4 Building Blocks 1: Oscillators

Peter Goodson, G4PCF

Oscillators are fundamental to radio and are used in almost all types of radio equipment. The purpose of an oscillator is to generate an output at a specific frequency. For most applications, an oscillator would ideally generate a pure sine wave. If a spectrum of voltage against frequency were to be plotted, it would consist of a single line at the required frequency. **Fig 4.1(a)** shows the spectrum of the 'ideal' oscillator.

No oscillator produces this ideal output, and there are always harmonics, noise and often sidebands in addition to the wanted output frequency. A more realistic plot is shown in **Fig 4.1(b)**. Tuning a receiver across this would show where the problems are. The noise floor (the background level of noise which is equal at all frequencies) is apparent, as is the '1/*f*' noise, which is an increase in noise close to the centre (wanted, or carrier) frequency. The harmonics also show these effects, as do the sidebands and spurious oscillations which often occur in synthesisers because of practical limitations in loop design.

Noise and sidebands are always undesirable, and can be very difficult to remove once generated, so one purpose of this chapter is to indicate ways of minimising them. Harmonics are less of a problem. They are a long way in frequency terms from the wanted signal, and can often be simply filtered out. In some cases, they may even be wanted, since it is often convenient to take an output frequency from an oscillator at a harmonic instead of the fundamental. Examples of this are shown below for crystal oscillators.

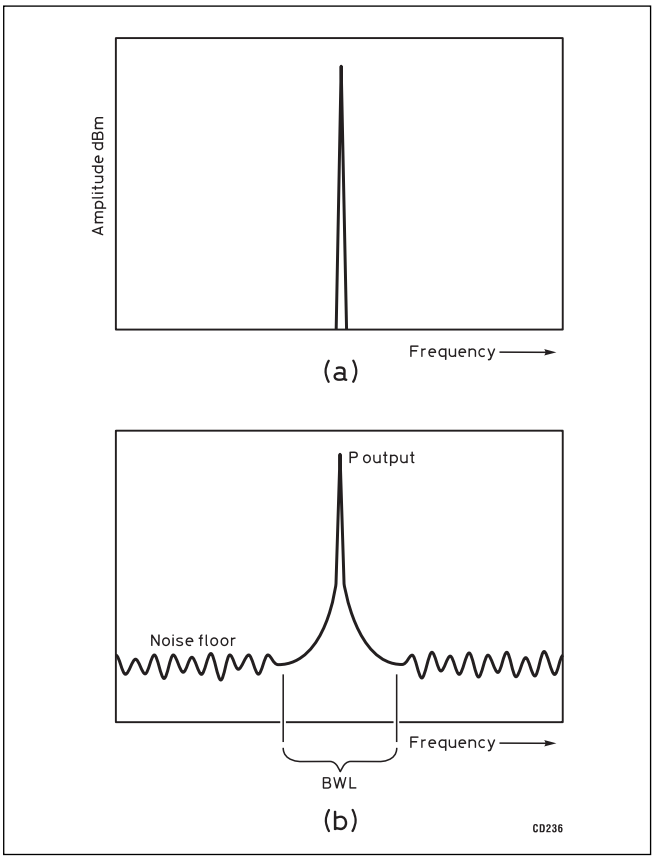

Fig 4.1: (a) Spectrum of 'ideal' oscillator. (b) Spectrum of local oscillator showing noise

Oscillators may be classed as variable frequency oscillators (VFOs), crystal oscillators (COs) (including the class of variable crystal and ceramic resonator oscillators (VXOs)), phase-locked loop synthesisers (PLLs), which include VFOs as part of their system and, more recently, direct digital synthesisers (DDSs). Any or all of these may be used in a particular piece of equipment, although recent trends commercially are to omit any form of 'free-running' (ie not synthesised) VFO. The name is retained, however, and applied to the sum of the oscillators in the equipment. Since these are digitally controlled, it is possible to have two (or more) virtual 'VFOs' in a synthesised rig, where probably only a single variable oscillator exists physically. It can be retuned in milliseconds to any alternative frequency by the PLL control circuit.

All oscillators must obey certain basic design rules. The first essential for an oscillator is an amplifying element. This will be an active device such as a bipolar transistor (sometimes called a 'BJT' for bipolar junction transistor), a junction field-effect transistor (JFET), a metal-oxide semiconductor FET (MOSFET) or a gallium arsenide FET (GaAsFET). There are lesser-used devices such as gallium arsenide bipolars, and gallium arsenide HEMTs (high electron mobility transistors), a variant of the GaAsFET with lower noise.

Operational amplifiers are rarely used at RF, since very few are capable of high-frequency operation. Some integrated circuits do contain oscillator circuits, but they are very often variants of standard discrete circuits. Valves, now long obsolete in the context of oscillator design, will be neglected in this text, but it is worth observing that the fundamental circuit configurations owe their origins to valve designs from the early part of the 20th century. The Colpitts and Hartley oscillators were both first published in 1915. The Colpitts oscillator is named after its developer, Edwin Henry Colpitts (1872-1949) and the Hartley oscillator was developed by Ralph Hartley (1889-1970). For those interested in valve design, reference should be made to earlier editions of this handbook, and to references [1] and [2]. A more recent source is Section 4 of reference [3].

The basic requirements for an oscillator are:

1. There must be gain over the whole frequency range required of the oscillator.
2. There must be a feedback path such that the product of the forward gain and the feedback attenuation still leaves a net loop gain greater than 1.
3. The feedback must occur in such a way that it is in phase with the input to the gain stage.
4. There must be some form of resonant circuit to control the oscillator frequency. This is most often an inductor and capacitor (L-C). More precise frequency control can be achieved using electromechanical structures such as crystals, ceramic resonators or surface acoustic wave (SAW) resonators. This latter category is only really feasible in specialist applications where the high cost can be justified.

At low frequencies, a resistor-capacitor (RC) or inductor-resistor (LR) network may be used for frequency control. These cannot be said to be truly resonant themselves but, with active gain stages, resonant peaks which are sometimes of very high Q can be achieved.

A potential successor to the crystal and ceramic resonator is the electrochemically etched mechanical resonator, usually in the same silicon chip as the oscillator circuit. These devices are still in the development stage, but may well come into common use soon.

An oscillator therefore needs gain, feedback and a resonant circuit. However, there are more criteria that can be applied:

5. For stable oscillations, the resonant circuit should have high Q. This tends to rule out RC and LR oscillators for high frequencies, although they are often used for audio applications. More information on low-frequency oscillators can be found in audio texts [4].
6. For stable oscillations, the active components should also show a minimum of loading on the resonant circuit, ie the loaded Q should be as high as possible. Although an oscillator could be said to be the ultimate in Q-multipliers, high inherent Q gives better stability to the final product.
7. The choice of active device is critical in achieving low-noise operation. This will be covered extensively below, since in almost all cases low-noise oscillations are required, and the difference between 'noisy' and 'quiet' oscillators may be only in the choice of a transistor and its matching circuit.
8. The output stage of an oscillator can be very important. It must drive its load adequately, without changing the loading of the oscillator and thereby changing the oscillator frequency (an effect known as load pulling).
9. The power supply to an oscillator should be very carefully considered. It can affect stability, noise performance and output amplitude. Separate supply decoupling and regulation arrangements are usually essential in good oscillator design.

EFFECT OF VFO PERFORMANCE ON RECEIVER PERFORMANCE

When an oscillator is used as the first local oscillator in a receiver, its performance has a bearing on the receiver performance. The most obvious of these is frequency drift, but other oscillator parameters include:

- The oscillator noise sideband performance will affect the ability of the receiver to resolve a weak signal when there are nearby strong signals.
- The oscillator noise floor may affect the ability of the receiver to resolve a weak signal when there are strong signals far removed (in frequency) from the wanted signal.
- A high level of harmonics from an oscillator may increase the spurious responses in the receiver.
- The value of VFO buffer output impedance may be important to enable some mixers to give their best performance.
- If the oscillator signal is picked up by the antenna it may cause a degradation of receiver performance. For this reason it may be necessary to screen the oscillator.

For a high performance receiver (a receiver which has excellent intermodulation performance) an oscillator with good noise performance will be required or else the work and cost put into achieving the intermodulation performance will be wasted.

However, if the receiver is one of fairly modest performance, it is not necessary to spend a lot of time or money on a high performance oscillator. This shows that the performance of the receiver signal path and oscillator need to be matched. This subject is covered in more detail in the chapter on HF Receivers.

VARIABLE-FREQUENCY OSCILLATORS

A VFO is a type of oscillator in which the oscillation frequency is adjustable by the operator. Normally this is done by the tuning control on the front panel of the equipment. This control operates through a variable capacitor or variable inductor to control the oscillation frequency. Most VFOs consist of a single active device with a tuned circuit and a feedback network to sustain the oscillation.

Two such examples are shown in **Figs 4.2 and 4.3**. Some comments on these diagrams will show the design compromises in action. Fig 4.2 uses a tapped inductor and is known as the Hartley oscillator. The diagram shows the basic circuit and does not include DC bias components. In the Hartley oscillator of Fig 4.2, the tuned circuit is designed for high Q. The active device is an FET, chosen for good gain at the intended frequency of oscillation. A dual-gate MOSFET could be used or, especially at UHF and microwave frequencies, a GaAsFET. The high input impedance of the FET puts little loading onto the tuned circuit, while the FET source provides feedback at low impedance. The fed-back voltage is transformed up to a higher impedance and appropriate phase, so the oscillation loop is complete. This is one example where the active device provides only impedance matching to the oscillator loop; no voltage gain or phase shift in the FET is involved.

Fig 4.3 shows the capacitively tapped version of basically the same circuit. This is known as the Colpitts oscillator. Capacitive transformers are less easy to understand intuitively but, if one thinks of the whole as a high-impedance resonant circuit, then a capacitive tap is a reasonable alternative. While in principle the two circuits are very similar, in practice the additional capacitors of the Colpitts make the circuit less easy to tune over a wide range. In the Hartley circuit, with careful design and a low-input-capacitance FET, a very wide range of total circuit capacitance variation (and hence wide frequency range) can be achieved. The Colpitts circuit is preferred at

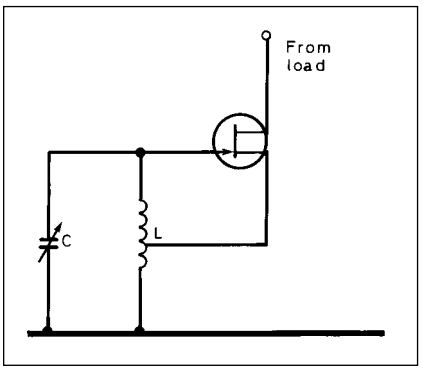

Fig 4.2: Hartley oscillator

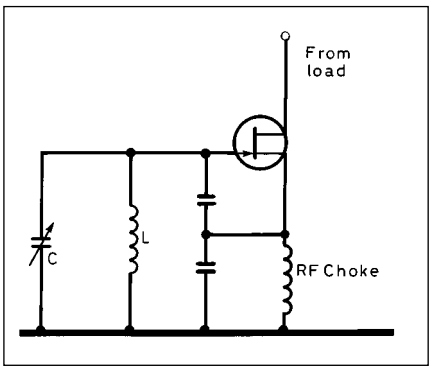

Fig 4.3: Colpitts oscillator

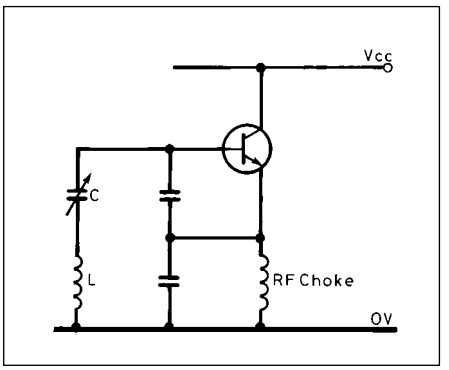

Fig 4.4: Clapp oscillator

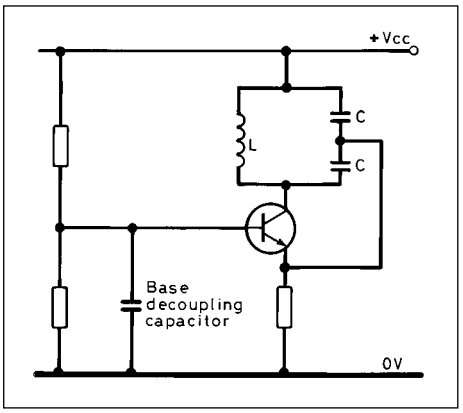

Fig 4.5: 180 degree phase-shift oscillator

higher frequencies, since the transformer action of the inductor is much reduced where the 'coil' is in the form of a straight wire, say at 500MHz or above.

Sometimes described as a variant of the Colpitts circuit is the Clapp oscillator: **Fig 4.4**. This shows a series-resonant circuit, which is more easily matched into the bipolar transistor shown. This circuit is especially suitable to UHF and low microwave work, where it is capable of operating very close to the cut-off frequency (Ft) of the transistor.

Variants of both the Hartley and Colpitts circuits have been derived for feedback around the active device including a 180° phase shift. An example is shown in **Fig 4.5** and this is again more suited to a bipolar transistor.

In the technical press, there has been much discussion on whether the best choice for the active device in an oscillator is a bipolar transistor or an FET. There are many factors involved in this choice, but some simple rules can be derived. First, since the loaded Q must be maintained as high as possible, an FET is attractive for its high input impedance. A bipolar transistor could be used as an emitter follower with an undecoupled emitter resistor but a noise analysis of the circuit would put all of that resistance in the noise path. As an example, **Figs 4.6 and 4.7** compare noise sources in oscillator (or amplifier) input circuits. The FET, provided it is operated well within its frequency range, has a high input impedance. Any input capacitance is absorbed into the tuned circuit fairly directly, since it looks just like a capacitor with at most an ohm or two of input series resistance and perhaps 1 or 2nH of series inductance from the bondwire on the chip. The output impedance at the source, ie in a follower, is $1/g_m$, (plus a small ohmic resistance term) where g_m is the mutual conductance at the operating frequency. As a gain stage, the output impedance of most FETs is very high, so the gain is determined by gm and the load impedance. An exception to this is the GaAs MESFET, where the output impedance is typically a few hundred ohms, so there is a serious limit to the gain available per stage. Some silicon junction FETs (JFETs) have low g_m, so it is worth choosing the device carefully for the frequency used. GaAsFETs tend to have much higher g_m, especially at high frequencies, and so are recommended for microwave work, with the proviso that they are rather prone to $1/f$ noise. This will be described below.

Turning now to the bipolar device, it can be seen that the input impedance is inherently low. At low frequencies, the input impedance is approximately h_{fe} (the AC current gain) multiplied by the emitter resistance ($R_e + r_e$). This term is comprised of R_e, the $1/g_m$ term as in the FET, and r_e, the ohmic series resistance.

This latter can be several ohms, while the former is determined by the current through the transistor; at room temperature, $R_e = 26/I_e$, where I_e is the emitter current in milliamps. Thus at, say, 10mA, the total emitter resistance of a transistor may be an ohm or two, and the input impedance say 50 to 100Ω. This will severely degrade the Q of most parallel-tuned circuits. There is no advantage in using a smaller emitter current, since this leads to lower F_t and hence lower gain; as the frequency approaches F_t, the AC current gain is degraded to unity by definition at F_t. For completeness, the bipolar noise sources are included. In this respect, the best silicon bipolar transistors compare roughly equally with the best silicon FETs. GaAs devices tend to be better again, especially above 1GHz, but the effect of $1/f$ noise has to be considered, so that in an oscillator (but not an amplifier) a silicon device has many advantages right into the microwave region. GaAs devices only predominate because most of them are designed for amplifier service; if sub-0.5 micron geometry, silicon discrete FETs became commercially available, they could become the mainstay of oscillators to 10GHz and beyond.

