A New Low-Noise VHF Crystal Oscillator Topology

©2008 Chris Bartram GW4DGU[1]

As we go higher in frequency, and as other users of the microwave spectrum become more numerous, we need better phase noise from our local oscillators and other signal sources. There are a number of reasons. On the lower frequencies, reciprocal mixing could be a major reason. Alternatively, people working on the bands above 50 or 100GHz need as coherent a signal as can reasonably be obtained. Multiplication, either by conventional multiplier strips, or by properly designed phase-locked oscillator frequency multipliers, from a very, very good VHF crystal source with the best possible close in phase noise performance, locked to a suitable frequency standard, is still one of the best ways of achieving this.

Much of the work on the new topology reported here has formed part of the uWSDR project.[2] One use of the oscillator within this project will be as part of a very high performance clock for present and future DDS-based microwave synthesisers.

The background to this work goes back a long way. For nearly three decades I've used a VHF crystal oscillator circuit based on one originally described by Driscoll for use with AT-cut overtone crystals at around 5MHz[3]. The original circuit can be easily redesigned to operate with overtone crystals in the 100MHz region. Driscoll published a paper on this, and a version using power FETs was published in the German amateur radio publication 'VHF Communications' during the late 1970s. Recently, I've come to realise that the Driscoll topology can be used as the basis for an even better oscillator.

An attraction of the Driscoll and its related topologies over many traditional oscillator circuits, is that, apart from providing very good performance indeed, it's easily possible to separate the several parts of the circuit, and optimise them more-or-less independently. The oscillator then ceases to be a complex, partly understood entity, designed by folklore, magic and mystic incantation, and can be broken-down and understood in relation to its system components.

Oscillator principles.

In an oscillator described in terms of a feedback loop[4], the condition for oscillation is that the loop gain is greater than unity, and the phase shift around the loop is an integer multiple of 360° at the frequency of operation..

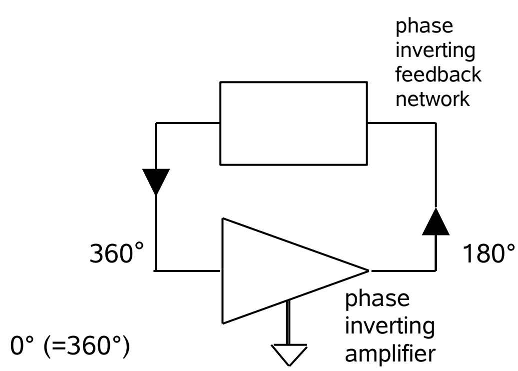

Fig 1. Oscillator with Inverting Amplifier

Phase Noise

Many years ago, Dave Leeson, W6NL, ex-W6QHS, published what must still be one of the most cited papers on oscillators[5]. He analysed the noise performance of a generalised oscillator, concluding that oscillator noise at a given frequency offset can be modelled by three main parameters:

- the loaded Q of the resonator
- the available power at the input of the sustaining amplifier
- the baseband noise power at the frequency offset, also measured at the input of the sustaining amplifier.

In fact it can get quite a lot more complex than that, and there have been subsequent papers which have introduced more complex models in order to describe oscillators more accurately, but as a brilliantly simple guide to understanding and optimising a low noise oscillator, Leeson's original paper is difficult to better.

Essentially, a noisy oscillator can be modelled as a perfect oscillator modulated by baseband amplitude noise. The amplitude noise is converted into phase modulation, and thus phase-noise sidebands.

Based on this understanding, four oscillator design desiderata become clear:

- maximise the loaded Q of the resonator
- use a maintaining amplifier with low 1/f (flicker) noise
- operate the maintaining amplifier at the highest possible signal power level.
- use a formal limiter in the oscillator feedback loop.

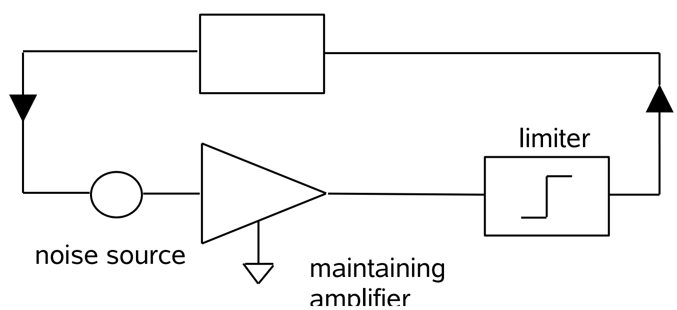

Fig 2 Oscillator Outline with Noise Source

Maximising Loaded Q

In an oscillator using a series resonator, maximising loaded Q is largely a matter of reducing the resistance in series with the resonator. Basic theory says that the loaded Q of a series resonant L/C circuit is a function of the source and load impedances.

