Understanding phase noise

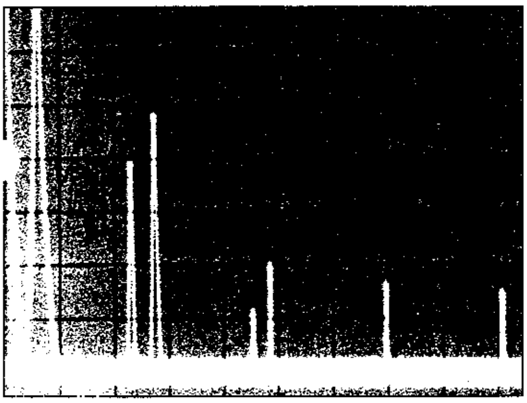

Oscillator purity is an increasingly important factor for designers of communications equipment. Ian Hickman has been investigating phase noise in rf oscillators and has discovered a problem that seems to have gone unnoticed.

Modern wireless communication often uses one or other of the various types of digital modulation.

The earlier, simpler forms, such as basic dpsk, or binary phase shift keying, are relatively robust, requiring only a modest signal-to-noise ratio at the receiver to guarantee successful reception. But shortage of spectrum space spurred the search for greater bandwidth efficiency. This lead first to the development of variations on the theme of qpsk – quadrature phase shift keying – which conveys two bits of information per signal element or 'symbol'.

Later, more exotic forms, such as 16psk, 64apk and even 256apk appeared, carrying four, six and eight bits per symbol respectively.

At the receiver, the demodulator must effectively measure the phase difference between successive symbols. At the transmitter, this starts out as 0 or 180° – in the case of asymmetrical dpsk – or only + or –90° in the symmetrical form. But by the time the signal is received, noise and interference will have eroded the available phase margin, possibly leading to bit errors.

With asymmetrical form qpsk, the phase change between symbols is 0, + or –90 or 180°, or +45, +135, –45 or –135° in the symmetrical case, sometimes called as '$\pi/4$ qpsk'. So a higher signal to noise ratio at the receiver is required for the same bit error rate.

With the advanced forms of modulation mentioned earlier, the phase change from one symbol to the next may be only $22^1/_2$ degrees or even less, so clearly an even greater signal-to-noise ratio is required for an acceptable bit error rate.

Noise in the receiver

Atmospheric noise and interference are not the only problems a digital data receiver faces. While an hf receiver with a reasonably efficient aerial is likely to be 'externally noise limited', at vhf and even more so at uhf and microwaves, external noise is so low that reception will usually be limited by the receiver's own noise.

In this context, one usually thinks of input stage noise. But in the reception of digital phase modulation, an important contribution to the factors eroding the essential phase discrimination, on which a low bit error rate depends, is the phase noise of the local oscillator.

Ideally, an oscillator produces an isolated spectral line, with zero energy output at any other frequency. Of course, there will be some harmonic content, but this is usually unimportant in a well designed receiver. Much more troublesome is energy at frequencies immediately adjacent to the oscillator output. This takes the form of noise sidebands, which can be quite large at very small offsets from the oscillator frequency, falling off at greater offsets, until at frequencies well removed from the carrier, their level bottoms out at the oscillator's far-out noise floor.

Why is phase noise important?

The sidebands consist of a mixture of amplitude-noise and phase-noise. In a receiver local-oscillator application, the amplitude noise sidebands are usually unimportant, since the local oscillator output is applied to the mixer at a high level: the local-oscillator input of the mixer thus operates in a heavily compressed mode. So minor level changes – even of a decibel or so – would have negligible effect. But the local-oscillator phase noise is quite a different story.

The intermediate-frequency signal reflects the phase difference between the rf signal input and the local-oscillator drive waveform. Thus local-oscillator phase noise adds linearly to phase disturbances of the wanted signal. These include noise, interference and multipath suffered in the over-the-air path, and front-end noise due to a marginal signal level.