So what is this $1/f$ noise, and how does noise affect an oscillator? Noise is more familiarly the province of the low-noise amplifier builders, but the principles are the same for oscillators. Where noise comes into the receiver context is that if the signal to be generated is the local oscillator (LO) in a receiver, then most or all of the local oscillator noise is modulated onto the wanted signal at the intermediate frequency (IF). There are several mechanisms for this, including the straightforward modulation of the wanted signal, and the intrusion in the IF of recip-

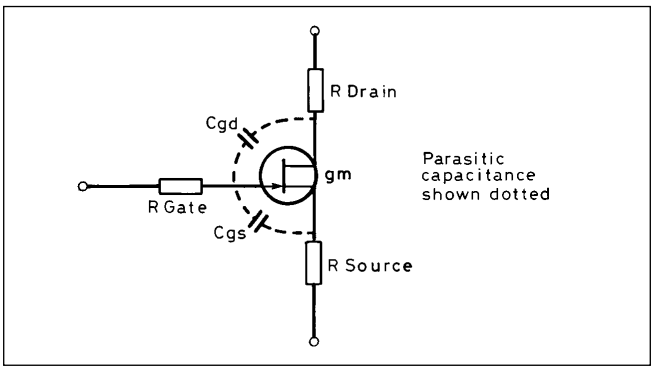

Fig 4.6: FET input characteristics. Input resistance is very high. Input capacitance is approximately equal to C_{gs} + (voltage gain) C_{gd}. Input equivalent noise resistance is approximately equal to R_{GATE} + R_{SOURCE} + $1/g_m$. Typically R_{GATE} is less than 10Ω and R_{SOURCE} is approximately equal to R_{DRAIN} and less than 100Ω. Power devices may be less than 1Ω

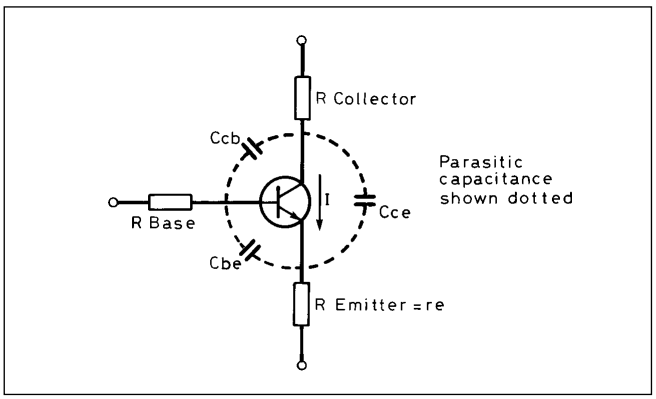

Fig 4.7: Bipolar input characteristics. Input resistance is approximately R_{BASE} + h_{fe} x (r_e + R_e) where R_e = $1/g_m$ and g_m = (q/kT) x I. Input capacitance is very approximately equal to C_{be} + V_{gain} x C_{cb} but transistor effects will usually increase this dynamically. Input equivalent noise resistance (R_{IN}) is equal to R_{BASE} + (R_e + r_e)/2. Typically, for a small device, R_{BASE} is about 100Ω, R_e is 26Ω at 1mA, and r_e is 3Ω. Therefore R_{IN} is approximately 129Ω

rocal mixing (see the chapter on HF receivers), where a strong but unwanted signal, which is close to the wanted one within the front-end pass-band, mixes with LO noise. A low-noise oscillator is therefore a major contributor to a low-noise receiver. On transmit, the effects may not be so obvious to the operator, but the transmission of noisy sidebands at potentially high powers is inconvenient to other band users, and may in extreme cases cause transmissions outside the band. In practice, an oscillator with low enough noise for reception is unlikely to be a problem on transmit.

All devices, active and passive, generate noise when not at absolute zero temperature. Inductors and capacitors only do so through their non-ideal resistive terms, so can be neglected in all practical cases. Resistors and transistors (FET and bipolar) are the real sources of noise in the circuit. The noise power generated by a resistor is:

$$\text{Noise power} = kTB$$

where k is Boltzmann's constant, a fundamental constant in the laws of physics; T is the absolute temperature (room temperature is usually approximated to 300K, ie 27°C); and B is the measurement bandwidth.

This equation is sufficiently accurate for most modern resistors like metal film and modern carbon film resistors. Note that some older types, notably carbon-composition types, do generate additional noise. Active devices are less so, and an equivalent noise resistance (R_{IN} in Fig 4.6 and 4.7) can be derived or measured for any active device which approximates the device to a resistor.

This may have a value close to the metallic resistance around the circuit, or it may be greater. For many applications, this resistance can be used to estimate the noise contribution of the whole circuit. However, there are other noise sources in the device, most noticeably in the oscillator context those due to $1/f$. Fig 4.1(b) shows $1/f$ noise diagrammatically. Noise sources of this type have been studied for many years in many different types of device. Classically, this noise source increases as frequency is reduced. It was thought that there had to be a turnover somewhere in frequency, but this need not be the case.

If we take a DC power supply, it has a voltage at the output, but it has not always been so. There was a time before the supply was switched on, so the 'noise' is infinite, ie $1/f$ holds good at least in qualitative terms. This gives rise to the concept of noise within a finite bandwidth, usually 1Hz for specification purposes. Practical measurements are made in a sensible bandwidth, say a few kilohertz, and then scaled appropriately to 1Hz. The goodness of an oscillator is therefore measured in noise power per hertz of bandwidth at a specified offset frequency from the carrier. A good crystal oscillator at 10MHz would show a noise level of -130dBc (decibels below carrier) at 10kHz offset. A top-class professional source might be -150dBc. There are several possible causes of $1/f$ noise and, when examining the output of an oscillator closely, there is invariably a region close to the carrier where the noise increases as $1/f$. This is clearly a modulation effect, but where does it come from?

The answer is that most (but not all) oscillators are very non-linear circuits. This is inherent in the design. Since it is necessary to have a net gain around the oscillation loop (which is the gain of gain stage(s) divided by the loss of feedback path), then the signal in the oscillator must grow. Suppose the net gain is 2. Then the signal after one pass round the loop is twice what it was, after a further pass four times and so on. After many passes, the amplitude should be 'infinite'. There clearly must be a practical limit. This is usually provided by voltage limitations due to power supply rails, or gain reduction due to overdrive and saturation of the gain stage(s). Occasionally, in a badly designed circuit, component breakdown can occur where breakdown would not be present under DC conditions. This can be disastrous as in a burn-out or, more sinister, as in base-emitter breakdown of bipolar transistors, which will lead to permanent reduction of transistor gain and eventual device failure in service.

So, almost all practical oscillators are highly non-linear and have sources of noise, be they wide-band (white noise) or $1/f$ (pink noise). The non-linearity thus modulates the noise onto the output frequency. In contrast, a properly designed amplifier is not non-linear, has very little modulation process, and is therefore only subject to additive noise.

The exceptions to this are the very few oscillators which have a soft limitation in the output amplitudes, ie they are inherently sine wave producers. This requires very fast-acting AGC circuitry, and is very rarely used because it adds much cost and complication. Most practical oscillators are inherently square-wave generators but a tuned circuit in the output will normally remove most harmonic energy so that the output looks like a sine wave on an oscilloscope. A spectrum analyser normally gives the game away, with harmonics clearly visible. The DDS devices are an exception to all this for reasons explained below.

Generally, a silicon device will have very much lower levels of $1/f$ noise than a gallium arsenide device. This is because silicon is a very homogeneous material, of very high purity, and in modern processes with very little surface contamination or surface states. Gallium arsenide is a heterogeneous material which is subject to surface states. The $1/f$ knee is the frequency below which the noise levels associated with the surface states start to increase, typically at 6dB/octave of frequency. In, say, 1980 this frequency would have been 100MHz or more for most gallium arsenide processes. In 1990 the figure was as low as 5MHz in some processes. Silicon typically has a $1/f$ knee of less than 100Hz, sometimes less than 10Hz, so $1/f$ noise is unimportant for most practical cases. It means that there is a possibility of increased close-to-carrier noise in GaAs-based oscillators. Whether this makes the oscillator better or worse than a silicon design depends on the exact choice of devices and circuits. More information on these topics can be found in [5].

PRACTICAL VFOs

A VFO design should start with a set of clear objectives, such as:

- The required frequency tuning range.
- The exact means of tuning.
- The requirement for frequency stability.
- The RF power output required from the oscillator. The RF level is important because most mixers have a minimum oscillator power requirement. This is the oscillator power that will give the specified mixer performance.
- For a high performance receiver it may be necessary to define the oscillator noise performance.

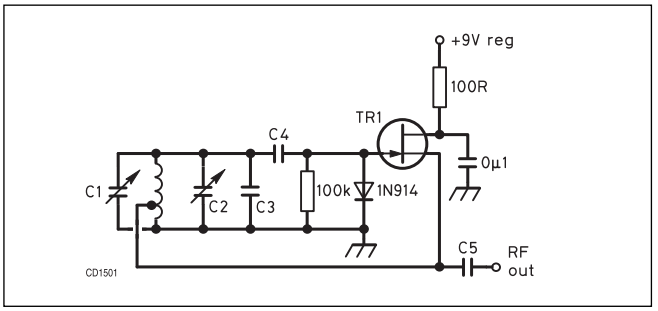

Fig 4.8: Hartley VFO (*W1FB's QRP Notebook*)

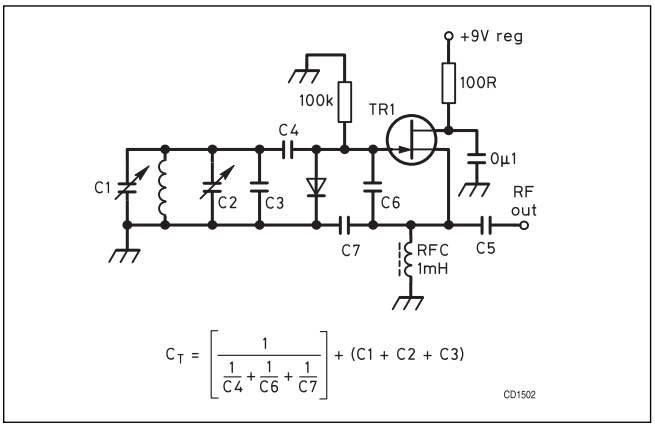

Fig 4.9: Colpitts VFO (*W1FB's QRP Notebook*)

A typical VFO design objective is as follows:

1. Frequency range 5.0 to 5.5 MHz. The VFO therefore tunes over 500kHz
2. The VFO is required to be tuned by a variable capacitor.
3. The VFO is required to tune an SSB / CW transceiver. Therefore this requires that the frequency should not drift by more than 100Hz in an hour, and should always be repeatable to within 200Hz.
4. RF Power output needs to be 10mW.

The frequency range required is determined by the band that is required to be covered, and the intended IF. In this example the use of a VFO covering 5 - 5.5MHz allows the 80m and 20m bands to be tuned, using an IF of 9MHz.

The third requirement is defined by the fact that an SSB receiver needs to be tuned to within about 50Hz of the correct frequency. This puts a difficult requirement on VFO stability because small changes in temperature caused by, for example switching on a heater in the room, can cause considerably more change than 50Hz.

Figs 4.8 to 4.10 show a range of practical VFO circuits, which include the DC bias components. In each circuit, C1 is the main tuning capacitor and C2 provides bandspread. The circuit in Fig 4.8 is a Hartley oscillator. The feedback tap should attach at about 25% of the coil from the earthy end; actually, the lower down the better, consistent with easy starting over the whole tuning range. A bipolar transistor could be used in this circuit with bias modification, but it will tend to reduce circuit Q and hence be less stable than a FET. The FET should be chosen for low noise, and JFETs are usually preferred. However this circuit will also work well with modern

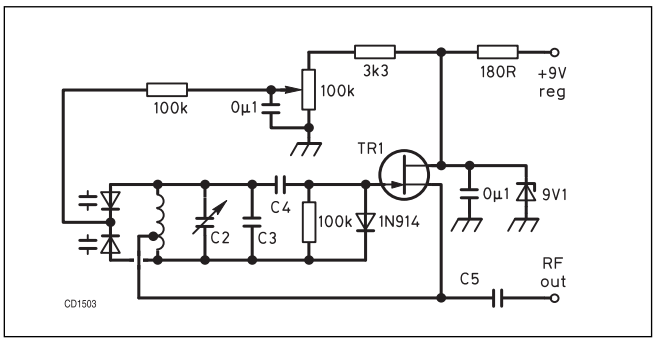

Fig 4.10: Varactor-tuned Hartley VFO. Suggested values for reactances are: X_{C1} 1200Ω; X_{C2} 3500Ω; X_{C3}, X_{C4}, X_{C5} 880Ω; X_{C6}, X_{C7} 90Ω; X_{L1} 50Ω (*W1FB's QRP Notebook*)

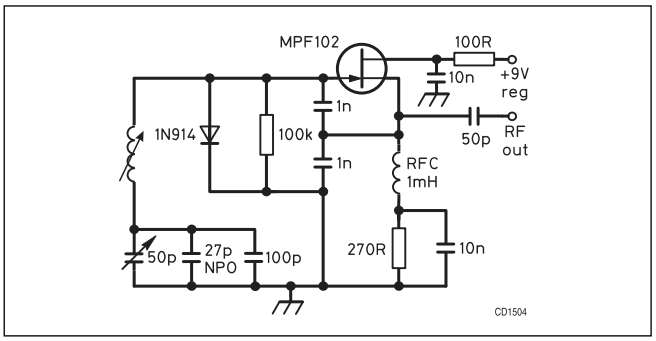

Fig 4.11: A JFET series-tuned Clapp VFO for 1.8MHz (*ARRL Electronic Data Book*)

dual-gate MOSFETs. Its second gate should be connected to a bias point at about 4V. Fig 4.9 shows the Colpitts version of the circuit. The additional input capacitance restricts the available tuning range but this is rarely significant. Again, a dual-gate MOSFET version is possible. A voltage-tuned version of the Hartley is shown in Fig 4.10. This is actually a VCO (see later) whose tuning voltage is derived from a potentiometer, so to the operator it operates like a VFO.

Many other variants on these basic circuits have been published over the years; some examples are shown in **Figs 4.11**, a Clapp oscillator and **4.12**, a Colpitts oscillator using a dual-gate MOSFET as the active device. In Fig 4.12, the biasing of the dual-gate MOSFET in oscillator service is typical; G1 is at source DC voltage level and G2 at 25% of the drain voltage.

As well as the Colpitts. Clapp and Hartley oscillator, other designs described in this chapter are the Franklin and Vackar oscillators.

Fig 4.13 illustrates a Franklin oscillator. This uses two active devices which in this example are GaAs MESFETs. As mentioned above, the MESFET tends to have relatively low gain per stage, so there is an advantage in using two gain stages in series to provide positive gain over a very wide frequency range. The

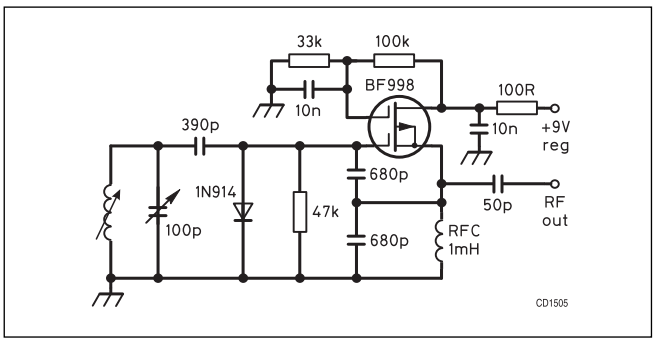

Fig 4.12: A dual-gate MOSFET in a common-drain Colpitts circuit (*ARRL Electronic Data Book*)

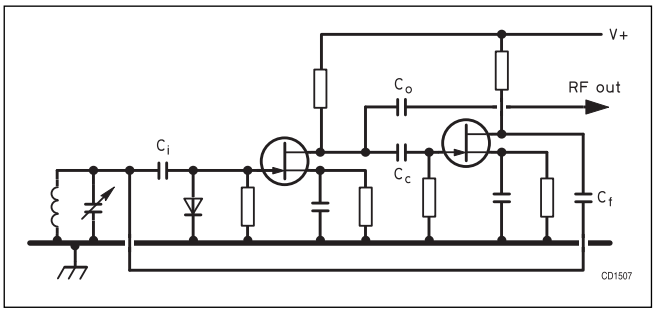

Fig 4.13: The Franklin oscillator

devices are shown running with self-bias on the gates. Ideally, they should be run close to I_{dss}, where the gain is usually greatest. The feedback path is arranged to give 360° phase shift, ie two device inversions. The feedback is then directed into the tuned circuit.

The output impedance of the MESFET is relatively low, but it is prevented from damping the resonant circuit by very light coupling; C_i and C_f typically are 1-2pF in the HF region, smaller still at higher frequencies. The high gain ensures oscillation even in this lightly coupled state, so that very high circuit Q is maintained, and the oscillation is stable and relatively noise-free. A particular feature of this configuration is the insensitivity of the circuit to supply variation; but the supply should still be well decoupled. A disadvantage of the circuit is the total delay through two stages, which limits the upper frequency, but with gallium arsenide devices this can still yield 5GHz oscillators.