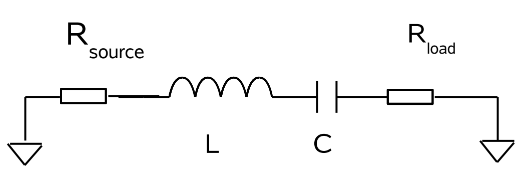

Fig 3 Loaded Q of Series Resonant Bandpass Network

Discounting losses in the reactive components, at resonance, the loaded Q of the network in Fig 3 is given by:

$$Q_L = X_L/(R_{source} + R_{load}) \quad \text{or} \quad Q_L = X_C/(R_{source} + R_{load}) \qquad \text{Eq.1}$$

From that it's clear that the loaded Q can be maximised by using the smallest possible values of source and load resistance. That's a very important concept in the design of series resonant oscillators.

Despite their appearance as a simple two, or three, legged component, crystals are complex electro-mechanical systems, and a full equivalent electrical circuit of that mechanical system would probably take-up most of this page! However a simple model is more than adequate to understand the behaviour of a crystal close to resonance.

At around 100MHz a well made AT-cut fifth-overtone crystal will have a simplified equivalent circuit looking something like this:

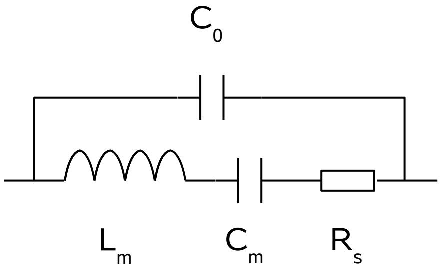

Fig 4. Simplified Equivalent Circuit of a Crystal

C_0 is the static capacitance across the crystal. Effectively, the electrodes sputtered onto either side of the quartz blank behave as a parallel-plate capacitor. It's usual to tune-out C_0 with a shunt inductor resonant at the operating frequency.

The series network of L_m, C_m, and R_s are the electrical consequence of the vibrating mechanical system, and are known as the motional parameters.

For a good fifth-overtone 100MHz crystal the static and motional parameters are typically around:

C_0	≈	6pF
C_m	=	500aF (500attoFarads = 500×10^{-18}F)
L_m	=	5.066059mH
Rs	≥	30Ω

The unloaded Q of that crystal is therefore in the region of 100000. However, be warned, I have measured crystals, (using a vector network analyser and a test jig with 12.5Ω source and load impedances made to the international IEC60444 standard) and seen Qs much smaller than this.

Crystals vary a lot. It is increasingly difficult to buy good crystals on non-standard frequencies, and good crystals aren't now cheap. Be enormously wary of anyone offering to make custom crystals in less than about a month. Crystals can be difficult to specify well, and, sadly there are manufacturers who will take advantage of this. It's important to specify the motional capacitance or inductance as well as C_0 and R_s. Sometimes you'll have to negotiate these parameters with the crystal manufacturer, who may also ask questions about ageing and burst or 1/f noise. If they do, it's probable that you've got a good supplier. The appropriate answer is 'as good as you can do within my budget!'

Optimising 1/f Noise

In any any high performance oscillator it is essential to minimise sources of noise in and around the maintaining amplifier. With very low noise oscillators, this is particularly important. Low-frequency amplitude domain noise from transistors, resistors, the resonator and power sources gets turned into phase noise by a modulation process.

Low 1/f noise from a maintaining amplifier is achieved by selection of a suitable device, operated at appropriate supply voltage and current, by presenting the device with suitable terminating impedances at both baseband and at the frequency of operation, and by ensuring that the amplifier operates as a linear device. Unfortunately, probably because it's difficult to guarantee, manufacturers of rf devices are coy about revealing 1/f noise data in their data sheets, however, it is sometimes contained in device SPICE parameters. Look for KF (1/f noise coefficient) and AF (1/f noise exponent). If these parameters are other than 0 and 1 respectively (the default values) then the 1/f noise model is probably valid at some level. Fig 5 shows a QUCS[7] low-frequency noise simulation of an NE68833. The 1/f noise corner frequency, defined as where the noise level rises by 3dB, is at around 300Hz in this case.

There are probably a large number of good oscillator devices. The NE688xx series is a favourite and it's use is cited in a number of professional papers on low-noise oscillator design. In the past I've used low-noise audio transistors, although they tend to have relatively low Ft and large capacitances which makes the RF design difficult. I have also used JFETS, such as the J310, but as they have less transconductance than typical bipolars, making the impedance at the source greater than that at the emitter of a bipolar, their use can impact

