Chapter 1

Frequency Sources

Crystal Oscillators

Microwave Oscillator Stability

What to consider in achieving "Best Performance"

John Hazell, G8ACE

This article discusses some of the considerations governing the stability of a quartz oscillator source.

Temperature dependence

Almost without exception, crystals are marked with their operating frequency. However, **operating temperature** is vitally important for stability and, for most crystals, this information is either missing or hidden within a manufacturer's code. The family of curves in the graph below show the turnover points for AT cut Crystals. The angle of cut governs the turnover or best operating temperature and this angle will typically lie between 34° 58' and 35° 13'.

The ideal condition is to operate the crystal on the flat part of its characteristic. Typically a crystal cut between 35° 0' and 35° 2' would be used at ambient temperature.

Over the wider temperature range of O°C to 60°C it might exhibit a stability of 5ppm or 100kHz variation at 10GHz. The need to control the temperature becomes obvious and to do this the crystal must be held above ambient in a Crystal Oven or Murata-style heater clip. Superior results can be obtained by using a crystal whose angle of cut results in a turnover point matching the heater. Finally, the oscillator circuitry may contain components whose characteristics vary with temperature and hence the frequency will still vary. One solution here is to also temperature stabilise this circuitry.

An ovened oscillator system, set to

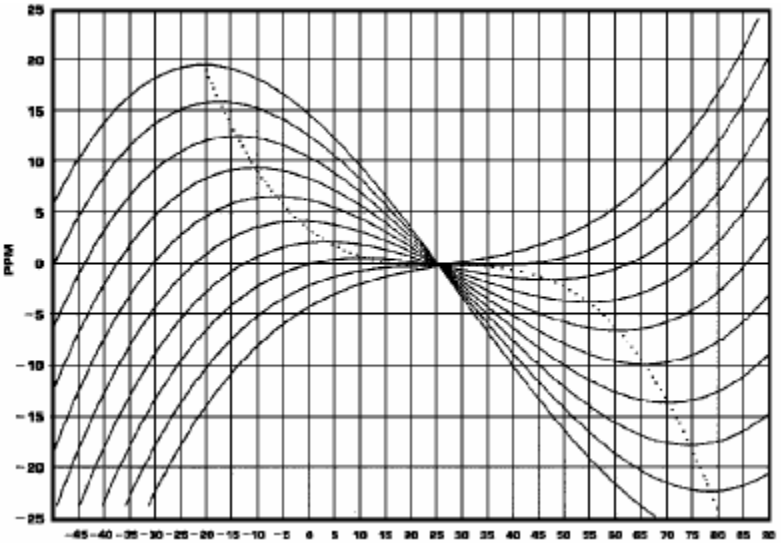

1°C of the crystal turnover temperature, would generally provide a high degree of stability. An ovened home-brew 10MHz reference set to 1° of turnover returns better than 1 part in 10**8**/°C stability after six months ageing, or 100Hz/°C at 10GHz.

Ageing rate
A new crystal will typically age at 0.5ppm/month (5kHz shift at 10GHz) with the figure declining substantially after the first few months of continuous operation. Cleanliness in processing, the type of mount and the seal type all affect ageing. The norm is for the frequency to increase in glass enclosures and decrease in metal.

Re-stabilisation
Once a crystal oscillator has been switched off it will take some time to return to a given ageing rate.

Frequency retracing
Again, once the oscillator is switched off, say at the end of an event, and re-powered some weeks later, there will be a frequency offset, even after warm up. This may help to explain frequency calibration errors between events.

Considerations
Microwave operators will know that the two most important factors in finding that distant signal are dish alignment and accurate frequency readout. The advantages of oven control on the oscillator are well worth while but, for full advantage, try doing some or all the following:

- In the case of a home constructed oven, measure crystal frequency against thermostat voltage in determining the turnover point of the crystal. Use this voltage as the set point for the oven comparator. Make sure the crystal turnover temperature is above the oven self-heating value so the oven is in control for all ambient operating temperatures.
- Age the oscillator for as long as possible and this generally means months, at least, to see stabilisation of the frequency.
- Avoid switching the unit off. Perhaps diode steering in the power connection arrangements is worth considering so that the oscillator can be moved from mains to battery power without breaking the supply.
- Adequately insulate the crystal oven. When large ambient temperature changes are encountered, an additional outer oven may be required.
- Short term stability ... some crystals exhibit large frequency fluctuations. One cause of this appears to be stress in the connections to the quartz. Heat treatment has been shown to improve this problem.
- Power supply voltage/frequency dependence, shock, vibration and humidity are all other factors to be considered.

Crystal Oscillators

Some ongoing discussion from the UK Microwave Reflector

This article is not the usual formal technical article, but an excellent example of how an internet reflector should be used …useful and informative discussion between three or four experienced microwavers, posted for all to read. We wish there was more of this kind of discussion on the various reflectors instead of the "back biting" that we sometimes see! So our thanks go to the authors of the posts you are about to read.

Kevin Murphy, ZL1UJG, posted the following on the UK Microwave internet reflector. It started off a most interest discussion as emails went back and forth over several days …..

I am working on some crystal oscillators for beacons. I have a G4DDK004 oscillator and also some Minikits EME65 PCB's, which both use the 2 transistor Butler crystal oscillator. I have looked at the DC across the V (base-emitter) on the 2 transistors (using a 10k resistor to isolate RF from the digital voltmeter and they both are below 0.5 Volts. (below 0.3 Volts in some cases). The voltage (b-e) on the transistors should be 0.6 to 0.7V so that constant source and load terminations are seen by the crystal. I remedied this by putting two 1N4148s across the tuned circuit, which gave 0.58 to 0.7V across each transistor base emitter, with minimal drop in harmonic output.

I added a choke across the crystal to cancel out the stray capacitance of the crystal holder (5 to 7 pF). I additionally raised the impedance of the 1st transistor, by putting a small resistor in series with the base (~ 30 ohms SMD) Measurements indicate that this raises the RF level and should increase the carrier noise, although it will degrade crystal operating Q. I used a resistive 20dB probe to make some comparative level measurements. Although the oscillators should run off a low noise regulator, these ones are run off a bench PSU. The complete oscillator will be run off a heavily filtered 7808 regulator, with the 2 transistor oscillator being run off an active finesse filter. (see www.wenzel.com). I measured the wideband noise and this measures around -160 dB below the carrier (referred back to the crystal oscillator frequency). (The instrument appears to have some margin left. (I can see the roll off of the wideband noise, due to a pipe cap filter used in a 13 cm LO output.)

I looked at the G8ACE website and the wideband noise floor of his oscillator, referred back to the crystal frequency, appears to be of the same order. Then I came across this link: **http://www.nitehawk.com/sm5bsz/linuxdsp/hware/lotest.htm** I understand what is being said and tried a 1uH in series with the emitter resistor of the 1st transistor of the Butler 2 transistor oscillator and, hello! ….I measured no difference!