DESIGNING / BUILDING A VFO

The following are recommendations for building variable frequency oscillators:

- For best stability, operate the VFO within the range 5 - 10MHz.
- Build the VFO in a sturdy box, eg a die-cast box. This will provide mechanical stability.
- Screen the oscillator to prevent the transfer of energy to/from other electronic circuits nearby. The screening box mentioned above will achieve this, but note that a box will not provide screening against magnetic fields.
- The VFO should be positioned away from circuits which generate high magnetic fields eg transformers.
- Protect the circuit from radiant and convected heat sources. Placing the VFO in a box will achieve this, but minimise the number of heat sources inside this box. For example, voltage regulators should be placed outside the box.
- Avoid draughts across the tuned circuit.
- Operate the VFO at a low power level. This will minimise the heating effect of RF currents in the oscillator components.
- Provide adequate decoupling.
- Use a buffer / amplifier stage to isolate the oscillator from the load.
- Provide regulated voltage supplies to the oscillator and buffer amplifier.
- Use a FET as the active device. The FET should have reasonably high gm (> 5mS). A J310 FET is a good example.
- Resistors should be metal film or carbon film types.
- Resistors should be rated at 0.5W or 1W dissipation. This reduces the temperature rise caused by currents flowing in the resistors.
- The active device will experience a temperature rise due to the operating current flow. This will cause a warm-up drift in the oscillator frequency. For this reason a conventional leaded package should be used, not a surface mounted package, and a small heat sink should be fitted.
- If the VFO is built on a PCB the board should be single-sided. This is because double-sided boards may show a capacitance change with temperature due to the dielectric material characteristics changing with temperature.
- Use single-point earthing of frequency-determining components in the tuned circuit.
- All components should be as clean as possible.
- The oscillator should be fitted with a good quality slow-motion drive to give the required tuning rate on the main tuning control. A typical figure chosen here might be 5kHz/turn of the tuning control for SSB mode.
- The receiver/transceiver operating frequency is most easily indicated using a digital counter/display. This should have a resolution of 0.1kHz or better, to allow easy re-setting to a certain frequency. The frequency display may use a software-based frequency counter (using a PIC) or it may be a all-hardware design using digital counters.

Oscillator Tuning Components

The oscillation frequency of an oscillator is primarily determined by the value of the inductor and capacitor(s) in the tuned circuit. Undesired changes in the reactance of these components (due to temperature change or some other factor) will have a direct effect on the oscillator frequency. It can be easily shown mathematically that if the inductor reactance changes by +100ppm then the oscillator frequency will move by -50ppm. A similar result is obtained if the capacitor reactance changes by +100ppm. This shows that there is a direct relationship between component value and oscillator frequency. For example, consider a 5MHz oscillator which employs a silvered mica capacitor with a temperature coefficient of +100ppm/°C. The capacitor would cause the oscillator to drift by -250Hz/°C. This is significant drift, and if the VFO was used in an SSB receiver, would require frequent retuning as the room temperature changed.

All components in the tuned circuit chosen should therefore be as stable as possible, and this is normally achieved by making the right choice of the type of component for each capacitor and inductor. A single poor quality component in an otherwise good oscillator design may ruin oscillator stability.

Capacitors

The following is a description of various capacitor types and their suitability for oscillator use.

- X7R dielectric ceramic capacitors must be avoided. These have low Q and are microphonic. They also exhibit capacitance change with applied voltage.
- COG dielectric ceramic capacitors are significantly better than X7R. They have higher Q and better temperature stability. These are available in surface-mounted (leadless) form and these are recommended. This is because they can be fitted directly onto a PCB and are then mechanically very stable.
- Silvered Mica. These are excellent for oscillators. They have high Q, and are temperature stable. However, the quantities used in the electronics industry are declining and they are becoming difficult to obtain. Typical temperature coefficient is -20 to +100 ppm/°C
- Polystyrene. These have a high Q at lower radio frequencies and are temperature stable. These are also becoming difficult to obtain. Typical temperature coefficient is typically -150 ppm/°C
- Porcelain capacitors. These are specialised capacitors, which are designed for use in high power microwave amplifiers. These offer very high Q and have typical temperature coefficient of +90 ppm/°C.
- Some trimmer capacitors show very poor temperature stability. Air dielectric trimmers on a ceramic base are preferred.
- The main VFO tuning capacitor should be a double bearing type to aid mechanical stability. Those with silver-plated brass vanes show better temperature stability than those with aluminium plates. The capacitor should be securely mounted to avoid vibration.
- If the Q of a capacitor is rather low, a useful technique is to make up the value required by putting several small

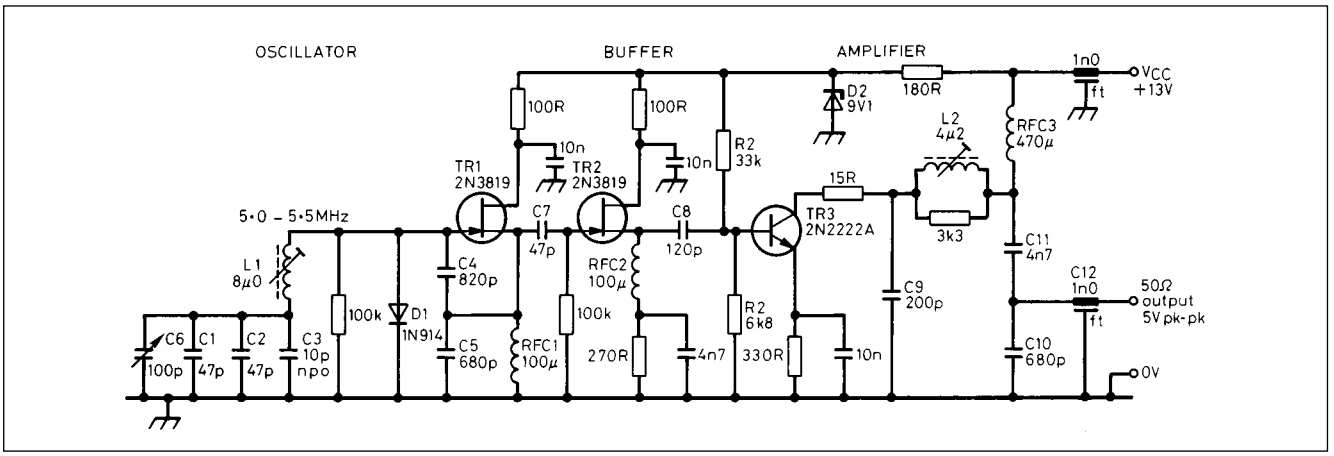

Fig 4.14: Clapp VFO with amplifiers for 5.0-5.5MHz. Reactance values are L1 = 265Ω, L2 = 140Ω, C1 = 690Ω, C2 = 690Ω, C3 = 2275Ω, C4 = 33Ω, C5 = 48Ω, C6 = 303Ω min, C7 = 690Ω, C8 = 227Ω, C9 = 152Ω, C10 = 48Ω, C11 = 4.5Ω, C12 = 23Ω, RFC1 = 4400Ω, RFC2 = 4400Ω

value capacitors in parallel. This usually produces a capacitor with a greater Q. It also reduces the RF current flowing in each capacitor, which reduces the heating effect of that current.

- Ceramic plate capacitors are available with positive and negative temperature coefficients as well as COG / NPO. Commonly available ceramic plate capacitors have designations such as N150, meaning it decreases its value by 150 ppm per degree Celsius. Conversely, a P100 capacitor increases the capacitance by 100 ppm for every degree Celsius. As the inductor in the VFO normally has a strong positive temperature coefficient the use of negative temperature coefficient capacitors (eg N150) in parallel with the tuning capacitor can reduce the drift with temperature to practically zero.

Inductors

If the guidelines (above) about capacitors are followed, the Q of an HF oscillator coil will usually be less than that of the capacitors. Therefore, the coil is the limiting factor on oscillator Q.

- The oscillator Q can usually be increased by raising the Q of the coil, and this can be done by winding the coil from Litz wire. Litz wire can be made by twisting a number of thin wire strands together. Each individual strand should have enamel insulation. Better Q may be obtained if the strands are plaited, rather than twisted.
- If the oscillator is operated on VHF, coil Q is usually high (typically 400) and the tuned circuit Q may be limited by the capacitors.
- Ceramic or fused silica formers are preferred for coil formers.
- Plastic formers should be avoided unless there is no other alternative.
- The coil former should not be too large. This is because a large coil will have a large magnetic field and it will be very susceptible to movement of objects which are within its field. 2 - 3 cm diameter is ideal
- If the coil is fitted with an adjustable screw slug the slug should be of a low permeability type. The number of turns on the coil should be adjusted so that the final adjustment position of the slug is just into one end of the coil.
- Toko coils are sometimes used in projects that must be easily reproducible. These are slug-tuned coils, and do not provide the best stability, but are commercially available as a ready-made part. For use in a VCO (corrected by a PLL) they offer good performance into the VHF range.
- If the coil is wound on a toroidal core the type of core material should be chosen carefully because some core materials have a high temperature coefficient of permeability. The resulting oscillator would show very high temperature drift. Powdered iron cores are more stable than ferrite.

Many of the specialised parts mentioned above may be obtained at radio rallies.

TWO PRACTICAL VFOS

Gouriet-Clapp Oscillator

The series-tuned Colpitts oscillator, often referred to as the Clapp or Gouriet-Clapp (devised by G C Gouriet of the BBC), has been a favourite with designers for many years. The series-tuned oscillator enables a higher value of inductance to be used than would normally be required for a parallel-tuned design, resulting in claims of improved stability. The circuit of a practical VFO, **Fig 4.14**, has component values selected for a nominal operating frequency of 5MHz, but reactance values are tabulated for frequency-determining components to allow calculation of optimum values for operation on other frequencies in the range 1.8 to around 10MHz. Simply substitute the desired VFO frequency into the reactance formula and the required L and C values can be deduced. Round off the calculated value to the nearest preferred value.

L1 should be wound on a ceramic low-loss former and the inductance is made variable by fitting a low-permeability, powdered-iron core. The series capacitors C1, C2 and C3 are the most critical components in the circuit; they carry high RF currents and the use of three capacitors effectively decreases the current through each one. A single capacitor in this location may cause frequency-jumping brought about by dielectric stress as a result of heat generated by the RF current. C3 is selected to counteract the drift of the circuit. In most circuits the tuning capacitor C6 will be located in series with the inductor and in parallel with C1-C3; a trimmer capacitor may also be placed in parallel with the main tuning capacitor. By placing the tuning capacitor in parallel with C5, a smaller tuning range is possible. This may be desirable for a 40m VFO where the frequency range required is only 200kHz.

The active device is a JFET and another JFET, TR2, acts as a buffer amplifier. This is lightly coupled to the source of TR1. Both

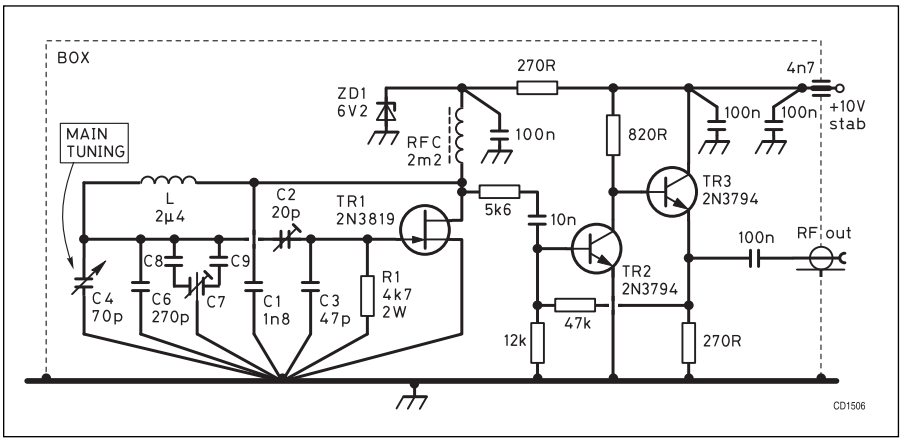

Fig 4.15: This Vackar oscillator covering 5.88 - 6.93MHz and two-stage buffering is based on a G3PDM design. A temperature compensation scheme consisting of C7, C8 and C9 is described in the text. A lower-noise option would be to derive the supply from a LM78L09 regulator with a liberal amount of electrolytic capacitors to kill any low frequency noise on the supply rail

transistors operate from a zener-stabilised 9V power supply. A better, lower-noise option would be to derive the supply from a LM78L09 regulator with a liberal amount of electrolytic capacitors to kill any low frequency noise on the supply rail.

TR3 is a voltage amplifier operating in Class A and uses a bipolar device such as the 2N2222A. Bias for Class A operation is established by the combination of the emitter resistor and the two base bias resistors which are also connected to the stabilised 9V supply.

The collector impedance of TR3 is transformed to 50Ω using a pi-network which also acts as a low-pass filter to attenuate harmonics. The loaded Q of the output network is reduced to approximately 4 by adding a parallel resistor across the inductor, and this broadens the bandwidth of the amplifier. For optimum output the network must be tuned either by adjusting L2 or by altering the value of C9. The 15Ω series resistor is included to aid stability.

The VFO described in Fig 4.14 is capable of driving a solid-state CW transmitter, and it is ideally suited to provide the local oscillator drive for a diode ring mixer. It also has sufficient output to drive a valve amplifier chain in a hybrid design.

Vackar Oscillator

This is an oscillator configuration which was first published in 1949 by Jiri Vackar. The circuit is a modification of the Colpitts oscillator, and is capable of extremely high stability. A notable example of the Vackar oscillator was a design by P G Martin, G3PDM (**Fig 4.15**) which is described more fully in [6]. This design has adjustable temperature-compensation and employed most of the measures described earlier to produce a stable oscillator. The design used a silver-plated copper wire for the tank coil (wound on a ceramic former), and included a two-transistor buffer stage.

The temperature compensation components are the 100 + 100pF differential air-spaced preset capacitor C7, the 100pF ceramic C8 with negative temperature coefficient and the 100pF ceramic C9 with positive temperature coefficient. Turning C7 effects a continuous variation of temperature coefficient of the combination between negative and positive with little variation of the total capacitance of about 67pF. Differential capacitors are hard to find but can be constructed by ganging two singles with semicircular plates, offset by 180 degrees. Adjustment for, and confirmation of, zero temperature coefficient of the whole oscillator is a time-consuming business. An updated version of this circuit would probably use a dual-gate MOSFET, which should offer marginally better performance. The buffer stages could use more easily obtainable devices such as the 2N3904.

VFOs are generally not used commercially because every VFO needs setting-up. This is costly and it is cheaper to use a synthesiser system, even if the synthesiser is fairly complex. However, for the amateur, building and adjusting a VFO for an SSB receiver/transceiver is a viable option to obtain variable frequency operation. It requires good mechanical skills to build a good quality VFO. It requires a significant amount of time and effort to build the VFO and adjust the temperature compensation. However, time spent will be rewarded with an oscillator which is essentially drift-free.

EVALUATION OF A NEWLY-BUILT VFO

It is essential that a VFO is properly tested for frequency stability before being used, particularly if it is to be used for transmitting. The process of testing the VFO is described below.

The oscillator should be located in the chassis in which it will finally be used and the output connected to a dummy load. The normal supply voltages supplied to the oscillator and buffer should be applied and the total current consumption checked. This should be no more than 30mA. Check that the circuit is oscillating by looking at the output on an oscilloscope, checking the frequency on a frequency counter, or listening on a receiver tuned to the expected frequency of oscillation. You may have to tune around on the receiver slightly because the oscillator probably won't be on the exact frequency you expected.

Adjust the oscillator to the required frequency range. This is normally done by adjusting the value of the inductor with the tuning capacitor set to mid-travel. The inductor is adjusted to set the oscillation frequency to the middle of the required tuning range. The padding capacitors are then adjusted to set the upper and lower frequency range required.

Frequency measurements should then be done to check the temperature stability. Temperature drift can be adjusted by exchanging fixed capacitors for ones having the same capacitance value, but different temperature coefficients (see the chapter on passive components).