The over-the-air path is outside the receiver designer's control; he can only concentrate on he other factors, of which – in a digital data ...eceiver – oscillator phase noise is a major component.

Phase noise from the local-oscillator

In special cases of fixed frequency operation, a receiver's local oscillator may be a crystal oscillator. Such an oscillator is characterised by extremely low levels of sideband noise – which is usually denoted by $\mathcal{L}(f_m)$ and defined as the noise power in a 1Hz bandwidth at an offset of f_m. But usually, the local oscillator will be an *LC* type, and these exhibit a higher level of sideband noise, extending out much further on either side.

To highlight the difference, note that a good crystal oscillator may show a level of sideband noise, $\mathcal{L}(f_m)$, which is already down to –140dBc at only 10Hz offset from the carrier. By contrast, a commercially advertised varactor-tuned voltage-controlled oscillator module, covering the range 100 to 200MHz, claims a typical $\mathcal{L}(f_m)$ of –105dBc at 10kHz ffset, and around –120dBc at 100kHz offset.

Where the *LC* oscillator forms the voltage-controlled oscillator in a phase-locked loop, its sideband phase noise within the loop bandwidth will be reduced by the loop negative feedback. Outside the loop bandwidth it returns to the level it would be were the voltage-controlled oscillator running open loop.

Clearly, even given a degree of phase-noise clean-up by the loop, one is better off starting out with a low phase-noise oscillator in the first place. A facility for measuring the phase noise of an oscillator is therefore an important item in any rf development laboratory, and can involve some very expensive equipment. For this reason, I was interested in an article which described such a measurement system using only standard laboratory instruments plus some inexpensive rf instrumentation[1]. The basic arrangement is shown in **Fig. 1**.

B is for bottom

I wanted to try and measure the phase noise of an oscillator, in order to settle a question which has interested me for some time. This question is, "is there an advantage in designing an *LC* oscillator in such a way that the transistor does not bottom at the negative-going peaks of the waveform?"

In fact, many *LC* oscillator designs do result in the transistor bottoming. This can be quite difficult to avoid in an oscillator with a wide tuning range, such as a three-to-one frequency ratio, given production spreads in transistor characteristics. The effects of bottoming in an rf oscillator had been explored in an earlier article[2], but equipment to measure phase noise was not available to me at that time.

I built an *LC* oscillator, operating at around 10MHz. Relative to vhf, this frequency is easier to measure. This, together with the other items needed for the Fig. 1 type set-up, is shown in detail in **Fig. 2**.

Tank circuit inductor L_1 was a Coilcraft

Working out sideband noise

In the lower trace of Fig. 5b), the measured level of sideband noise at 2.5kHz offset from carrier, with the circuit of Fig. 2, is –108dBV in a 30Hz measurement bandwidth. To work out $\mathcal{L}(f_m)$, the value in a 1Hz bandwidth is needed.

The analyser's intermediate-frequency filters consist of five synchronously tuned crystal filter stages, providing a Gaussian response. This characteristic is optimum for rapid settling to the true value of a swept signal. The noise bandwidth of such a filter is 12% greater than the actual –3dB bandwidth. The nominal 30Hz bandwidth is subject to a ±15% tolerance, so the actual –3dB bandwidth was measured, using the 1dB/division scale. It turned out to be 27Hz, giving a noise bandwidth of 30Hz, as near as makes no odds. Thus the level of –108dBV in 30Hz translates to –123dB in a 1Hz bandwidth. This represents the sum of the noise energy in both upper and lower sidebands, giving a figure of –126dBV or 0.5μV for the single sideband noise.

Given the measured sensitivity of the frequency discriminator of 6.1μV/Hz (see below), the rms frequency deviation f_d is 0.082Hz. For sinewave modulation at a frequency f_m, the modulation index $m=f_d/f_m$ equals the peak phase deviation in radians. Now $0.082/2500=3.3.10^{-5}$ radians, and for such a small phase deviation, only the first order fm sidebands are significant.