The comments indicate that the KD6OZH oscillator is really good (-178dBc for wideband noise) (-172dBc in the original article). Have I missed something? Should I try a FET as the 1st device, since it apparently has less noise? Is the noise generated by the BFS17 (which I have

used for the 2 oscillator transistors) that much higher? Why is the G8ACE oscillator apparently not as good as the KD6OZH oscillator? Instrument limitations? Is there another noise source in the oscillators? I know that PSU noise from 78## and 78L## regulators is an issue. I notice that these 2 transistor Butler oscillators, have an extraordinary tuning range, pulling the output frequency +/- many 10s of kHz on the 13cm band, before the crystal oscillator stops. Is this normal? I expect the FET/bipolar version from G8ACE, would have less tuning range as the overall gain in the oscillator circuit is less but is the sensitivity to tuning similar? I know there are probably better oscillators, but I haven't seen any schematics/circuits.
It has been an interesting exercise taking measurements and learning a little more about oscillators.
Comments and discussion please!

Following Kevin's post, the following replies were received…..
From SAM JEWELL, G4DDK, <jewell@btinternet.com>
KD6OZH states in his QEX article that it is important to achieve and retain low noise figure and high linearity in the second stage of the oscillator to obtain the very low noise performance he claims. The use of a cascode second stage can sometimes make this difficult to achieve because the available supply volts is shared between the two devices in series, which might lead to less than optimum performance. My experience is that you need to use a transformer coupled output or a tuned output although the use of low voltage silicon/germanium devices in the cascode might also work.

From John Hazell, G8ACE: The importance of Crystal Temperature Control
The debate on oscillator stability has to date been centred around the Q of the crystal and other parameters of the RF circuit used. Crystals are unstable with temperature and this stability aspect needs addressing just as carefully as the RF electronics. The following comments relate to AT cut overtone crystals as commonly used in microwave equipments. The cutting angle of the quartz controls the knee point. That is the operating point at which the frequency shift is least with temperature change. Common cut angles give knee points of 25C, 40C and 60C although since the angle is so critical some variation from the expected point will be found. 25C is satisfactory for low frequency equipment with little internal heating and has the advantage that the knee slope angle is shallow with temperature changes. 40C is useful for clip heaters but if the equipment temperature rises above the clip temperature then the stability is lost. 60C ensures that for the majority of the time the crystal is above the influence of ambient temperature changes within the equipment. In the G8ACE MKII OCXO part of the setup procedure is to explore for the knee point during the alignment. This is for two reasons: firstly, suppliers do not necessarily supply what is required and prefer their own spec of quoting stability in ppm over a given temperature range. Most cut angles will easily fit into this spec.

It is most important when ordering a crystal to quote for OCXO use and the required turnover point. Secondly, as mentioned above, the precision required for the cutting angle leads to some variation anyway in the actual knee point and only by individual oven adjustment will this point be found with certainty. A guide to the enormity of the frequency change with temperature is that a commonly used 106.5 MHz/60C crystal will move some 800Hz during the heating phase, starting from 15C. Multiplied by 96 for 10GHz this

represents some 75kHz of system calibration shift. To confirm this point consider a commercial 5 or 10MHz OCXO where large shifts in frequency occur during warm up but then is quite stable. Once the knee point has been determined, the stability can be good but some crystals do move in frequency very slowly. This is assumed to be part of the ageing process. However, it has been observed that small adjustments to the oven temperature around the knee point affect the amount of frequency creep. For a crystal operated for say 8 hours daily the crystal will retrace overnight and each subsequent day will repeat the process. Powering the OCXO 24 hours a day the frequency creep will slow down but there will be variations between samples. It has been noted that the creep amount is dependent on exactly where the oven temperature is set about the knee point. Nothing has been found in documents relating to this unless the writer has misunderstood that published.

Eric, F1GHB has spent some time also looking into the finer aspects of stability and has kindly sent his results which are shown in page two of his pdf document. For the advanced setup Process, access to either GPs or Rubidium or similar source is necessary. Once the creep rate of the OCXO has been determined the OCXO temperature is adjusted in increments to +/- 0.5v either side of the knee point and the change in the frequency creep noted until the flattest time/frequency response is achieved. In the absence of a high stability source for advanced setup a beacon could be used in the 3cm instance but it must be chosen with care. GB3MHX, whilst free running, is quite stable. GB3SCX is also very stable with Hz errors quoted by Andy G4JNT from time to time. Other beacons driven by Adret sources could also be suitable. This document plus the additional information from Eric F1GHB can be found at these sites: **www.microwaves.dsl.pipex.com/mk2/advanced.pdf** and **www.microwaves.mcmail.com/mk2/advanced.pdf**

From Chris Bartram, GW4DGU

I agree completely with John about finding the inflection/turnover/knee point being the key to good OCXO performance. Trying to significantly temperature compensate oscillators with compensating capacitors is a bit of a nightmare - I know, I've been there! For amateur usage, over a limited temperature range, it's probably possible but it's a slow process, particularly as you have to wait for the oscillator to reach thermal equilibrium every time you change a capacitor. One of the problems with trying to temperature compensate a crystal oscillator is that the frequency/temperature curve isn't linear. Also, you really need some form of controllable thermal environment. Very many years ago, Dave Tong (of Datong fame) wrote a very nice article about compensating VFOs using that technique in 'Wireless World' which very succinctly details how to do it. That can be applied to crystal oscillators. Julian, 'YGF, described a circuit using a thermistor and varactor in Radcom about 15 years ago. There's also a good book on the subject by a guy called Marvin Frerking, if anyone is seriously interested.

For low-power portable operation, I'd probably find a good (0.5ppm) TCXO at a rally (or even buy a new one!!) and lock the DB6NT oscillator to it using one of the techniques that Andy, 'JNT, has recently described. That will give at least an order of magnitude better frequency stability than a simple crystal, and I'd be prepared to lay good odds that it could be done with a current drain of <25mA at 12V. Kevin's measurements of his

oscillator are very interesting. He's seeing better performance than I would have expected from my simulations, but, from memory, he's breaking the loop differently, and we may well have used different component values. Mine were lifted from the RSGB Microwave Manual. My simulations of three years ago are now stored on a deep backup CD somewhere.

The important points to note are:
a. He's clearly got a good crystal! But then Rakon (a NZ company) is a VERY good crystal supplier, a world leader, in fact. I've specified them for work projects where excellent crystal performance was required.
b. He shows the reduction in loaded Q with drive-level very clearly.
c. >15dB open-loop gain is grossly excessive and will cause the oscillator maintaining amplifier, operating with the loop closed, to be cut-off most of the time, resulting in very severe degradation of the loaded Q.