VOLTAGE-CONTROLLED OSCILLATORS

These are a class of oscillator which are similar to a VFO. However a tuning capacitor or tuning inductor is not used and the oscillation frequency is controlled by an applied voltage. This is achieved by applying the control voltage to a component (or components) in the oscillator tuned circuit whose reactance varies with the applied voltage. Having the oscillator frequency controlled by a voltage allows the oscillator to be used in applications where electronic control of the frequency is required. The usual element used for this is a tuning diode. This is a junction diode which is operated in reverse bias, and exhibits a capacitance which varies with the applied bias voltage. Diodes are available which are designed specifically for operation as a tuning diode. These are optimised for wide capacitance variation with applied voltage, and good Q. Other means of electronic tuning are possible in some circuits, but

the use of tuning diodes is most common. The important parameter of the VCO is its tuning rate in MHz/volt and a typical design figure might be 1MHz/volt. VCO design is not without problems, particularly if the design needs a high tuning rate.

A VCO is normally used in a closed-loop control system. The system is designed to control the frequency of oscillation by means of the DC input control voltage. The requirements for short term stability (ie phase noise) are similar to VFOs. Therefore, most of the design rules presented above are equally applicable for VCOs as for VFOs. There is one exception to this: Because of the closed-loop nature of the VCO control system, the design rules for long-term stability of the oscillator no longer apply or can be significantly relaxed. VCOs can be built for virtually any frequency range from LF to microwave.

CRYSTAL OSCILLATORS

A crystal oscillator is a type of oscillator which is designed to operate on one frequency only. The frequency of oscillation is determined by a quartz crystal and the crystal must be manufactured specifically for the frequency required. Quartz shows the property of the Piezo-electric effect and can be made to oscillate at radio frequencies. The main virtue of quartz is that it has a very low temperature coefficient so the resulting oscillator will be very temperature stable. It also has a very high Q. The types of crystal are similar to the ones used in quartz crystal filters.
Crystal oscillators exhibit the following characteristics.

- Good temperature stability.
- The high Q gives good oscillator phase noise.

They are used where an accurate frequency source is needed. Examples of this are in electronic instruments, frequency synthesisers and RF signal generators.

For a description of the equivalent electrical circuit of a crystal see the next building blocks chapter.

Crystal oscillators employ one of two modes of oscillation:

- Fundamental mode, which is normally used up to about 24MHz.
- Overtone mode, which is normally used above about 24MHz.

There really are no user serviceable parts inside a crystal, but it is interesting to open up an old one carefully to examine the construction. The crystal plate is held in fairly delicate metal springs to avoid shocks to the crystal. The higher-frequency crystals are very delicate indeed.

In the last few years, crystals have become available which are intended to be operated in fundamental mode up to about 250MHz. They are manufactured using ion beam milling techniques to produce local areas of thinning in otherwise relatively robust crystal blanks. This maintains some strength while achieving previously unattainable frequencies. These are intended for professional applications and are generally too expensive for the amateur, and not easily available.

Crystals intended for fundamental-mode operation are generally specified at their anti-resonant or 'parallel' frequency with a specific 'load' capacity, most often between 12 and 30pF. When a crystal is operated in fundamental mode, it is possible to tune, or trim, the oscillation frequency slightly by adjustment of a reactance in the oscillator circuit. This is usually done with a variable capacitor and is done to adjust for the initial tolerance of the crystal frequency and can be done subsequently to adjust the frequency as the crystal ages.

When a crystal is operated in the overtone mode, it is in a series resonant, mode. Overtones available are always odd integers and the normal overtones used are 3rd and 5th. Overtone operation allows an oscillator to operate up to about 130MHz (5th overtone). A trimmer capacitor can be used to effect a slight adjustment of the frequency but the ability to adjust the frequency is much less than with the fundamental mode of operation. The 'pullability' of an overtone crystal is approximately $R = 1/N^2$, where R is the ratio in number of Hz it can be pulled and N is the overtone number eg 3 or 5 etc compared to a fundamental (parallel resonant) crystal of the same frequency. Hence, a 10MHz fundamental might be able to be pulled, say, 200Hz and a 30MHz third overtone only about 70Hz at 10MHz and hence a total of 210Hz at 30MHz. The 10MHz fundamental when multiplied up to 30MHz would allow about 600Hz variation. For FM transmitters a fundamental crystal oscillator provides a more linear deviation method than overtone oscillators.

Crystal Oscillator Circuits

The operation of a crystal oscillator is basically the same as the principles of any oscillator (VFO or VCO) described so far in this chapter. The best option for the amateur is to follow one of the circuits described here. Some crystal oscillator circuits do not need to use inductors, so the Hartley circuit is less common than the Colpitts type of circuit. Most circuits use a single active device.

For fundamental-mode crystal oscillators, any low power HF transistor is adequate, even a BC109. For overtone oscillators any RF type is recommended, eg NPN transistor types 2N2222A (metal can), BF494 or 2N3904 (plastic); N-channel plastic JFETs include J310, MPF102 and 2N3819. Virtually any dual-gate MOSFETs can also be used. A good rule of thumb is to choose a device with an Ft that is at least three to five-times the oscillator frequency. The BC109 and 2N2222 both have an Ft of about 100MHz so would be suitable for oscillators up to about 30MHz.

Excessive feedback should be avoided, primarily because it tends to reduce stability. Excessive oscillation levels at the crystal can lead to damage, but this is very unlikely in circuits operating on a supply of 12V or less. Sufficient feedback is needed to guarantee starting of the oscillator. Some experimentation may be needed, especially if the crystal is of poor quality. Pressure-mounted crystals and those in epoxy-sealed aluminium cans, eg colourburst crystals from early colour television sets, are often poor.

This section describes a range of fixed-frequency crystal oscillators with practical circuit details. Crystal oscillator circuits are not hard to build, and most reasonable designs work fairly well since the crystal itself provides extremely high Q, and in such a way that it is difficult to degrade the Q without deliberately setting out to do so.

Note that an overtone crystal is more active on its fundamental frequency than the overtone. If no method is used to suppress the

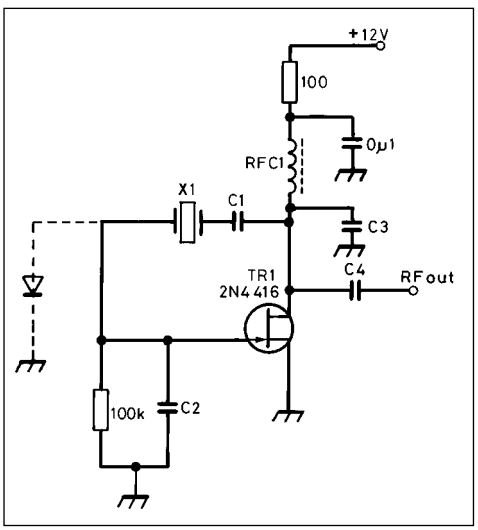

Fig 4.16: Pierce oscillator. Reactance values are: X_{C1}, X_{C3} 230Ω; X_{C2}, X_{C4} 450Ω. (*W1FB's QRP Notebook*)

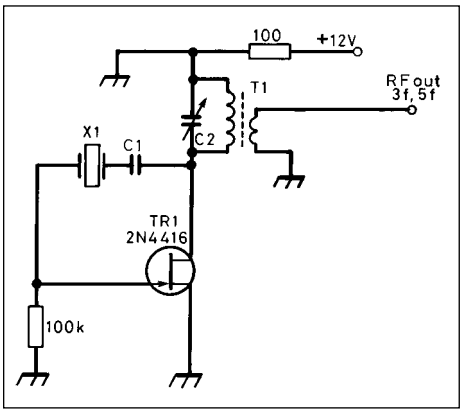

Fig 4.17: Overtone oscillator. Reactance values are: X_{C1} 22Ω., X_{C2} 150Ω. at resonance. (*W1FB's QRP Notebook*)

fundamental mode then spurii can be generated. A common method with overtone crystal oscillators is to use either an inductor in shunt with the crystal resonating with the holder capacitance or a low value resistor, the value should be about ten-times the ESR of the series resonant crystal. A third overtone crystal at 50MHz has an ESR of about 40 ohms worst case, so a 390 ohm resistor or slightly higher would be suitable.

Figs 4.16 to 4.19 show four basic crystal oscillator circuits. Fig 4.16 is a Pierce oscillator using a FET as the active device. The circuit operates as a fundamental oscillator and the crystal is connected in the feedback path from drain to gate. The only frequency setting element is the crystal. For this to be the case, the RFC must tune with the circuit parasitic capacitances to a frequency lower than the crystal. Most small JFETs will work satisfactorily in this circuit. Note that the FET source is grounded, which gives maximum gain.

Fig 4.17 shows an overtone version of Fig 4.16. In this, a tuned circuit is inserted in the drain of the FET. The tuned circuit is tuned to the overtone frequency required and hence the FET only has gain at that frequency. This causes the circuit to operate only on the wanted overtone.

Fig 4.18 shows a Colpitts fundamental oscillator circuit, which uses a bipolar transistor. The chief advantage of this configuration is that the crystal is grounded at one end. This facilitates switching between crystals, which is much less easily achieved in the Pierce circuits.

Fig 4.19 shows an overtone oscillator using a bipolar transistor. One side of the crystal is earthed, which makes switching crystals easier. However, switching crystals in an overtone oscillators is not recommended, since the stray reactance associated with the switching can lead to loss of stability or moding of the crystal. This causes some crystals to oscillate at an unintended overtone, or on another frequency altogether. If switching of overtone crystals is required, it is better to build a separate oscillator for each crystal and then switch oscillators.

Oscillators can also be built using an integrated circuit as the active device. **Fig 4.20** shows an oscillator using a standard

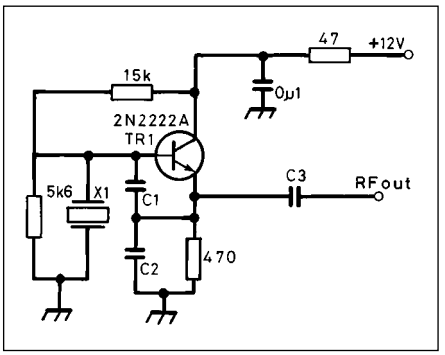

Fig 4.18: Colpitts oscillator. Reactance values are: X_{C1}, X_{C2}, X_{C3} 450Ω.. (*W1FB's QRP Notebook*)

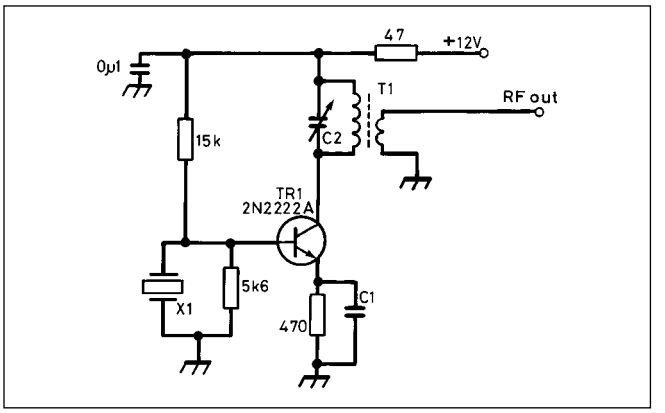

Fig 4.19: The tuning of the LC collector circuit determines the mode of oscillation: fundamental or overtone

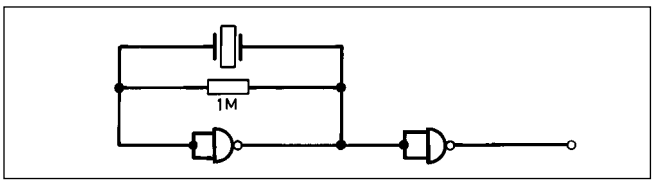

Fig 4.20: CMOS oscillator circuit. With modern 'HC' CMOS, this will oscillate to over 20MHz with a good square-wave output

CMOS gate package. Almost any of the CMOS logic families will work in this circuit. The primary frequency limitation is the crystal. This oscillator will generate a very square output waveform, and with the faster CMOS devices can be a rich source of harmonics to be used, for example, as marker frequencies. Note that a square wave contains only odd harmonics.

Self-contained crystal oscillators can be purchased which are intended for use as clocks for digital circuits. They frequently contain the crystal plate, the gates and any resistors on a substrate in a hermetic metal case. These are unsuitable for most amateur receiver or transmitter applications because there is no way of tuning out manufacturing frequency tolerance or subsequent ageing. Also, the phase noise performance is usually poor.

Where an oscillator with very high temperature stability is required, a temperature-compensated crystal oscillator (TCXO) can be used. These are normally supplied in a sealed package and they were developed to maintain a precise frequency over a wide temperature range for long periods without consuming the stand-by power required for a crystal oven. These are suitable for use as a frequency reference in frequency synthesisers and are used by the mobile communication industry. They are now included in top-of-the-line amateur transceivers and offered as an accessory in the mid-price field.

A TCXO is pre-aged by temperature cycling well above and below its normal operating range. During the last cycle, the frequency-vs-temperature characteristic is computer logged; the computer then calculates a set of temperature compensation values, which are stored in an on-board ROM. A good TCXO might have a tolerance of 1ppm over 0 - 60°C and age no more than 1ppm in the first year after manufacture, less thereafter. There is normally a provision to make small adjustments to the output frequency so that it can be adjusted from time to time.

Buying Crystals

Crystals must be bought for the frequency and mode that you require. Many common frequencies can be bought ready-made (eg exact frequencies like 5MHz or 10MHz). If you need to get a crystal made you will need to specify (as a minimum):

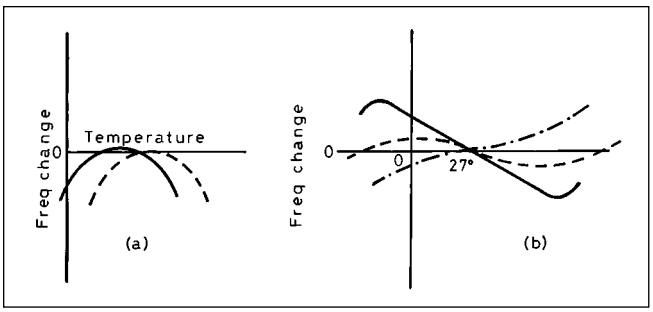

Fig 4.21: Crystal temperature coefficients. Frequency / temperature curves of 'zero-temperature coefficient' crystals. (a) Typical BT and DT-cut crystals. (b) Typical AT-cut crystals

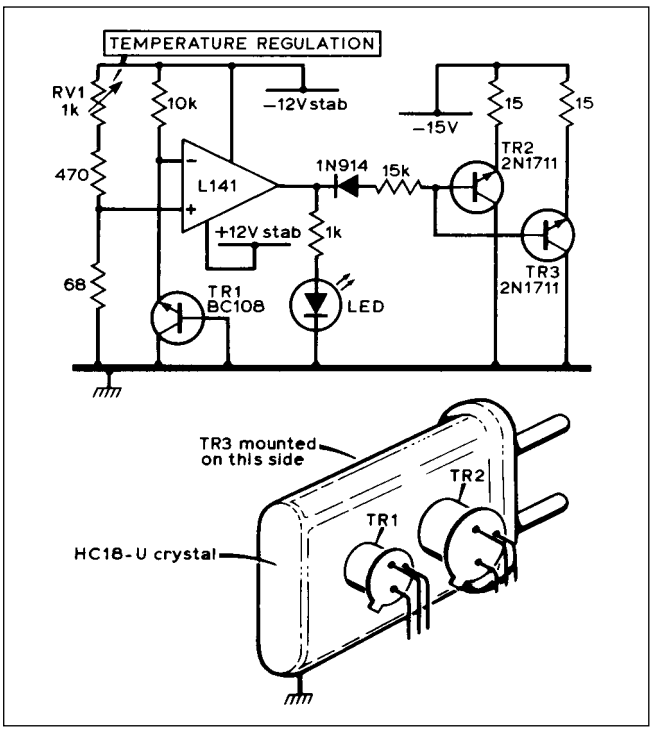

Fig 4.22: Proportional temperature control system for HC/16U crystals as proposed by I6MCF and using a pair of transistors as the heating elements and a BC108 transistor as the temperature sensor

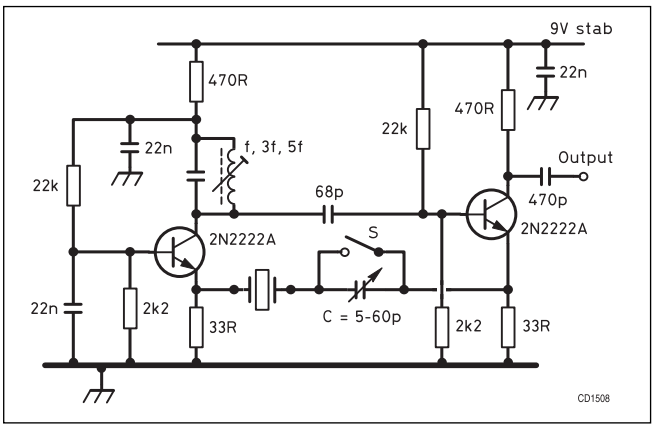

Fig 4.23: A crystal test oscillator using the Butler circuit. With S closed, the resonant frequency is generated. With S open and C set to the specified load capacitance, the anti-resonant frequency can be measured

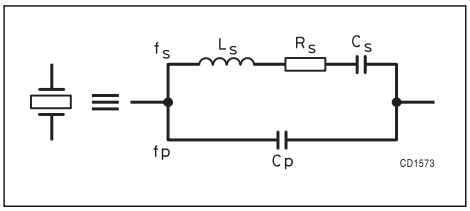

Fig 4.24: Crystal equivalent circuit

- Frequency
- Mode
- Loading
- Operating temperature

Crystals are normally available in HC6/U, HC18/U, HC25/U or similar holders. If crystals are bought at a rally, much older, larger packages may be found.