So if the modulating frequency f_m were a 2.5kHz sinewave rather than narrow band noise, the first order sidebands would each be $(3.3.10^{-5})/2$ in amplitude relative to the carrier, since for small angles, $\arctan\theta=\tan\theta=\sin\theta=\theta$, with negligible error. So the sinewave single sideband amplitude would be simply $20\log(1.65\ 10^{-5})$ relative to the carrier, or –96dBc. This may be taken as a first order approximation to the value of $\mathcal{L}(f_m)$ at 2.5kHz offset, for the circuit of Fig. 2.

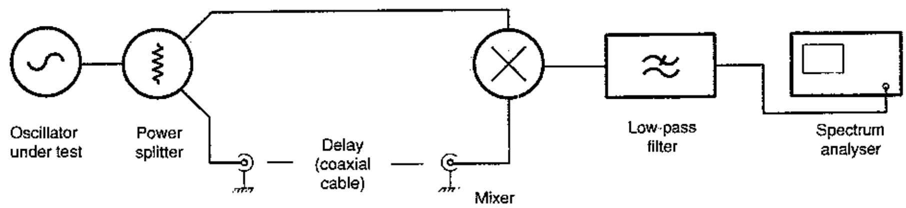

Fig. 1 Block diagram of a set-up to investigate oscillator phase noise.

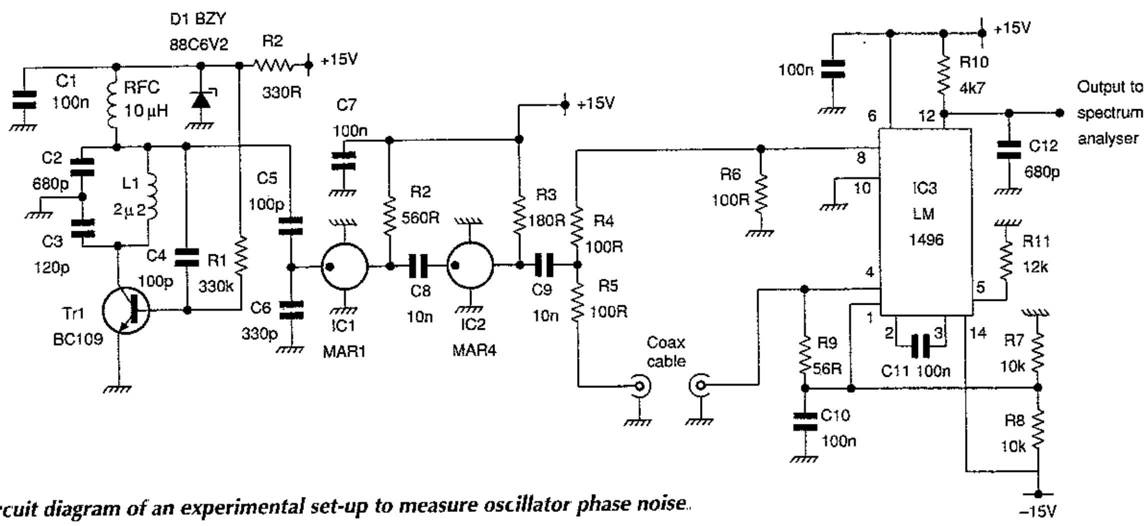

Fig. 2. Circuit diagram of an experimental set-up to measure oscillator phase noise.

SLOT-TEN-1-03 unshielded inductor with a carbonyl E core, having a quoted nominal inductance of 2.2μH and Q of 56 at 7.9MHz. I chose a Colpitts oscillator, as the inductor was untapped, arranged so that the transistor could be operated with the emitter connected directly to circuit ground.