Typically, I'd design a crystal oscillator for about +6dB open-loop gain. That would ensure reliable starting. I'd also use a separate limiter, such as a pair of 'crossed' schottky diodes, so that the amplifier is forced to operate in class-A. This could also be used to control the crystal dissipation. There are other ways of making a suitable limiter, such as using a 'long-tailed-pair' (I did work on this many years ago using a Plessey (remember them?) IC limiter, and the idea found its way back into the device data sheet!) but 'crossed' Schottkys are a good solution and much simpler to implement.
There are a few points to note about the use of formal limiters in an oscillator design. Firstly, because the amplifier is operating in class-A, large levels of harmonics are not available. I like balanced schottky diode frequency multipliers, as they add almost no additional phase noise.
Alternatively resistive mode FET or BJT multipliers can be good but beware of highly efficient BJT multiplier, as they are probably at least partially employing some form of parametric effect, and can act as highly efficient phase modulators, transferring supply-line noise to the carrier. Secondly, it's quite a good idea to take the oscillator output across the limiter, as the amplitude domain noise will be 'squashed'. Don't try to take the output from in series with the crystal - you'll just reduce the loaded Q. There's no free lunch there!
A third point to note is that the use of AGC instead of a limiter to keep the amplifier in class-A can easily lead to the generation of phase noise due to the amplitude-to-phase conversion in the gain control device. I hope that's useful to someone.

For what it's worth, it doesn't come as a surprise that changing a crystal in an oscillator to one on a frequency not so different from the previous one results in the crystal coming-up on frequency! If the oscillator had been set-up to put the frequency of the first crystal on its nominal frequency and the loop phase response of the maintaining amplifier was essentially unchanged, then I'd expect the second crystal to also come-up on frequency, give or take the crystal calibration.
There's no reason why the Driscoll oscillator can't be turned into a VCXO. It's entirely possible to tune it, and in an more controlled way an the current alternative. The balloon-board 'Butler' works quite well as a VCXO in a rather uncontrolled manner because the loaded Q of the resonator is relatively low. But, it has disadvantages. In particular, both ends of the varactor diode float above RF ground, and it's quite difficult to get good loaded Q, which means that you won't get the best phase noise, and short/medium term frequency stability. Why do I call the emitter-coupled oscillator, known to the amateur microwave community as

a 'Butler' oscillator, the 'balloon board' oscillator? The ancient history is that in the early days of narrowband microwave operation the blessed Mike Walters, G3JVL, was able to procure a significant numbers of a meteorological balloon transmitter which was designed to send data as FSK at about 400MHz. This was an empirical design produced before even most professionals had access time but it was designed to accommodate relatively wide shift direct FSK at UHF, not as a high quality source. It was never designed as an oscillator intended for multiplication to microwave frequencies! Nonetheless, as a convenient source of 384MHz it was used in most of the early 10GHz transverters as a driver for a step-recovery diode multiplier, and formed the basis of a number of subsequent 'designs'. I know, as the result of simulation and analysis, that the 'balloon board' oscillator isn't terribly good. That was clear even with the relatively crude software tools we had 20+years ago.

I've been using both FET and BJT Driscoll-type oscillators for the highest performance applications since the early 1980s, and I've yet to find a better circuit topology for a VHF overtone crystal oscillator. My transverter designs of 20+ years ago for muTek used a single JFET Butler oscillator. (I think DB6NT still uses something similar.) This was acceptable at VHF, but I wouldn't use it at higher frequencies, as it's quite difficult to control crystal dissipation, so it drifts and the loaded Q isn't **that** good. The phase shifts, and thus frequency changes, due to temperature dependent changes in maintaining amplifier component values make it unsuitable for applications which have to meet commercial approvals even at VHF. In other work applications I've used the impedance inverting overtone oscillator with good results and a single transistor version of the Driscoll.

So you think you know all about crystals?

Part 1

From the WA1MBA Microwave Reflector

The following discussion appeared on the WA1MBA Internet Microwave Reflector. We are very grateful to Bob, K3VOT, for allowing us to publish his email information in the form of this article. Bob really knows what he is talking about! While some of his comments are more relevant to American microwavers than to us in the UK, we have not omitted them as they still make interesting reading.

From: K3VOT [Bob@aurand.com]
My qualifications? 25+ years working in the crystal oscillator and microwave oscillator business, including custom microwave bricks, high stability OCXO's and TCXO's, synthesizers, etc. I currently work as engineering manager for a large international corporation in their crystal oscillator division. In a typical year we buy, oh .....400,000 to 500,000 crystals.

Crystal Suppliers:
First of all, I sometimes use International Crystal for home projects, consulting jobs and occasionally for projects with my current employer. I've had excellent success with them, their major shortcoming being lack of capacity—they just can't deal with the kind of large orders we place. Their chief crystal engineer and a good contact there is Darrell Brehm, WA3OPY. Darrell has been around, having worked at McCoy and Reeves-Hoffman before ICM. Some (all?) of the other crystal makers mentioned on this reflector are small outfits and I have no experience with them except Colorado Crystal. Colorado is probably in the top tier of crystal makers, worldwide, and they have an excellent hi-precision process. Contact is Tom Schulyer, KD0JP I believe. In the last 5 years or so, much consolidation and acquisition has occurred in the US crystal industry. Colorado has been one of the survivors and, as a result, has more work than they need or want. Lead times are very long ... for work, we are regularly quoted 30+ weeks.

A good resource for finding crystal and oscillator vendors is the IEEE Frequency Control Society. They have hotlinks on their web site to most of the major suppliers worldwide. Lots of good technical papers here too. I recommend John Vig's tutorial on crystals and oscillators as required reading for all experimenters. The IEEE web page that links to crystal and oscillator related information is:
http://www.ieee-uffc.org/index.asp?page=freqcontrol/freqmain.html&Part=5#top
Crash around in there and you will find lots of good stuff!

Crystal accuracy and behaviour:
Firstly, virtually no one makes solder seal crystals anymore. It's a dirty process, leaving splatters of solder and flux inside the can. Any contamination such as this will move around, eventually landing on the electrodes, changing the frequency. In the biz we call this the quartz to crud ratio. For various reasons the crud can subsequently leave the electrode area only to return later, changing the frequency every time. Additionally, solder sealing is intrinsically an "atmospheric" sealing operation. So, what ever contaminate (or moisture) was inside the building at the time gets sealed inside the can to do

it's damage over time. Modern crystals are invariably either resistance welded or cold welded. Resistance welding is far more common for several reasons: commercial welding equipment can be purchased -- cold welders are usually custom built. R-weld is faster and therefore more suitable for production environments. R-welders are easier to operate and maintain. Among oscillator engineers, it's generally believed that cold welding results in better performance, but some crystal engineers swear resistance welding can achieve the same performance. The jury is still out on that one, as far as I'm concerned. All UM-1's are resistance welded. TO-3 and other 2 leaded crystals can be either R weld or cold weld. Look at the base ... if it's one large piece of glass with the crystal leads poking through, it's cold weld. If the base is mostly metal with matched glass eyelets for the leads, it's R weld. No crystal maker can get the calibration right unless you, the user, supply the right information.