Crystals made for commercial use are similar in specification to amateur devices; they may be additionally aged or, more commonly, ground slightly more precisely to the nominal frequency. For most purposes this is not significant unless multiplication into the microwave region is required.

Occasionally, crystals are specified at a temperature other than room temperature (see **Fig 4.21**). Ideally, the crystal should have a zero temperature coefficient around the operating temperature, so that temperature variations in the equipment or surroundings have little effect. In some cases, the crystal is placed in a crystal oven at 70ºC. It is important to note, however, that the best result will only be obtained if the crystal is cut for the oven temperature. Although any temperature stabilisation is useful, a room-temperature-cut crystal may have little advantage in an oven. A scheme for thermal stabilisation of crystals is shown in **Fig 4.22**.

Fig 4.23 shows a Butler oscillator which is a favourite for testing oscillator crystals. With the switch S closed, the circuit will oscillate at the crystal's resonant frequency; the output amplitude is then measured. Next, a 100Ω variable composition resistor is substituted for the crystal and adjusted to produce the same output. Its final resistance value is equal to the crystal resistance, R_s in **Fig 4.24**.

VARIABLE CRYSTAL (VXOS) AND CERAMIC-RESONATOR OSCILLATORS

A VXO is a type of crystal oscillator which is designed to have a much greater frequency adjustment range than a standard crystal oscillator. A typical use for a VXO is to obtain crystal control on an HF amateur band. The frequency can be adjusted by several kilohertz to avoid causing interference to other band users.

VXOs are used frequently in portable equipment where the virtues of simplicity and stability are particularly useful. Fundamental mode oscillators should be used because these can be pulled in frequency much more than overtone oscillators. The principle can be understood if one replaces the piezo-electric crystal by the electrical equivalent circuit shown in Fig 4.24. If a variable inductor is placed in series with that equivalent circuit, the resonant frequency of the combination can be 'pulled' somewhat below that of the crystal alone; for convenience, tuning is usually carried out with a variable capacitor in series with this coil which is over-dimensioned for that purpose.

When building VXOs the precautions described for VFO design should be used, but VXOs are less sensitive to external influences since the oscillation frequency is controlled primarily by the crystal. The two problems of VXOs are the very limited tuning range,

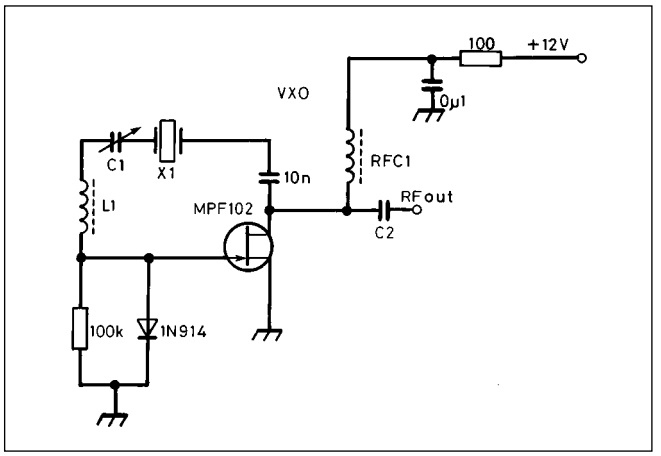

Fig 4.25: A Pierce VXO for fundamental-mode crystals. The diode limits the output amplitude variation over the tuning range. Reactances at the crystal frequency are: C1 = 110Ω, C2 = 450Ω, L = 1.3Ω, RFC = 12.5kΩ

typically of 0.1%, and the tendency to instability if this range is exceeded. A good VXO should be very close to crystal stability, with the added ability to change frequency just a little. Of course, multiplication to higher frequencies is possible, and several of the older types of commercial equipment successfully used VXOs on the 2m band. Typical VXO circuits are shown in **Figs 4.25 and 4.26**. Fig 4.25 is a Pierce oscillator using an FET as the active device. Stray capacitance in the circuit should be minimised. Ideally, the VXO should be able to shift from below to just slightly above the nominal crystal frequency. Fig 4.26 shows a two-transistor Butler oscillator.

Crystal switching in VXOs should be avoided and it is better to switch oscillators. Non-linear operation of capacitance against frequency should be expected, so the tuning scale will be non-linear. Note that some linearisation is possible using combinations of series and parallel trimmer capacitors with the crystal. Instead of a tuning scale, a better approach is to fit an LCD frequency readout or to use the station digital frequency meter, thus avoiding altogether the need for a tuning scale.

One way to achieve a wider tuning range is to use two VXOs ganged to shift in opposite directions at two high frequencies differing by the desired much lower frequency, as proposed by G3MEV in reference [7]. **Fig 4.27** shows G3MEV's VXO circuit.

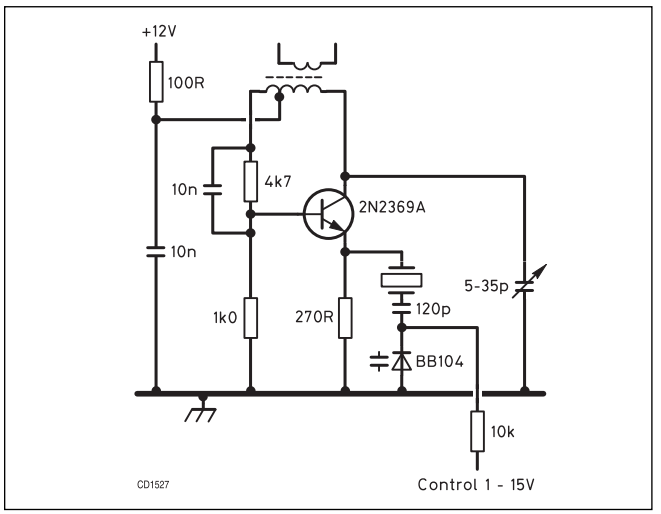

Fig 4.27: Voltage-controlled 21.4 and 20MHz oscillators used by G3MEV in his 'heterodyne' 1.4MHz VXO system

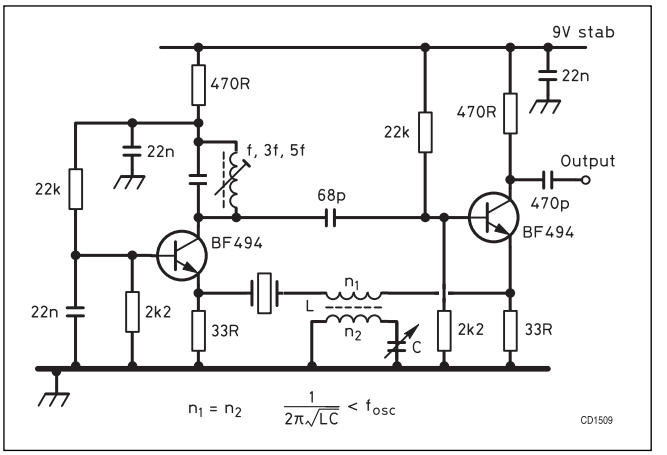

Fig 4.26: A Butler VXO for fundamental or overtone crystals

Another idea is the use of one VXO to interpolate between closely-spaced multi-channel synthesiser frequencies; The difference frequencies between a VXO tuning 24.10125 to 24.11125MHz (only 10kHz in 24MHz) and the 40 synthesised channels of a CB transmitter will cover 3.5 - 3.9MHz.

Ceramic resonators are now inexpensively available in many standard frequencies. They permit the construction of stable variable-frequency oscillators with a wider pulling range than VXOs. See reference [8].

Like quartz crystals, they are piezo-electric devices; they have a worse temperature coefficient than crystals and a lower Q, but for most HF applications it is adequate and the lower Q permits pulling over a wider frequency range. The oscillator circuits are similar to those for VXOs. **Figs 4.28 and 4.29** are examples.

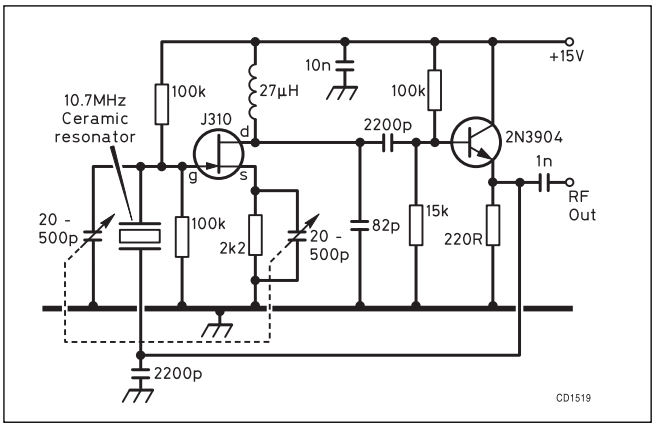

Fig 4.28: K2BLA's variable ceramic-resonator oscillator covers a 2% frequency range

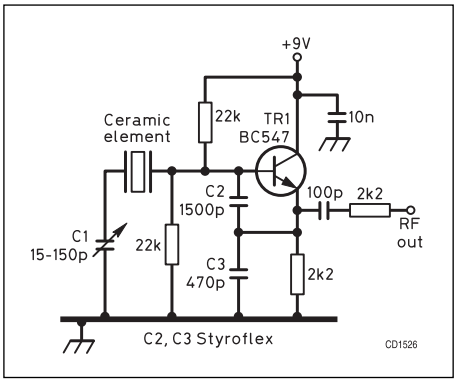

Fig 4.29: LA8AK's oscillator can tune 3.5-3.6MHz with a 3.58MHz ceramic resonator

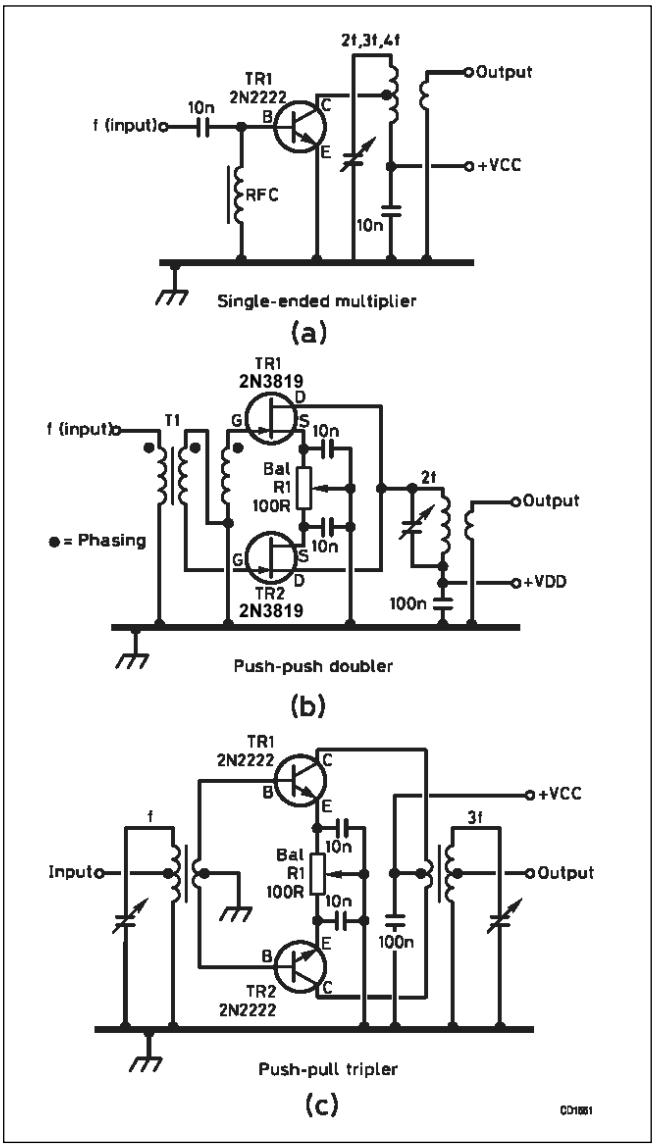

Fig 4.30: Frequency multipliers. (a) Single-ended. (b) Push-push doubler. (c) Push-pull tripler

USING CRYSTAL CONTROL ON VHF

To obtain a crystal controlled source on VHF an overtone oscillator could be used. For example, if a crystal-controlled AM or CW transmitter is to be built for the 50MHz or 70MHz amateur bands, then there is significant benefit in using an overtone oscillator to directly generate the frequency required. Compared to using a fundamental oscillator with frequency multiplier stage(s), the overtone circuit is much simpler, and reduces the likelihood of spurious outputs occurring from the transmitter.

Crystals for fundamental use are, however, much cheaper than those intended for overtone use. Also if the source is intended for use in an FM transmitter, it will be easier to achieve the frequency deviation required if a fundamental oscillator is used. It is necessary to multiply the frequency and this can be achieved using frequency multiplier stages.

Normally an RF amplifier is biased and driven so that it does not generate a high level of harmonics at its output. However, when a frequency multiplier is designed, its bias and drive level are set so that it will deliberately generate harmonics of the the input frequency.

A frequency multiplier has its output tuned to a multiple of its input frequency, eg x2, x3 or x4. These multipliers are known as doublers, triplers and quadruplers, respectively. VHF/UHF bipolar transistors work well as frequency multipliers. FETs do not generally make good frequency multipliers, except where a frequency doubler is required. The square-law characteristic of FETs means that they make very effective frequency doublers. **Fig 4.30(a)** illustrates a typical multiplier.

If a higher multiplication factor is required, it is better to use two or more multipliers, eg for x9 multiplication use two tripler stages. The multiplier may be operated with a small amount of forward bias dependent on the transistor forward transfer characteristic but the stage must be driven hard enough (with RF) for it to function correctly. The bias is adjusted for maximum efficiency (output) after the multiplier has been tuned to the desired output frequency.

The multiplier can be single-ended, as in **Fig 4.30(a)**, or with two FETs configured as a push-push doubler as in **(b)**, or two transistors as a push-pull tripler as in **(c)**. The efficiency of **(b)** and **(c)** is typically higher than (a). The push-push doubler discriminates against odd-order multiples and the push-pull tripler discriminates against even-order multiples, so additional attenuation of unwanted outputs is much improved over the single-ended multiplier. R1 in (b) and (c) is adjusted for optimum dynamic balance between TR1 and TR2, ie maximum multiplier output.

It is always good practice to start with a reasonably high oscillator frequency in a receiver LO multiplier chain, thus requiring fewer multiplier stages. This is because there is the distinct possibility that one of the unwanted multiples will reach the mixer and cause a spurious response (or responses) within the receiver tuning range. The magnitude of any unwanted injection frequencies depends on the operating conditions of the final multiplier, including the working Q of its output circuit. To minimise this problem it is advisable to include a buffer amplifier stage with input and output circuits tuned to the LO frequency between the final multiplier and the mixer.

These circuits may consist of a single high-Q tuned circuit called a high-Q break or two loosely coupled circuits. Helical filters provide an ideal solution to this problem. Two filters may be cascaded to give greater attenuation.

Adequate filtering is essential in transmitter frequency multipliers, particularly between the final multiplier and the first RF amplifier, to attenuate all unwanted multiplier output frequencies from the oscillator to the final multiplier. In order to establish the crystal frequency it is useful to draw a chart such as **Fig 4.31**. This shows the combination of frequencies and multiplication factors which could be used to obtain an output on 118MHz.