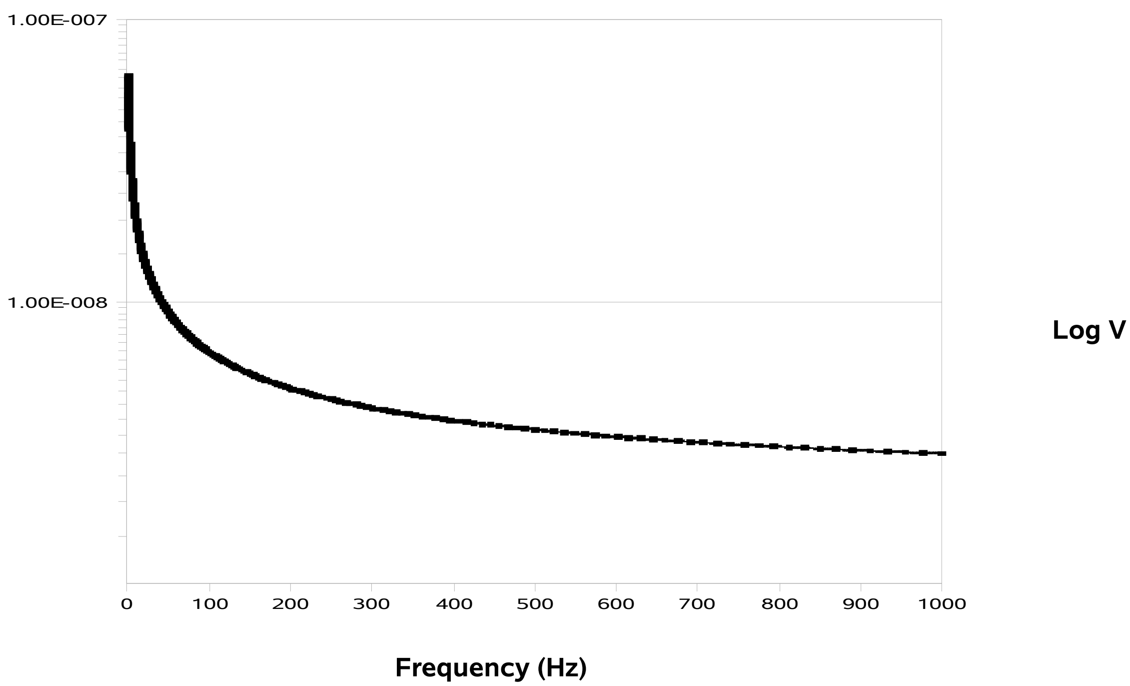

Fig 5 Simulated 1/f Noise Voltage (in 1Hz Bandwidth) with Frequency: NE68833 in Test Circuit

negatively on the obtainable loaded Q. Conventional bipolar junction transistors still seem to be the best solution at VHF.

Crystals can also exhibit 1/f noise characteristics, although with good parts, noise from the amplifier is likely to be dominant.

Another low-frequency noise phenomenon in both amplifiers and crystals used to be called 'burst noise' and seems to increasingly be known as 'random telegraph signals' (RTS). That is getting a bit beyond the scope of this note, and *may* form the subject of a future investigation...

Bias networks are frequently underestimated as potential noise sources. Resistor Johnson noise can be very significant. Close to minimum 1/f noise can be obtained by driving a suitable bipolar transistor with a very low impedance at baseband. This is an area which is addressed by the new topology.

Operating the Maintaining Amplifier at the Highest Possible Signal Level

There is a compromise here, determined mainly by allowable crystal dissipation. Typically 250µW is typically the maximum which should be allowed for AT-cut VHF overtone crystals. The choice of amplifier topology is important in maximising the level at the input of the maintaining amplifier while keeping the crystal dissipation to a minimum.

The Limiter

All oscillators must have some form of limiter or gain control. Most simple oscillators use self-limiting. This leads to a number of complications. With many series resonant oscillators, the primary problems are that as the maintaining amplifier saturates, the active device is cut-off for part of the cycle. That raises the impedance(s) seen by the resonator. With a series resonant circuit, high Q equates with low source and load impedances. Secondly, informal limiting schemes are usually very asymmetric. That leads to a poor amplitude-to-phase conversion characteristic, which is undesirable in a low phase-noise oscillator.

Also, an amplifier which is in cut-off for a large part of its cycle can't be expected to have good noise performance...

As I've already noted, it's essential to keep the power dissipated in the crystal within safe limits. If it is too high, the crystal will drift due to heating, and will age rapidly, due to changes in the crystal lattice. In extremes it can even shatter! I plead guilty, in the paleothermionic era, to having physically shattered a massive WW2-era 10X crystal. A misguided, youthful attempt to use an 807 as a single tube 50W 80m transmitter had something to do with it! Quite recently I was able to permanently damage a good VHF crystal, moving its zero-phase frequency by over 10kHz, and significantly reducing its Q when I forgot to include the limiter in a prototype oscillator...

Another very good reason for including limiting in a crystal oscillator, is that crystals exhibit a marked tendency to 'drive level dependence' of a number of parameters. Usually DLD is only considered with respect to Rs, but it is also present in the frequency stability of the crystal. I've seen a figure of 1part in 10^9 quoted for a 1% change in drive level. In a high performance oscillator, that starts to become significant.