To minimise loading and to maintain a reasonably high working Q, the output was taken from the base end of the tank circuit. This is a much lower impedance point than the collector end, and loading was further reduced by the using a capacitive divider, C_5 and C_6, to buffer the 50Ω input of IC_1. Together with IC_2, IC_1 provides a total gain of 26.7dB nominal, providing a level of –8dBm into 50Ω at the coaxial socket connected to R_5.

Frequency discriminator

The output of IC_2, which sees approximately 50Ω loading, is applied to the local-oscillator port of an active double balanced mixer, IC_3, an *LM1496*. **Figure 3a)** shows the internal circuit of IC_3. The 'carrier' or local oscillator is applied between pins 8 and 10, to four transistors connected in an arrangement often referred to as a Gilbert Cell. The signal input is applied between pins 1 and 4, the signal being steered in phase or in antiphase to the outputs at pins 6 and 12 (note the pin numbers quoted refer to the DIP packaged version of the *LM1496*). The transconductance of the signal long-tailed pair is set by the value of a resistor connected between pins 2 and 3. The magnitude of the tail currents is set by the current injected into the bias port, pin 5.

Figure **3b)** shows how the output at pin 12 is at its maximum positive level if the local oscillator and signal are in phase, is at zero (relative to its level in the absence of a signal input) when they are in quadrature, and at maximum negative level when in antiphase.

If pins 2 and 3 are shorted, so that both signal and local oscillator ports are overdriven – equivalent to squarewave drive in each case – the input phase to output voltage characteristic is linear, as in Fig. 3b), right hand side. If the signal port is operated in a linear manner, the characteristic is cosinusoidal, also shown in Fig. 3b).

In Fig. 2, the signal is applied to the signal input port via R_5 and a length of coaxial cable. The latter provides a fixed time delay, independent of frequency. Therefore if the oscillator frequency is varied, the electrical length of the cable varies, and so the phase of the signal applied to pin 4 of IC_3 will vary. So although IC_3 is a phase sensitive detector, in conjunction with the delay cable it forms a frequency discriminator.

The delay was provided by a reel of miniature polythene insulated coaxial cable, unearthed from my stock of handy bits and pieces. This cable had a silver on copper on steel inner, and could well have been UR94.

Monitoring pin 4 of IC_3 with one 'scope probe and the junction of R_4 and R_5 with the other, the waveforms were found to be in quadrature at 10.377MHz and in antiphase at 9.726MHz.

From these results, and assuming the velocity of propagation in the cable is two thirds that in free space, some simple algebra gives the length of coaxial cable as 14.95 quarters of a wavelength at 10.377MHz, say $3^3/_4$ wavelengths, allowing for experimental error. Thus T_d is 361ns and the physical length of the cable turns out to be 108.4m. I took this figure on trust, rather than unreeling the cable to find out.

Frequency discriminator sensitivity

Maximum sensitivity is ensured by C_{11} which provides an ac short between pins 2 and 3. As the dc resistance between these pins is infinite, in the absence of a signal input, the output sits at the midpoint of the characteristic, despite any small input offset voltage that there might be between pins 1 and 4.

By varying the tuning with the core of L_1, measuring the frequency with a digital frequency meter and the output level at pin 12 of IC_3 with a digital voltmeter, the frequency discriminator characteristic was measured. This is shown plotted in Fig. **3c)**. Due to the limited available tuning range, for the most part, only one side of the characteristic could be plotted, as shown.

The considerable length of coaxial cable used achieved a high sensitivity in the frequency discriminator, but introduced some inevitable attenuation. Consequently the signal voltage swing available at pin 4 was less than the local-oscillator input at pin 8, the attenuation in the cable being some 7dB. The result is that the discriminator characteristic is intermediate between those shown in Fig. 3b).