It is important to understand that any frequency change in the final multiplier output is the frequency change (or drift) of the oscillator multiplied by the multiplication factor. As an example, if the frequency change in a 6MHz crystal oscillator was 50Hz and the multiplication factor is 24, then at 144MHz the change will be 1.2kHz. This change is not too serious for NBFM but would not be acceptable for SSB equipment.

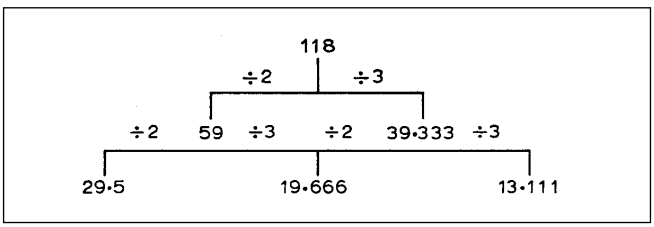

Fig 4.31: Type of chart for determining the fundamental crystal frequency required in an oscillator multiplier chain

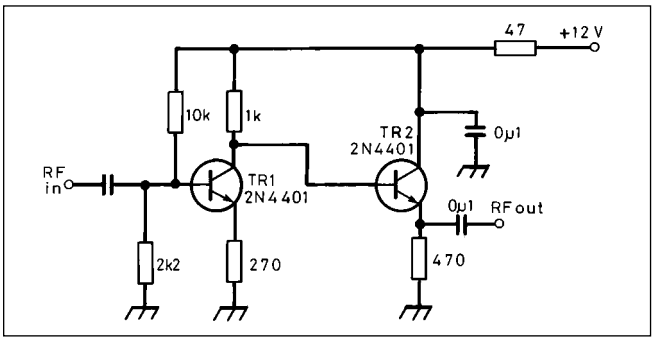

Fig 4.32: Direct-coupled buffer circuit (*W1FB's QRP Notebook*)

BUFFER STAGE

The output of an oscillator is not normally connected directly to the load. This is because changes in load impedance will affect the oscillator frequency. Changes in load impedance may be caused by:

- Changes in temperature
- Switching from receive to transmit
- Changes in supply voltage

A buffer amplifier must always be used to isolate the oscillator from the load. A buffer should be used with VFOs, VCOs and crystal oscillators. It is an RF amplifier which is designed for high reverse isolation. This greatly reduces the effect of load changes on the oscillator frequency.

The following needs to be known when designing a buffer amplifier:

- The expected load impedance.
- The gain required.
- The output level required.

Some of the oscillator circuits in this chapter show a buffer amplifier. Examples are:

- Fig 4.15. In this circuit the buffer is formed by TR2 and TR3.
- **Fig 4.32**. In this circuit the buffer is formed by TR2 which operates as an emitter follower.
- **Fig 4.33** shows a two-stage buffer amplifier which is designed to drive a 50-ohm load.
- **Fig 4.34** shows a buffer using a CMOS Schmitt trigger invertor. This is suitable for mixers that require a large voltage swing (5V p-p). There is also an alternative low-level output that can be used to drive Gilbert Cell type mixers.

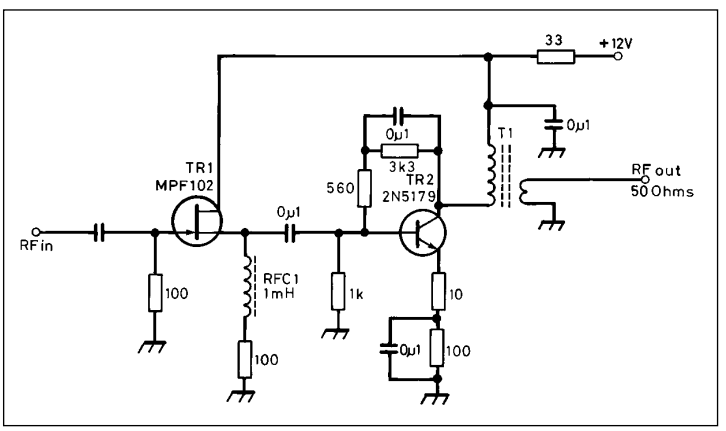

Fig 4.33: Transformer output buffer circuit. (*W1FB's QRP Notebook*)

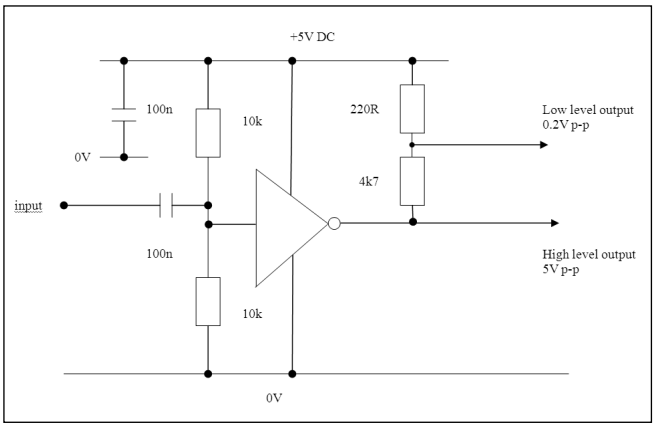

Fig 4.34: Buffer using a Schmitt trigger

Amplifiers employing feedback should not be used on their own as a buffer because these have poor reverse isolation. Buffers should be built using high quality components and should be solidly built like a VFO.

FREQUENCY SYNTHESIS

The description 'frequency synthesiser' is generally applied to a circuit in which the output frequency is non-harmonically related to the synthesiser reference frequency. Therefore, this excludes circuits which use a crystal oscillator and frequency multiplier stage to generate the output frequency. This section contains a description of the operation of analogue synthesis, PLL frequency synthesis and Direct Digital synthesis (DDS). Early forms of synthesis are no longer used commercially, but are described here because they are still of relevance to the radio amateur.

Analogue Synthesis Systems

Analogue (or Direct) synthesis was the first form of synthesis. It involves the production of an output frequency, by the addition of several (sometimes many) oscillators, usually by a combination of mixing, multiplying, re-mixing etc with many stages of filtering. This was difficult and expensive to set up really well. Examples were primarily limited to specialist, often military, synthesisers in the 'fifties. The output spectrum could be made very clean, especially in respect of close-to-carrier phase noise, but only by significant amounts of filtering. Direct synthesisers are still used in specialist areas, because they can give very fast frequency hopping and low spurious levels well into the microwave region.

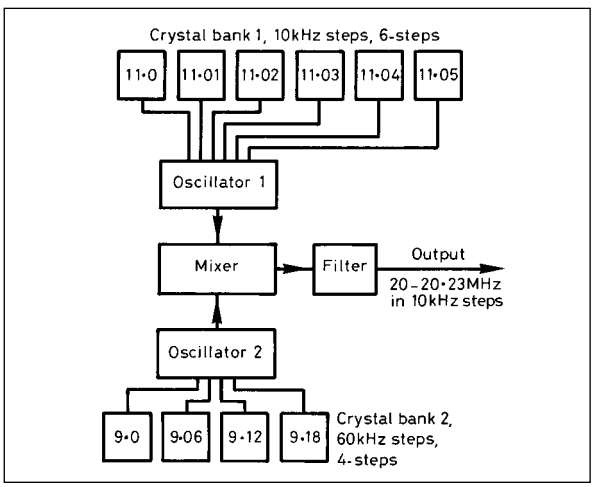

Fig 4.35: Simplified crystal bank sythesiser

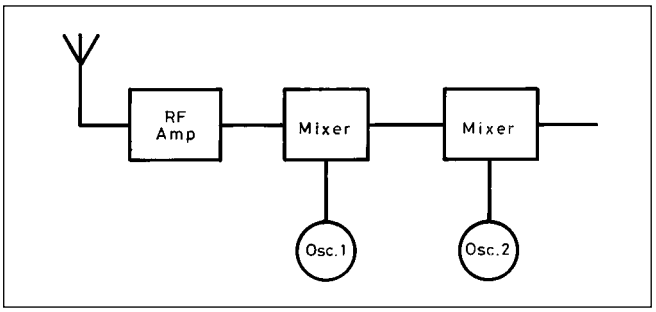

Fig 4.36: Multiple oscillators/mixers, an alternative to frequency synthesis

A simple version did go into quantity production, in the early US-market CB radios. This was the crystal bank synthesiser shown in **Fig 4.35**. This arrangement used two arrays of crystals, selected in appropriate combinations to give the required coverage and channel spacing. It worked very well over a restricted frequency range, could offer very low phase noise (because crystal oscillators are inherently 'quiet'), and could be made acceptably (for the time) compact and low power. An example on the UK market in the 1970s, was the Belcom Liner 2, which was an SSB transceiver operating in the 2m (144-146MHz) band. This used crystal bank synthesis with a front-panel-tuned VXO on a further crystal oscillator for the conversion up and down from the effective 'tuneable IF' at 30MHz to the working frequency of 145MHz.

The need for the VXO illustrates the weakness of this type of synthesiser; close channel spacing is essentially prohibited by the very large number of crystals which would be needed. Triple bank versions were described in the professional literature, but were not produced for amateur purposes. The other disadvantage of this class of synthesiser is the high cost of the crystals, but it remains an interesting technique capable of very high performance with reasonable design.

A popular alternative to synthesis for many years used a carefully engineered VFO and multiple mixer format. Examples are in the KW range, the Yaesu/Sommerkamp FT-DX series and, in homebrew form, the G2DAF designs (see earlier editions of this *Handbook*). This technique survived in solid-state form, with among others, the FT101 [9]. The technique is illustrated in **Fig 4.36** [10].

Phase-Locked Loop (PLL) Frequency Synthesisers

In parallel with the analogue synthesis systems described above, came early attempts at combining the best VFO characteristics with PLL synthesis. A PLL synthesiser employs a closed-loop control system, which uses analogue and/or digital circuits to generate the required output frequency. This allows a single-conversion receiver, and is able to achieve greater receiver dynamic range. An example of this was G3PDM's classic phase-locked oscillator, as used in his receiver and described in earlier editions of this *Handbook* [6]. This was an analogue synthesiser which used mainly valves, but with a FET VFO. This design is still regarded as the standard to beat. The FET VFO was discussed earlier in this chapter (Fig 4.15). A demonstration of this receiver was given where a signal below 1µV suffered no apparent degradation from a 10V signal 50kHz away. This amazing >140dB dynamic range was made possible by the extremely low phase noise of this synthesiser, and the superb linearity of a beam-deflection mixer valve.

Most modern PLLs contain digital circuits and the PLL output frequency is normally driven by digital control inputs. The first digital synthesisers appeared in about 1969, chiefly in military equipment. An outstanding example was the UK-sourced Clansman series of military radios.

A basic PLL digital synthesiser block diagram is shown in **Fig 4.37**. The VCO output is divided in a programmable divider down to a frequency equal to that of the crystal-controlled reference source; this may be at the crystal frequency or more usually at some fraction of it. The phase comparator compares the two frequencies, and then produces a DC voltage which is proportional to the difference in phase. This is applied to the VCO in such a way as to drive it towards the wanted frequency (negative feedback). When the two frequencies at the phase comparator input are identical, the synthesiser is said to be in-lock. Ideally, the phase comparator is actually a phase and frequency comparator, so that it can bring the signals into lock from well away. The 'DC' term referred to is a varying voltage determined by the frequency and phase errors. The bandwidth available to this voltage is said to be the loop bandwidth and it determines the speed at which lock can be achieved.

The programmable divider is a frequency divider using digital counters. Fully programmable dividers of appropriate frequency range were not immediately available to the early PLL designers. Two solutions were proposed for this:

- The dual-modulus pre-scaler.
- A variant of the crystal mix scheme.

The dual-modulus counter (**Fig 4.38** uses a dual-modulus prescaler and is a particularly clever use of circuit tricks to achieve the objective of high-speed variable ratio division. Dividers were

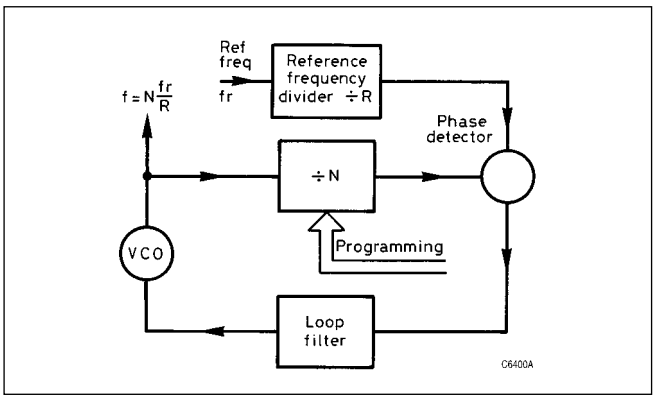

Fig 4.37: Basic PLL synthesiser

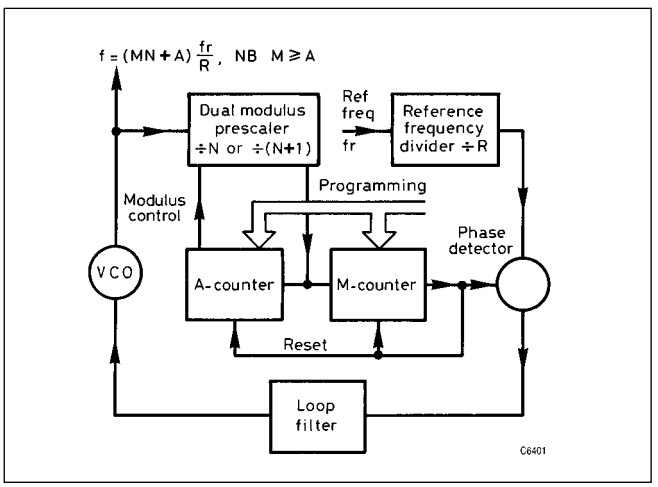

Fig 4.38: Dual modulus prescaler

designed which could be switched very rapidly between two different division ratios, for example 10 and 11. If the divider is allowed to divide by 11 until the first programmable counter (the A-counter) reaches a preset count (say, *A*) and then divides by 10 until the second counter reaches *M*, the division ratio is:

$$(11 \times A) + 10 \times (M - A) \text{ or } (10M + A)$$

The advantage in the system is that the fully programmable counters M and A need only respond, in this example, to one-tenth of the input frequency, and can therefore use lower cost, slower logic. The disadvantage is that there is a minimum available count of A x (11), where *A* is the largest count possible in the A-counter. This is rarely a problem in practice.

Fig 4.38 shows the complete loop consisting of the VCO, the dual-modulus counter, a main counter, a reference oscillator and divider chain, a phase detector and a loop amplifier/ filter.

This type of synthesiser is relatively easy to design when a synthesiser is needed with large step sizes. It has been described many times in the professional and amateur radio fields [11]. An example of this is a synthesiser for an FM receiver, which is required to step in 25kHz or 12.5kHz steps. This is also used in TV receivers, satellite receiver and Band II FM tuners. Most commercially available PLL synthesiser chips contain the digital parts of the PLL. That is the main divider, the reference divider and the phase detector, and the counters are normally programmed by serial input. Most modern PLL chips will operate up to VHF without the need for a dual-modulus pre-scaler. They also provide two (or more) phase detectors to cover the far-from-lock and locked-in cases.

The second method, that of crystal mixing with a digital synthesiser, is probably the most widely used technique in amateur equipment, where wide operating frequency range is not needed. It offers advantages of simplicity and low component count, especially when the digital functions of the system can be contained in a single chip; see **Fig 4.39**. Unlike the dual modulus scheme, no pre-scaler is needed, but a high-frequency mixer is necessary, usually with two crystals, one for mixing down and one for the reference.

The PLL is a closed-loop system, which employs negative feedback. Like any feedback system, the loop has a finite speed at which it can operate and this is defined by the closed-loop bandwidth. There is a (usually) wider bandwidth called the lock-in bandwidth, which is the range of starting frequencies from which lock can be achieved.