A number of approaches have been used to make a symmetrical limiter. Twenty-five years ago, Plessey manufactured a very low amplitude to phase conversion limiter chip based on a transistor differential pair. I used that in a crystal oscillator at the time, but at that time I didn't have access to the test equipment needed to measure it properly – neither did Plessey - so I've no idea just how well it worked. One potential problem with that part was that the 1/f noise of the device was unknown – in retrospect, it was probably not that good. A matched pair of discrete devices could be used, but they would need to have good 1/f noise performance. I favour monolithic high-barrier dual schottky diodes as 'crossed diode' limiters, as these devices have intrinsically low 1/f noise. The circuit is also very simple!

The use of gain control loops is not necessarily a good idea. The introduction of an amplitude modulator, which will inevitably be a source of amplitude noise, is also likely to have significant amplitude to phase conversion and can be a source of difficulties within the loop bandwidth of the control loop.

A limiter is also a good place to take power out of the oscillator, as the signal at that point, by definition, has very little amplitude noise!

Oscillator Topologies

There are two basic ways of making a series-resonant crystal oscillator.

One way of achieving that is to define the bandwidth of the feedback loop by a crystal acting as a bandpass filter.

The Butler oscillator is one example of this approach. A crystal, series resonant at the desired frequency of operation, is placed in the feedback loop. With proper design, the oscillator frequency will be determined by the frequency at which the crystal has zero phase shift. Unfortunately, this method has some disadvantages.

Most important is that it's quite difficult to operate the crystal in a low enough impedance environment to maximise its loaded Q. There's another related problem inherent in the two-transistor Butler oscillator, widely used in amateur microwave equipment: the input impedance of the common-base amplifier sums with the output impedance of the common collector stage, and the sum of these impedances is then placed in series with the crystal, reducing its loaded Q. A further problem is that the topology also tends to result in relatively large crystal dissipation.

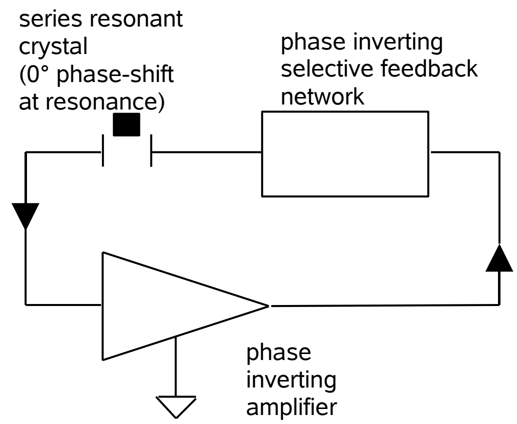

Fig 5 Oscillator with Crystal in Series with the Feedback Path

A second technique is to effectively turn the maintaining amplifier itself into a narrow bandpass filter. This was a technique used in the oscillators employed by Microwave Modules in their transverters and converters. Unfortunately, they used a simple common-base Colpitts oscillator with a series resonant crystal in place of a base decoupling capacitor. This results in quite noisy oscillators, as the crystal operates in a relatively high impedance environment, reducing its loaded Q.

The Driscoll oscillator effectively turns this oscillator around, and uses a common-emitter oscillator with the crystal connected between the emitter to ground, putting the crystal in a low impedance environment, and

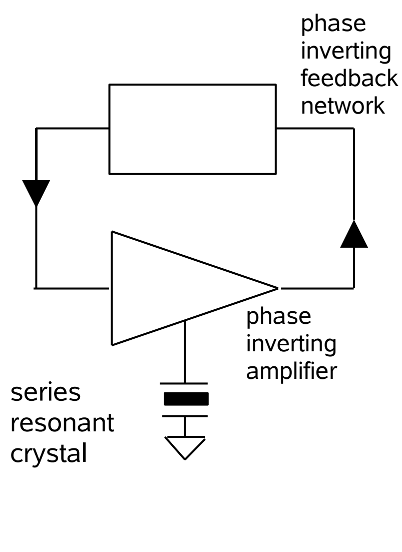

Fig 6 Oscillator with selective amplifier

increasing its loaded Q to a substantial proportion of its unloaded Q. As there is only a single stage maintaining amplifier, the crystal effectively has only around half of the resistance in series with it when compared with a two transistor Butler oscillator. It is therefore possible to obtain higher loaded Q, and thus lower phase noise at a given offset.

In Driscoll-like topologies, the resonator loaded Q has little to do with the loaded Q of the feedback network. Changes in the phase response of the feedback loop will cause small variations in frequency, but providing the loaded Q of the feedback network is kept low, the crystal will be firmly in charge. The best way to adjust the phase of the feedback network, and to trim the frequency of the oscillator, is to vary either the inductor or capacitor in the series arm of the feedback network. Changing the shunt capacitors will have a significant effect on the impedance matching, which is not desirable... One useful indication that a crystal oscillator is properly designed is that the oscillator will operate on a frequency close to the zero-phase frequency of the crystal. Another is that it shouldn't be possible to pull the frequency of the oscillator very far. The pulling range of a good crystal in a well designed oscillator at around 100MHz should perhaps be around ±600Hz. Much more than that and you are throwing away noise and short-term stability performance, although it may be necessary, if the oscillator forms part of a phase-locked loop to allow greater tuning range by reducing the loaded Q a little, ensuring that the oscillator remains locked under temperature extremes and with ageing.