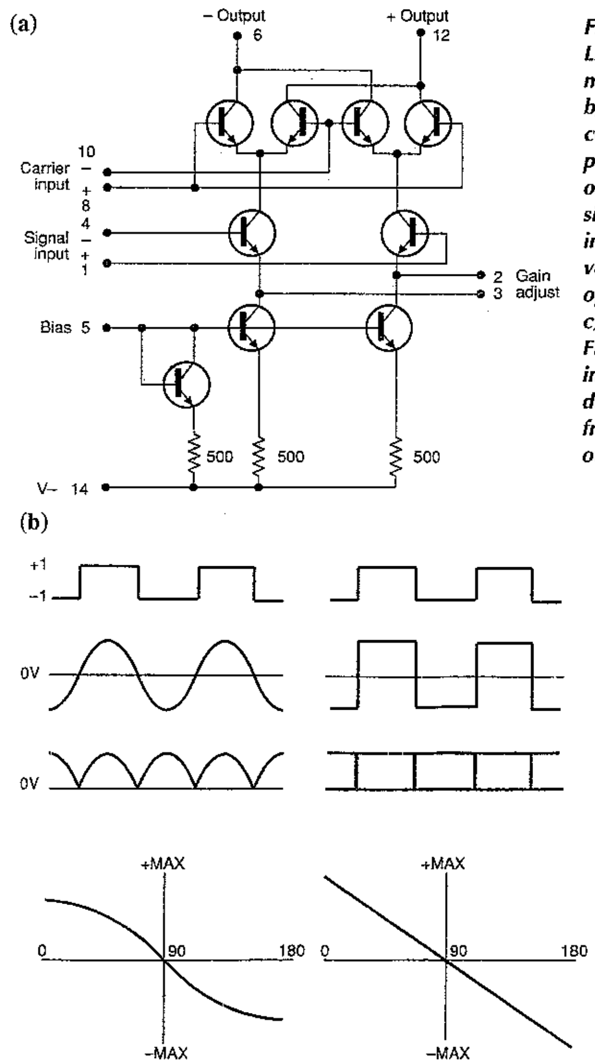

Fig. 3a) Internal circuitry of the LM1496 active double balanced mixer.
b) Showing response of the dc component of output voltage to phase changes between local oscillator and signal inputs, for sine and square wave signal inputs of equal peak-to-peak voltage, assuming linear operation of the signal port.
c) Showing the response of the Fig. 1 frequency discriminator, as implemented in Fig. 2. The black dot shows the measured centre frequency response, crosses show other measured points.