The closed loop bandwidth is under control of the designer and is set by the open-loop parameters. These are:

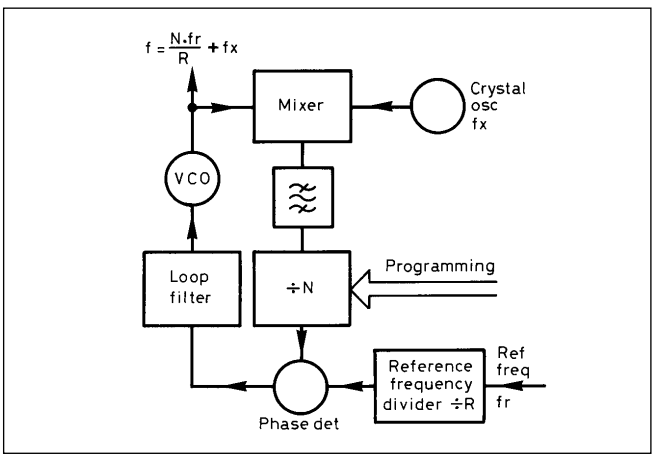

Fig 4.39: Mixer-loop PLL

- VCO gain
- Divider ratios
- Phase detector gain
- Loop filter amplitude/frequency and phase/frequency responses

The control loop is able to control and hence reduce the effects of VCO noise within the closed-loop bandwidth, this includes long term drift in the VCO. Outside the loop bandwidth, VCO noise is unaffected by the loop and the closed-loop bandwidth is normally set to about one tenth of the reference frequency. The closed loop is a second-order system and so it is possible for the loop to exhibit loop instability. The designer controls the loop stability primarily by controlling the response of the loop filter. The lock-up time is inherently related to the closed-loop bandwidth. For an FM-only rig, with say 12.5kHz minimum frequency step, lock-up time can be very quick.

For SSB receiver and transceivers, there are strong commercial pressures to replace the VFO with a full digital synthesis system because this avoids many of the problems and costs of building and setting up a good VFO. However, trying to reproduce the action of a good VFO in digital synthesiser form is difficult. A PLL synthesiser can only tune in steps and so to reproduce the VFO action for SSB requires very small steps, typically 20Hz. This cannot be done in a single loop because the loop bandwidth would have to be very narrow, and the loop would be very slow. Professional systems overcome this by complex, multiple-loop synthesis, which can be expensive to design. They can also be expensive to manufacture and are sometimes bulky.

The operating frequency of a digital synthesiser is normally controlled by pulses from a shaft encoder system. This allows control of the frequency by a tuning knob in the same way as a VFO. Frequently, the tuning rate of the knob is programmable so that the tuning rate (in kHz per turn) is adjustable for different modes of signal. AM or FM modes can use faster tuning rates than SSB. The equipment uses a digital frequency display and this normally has resolution of 100Hz or 10Hz. It should be noted that this does not mean that the frequency is accurate to 100Hz or 10Hz. (Resolution doesn't equal accuracy). The frequency accuracy of the synthesiser is determined entirely by the accuracy of the reference oscillator.

PLL synthesiser design is a very demanding and essentially analogue task. The VCO is particularly difficult to design because very small unwanted voltages induced onto the tuning voltage will generate noise or spurious outputs from the synthesiser. Nevertheless, truly excellent performance can be achieved by the professional engineer in factory-built equipment. Properly designed, PLL synthesis works very well indeed, and most commercial amateur equipment uses some form of PLL synthesiser. However, the problems of providing narrow frequency increments, with the demand for complete HF band coverage (at least on receive) have lead to rather compromised synthesiser designs. This has been seen in particular in reviews of otherwise excellent radios limited in performance by synthesiser noise and spurious outputs causing reciprocal mixing.

A PLL digital frequency synthesiser can be designed to cover a much wider frequency range than a typical good VFO, but still achieve good frequency stability. It offers an alternative to the amateur who doesn't have the skills or tools to build a VFO. Some synthesisers (eg those that use serial control) require software to enable them to operate, but synthesisers can be built using hardware only, and hence require no software. The frequency display is normally digital, and it is generally easy to obtain the frequency display information from the synthesiser control circuits.

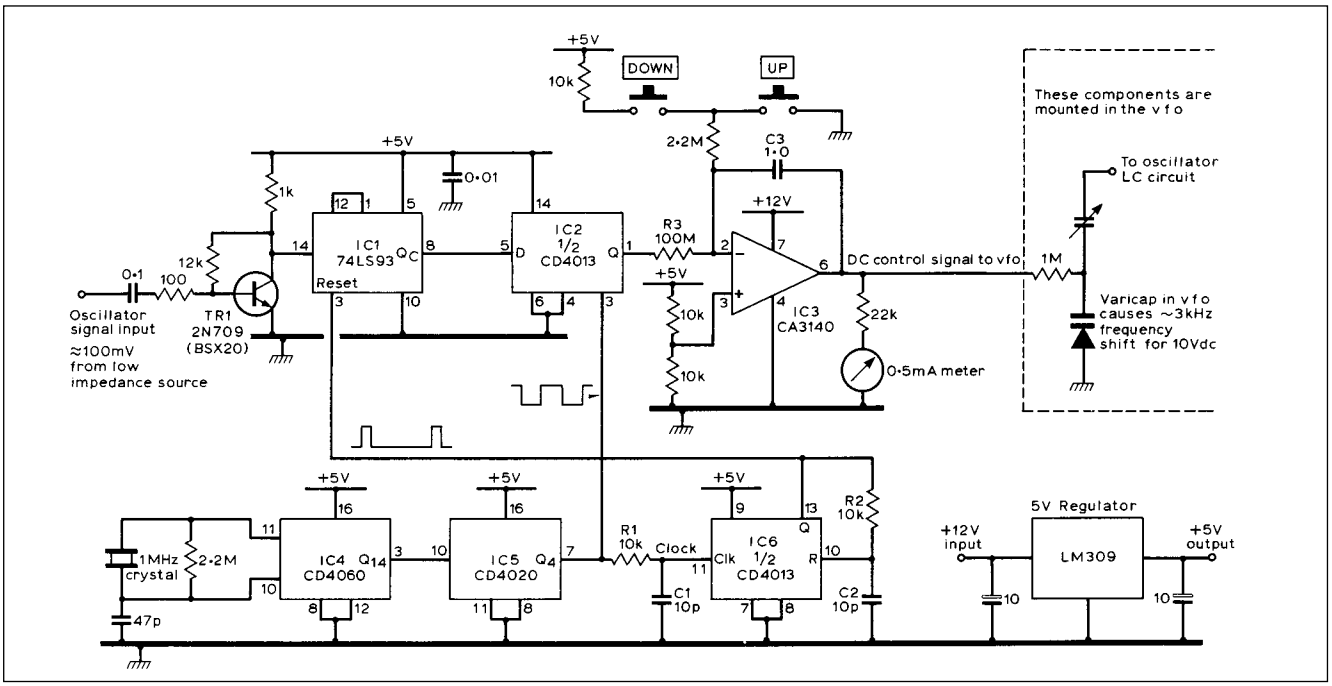

Fig 4.40: Circuit details of the CMOS form of the 'huff-and-puff' VFO stabilising system - a version of PA0KSB's system as described in *Ham Radio*

It is possible for the very dedicated and well-equipped amateur to design a PLL system for SSB use, particularly as he does not cost his time into the project. It should be said that a spectrum analyser, or at least a continuously tuneable receiver covering the frequency of interest and a harmonic or two, is essential for any synthesiser work. Loop stability and noise should be checked at several frequencies in the synthesiser tuning range.

Amateur equipment practice has tended to take a different route. A synthesiser is employed which has a minimum step size of typically 1kHz. Increments between the 1kHz steps are achieved either with a separate control knob, on older equipment, or by a digital-to-analogue converter (DAC) operated by the last digit of the frequency setting control. The analogue voltage tunes a VXO in the rig, which may either be the conversion crystal in the synthesiser or the reference crystal itself. Care must be taken to ensure that the pulling range is accurate or the steps will show a jump in one direction or the other at the 1kHz increments. Actually, many rigs do show this if observed carefully; the reason is that until recently the DAC was a fairly simple affair consisting of a resistor array and switching transistors.

More recently, commercial DACs have been used, but this is not a complete solution, since the VXO is unlikely to be linear to the required degree, so some step non-linearity is inevitable. At least if the end points are not seriously wrong, this should not be a problem. Some rigs do tune in 10Hz steps, especially on HF; this is an extension of the technique to 100 steps instead of just 10. Again, there is the issue of the 1kHz crossover points; but with 10Hz steps this can be made less noticeable on a well-designed and adjusted rig.

One variant on PLL synthesis which appeared some years ago, and which is particularly suitable to amateur construction, is the so-called huff-and-puff VFO [12]. This could be described as a frequency-locked loop and consists (**Fig 4.40**) of a VFO, and a digital locking circuit. This has the same frequency reference chain and phase-sensitive detector, but the divider chain is largely omitted. The VFO is tuned using a variable capacitor, and the loop locks the VFO frequency to the nearest multiple of the reference frequency possible. The reference is typically below 10Hz, so that analogue feel is retained in the tuning. To the operator, this behaves like a PLL. The loop can take a significant time to lock up, so various provisions to speed this up can be employed. The loop must not jitter between different multiples of the reference, hence the need for very high initial stability in the VFO. As a commercial proposition, this scheme is unattractive compared to a true PLL, but as an amateur approach it can overcome the VFO tuning problem in an elegant way.

Direct Digital Synthesis (DDS)

Direct Digital Synthesis is a method of frequency synthesis which differs from the methods described so far. A DDS is virtually all-digital and the only analogue part is a digital-to-analogue convertor (DAC) at the DDS output. In fact many modern DDS chips now contain two DACs, supplying two separate outputs. A DDS does not contain a negative feedback closed-loop system like a PLL. The DDS offers the possibility of generating an output which has very small frequency steps but, unlike a PLL, very rapid tuning to new frequency. Compared to a digital synthesis PLL type system, a DDS is relatively easy to design because the complex, critical electronics are contained within the DDS chip. Construction is relatively easy because it is built up as a conventional logic PCB and there are none of the problems of setting-up which occur with a conventional VFO. The DDS approach to synthesis is most suited to the amateur who is familiar with building digital systems. It may require some software to be written to drive the DDS chip. The output frequency stability (both short-term and long-term) is determined by the characteristics of the reference oscillator. For the amateur, the best option for the reference oscillator is to use a quartz crystal oscillator and it is worthwhile using a high quality oscillator circuit to obtain low phase noise and good frequency stability. The alternative is to purchase a ready-made high- performance crystal oscillator module.

Direct digital synthesis has been frequently mentioned in *RadCom*, and is a feature of current amateur equipment [13, 14]. The basic direct digital synthesiser consists of an arrangement to generate the output frequency directly from the clock

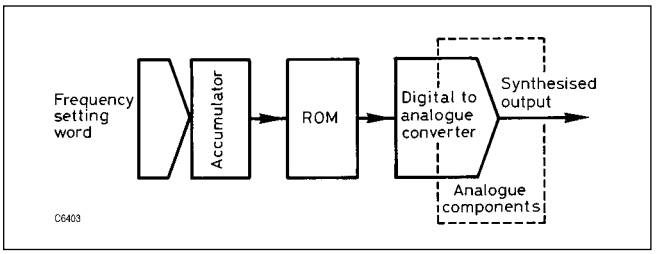

Fig 4.41: The DDS concept

and the input data. The simplest conception is shown in **Fig 4.41**. This consists of a digital accumulator, a ROM containing, in digital form, the pattern of a sine wave, and a digital-to-analogue converter. Dealing with the accumulator first, this is simply an adder with a store at each bit. It adds the input data word to that in the store. The input data word only changes when the required frequency is to be changed. In the simplest case the length of accumulator is the clock frequency divided by the channel spacing, although it is usually calculated the other way round, ie if a 5kHz channel spacing up to 150MHz is needed, and since at least two clock pulses are required per output cycle, a clock frequency of at least 300MHz will be required. More conveniently, a 16-bit accumulator gives 65,536 steps. If the step size is to be 5kHz, then a clock frequency to the accumulator of 65,536 x 5kHz is needed, ie 327.68MHz. Such a DDS would produce any frequency in the range covered, ie 5kHz to over 100MHz in a single range without any tuned circuits.

Other frequency increments are available; any multiple of 5kHz by selection of input data, and others by choice of clock frequency, eg 6.25kHz requires a clock at 204.8MHz with a two-channel (2 x 3.125kHz) program word. Of course, frequency multiplication or mixing can be used to take the output to higher frequencies. Programming is very easy, DDS devices have either a parallel-input data format or are designed to be driven from a serial input. The control circuits which drive the DDS can also be used to drive a digital frequency display.

A more advanced device is shown in **Fig 4.42** [15]. This has on-chip DACs and facilities for square, triangle and sine outputs. Two DACs are used, because both phase and quadrature output signals are available in true and complement form. The most significant bit (MSB) from the accumulator feeds the square-wave output buffer direct. In parallel, the next seven bits from the accumulator, which digitally represent a sawtooth waveform, feed a set of XOR gates, under control of the MSB, so that a triangle output is generated, in digital form, at this point in the circuit. Actually, two triangles in quadrature are generated. These can be digitally steered to the output DACs or can be used to

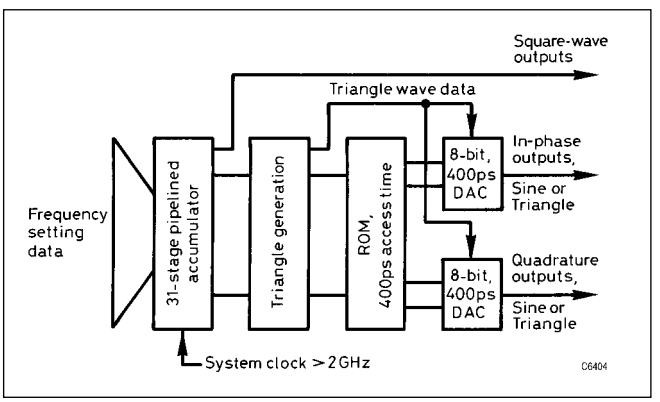

Fig 4.42: More complex DDS

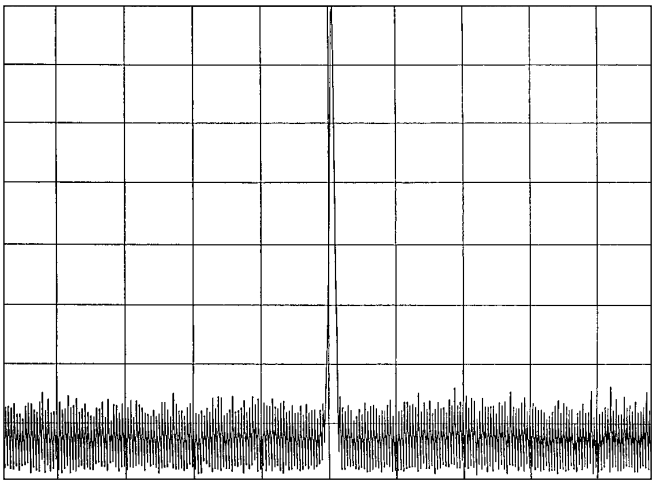

Fig 4.43: DDS spectrum at 250MHz with a 1GHz clock. 10MHz/div, X axis; 10dB/div, Y axis

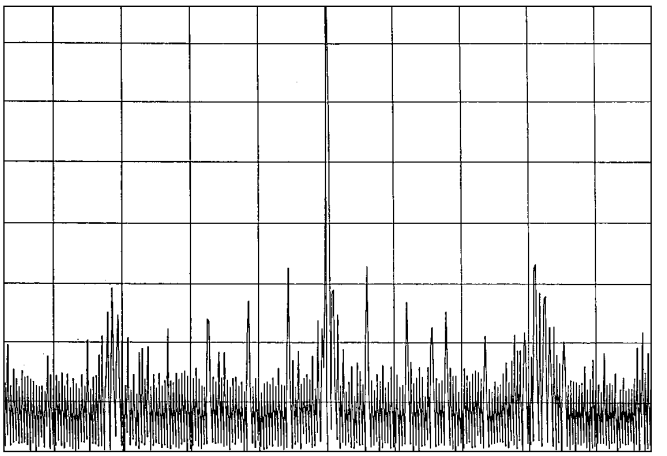

Fig 4.44: DDS spectrum at 225MHz with a 1GHz clock. 10MHz/div, X axis; 10dB/div, Y axis

address a ROM containing data for sine and cosine waves; only 90° is needed, since all four quadrants can be generated from one. Finally, if selected, the digital sine/cosine is fed to the DACs for conversion to analogue form.