For fixed frequency oscillators (not voltage controlled) this is largely the load capacitance, sometimes called "finish point". "Series" is a legal load corresponding to infinite load cap. Series crystals vs. load crystals are all made the same way. The difference is made up in the finish plating, what someone on this reflector referred to as the final electro deposition of electrode material and is strictly a matter of where on the reactance curve to put the frequency. All crystal makers need some kind of tolerance on load cap (if used) and frequency tolerance. +/- 0.5 pF is typical for load and +/- 10 ppm for frequency. It's up to you to adjust your circuit to frequency (try **THAT** on 400,000 oscillators a year!). Remember, at 1296 MHz, 10 ppm is 12.96 KHz! If the crystal is used in a oven, the operating temperature **must** be specified and given a tolerance. +/- 5 degrees Centigrade is typical. Usually, with OCXO's, the OSCILLATOR maker adjusts the temperature of each oven to match the turnover temperature of the crystal and each crystal in a batch will have slightly different turnover within the tolerance. Don't expect a well-centred process with a Gaussian distribution unless you are buying a VERY LARGE number of crystals. In practice, this means, when you are replacing a crystal in an oven, say to change frequency of a brick, and the documentation says maybe 75 degrees, that's the NOMINAL design temperature for the crystal. Maybe the old crystal came in at 73.4 and the brick manufacturer set the temp on that individual oven there. Now you buy a crystal from say, International, and give them the 75 degree spec but they end up at 78 degrees. When this crystal is installed in the 73.4 degree oven, it will run high and on a slope. If the oven temperature creeps up over time, due to ageing of the thermistor or other components, the frequency will drop towards the turn and this change in frequency will be indistinguishable from ageing.

By the way, this oven setting information applies to AT cut crystals, only. SCs behave differently but are unlikely to show up in ham equipment. NEVER expect a crystal calibrated at room temperature to be on frequency when used in an OCXO. Expect a 20 to 30 ppm drop in frequency (for the same load) when the crystal gets to oven temperature.

Crystal Oscillators circuits:

No oscillator circuit is perfect but by far my favourite is the Butler. It can easily work from low MHz frequencies to VHF. It's easy to set-up, accommodates both fundamental and overtone crystals and can deliver very good phase noise. We run thousands of these using 155MHz fundamental

crystals each week ... no problems.

Here's how to do it:
Select a high Ft transistor ... my favourite is the NE856 from NEC and available from Digikey. This transistor has low junction capacitances, low Rbb and good 1/f noise. In fact, I use this transistor in low level stages everywhere from audio to microwave. If you are running from a +12 volt or thereabouts supply, bias the base with 1K ohm to ground and 3.3 K ohm to +V. Chose the emitter resistor to be 1K at low frequencies and maybe as low as 330 ohm at 150MHz. Considerable variation in bias is possible. Bypass the base with a 0.1 uF chip cap close to the base lead. For the tank, select the inductor value to be between 50 and 150 ohms reactance at the oscillation frequency. Wind the coil on a toroid core. Micro Metals T-25-6 or T-30-6 are good choices at 100MHz. Use a larger core at lower frequencies.
Calculate the capacitance to resonate with the chosen inductor. Split the cap in a 1:3 to 1:5 ratio. Again considerable variation is possible. You might need to go back and adjust the inductor value slightly to accommodate standard value caps. Use C0G chip caps at VHF. At lower frequencies where C0G are not practical you will need to go to a higher dielectric constant like X7R for at least one of the caps. Put the larger value cap (smaller reactance) to ground and the other cap to the collector. Initially, take output from a one turn link wound on the toroid.

To set up, connect a jumper or small value resistor from the emitter to the junction of the two tank caps to make the circuit free-run. If necessary, adjust the tank to set the frequency roughly to value. Remove the jumper or resistor and connect the crystal. Observe the output frequency and level and squeeze or stretch the toroid windings if necessary to peak the output level.

Now adjust the output coupling by adding one turn at a time to the previous one turn link, each time squeezing or stretching the toroid winding if needed. While adding turns you will find the output power to increase, level off and then decrease. Proper coupling is fewest number of turns to reach the peak output power. At VHF, one turn may be sufficient. If you must adjust the oscillator frequency, put a setting capacitor in series with the crystal.

If you need a VCXO, put a varactor diode in series with the setting capacitor and the crystal. In the case of a VCXO, the varactor anode goes to the tapped capacitors. A 10K resistor from this point to ground establishes the anode voltage at zero. Another 10K resistor from the cathode provides a tuning voltage input port. Resistor values are not critical ... just don't go too low or too high. The Zetex ZC8xx series varactors, available from Digikey, work pretty well. If pullability is too low you can try some coil in series with the crystal and varactor. Again, wind it on a small toroid and you can have almost infinite adjustability.

The general idea is to make the coil + varactor + setting cap combination approximately series resonant at the oscillation frequency. Pull range is also very dependent on crystal parameters and will ultimately limit the range. If reactances in the series leg get too large loop gain will drop below 1 and oscillation will cease. Make this circuit small, shove it into a piece of 1" diameter copper water pipe, wind the outside of the pipe with insulated resistance wire and you have the basics of an OCXO. I hope this makes sense. A two transistor Butler circuit is possible and has some advantages at low frequencies (like 10MHz). That will have to wait for my web page and real schematics!

So you think you know all about crystals?

Part 2

The reply from Chris Bartram, GW4DGU

While I largely agree with the first 60% of the notes, I'd very strongly suggest that the section about crystal oscillator circuits be renamed 'Here's how NOT to do it', at least for microwave applications!

I've designed a large number of crystal oscillators in the course of my job over the years, many of which my clients have put into large scale production. As part of the design process, I've looked analytically at most of the common designs of vhf overtone oscillator, and come to some conclusions based on hard-headed analysis.