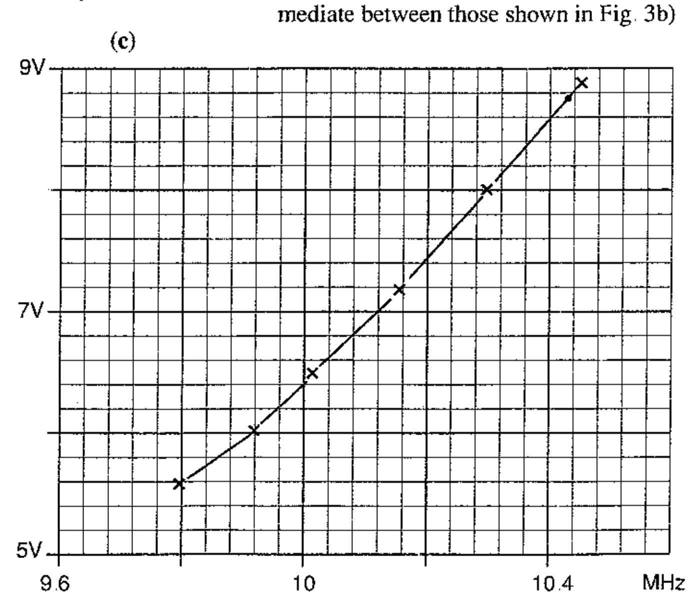

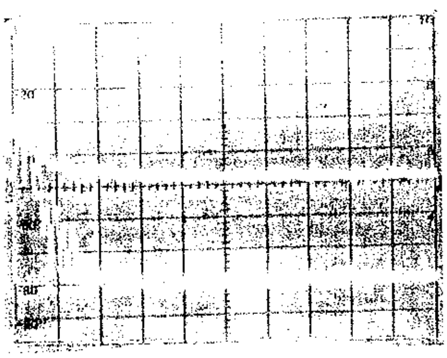

Fig. 4a) Spectrum analyser sweep, 0 to 5kHz, reference level (top of screen) –60dB, 10dB/division vertical, intermediate-frequency bandwidth 30Hz, smoothing maximum, 100s/div sweep speed. Lower trace, with + and –15V supplies off. Upper trace, supplies on, circuit as in Fig. 2.

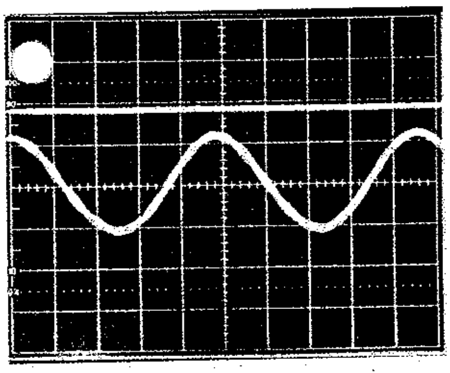

b) Oscilloscope traces; horizontal, 20ns/division. Upper trace, IC_3 pin 12, 5V/division, 0V at centre line, with coaxial cable disconnected. Lower trace, IC_3 pin 4, 50mV/division ac coupled, coaxial cable connected.

Over the central linear portion, the characteristic sensitivity is 164kHz/V or 6.09µV/Hz

Measured results

Output of the frequency discriminator, at pin 12 of IC_3, was connected to an *HP3580A* lf spectrum analyser, via the low pass-filter shown in Fig. 1. Figure 2 shows that the filter consisted simply of the 4.7kΩ phase detector output resistor R_{10}, in conjunction with some 800pF or so. This consisted of C_{12} plus about 100pF due to a screened input lead and the analyser's input capacitance. The cut-off frequency of this filter is a little over 40kHz, well clear of my range of interest, which was in noise sidebands up to 5kHz.

First of all, to establish a measurement noise floor, a spectrum analyser sweep from 0 to 5kHz was recorded with the power supplies switched off, **Fig. 4a)**, lower trace. This shows a measurement noise floor of about 80dB below a top-of-screen reference level of –60dBV, or some –140dBV. At this level, it is difficult to avoid some response from supply rail residual hum, visible as 100Hz and its harmonics at the left hand side of the trace.

Next, the circuit was powered up, but with the coaxial cable disconnected. Figure **4b)** upper trace 5V/division, 0V at centre line, shows the standing voltage at the frequency discriminator output, IC_3 pin 12. This was +8.75V, corresponding to the discriminator centre frequency.

Next, the coaxial cable was reconnected and the lower trace (50mV/division, 20ns/division ac coupled) shows the delayed signal applied to IC_3 pin 4. Some modulation of the trace is visible, but this was still there when the supplies were turned off – it turned out to be pick-up of the local fm radio station. As the frequency is unrelated to the local-oscillator waveform at IC_3 pin 8, it will not affect the result and can be safely ignored.

With the coaxial cable reconnected, the frequency was adjusted to 10.377MHz, by means of the core in L_1. At this frequency the signal input at pin 4 of IC_3 was in quadrature to the local-oscillator input at pin 8, corresponding to zero deviation from the discriminator's centre frequency. This is indicated by the black dot in Fig. 3c).

The oscillator's phase-noise sidebands (on both sides of the carrier) are translated by the frequency discriminator to baseband – from zero hertz upwards. The result is displayed in Fig. 4a), upper trace. This is over 30dB clear of the measurement noise floor, due to the high system sensitivity ensured by the generous length of coaxial cable used.