This device was designed to operate at up to 500MHz output frequency with 1Hz steps, so the full range is 1Hz - 500MHz, again in a single 'range' with no means of or requirement for tuning. The clock frequency is necessarily very high; the quadrature requirement adds a further factor of two, so nominal clock frequency is 2^{31} Hz, ie 2.147483648GHz. This must be supplied with crystal-controlled stability, although since phase noise is effectively divided down in the synthesiser and the frequency is fixed, it can be generated by a simple multiplier from a crystal source.

Output spectra are shown in **Figs 4.43 and 4.44**. Fig 4.43 shows a clean output at exactly one quarter of the clock frequency. Fig 4.44 shows a frequency not integrally related to the clock frequency. Here the spurious sidebands have come up to a level about 48dB below the carrier. This is the fundamental limitation of the DDS technique, and where most development work is going on. The limit comes from the finite word size and accuracy of the DAC, and incidentally the ROM, although this could have been made bigger fairly easily.

Fast DACs are difficult to make accurately, for two reasons. The first is technological. IC processes in general have a limit to a component matching accuracy of about 0.1%, ie 9 or 10 bits accuracy in a careful design. Other techniques, such as laser

trimming of the resistors, can be used on some processes. However, this tends to use older, slower processes, and especially requires large resistors for trimming which slow the DAC settling time. The second problem is simply the requirements on fast settling: the DAC is required to get to the final value quickly and this will in general happen with the smallest number of bits. An 8-bit system as in this example was chosen to fit the process capabilities and the device requirements, but this leads to a limitation in the high level of spurious signals present.

Basically, although some frequencies are very clean, in the worst case the spurious level is 6*N* dB below the carrier, where *N* is the number of effective DAC bits. For an 8-bit system, this gives -48dB. Oversampling, ie running the clock at more than twice the output frequency, gives some improvement (at 6dB per octave) by improvement of the DAC resolution, but only up to a limit of the DACs accuracy. Typically this is about 9 bits or -54dB. In some applications, this may not be serious, since filtering can remove all but the close-in spurs. Phase-locked translation loops into the microwave region also act as filters of relatively narrow bandwidth, while retaining the 1Hz step capability.

In the amateur rigs using DDS, the synthesisers operate at relatively low frequencies and are raised to the working frequency by PLL techniques. This is complicated, but makes fine frequency increments available without the compromise of PLL design. The devices used are CMOS types, with DAC accuracies of 8 and 10 bits. This gives spurious signals of theoretically -60dB referred to the carrier, which is adequate with filtering from the PLL.

Direct Digital Synthesiser for Radio Projects

This project was written for *RadCom* by Andy Talbot, G4JNT [16]. In this article, a DDS chip is included in a small stand-alone module controlled by straightforward text-based messages from a PC. The module has been designed to be used as a drop-in component in larger projects such as receivers or transceivers.

An onboard PIC microcontroller allows control of the DDS chip from a standard PC via the serial port, and a straightforward command syntax has been developed so that standard software commands, written in any language, can be used to set the operating parameters. This PIC can easily be reprogrammed to suit any user's requirements in a stand-alone project and enough spare Input / Output lines are provided to allow for this.

This article is intended to provide an overview of the DDS module as a component for larger projects and, therefore, only limited details are included.

DDS module design

The module is based around an Analog Devices AD9850 DDS chip, full details of which are available from Analog Devices [17]. The device will accept a clock signal up to 120MHz, although a suitable source is not provided within the module. This is best left to individual constructors.

The DDS can generate an output up to approximately one third of the clock frequency with a resolution of over 4 billion so, for a 120MHz clock, frequencies from DC to 40MHz can be produced in steps of approximately 28 millihertz (mHz), and so the actual frequency of the clock is unimportant, provided of course it is known with reasonable accuracy. The RF output is at a level of 1V p-p (0.35V RMS) and the output impedance is 100Ω.

The supplied PIC controller translates text based messages received from the serial port into command codes for the AD9850, and (optionally) stores these in non-volatile RAM for immediate setting at switch-on. Spare memory in the PIC allows user information, such as the exact clock frequency, to be readout on request, allowing common software to be written that can drive individual modules, each having different clock frequencies.

Another option available in the firmware is a 'times-four' output, where the output from the module is designed to drive a

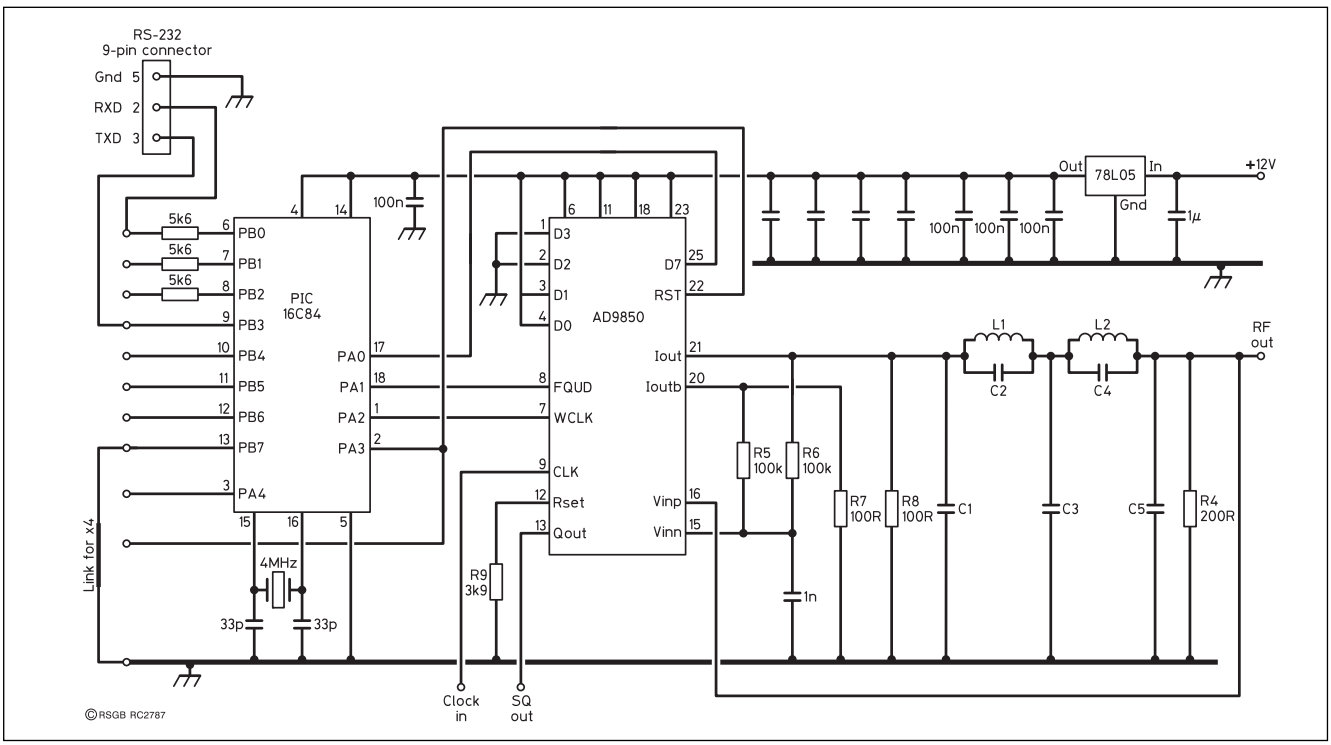

Fig 4.45: Circuit diagram of the AD9850 DDS module. Filter components and some decoupling components are dependent upon clock frequency

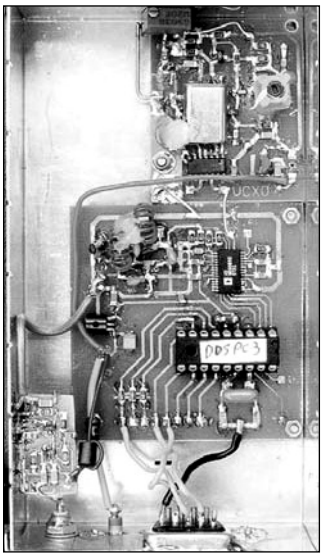

Fig 4.46: A prototype version of the G4JNT DDS board

quadrature frequency generator which performs the final frequency division.

A single hex-digit module address is included as part of the command syntax, to allow multiple modules to be driven in a multi-drop arrangement from a common controller with one COM port.

Output spurii

Spurious output signals from a DDS are complex in their nature and are not harmonically related - the device data sheet [17] gives more details.

All spurii from the design here are at levels of -60dB or better. This figure is obtained from the manufacturer's specification and is affected by the design of the output filter. One useful facet about DDS circuitry is that the spurious levels below the output filter cut-off frequency are inherently dependent on the circuits internal to the DDS chip, and are not affected by any poor circuit layout; this aspect is often trouble-some in other synthesiser designs, as filtering cannot remove close-in products.

60dBc spurious levels may be considered a bit excessive if this module were to be used alone as the local oscillator of a high-performance wide-band receiver, but there are additional techniques such as Phase Locked Loops that can clean up this signal.

Construction

The circuit diagram is shown in **Fig 4..45**, from which it can be seen that the module actually has very few components, most of them being for decoupling and output filtering. Apart from the integrated circuits, all other components are either of the surface-mount 1206 or 0805 style; for the output filter, wire-ended components are used. The PIC controller is mounted in a socket for easy re-programming.

The photograph **Fig 4.46** shows a prototype version of the DDS board, together with a high-stability 94MHz source on the PCB top right and an output buffer amplifier bottom left.

Values of components for the output filter depend on the clock frequency chosen, as do the values of decoupling capacitors. With a clock input that can range from a few kHz up to 120MHz, the optimum values of these can vary over a wide range.

Software control

Commands are sent using ASCII / Hex characters over a bi-directional RS-232 link with no handshaking. Parameters are 19,200 baud, 8 data bits, no parity, 1 stop bit. A simple terminal programme such as HYPERTERM (Windows®) or PROCOMM (DOS) can be used to command the frequency source. Set this to 19200 N 8 1, full-duplex, no flow control and all start up, and modem commands set to null.

Alternatively, custom software can be written to drive the COM port with the commands.

The first character sent is a board address which precedes all commands. This is a single Hex character sent as ASCII 0 - F and potentially allows up to 16 modules to be driven from the same COM port.

The next character is a command which may have hex data following it.

Q	followed by eight hex digits for the frequency command word terminated by a carriage return [CR]
P	followed by two hex digits for phase word and [CR]
U	writes the data sent above to the AD9850 DDS chip
W	as for U, and also stores all data in the PIC's non-volatile EEPROM memory for switch-on next time
Y	followed by one Hex digit, changes the board address and stores in EEPROM. No [CR] needed
K	followed by 10 hex digits and [CR]. User data, not used for driving the DDS. (In practice, read as decimal number for user data, typically clock frequency)
R	read back current data values - not necessarily those in EEPROM

The 32 bit or 8 hexadecimal character, value N (required for frequency-setting), can be derived from:

$$N = F_{out} / F_{clock} \times 2^{32}$$

Phase can be set to any one of 32 values in increments of 11.25 degrees. These form the five highest significant bits of the phase word Pxx. The lowest three bits are ignored.

Data is sent back from the DDS in text strings which can be read directly by application software.

An example of a command to set the output frequency is:

5Q03D70A3D [CR]	Board address 5, set frequency word N = hex 03D70A3D
5U	Programme the DDS to this value (with a 120MHz clock, this gives an output at 1.8MHz).

Applications

Any experiments or testing that needs an agile frequency source is a candidate. Just about any programming language that includes commands to drive the serial port can be employed, and does not even have to be PC-based. The only requirement is that you can actually write suitable software!

One application written to demonstrate the functionality of the module is for generating a narrow-band Multi-Tone Hellschreiber signal for LF use. SMT Hell transmits visible text as an image, and can be received on a frequency / time plot, commonly known as a spectrogram or waterfall display, using public domain commonly-available audio analysis software. The software generates SMT-Hell signals by directly commanding the DDS module (in real time) to set the frequencies that make up the vertical elements of each character sent. The horizontal components are made up by appropriately setting software delays. Transmissions as narrow as 2.5Hz bandwidth have been sent on 137kHz and successfully decoded as visible letters even where the signal is completely inaudible below the noise.

The ability to set the output signal phase to one of 32 values means that slow Phase Shift Keying is also possible by direct command. An additional command code (T) allows frequency and phase updates to be synchronised to an external trigger input on port B1 such as that from a GPS receiver. The DDS chip is updated within 3µs of the trigger signal rising edge and will allow, for example, precisely-timed low data rate signalling experiments.

By using the DDS output to drive the reference input to a conventional Phase Locked Loop synthesiser, the best of both worlds becomes possible.

The high frequency capability of PLL synthesisers, up to many GHz, can be coupled with the tiny step size of the DDS. For example, a PLL operating with an output at 2.4GHz could be made with a step size of 0.55Hz.DDS frequency resolution is considerably better than the stability of most crystals will allow. In fact, the actual crystal frequency can be measured, stored in the user data area of memory and then used in subsequent high-accuracy output frequency calculations.

DDS chips are evolving rapidly, to operate at higher frequencies and with more bits, but at reasonable cost. This trend will continue, and will make DDS more and more attractive for amateur applications. Web sites are the best source of up-to-date information on DDS chips.

At the time of writing, the fastest devices available from Analog Devices are the AD9858, AD9910 and AD9912. These will operate with a clock frequency up to 1000MHz, and can operate with either series or parallel control (except the AD9912 which uses serial control only). The AD9858 is a 10-bit device and the AD9910 and AD9912 are 14-bit devices.

For other examples of the use of DDS techniques, see [18] and the chapter in this handbook on software defined radios.

REFERENCES AND BIBLIOGRAPHY

[1] *Electronic and Radio Engineering*, F E Terman, McGraw-Hill, any older edition, eg the fourth, 1955.

[2] *Radio Receiver Design*, K R Sturley, Chapman and Hall, 1953.

[3] *Amateur Radio Techniques*, 7th edn, Pat Hawker, RSGB.

[4] *Special Circuits Ready-Reference*, J Markus, McGraw-Hill.

[5] *GaAs Devices and their Impact on Circuits and Systems*, ed Jeremy Everard, Peter Peregrinus.

[6] 'A receiver with noise immunity and frequency synthesis', Peter Martin, G3PDM, *Radio Communication Handbook*, 5th edn, RSGB, 1976, pp10.104-10.108.

[7] 'Technical Topics', *Radio Communication* May 1991, p31.

[8] See 'Using ceramic resonators in oscillators', Ian Braithwaite, G4COL, *RadCom*, February 1994, pp38-39.

[9] *Maintenance Service Manual*, FT-101 series, Yaesu Musen Co Ltd.

[10] 'An easy-to-set-up amateur band synthesiser', Ian Keyser, G3ROO, *Radio Communication* December 1993, pp33-36.

[11] The Plessey Company, part of GEC, have, when they were still making telecommunications ICs, published several applications books describing PLL synthesis; any of these are well worth acquiring, although they are only obtainable second-hand. Examples are *Radio Communications Handbook*, 1977; *Professional Radio Communications*, 1979; R*adio Telecoms IC Handbook*, 1987; *Professional Data Book*, 1991; and *Frequency Dividers And Synthesisers IC Handbook*.

[12] Klaas Spaargaren, PA0KSB, in the ARRL publication *QEX* (Feb 1996, pp19-23), Chas F Fletcher, G3DXZ in *RadCom* December 1997, as well as corrections, improvements, additions and comments in 'Technical Topics' in *RadCom*, Jul, Aug, Sep 1996, Feb, Dec 1997, and Feb, Apr 1998.

[13] 'Technical Topics', Pat Hawker, G3VA, *Radio Communication* Dec 1988, pp957-958.

[14] 'Direct digital synthesis, what is it and how can I use it?', P H Saul, *Radio Communication* Dec 1990, pp44-46.

[15] Plessey Semiconductors Data Sheet *SP2002*.

[16] 'Direct Digital Synthesis for Radio Projects'. By Andy Talbot G4JNT. RadCom Nov 2000.

[17] Analog Devices web site. http://www.analog.com

[18] 'CDG 2000 HF Transceiver'. *RadCom* Jun 2002 to Dec 2002.

About the Author

Peter Goodson is qualified with a Hons Degree in Communication Engineering. He began his career working on the component-level design of HF (SSB) and VHF (FM) transceivers. He now works on the system level design of RF systems, which includes the issues of communications security, EMC, co-siting of radio equipment, RF hazard assessments and Immunity to lightning strikes.