and Applications*, 5th Edition, McGraw-Hill, 2003.

Dacus, F., "Design Methods of Modern Ultra-Low-Noise Synthesizers," *Microwaves & RF*, Dec. 2018, pp. 36 – 42 (multi-part series).

Lam, C., Razavi, B., "A 2.6/5.2 GHz Frequency Synthesizer in 0.4um CMOS Technology," *IEEE JSSC*, Vol. 35, May 2000, pp. 788 – 79.

Lee, J., "A 3-to-8 GHz Fast-Hopping Frequency Synthesizer in 0.18-um CMOS Technology," *IEEE Journal of SSC*, Vol. 41, No. 3, Mar. 2006, pp. 566 – 573.

Natrajan, A., Komijani, A., and Hajimiri, A., "A Fully 24-GHz Phased-Array Transmitter in CMOS," *IEEE Journal of SSC*, Vol. 40, No. 12, Dec. 2005, pp. 2502 – 2514.

Rahajandraibe, W., Zaid, L., de Beaupre, V., and Bas, G., "Frequency Synthesizer and FSK Modulator For IEEE 802.15.4 Based Applications," *2007 RFIC Symp. Digest*, pp. 229 – 232.

Sandner, C., Wiesbauer, A., "A 3 GHz to 7 GHz Fast-Hopping Frequency Synthesizer For UWB," *Proc. Int. Workshop on Ultra Wideband Systems* 2004, pp. 405 – 409.

Staszewski R., et. al., "All Digital PLL and Transmitter for Mobile Phones," *IEEE J. Solid-State Circuits*, Vol. 40, No. 12, Dec. 2005, pp. 2469 – 2482.