The corresponding value of $\mathcal{L}(f_m)$ at 2.5kHz offset was calculated as shown in the panel. The result seems plausible, even if only an approximation. However, for the purposes of comparing phase noise with the transistor bottoming, or not bottoming, comparative measurements suffice, and proved revealing.

I needed to know whether the oscillator was bottoming or not. An *HP8558B* spectrum analyser was used to sample the output at the base end of L_1. To avoid excessive loading of the circuit, the 50Ω coaxial lead to the spectrum analyser was connected via a 4.7kΩ resistor.

Figure 5a) shows the spectrum of the oscillator, with settings of 10dB/division vertical, reference level –10dBm, 5MHz/division horizontal, 30kHz intermediate-frequency bandwidth, video filter on maximum. The illustration is a double exposure, showing the output of the circuit as in Fig. 2 (0Hz marker at extreme left), with the fundamental at just over 10MHz and its second harmonic nearly 30dB down. The second trace, with increased Tr_1 base current (offset half a division to the right), shows a larger fundamental and prominent third and fourth harmonics in addition to the second.

The second trace is the result of connecting a 56kΩ resistor in parallel with R_1. Thus the base current was increased by a factor of over six, while the output amplitude increased only by some 8dB or times 2.5. This, together with the marked level of higher harmonics, shows that with the additional base current the circuit was bottoming, but without it was not.

Figure **5b)** shows (upper trace) the 0 to 5kHz baseband spectrum, with the increased base current, resulting in the transistor bottoming. The lower trace is a repeat of the upper trace in Fig. 4a), for comparison.

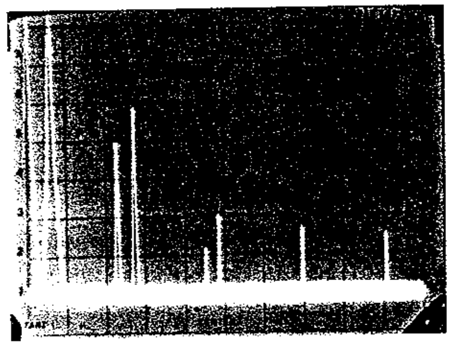

Fig. 5. a) Spectrum of the oscillator, with settings of 10dB/division vertical, reference level –10dBm, 5MHz/ division horizontal, 30kHz intermediate frequency bandwidth, video filter on maximum. Double exposure. Circuit as in Fig. 2 (0Hz marker at extreme left) shows the fundamental at just over 10MHz and its second harmonic nearly 30dB down. Trace with increased Tr_1 base current (offset half a division to the right) shows larger fundamental and prominent third and fourth harmonics.

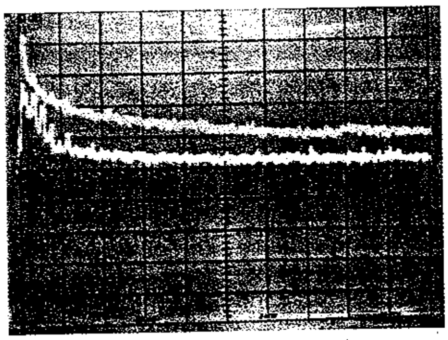

b) Upper trace, 0 to 5kHz baseband spectrum, with increased Tr_1 base current, transistor bottoming. Lower trace, repeat of the upper trace in Fig. 4a), for comparison. Both use same settings as Fig. 4a).

Both traces were recorded with the same settings as Fig. 4a). For this test, care was taken that the signal applied to the frequency discriminator was the same as it was without the increased base current. To this end, after adding the 56kΩ resistor in parallel with R_1, the 100pF capacitor C_5 was replaced by a 5 to 65pF trimmer. This was adjusted to give the same amplitude inputs at the local oscillator and signal ports of IC_3 as previously.

The resultant small shift in oscillator frequency, due to the slightly reduced loading on the tank circuit, was removed by readjusting the core of L_1.