Fig 3. shows that a crystal will have a significant shunt capacitance. It's very desirable indeed to tune that out with a parallel inductor. To frequency modulate the oscillator eg. to tune it for phase-locking, a series L/C circuit resonant at the oscillator frequency in series with the crystal is just about as good as it gets.

Crystal dissipation will usually limit the power level at which the maintaining amplifier can be run.

Typically, with AT-cut VHF crystals, the dissipation shouldn't exceed about 250µW. A pair of 'crossed' schottky diodes across the collector load will result in an output voltage of ~0.7V peak-to-peak. This will be stepped-down to about 0.3V p-p at the base of the bottom transistor of the maintaining amplifier, and if the impedance at the emitter of that transistor is of the order of 30Ω, about half the voltage will appear across the loss resistance of the crystal, resulting in a crystal dissipation of about 100µW.

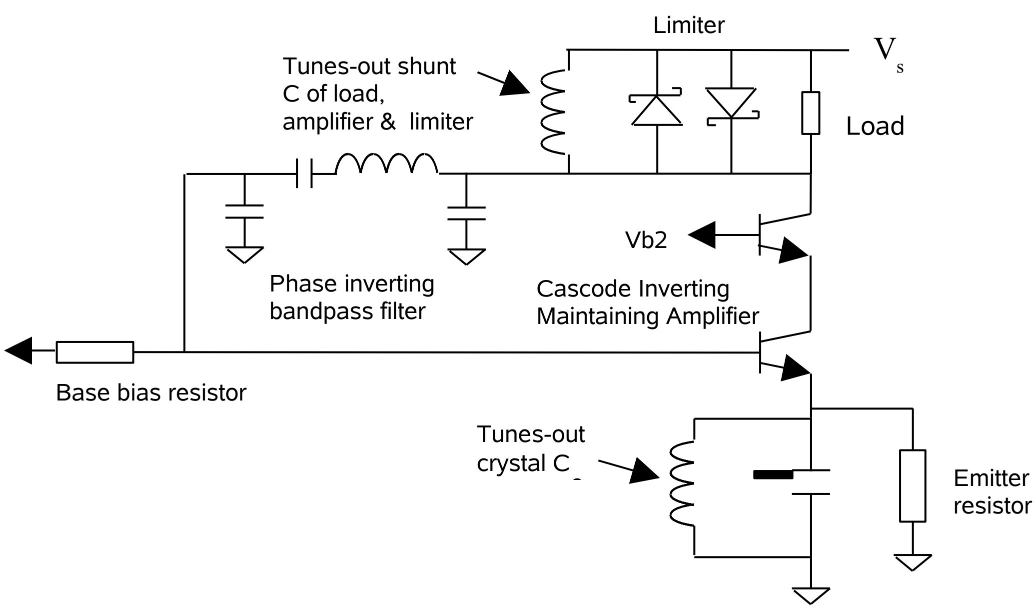

Fig 7. Outline Schematic of Driscoll VHF Overtone Crystal Oscillator.

It is also important to keep the maintaining amplifier operating as a low-noise class-A amplifier. To do this, not only is it important to keep the drive to the amplifier within reasonable limits, so that the noise performance of the amplifier isn't compromised by driving the amplifier into saturation, but the amplifier should present a stable instantaneous impedance throughout each cycle. If the amplifier cuts-off for part of each cycle, the impedance presented by the emitter will rise, and result in a reduction of loaded Q.

With suitable transistors, and clean power supplies, the most difficult noise source to eliminate is Johnson noise generated by biasing resistors. The classic Driscoll oscillator topology uses series resistors to supply base bias. This is more important in the case of the lower transistor in the cascode maintaining amplifier, than in the upper device.

A New Oscillator Topology?

Essentially, the new topology replaces the shunt-C, series-L, series-C, shunt-C bandpass phase-inverting feedback network used by Driscoll, with a shunt-L, series-C, series-L, shunt-L, phase-inverting bandpass network. This has the advantage of eliminating series resistor feeds, and thus the Johnson noise generated by them. It also provides some circuit simplification. A further advantage is that the source impedance seen by the lower transistor at low frequencies approaches zero, which leads to good 1/f noise performance. The combination of these changes leads to improved close-in phase-noise/short-term frequency stability performance.

It is possible to eliminate the series inductor in the phase inverting bandpass filter. This can result in a 'peaky' highpass filter with the correct phase response. This approach is fine for most 3rd overtone crystals, although the extra selectivity of the bandpass network will probably be needed to reduce the risk of higher order overtone crystals oscillating on unwanted overtones.