One of the most useful indicators of crystal oscillator performance is the loaded Q of the crystal. The more energy retained in each cycle by the crystal (the greater the loaded Q), the smaller the effect of the maintaining amplifier, and the greater the stability and the smaller the level of the phase noise. Of course it's not quite as simple as that, and the hf noise figure and 1/f noise performance of the maintaining amplifier as well as the power levels and limiting strategies within the oscillator all play their part. To a first approximation, these are all related by simple linear maths. For those who want to go further there are books: Randy Rhea's 'Oscillator Design and Computer Analysis' is a good accessible start. What is clear is that some very well respected designs aren't actually that good, and there are very large differences in performance between the good and the bad. I'm currently working on a very high performance local oscillator/multiplier strip for my own use, and I've run simulations using Genesys and Spice (and, in a couple of cases, Serenade) of oscillator designs using a model of a good quality 128MHz 5th overtone crystal (Qu = 64000) and NE856xx transistor(s) (a highly recommended device BTW) in the maintaining amplifier. These are my results after optimisation:

Single transistor Butler oscillator (like K3VOT's): Qu=2400

'Balloon board' Butler oscillator (typically used in many UK microwave designs): Ql=4600

Single transistor self-limiting impedance inverting oscillator: Ql=16000

Cascode Pierce oscillator with separate limiter: Ql=24000

Cascode Driscoll oscillator with separate limiter: Ql=33000

The Driscoll circuit was originally designed for use in 5MHz frequency standards, but it can be made to work very well at VHF. It's capable of extremely good noise performance, while not over dissipating the crystal.

My oscillator strip uses a cascode Driscoll with some additional circuitry to further improve the very close-in noise. It has facilities to phase-lock the VHF oscillator to an external standard. The design uses mainly 0603 parts and has 'no tweak' filtering. Currently the output is in the 400MHz region, but I'm working on multipliers into the microwave region.

Improving Crystal Oscillator Signal Purity

John Hazell, G8ACE

Many microwavers purchase low cost crystals from providers such as QuartSlab, Klove and Eisch. At 24GHz and above it sometimes happens that the received signal using a new low cost crystal can sound full of LF jitter or perturbations. This makes copying of weaker SSB signals quite difficult. The blame for this effect can be laid directly at the door of the crystal itself.

I have been pursuing a cure for this problem for some time, it being quite difficult to get suppliers to discuss the problem and therefore track down a solution. The crystal blank quality on low cost crystals appears to drop to very poor quality at times, possibly quality that is adequate only for computer crystals. Both G4BRK and myself purchased 100.2 MHz crystals from Klove recently and both these showed large amounts of jitter. Other crystals purchased at the same time from Klove are perfectly good but because the frequency was different they where almost certainly made from different quartz blanks of better quality.

So the first unknown is the **crystal blank quality**. It seems likely that once a bad quality crystal frequency appears its likely to persist until all that stock of blanks is exhausted.

Additionally there appears to be a likely hood that variable stresses are contained in the bond wires, the crystal element within the can adding to this jitter problem. QuartSlab suggested some time ago that holding a crystal vertically and using a soldering iron to heat the lead out wire for ten seconds to allow the heat to travel up inside the unit would reduce the jitter.

This was tried and an improvement in the amplitude of the jitter of perhaps 2:1 was obtained. I also understood that professional users often store new crystals at around 80º C until they are required. I had kept some crystals under the jacket of my hot water cylinder for several years but on finally using these the jitter was no different to a new crystal. This warm storage may help with ageing though I have no measurements to support this aspect. Another crystal supplier then suggested cycling the troublesome crystals between around +80º C and -10º C for one hour at each temperature for 48 hours. A simple arrangement to do this cycling was to use my OCXO[1] unit itself to do the crystal heating. A resistor of 4K7 was added from TP1 on the circuit diagram to the +10.5v regulated supply. This allows the upper oven temperature range to be extended. Additionally if the heating is done using the complete OCXO it is important that high temperature Epoxy was used to bond the heater plate to the PCB board. For the low temperature simply placing the bare OCXO module in the freezer was the simplest solution. Only the crystal itself needs to be cycled so if you have a means then it can be done outside the OCXO.

An evaluation method needed to be established to verify any improvements that might be made. The **system diagram** (seen here on the following page) shows the arrangement employed. Two Adret 5104 synthesizer units are employed to establish a reasonable quality reference. The Adrets themselves will not be perfect but the results show them to be considerably better than a 'bad' crystal. The IF output from the transverter is connected to the FT290 Rx and the resulting CW audio note connected to the sound card in the PC and the DL4YHF[2]

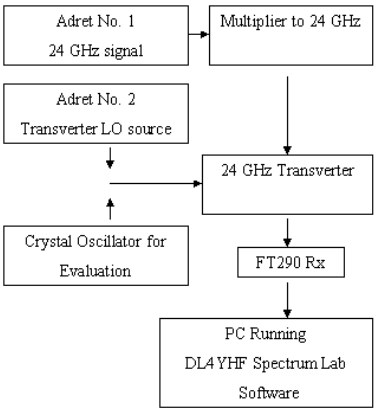

Spectrum Lab software.

Figure 1, shows the frequency perturbations of the test set-up. Vertical scale is 10Hz/div. Horizontal divisions are 10 seconds. The amplitude of the jitter is around 10-15Hz pp. This amount sounds perfectly satisfactory on CW.

test setup. Vertical scale is now 25 Hz/div. The jitter amplitude being as high as 300Hz pp. This is very audible as a nasty wobbly note. The method of heating and cooling the crystal is rather arduous and was done as time permitted. The crystal being re-measured at intervals for signs of improvement. Temperature cycling during the daytime over a period of around five days gave the result in **Figure 3**.

Here the vertical scale is 20 Hz/div. So it's possible to see around 30Hz pp. jitter and a result about 2:1 worse than that of the Adret but an overall improvement of about 10:1 for this particular crystal.

Another crystal, for which a graph is not available, improved to match the Adret result so it's possible that the Adret jitter would be the limiting measurement factor for this case.

The improvement is not necessarily as good for all crystals. If you have the jitter problem then only by doing the

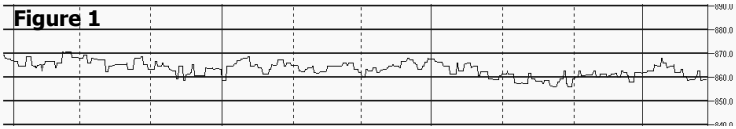

Substituting the test crystal in its OCXO for Adret No. 2 a measurement of that can be made.