Findings

You can see from Fig. 5b) that in the range above 2.5kHz offset, the magnitude of the phase noise relative to the carrier is nearly 10dB lower when the transistor is not bottoming, than when it is. Note particularly, that the gap widens at lower offsets. This is presumably because bottoming involves higher

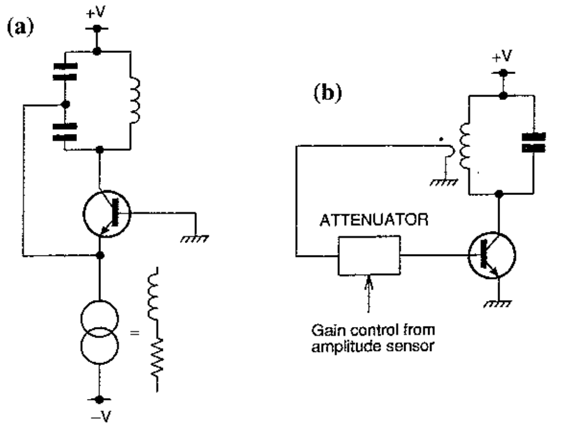

Fig. 6a). Defining the transistor's collector current. By means of a long tail as here is just one of many ways. The resistor may be replaced by the output of a d-to-a converter, permitting adjustment of the tail current under program control.
b) Separating the amplitude control mechanism from the oscillator should permit operation of the transistor in a linear regime. This should result in much reduced phase noise sidebands, by preventing the transistor's $^1/_f$ noise cross-modulating onto the carrier.

order non-linearities, resulting in the transistor's $^1/_f$ noise, cross modulated onto the carrier, effectively extending further out into each sideband.

By 5kHz, the noise, as measured with a frequency discriminator, has clearly flattened out. This corresponds to phase noise falling at 6dB/octave of offset frequency, or the f^{-1} region of phase noise, which continues until the far out noise floor is reached.

At smaller and smaller offsets, the slope becomes greater, f^{-2}, f^{-3} and at very small offsets f^{-4}. This tendency is visible in both traces in Fig. 4a), though setting in at a higher frequency when the transistor is bottoming. As the offset reduces to zero, the amplitude increases, up to the value of the carrier output. The trace in Fig. 4a) does not show this below 5Hz, as this is the low frequency limit of the *HP3580A* spectrum analyser. In any case, the output due to the carrier itself is (near) zero, since the local oscillator and signal inputs are in quadrature.

So when an oscillator with low phase noise is required, a circuit design should be selected which avoids bottoming of the collector. This can be achieved in a number of ways, for instance using a 'long tail' to define the emitter current, **Fig. 6a)**

Where a large tuning range is involved, it may be advantageous to vary the tail current. Assuming capacitive tuning, the dynamic resistance of the tank circuit will increase with frequency. So to maintain a constant amplitude of oscillation the tail current should be varied inversely as the oscillator frequency.

Of course, even when not bottoming, the transistor is still operating non-linearly, the collector current being cut off for part of each cycle. If amplitude control could be implemented independently of the transistor, as indicated in Fig. **6b)**, it should be possible to operate the transistor entirely in a linear mode, preventing the cross modulation of its $^1/_f$ noise onto the carrier output. This is an interesting possibility and one which I intend to pursue. Doubtless this has been done many times already, but I don't recall having seen any published results.

An alternative to Fig. 6b) would be to use a variable gain amplifier as the maintaining amplifier. A suitable candidate would seem to be the recently announced *CLC5523* from National Semiconductor, which I am trying to obtain a sample of. It has a 250MHz bandwidth at 135mW power consumption ■

References

1. Suter, W. A., 'Phase Noise Measurement for Under $250 *RF Design*, Sep. 1995, pp. 60-69
2. Hickman, I., 'The ins and outs of oscillators ' *Electronics World*, July 1994, pp 586-589.