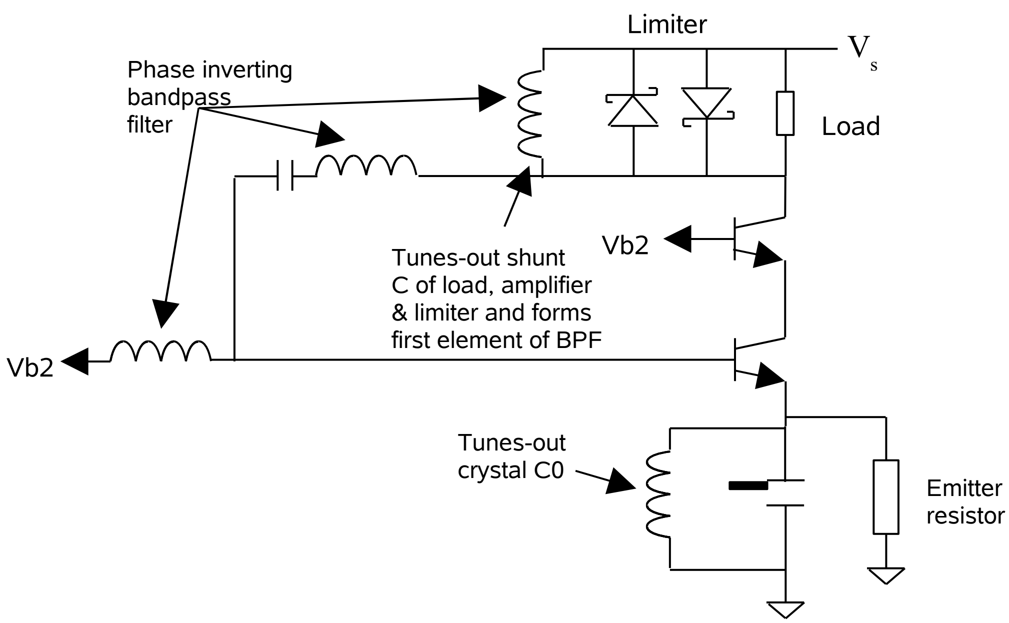

Fig 8 Improved Overtone VHF Crystal Oscillator

A Real Oscillator

The schematic of a prototype of my new oscillator is given in Fig 8. In this I initially used a crystal which was manufactured about a decade ago, and which had lain unused. Like any new crystal it took a little while to settle-down, but after a month at room temperature was giving frequency stability of significantly better than 10^{-10} (root Allan Variance) over 1s, when medium-term thermal drift effects were removed statistically. My measurements have been made with an Agilent 53181A reciprocal counter with ~1mHz (yes, millihertz...) frequency resolution for a 1s gate time, clocked with a Rubidium source.

Direct measurements of phase noise at the level expected from this oscillator aren't easy. I don't have ready access to specialist test equipment, and have been slowly building my own test set-up using general purpose lab equipment. Initial measurements of a 125.000MHz oscillator, using a crystal discriminator, currently show <-130dBc/Hz at 1kHz offset, although that appears to be limited by the baseband LNA I am currently using. Judging both by the baseband noise spectrum I currently see, and from calculations based on a knowledge of the performance of the individual system components, it seems very probable that this will measure at least 10dB better, once I've completely debugged my test equipment. If that is true, the phase noise performance at 1kHz will be 10dB better than measurements reported of current highly respected Butler oscillator designs. This advantage is likely to increase closer to the carrier.

'Pushing' or the change of frequency with supply voltage has been measured on my prototypes as of the order of 0.15ppm/volt. Again, a very good result.

The performance of this circuit is partially limited by the quality of the crystal employed. I believe the basic design would be usable with SC-cut crystals, and could potentially give even better performance than with the AT-cut crystals I've used so far. Another possibility which I plan to investigate is to use a coupled pair of crystals in order to increase the effective loaded Q.

In my prototype, all of the capacitors were NPO/COG 0603 parts at 2n2 and below. Larger capacitors had X7R dielectrics at 100n and below and X5R above that value. The electrolytics were conventional aluminium parts. All of the inductors were Coilcraft 1008CS wound surface mount components. The ferrite beads were standard 0603 600ohm impeders. The choice of oscillator transistors is important. Substitution with devices such as the BFR92/BFR93 could seriously damage the potential performance. The BGA616 is a low cost HEMT MMIC by Infineon which costs about US$0.35 and is useful all the way to 10GHz.

One of the early prototypes of the oscillator was donated to the builders of a new 24GHz beacon. This generated enthusiastic responses which commented on the purity of the note at the output frequency. While no measurements have yet been made, simulations suggest that the phase noise at 100Hz offset could show an improvement of 10 or even 20dB over existing amateur designs.

This project is still very much a work in progress, and is presented here as such. I feel that the oscillator is a significant step forward from traditional designs, and is likely to be very useful to the amateur microwave and mm-wave community. The information presented here isn't a step-by-step construction project, and unfortunately I don't have the time to provide 'hand-holding' support, as I do like to get on the air sometimes! However, I'd be happy to hear of the experiences of anyone trying to duplicate the design. I'll try to answer reasonable questions, and I'll place any new information, and the PCB Gerber files on my website.