Figure 2 shows the new 100.2 MHz crystal as considerably worse than the

temperature stressing will you find how good the crystal is capable of becoming. The sharp frequency transitions shown in Fig. 2 are assumed to be stresses in the bonding wires which presumably are

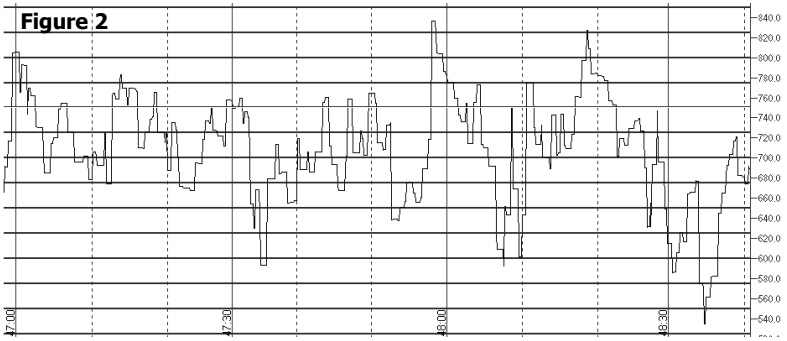

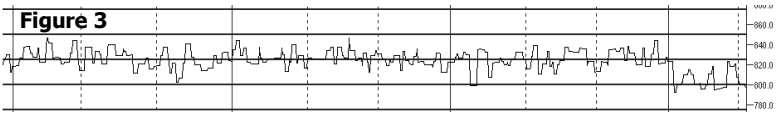

annealed in the temperature cycling process along with some stresses within the crystal resonator.
Luis Cupido, CT1DMK, on a recent visit suggested increasing the crystal drive level. This I assume will help accelerate the shedding of any rubbish on the crystal itself. The design drive level in my OCXO is nominally 750uW to the crystal. By shorting R2, the effect of the amplitude limiter diode D1 is removed and by shorting R8 the crystal drive level is raised to 10mW whilst maintaining around +4dbm output level. A really high power oscillator is possibly needed but there is some evidence that running the OCXO for several days at this higher drive level has some improvement effect. The tests have been done at 24GHz because the writer's portable transverter design allows easy testing. Testing at 10GHz would be possible but the amplitude of the measurements would be around 2.3 times smaller. It is assumed from the crystals evaluated so far that this jitter only manifests itself as an operating problem at 10GHz and higher. Regrettably the degree of improvements that can be obtained can only be found with the extended cycling and re-testing. So far all crystals have improved to a greater or lesser extent.

It has been reported and observed that the problem does improve on its own over a very long period but I assume this is due to the slower action of a crystal being heated and cooled in its application at a much lower cycling rate. If you have a problem crystal then testing and cycling should reward you with an improved LO note purity.

References:

1. http://
www.microwaves.dsl.pipex.com/
2. http://www.qsl.net/dl4yhf/
spectra1.html

Synthesisers
PLOs and DDSs

Experiences with the San Diego Microwave Group 1152MHz Synthesiser

Peter Day, G3PHO

This article outlines my own experiences with the 1152MHz Qualcom synthesiser board. The board, now used as my primary PORTABLE microwave frequency marker.

The pc board is quite large! It measures some 6 by 9 inches and so requires, by some standards, a bulky enclosure. I housed mine in a plastic box obtained from Maplin. A plastic box was deliberately chosen as the module is used as a multi-microwave band frequency marker when I am out portable. It radiates strong, very stable signals on all bands 2.3GHz to 24GHz (and maybe higher). It actually includes a great deal more circuitry than the basic synthesiser — a 1GHz IF amplifier and receive/transmit amplifiers for example. Originally, the VCO operates on a frequency of 750MHz to 1000MHz but a simple modification easily gets it onto 1152MHz. A high quality, 10MHz TCXO reference oscillator is divided down for the internal PLL reference frequency of 1MHz. The synthesiser chip is a 40 pin Qualcom 30361-16N or 32161-16N.

If one of the onboard MMICs is enabled as an amplifier, some +10dBm RF output at 1152MHz is available … an excellent LO source, especially for a beacon. In my case I did not do this modification, preferring instead to have the rich harmonic output which read 4mW on my power meter (probably at a multitude of frequencies at the same time!).

There are five simple modifications to be made to the board to get it to operate as a microwave marker. These are detailed on notes provided with the pcb. One of the modifications, adjusting the 10MHz TCXO, was not possible with my board as the 10MHz module was in a sealed can and was marked as having zero kHz offset. In reality it turned out to have a slight offset at 4.8Hz at 10MHz resulting in the 10GHz marker being at 10368.005MHz. However this is no real problem as the frequency is very stable, providing a constant, accurate marker just inside the narrowband portion of the band. On 24GHz it is just 12kHz up from the bottom of the band, no big deal with an IC202 as the RX IF!

Most of boards obtained from Chuck have an adjustable TCXO so purists can get theirs spot on all the zeros! Unlike other frequency sources I have, no multiplier diode was required at the RF output connector to produce harmonics above 1152. The module radiates plenty of those with no help at all! In fact the connector is left unloaded. When a Sat TV multiplier diode was connected at this point (in the manner described in the Newsletter last year) the output broke up into an unstable mess! One minor disadvantage with the module is the power supply requirement. Two voltages are need…. +15V DC and +5V DC. The latter draws around 0.5 amps and is thus a potential battery "zapper" when out portable over long periods. The +15V supply is only required to produce around 60mA or so. It actually feeds a 12 volt regulator which then supplies the 10MHz reference oscillator and so it would be possible to run the whole board from a 12V battery and use a 3 pin, step-down regulator (i.e. a 7805) for the +5V circuitry.

In my case I decided to keep the 15V input and built up a small, external 12V DC input power supply that pro-

duces the +15V from an LM2577T-Adj "switcher" IC (see Microwave Newsletter, May 1997) and the +5V from a 7805 regulator. The latter requires good heat sinking but the switcher IC runs very cool with no heat sink. A small die cast box houses the PSU. Leads in and out of it are thoroughly filtered to prevent any radiation from the switch mode section. The heavy current consumption at 5V was solved by having the inputs to each IC separately switched. Provided the 15V supply is left running, no drift occurs if the 5V regulator is switched off until a marker signal is required. Instant lock is achieved on switch on from cold. A surface-mount LED on the Qualcom board just flicks on for a second and then goes out, indicating lock.

Unfortunately there is no circuit diagram available with the board, as received from San Diego, so it was not possible to decide what to cut out of the board that was of no use, in an effort to reduce overall current consumption. Under portable conditions, I switch on the 15V supply to the TCXO just before I set off to the /P location. By the time I am there (usually 30 minutes) the module is drift-free. Once the microwave receiving equipment has settled down I then switch on the +5V supply and the marker is instantly available as an accurate frequency reference, whenever I need it. Prior to using this system I had an old Microwave Module 384MHz, 500mW source (suitable re crystalled) driving a G3JVL type multiplier using MD4901 varactor. I had already fixed a Murata crystal heater into the MM module but even then the stability over, say, 10 hours was not really good enough for me to be confident of "spotting" frequencies on 10GHz within +/- 1kHz accuracy. The Qualcom module can sit on the car dashboard, away from the vagaries of the weather outside and thus remain on frequency throughout the day's activities. Short of taking out my precious 90-120MHz Adret synthesiser or locking my transverter LO to a TV-locked or MSF-locked source, I cannot think of a more accurate way of keeping on frequency under portable conditions. Of course, the situation at home, where a good frequency counter and the Adret are on 24 hours a day, 365 days a year, is completely different! The Qualcom 1152MHz synthesiser is therefore recommended as a valuable tool for portable microwave work.