Notes:

(1) Blaenffôs, Pen-y-Bont, CAERFYRDDIN, Cymru. SA33 6QG. UK. email: <gw4dgu@blaenffos.org>. Web: <www.christopherbartramrfdesign.com/blaenffos/indexgw4dgu.html>

(2) <uwsdr.berlios.de>

(3) Michael M.Driscoll, Two-Stage Self-Limiting Series Mode Type Quartz Crystal Oscillator Exhibiting Improved Short-Term Frequency Stability, IEEE Transactions on Instrumentation and Measurement June 1973, pp130 – 138.

(4) The other way of describing an oscillator is in terms of a resonator excited by a negative resistance generator. Both negative resistance and feedback analyses are two sides of the same coin, and the

choice really depends on the available data. Negative resistance analysis tends to be more appropriate, for instance, for a microwave VCO, whereas the feedback loop model is my usual choice at VHF and below.

(5) D.B.Leeson, A Simple Model of Feedback Oscillator Noise Spectrum, Proceedings of the IEEE, February 1966, pp.329 - 330

(6) Leeson describes the carrier to noise power ratio (in dB) in one sideband at a given frequency offset from the carrier:

$$L(fm)=10\log\left[\frac{1}{2}\left(\left(\frac{f0}{(2\,Ql\,fm)}\right)^{2}+1\right)\left(\frac{fc}{fm}+1\right)\left(\frac{FkT}{Ps}\right)\right]$$

where:	L(fm)	the relative noise power in dBc/Hz at the given frequency offset
	Ql	the loaded Q of the resonator
	fm	the frequency offset
	f0	the oscillator frequency
	fc	the flicker noise (1/f noise) corner frequency of the sustaining amplifier
	F	the noise ***factor*** of the sustaining amplifier at the offset frequency
	k	Boltzmann's constant ($1.38.10^{-23}$)
	T	the physical temperature of the amplifier (°K)
	Ps	the oscillator power available at the input to the sustaining amplifier

(7) <qucs.sorceforge.net>

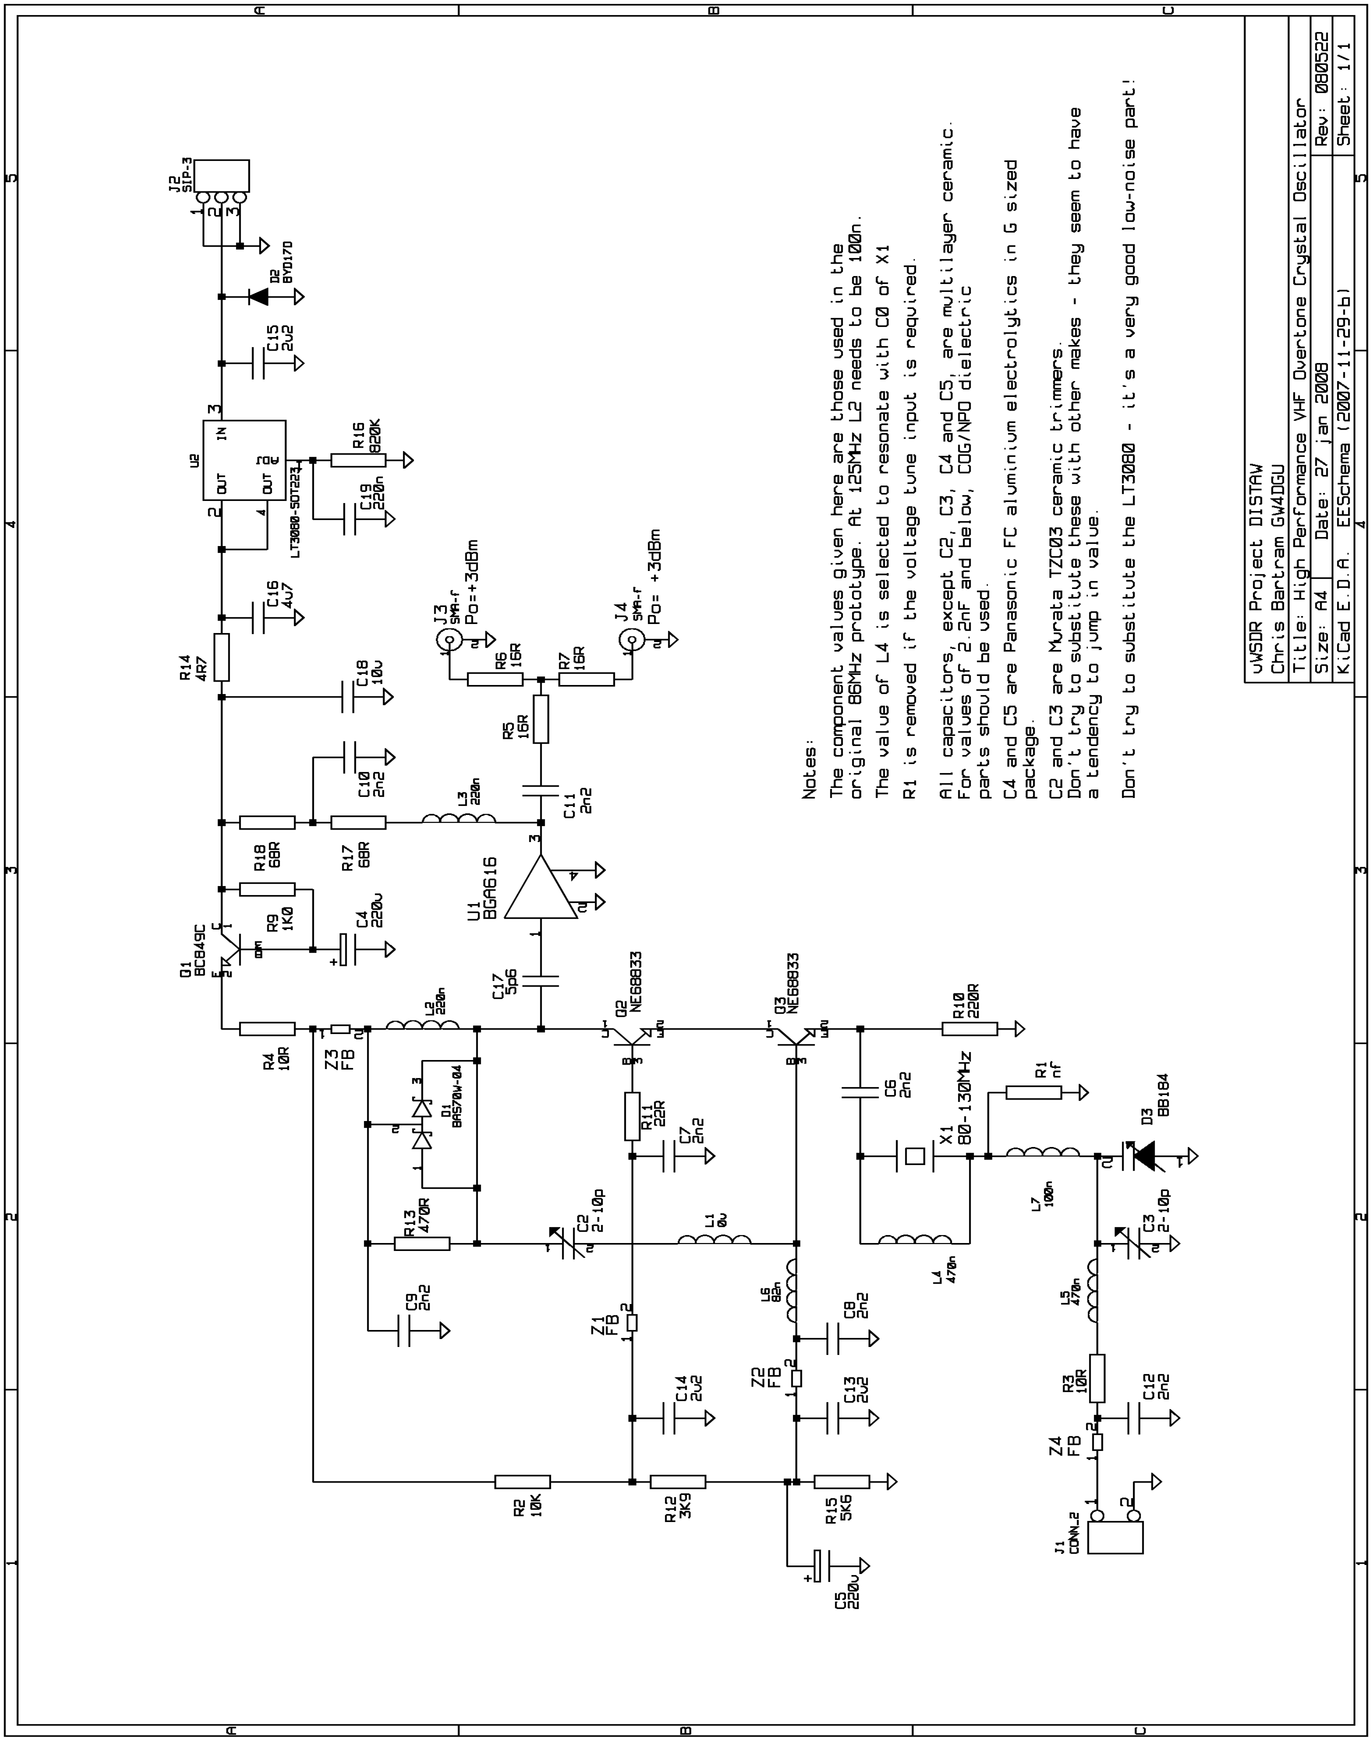

Fig 7 Schematic of Prototype