Oscillator Sources for Microwave Use

Andy Talbot, G4JNT

For some time I have being looking at ways of locking microwave source oscillators to a master frequency standard to allow high stability operation for propagation monitoring of beacons and to remove frequency drift with temperature and ageing. Most good quality frequency standards operate at either 10MHz or 5MHz, whereas microwave sources generally start from a crystal in the 100MHz region and are multiplied up. Going from 10MHz to some arbitrary 100MHz value is not straightforward. For local oscillators which generally operate at multiples of exact MHz, it is sometimes possible to find frequencies that can be generated relatively easily from 10MHz (e.g. 106.5MHz for a 10GHz LO source when multiplied by 96) but a beacon on 10368.905 (GB3SCX) needs 108.0094270833MHz - not nearly so easy!

The Conventional Approach

The approach taken by WA6CGR (www.ham-radio.com/wa6cgr) in his phase locked microwave source is to use a Phase Locked Loop with arbitrary values of division for both the reference and VCO divider chains, as shown in Figure 1. The output frequency is given by :

$$F_{rf} = M * N / R * F_{reference}$$

where M is the RF multiplication, N is the PLL VCO division and R is the reference division value. Is usually possible to find some pair of values of N and R where the resulting locked frequency is "near enough" to the wanted value to be acceptable. N and R should not be made too high as otherwise the divided down comparison frequency becomes unmanageably low, and some arbitrary value has to be decided on when working out R and N values, which usually have to be derived by trial and error, or more easily by a computer search of all possible values, given the dictates.

For the example above comparison frequency was specified as being 10kHz minimum (see below for justification) and a computer search came up with values of R = 849 and N = 9170 with an RF multiplication M of 96 for a result which is 406Hz low - probably good enough in practice. The comparison frequency ended up at around 11.79kHz. For our purposes on microwaves, we probably have one of the most stringent stability and phase noise specifications of any user on these frequency bands. Microwave amateurs are (probably) the only group who are concerned with phase noise performance at GHz that is only tens or even hundreds of Hz away from the carrier centre. Poor phase noise performance in this region manifests itself as an annoying rough tone to recovered SSB or CW signals, and even closer in to the carrier as a frequency jitter or wandering. So the basic microwave oscillator source at 100MHz has to have a very good performance so that when multiplied up many times its own phase noise (which is multiplied by the square of the multiplication factor) is acceptably low. 100MHz overtone crystal oscillators are usually quite adequate in this respect, although their frequency setting accuracy, temperature drift and ageing leave a lot to be desired. Furthermore, unless they are designed and cut for operation at elevated temperature, putting them in an oven or adding a clip on crystal heater hardly helps in temperature stability which can cause

several parts per million change even over day / night inside a house. A varicap diode in series with the crystal can easily be added to and allows the crystal frequency to be pulled a few parts per million - a few hundred Hz - by varying the control voltage applied to the varicap. This item will be referred to as the VCXO, or Voltage Controlled Crystal Oscillator.

10MHz reference oscillators on their own, free running, make use of the very best quality crystals there are, and the inherent phase noise of these if they could be multiplied up directly is usually more than good enough for our final requirements. However, being forced to use a phase locked loop in the frequency determining process can undo all the oscillator manufacturers good work! So if we can phase lock the crystal to a divided and multiplied version of a reference we're home and dry. The main problem here is the restricted pulling range of the 100MHz overtone crystal oscillator. As a rule of thumb, the bandwidth of a phase locked loop - which dictates the time required to lock up and the phase noise performance - has an absolute maximum value given by the product of the Voltage Controlled Oscillator control constant (Kv) defined in Hz/Volt and the Phase Detector constant (Kd) defined in Volts per radian. Where the output of the VCO is divided down, the Kv value has to be divided by N.

To a first approximation the absolute maximum PLL bandwidth is given by:

$$BW = Kv \cdot Kd / N$$

Phase Locked Loops in other applications often filter to narrower bandwidths, but for synthesisers we always want the absolute maximum bandwidth that can be achieved to minimise any noise added onto the VCO from external or internal sources.- this can make PLL design both easier - in the area of loop filtering, and also more difficult - in the area of high speed phase comparator performance. For the GB3SCX example above, Kv of a typical VCXO Butler oscillator was measured at 250Hz/Volt (divided by 9170), and a standard edge triggered phase detector in a 5V system usually has Kd = 1.6 V/radian, giving a maximum loop bandwidth:

$BW_{(Max)}$ = 250 x1.6 / 9170 = 0.04Hz or 1 cycle per 23 seconds

Since loop lock and stabilisation from switch on up can take several cycles we are looking at a period of several minutes before an acceptable output stability is achieved. Furthermore, and even more serious for our purposes, the inherent stability of the VCXO has to be stable in its own right within this bandwidth period, as otherwise the loop cannot correct it. Asking an non-ovened crystal oscillator to hold a few parts in 10^{-8} several tens of Hz at the final 10GHz microwave frequency) within even one cycle of loop bandwidth is not really on with temperature shifts and draughts. So it looks as if even the VCXO has to be oven controlled itself if taking this route. I first encountered this VCXO / PLL problem not originally on the microwave bands, but at LF when playing with ultra narrow band processing where the frequency stability requirements, (in relative terms anyway) are similar. A 24.576MHz clock oscillator had to be locked to a 5MHz reference and the only common frequency for comparison was 8kHz. It was necessary to allow 30 minutes for satisfactory PLL stabilisation and this was using a fundamental mode crystal with inherently higher pulling ability.

Direct Digital Synthesis

A few years ago DDS devices became widely available to amateurs and it became easy and straightforward to construct a source of virtually any arbitrary frequency up to a value of about 40% of the clock applied to the DDS -

which is typically in the 100 to 200MHz region for low cost devices. See my DDS module design published in Rad-Com November 2000. This used an AD9850 DDS with a maximum clock of 120MHz, allowing frequencies up to around 40MHz to be generated with a resolution, or tuning step, of $F_{clock} / 2_{32}$. A later device, the AD9851, which is allowed a clock frequency up to 180MHz, and also had an internal X6 PLL multiplier on board the chip itself so the actual clock source need only be 20MHz maximum. The beauty of DDS solutions is that there is no inherent increase in phase noise due to the DDS process - phase noise in is scaled exactly by the frequency division process.

I did wonder if the internal X6 in the AD9851 added any phase noise, but it is a very high bandwidth PLL and the data sheet states that his is "insignificant". Tests by multiplying a DDS output up to 10GHz did tend to agree with this statement. By driving an AD9851 with 10MHz from a frequency reference, a 60MHz clock results which can generate any frequency from 0 to around 24MHz in steps of 0.014Hz The downside to DDS sources is that they are not very clean, 50 - 60dBc spurii are common, and these seem to bear no relationship to the carrier frequency (they do, but the relationship is complex and is based on harmonics, alias products, harmonics of alias products and intermods). Suffice to say, the sprogs are troublesome, and some low level 'sprogs' are quite close to the carrier, albeit many dB down. I initially thought the DDS output at 21.6MHz could be multiplied by 5 directly to 108, and did try a breadboard of this, generating a test signal multiplied up to 10368GHz. It sounded terrible!

The centre frequency has a 'pure' enough tone, but it was surrounded by whiskers and rubbish that sounded like very high level of phase noise. What was happening was that all the close in spurii, that may well have been 80dB down and undetectable on the source, when multiplied by 540 were increased by the square of the multiplication factor and were now very significant. This idea was very quickly abandoned. **Next I considered the PLL route - see Figure 2.** Since the DDS can generate tens of MHz directly, the loop bandwidth could be made very high—certainly in the tens of kHz region and ought to be capable of stabilising a poorer quality oscillator such as an LC design. Various attempt were made to do this, but a sufficiently high bandwidth PLL never materialised. Edge triggered phase/frequency detectors require high speed logic to make them operate and I was losing patience at this point, quite apart from the fact that the DDS spurii were still proving difficult to filter out with loop bandwidths sufficient to clean up an LC source. More perseverance with good LC oscillators and fast logic may have come up with something, but I could be bothered.

So, back to a crystal oscillator as the source - at least these are always clean. Now a PLL using this as the VCXO does not have to divide down by a large factor, as the comparison can be made at tens of MHz. By using a simple divide by 8 prescaler the VCXO can be brought down to 13.5MHz for comparison with a reference generated from the DDS. The exact division ratio is unimportant as the DDS can generate any frequency, and a divide by 10 would work just as well. A simple mixer type phase detector was implemented using an NE612 receiver chip followed by a low noise op-amp, rather than trying to make a complex high speed edge triggered D-type flip flop design. Kv of the VCXO was 250Hz/V now only divided by 8, and the NE612 / amplifier combination had a Kd of 2 V/rad - or more depending on the amplifier gain.

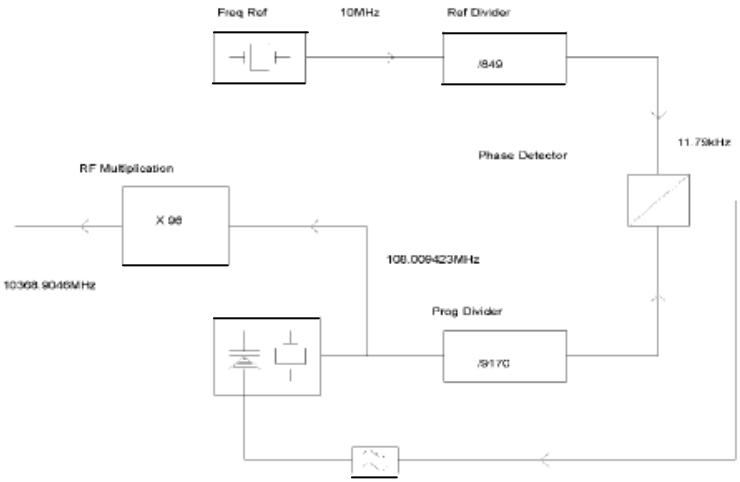

Now, maximum loop bandwidth = 250 * 2 / 8 = 63Hz ….. much more acceptable.

The circuit of Figure 3 was built up and tested. It locks up virtually immediately, the note when multiplied up to 10GHz sounds 'perfect' and all appears to work as required. No particular effort was made to optimise loop filtering - as the inherent loop bandwidth is dictated solely by the Kv.Kd product this was just left to cope as best it could and give the highest possible bandwidth, with no attempt to control overshoot or anything. As with any PLL based on a mixer rather than an edge triggered phase detector, the loop will only lock up if the initial frequency error at switch on is within the loop bandwidth. So the crystal oscillator has to be within 63*8 or around 500Hz of the wanted frequency. In practice, and to cope with temperature shifts, a starting frequency of significantly less than this should be aimed for - say within 200Hz which corresponds to a couple of PPM. This is getting a bit tight for an non-ovened crystal to be used over a wide frequency range, but, with decent overtone devices, it can usually be achieved. The divide by 8 prescaler chip used here, the SP4908 is probably obsolete now, but I had one Any suitable device from the junk box will do, and doesn't have to divide by exactly 8 - so long as the final frequency is within the range of the DDS source. VHF prescalers are not too common now as synthesiser chips themselves operate at 1 - 2GHz; the prescalers that are around tend to operate at many GHz, but even one of these can be pressed into service as they usually operate down to DC. An incidental advantage of the loop locking up quickly it that is just about possible to generate frequency shift keying by reprogramming the DDS. Provided the frequency shift is small - no more than around 1 to 2kHz - the loop can track this change within a few milliseconds. If CW at around 12WPM is used then the keying just sounds a little soft, as the tone changes from mark to space. Much faster than this, and the edges are smeared out too much.

The DDS frequency resolution is $F_{clock} / 2_{32}$ which for a 60MHz clock is 0.014Hz. When multiplied up to 10GHz by the X8 prescaler and X96 RF Multiplication, this step size translates to a little over 10.7Hz - probably accurate

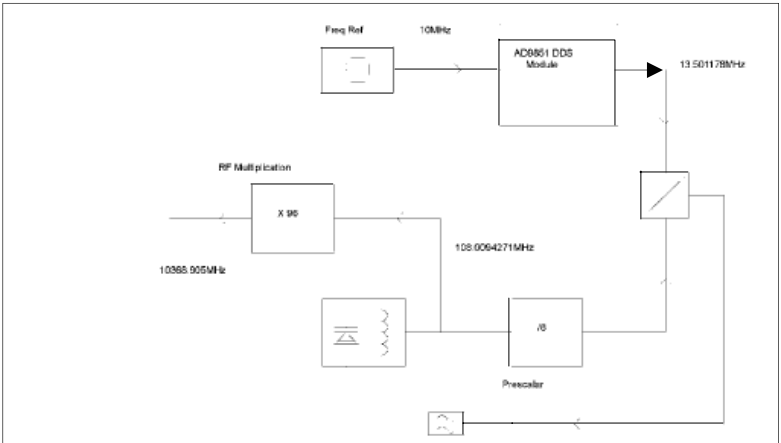

Fig 2 DDS Source for generating the PLL reference

enough for most purposes!

An afterthought that occurred to me as this article was being written ...

Instead of using a prescaler to divide the VCXO down to a frequency that could be generated by the DDS, why not use the VCXO directly as the DDS clock (without the X6 PLL option enabled)? Then by thinking backwards, programme a value into the DDS that generates exactly 10MHz from the chosen oscillator frequency. As the internal PLL is not needed, and the device is only operating at 100MHz the original AD9850 can be used in this adaptation. This concept will form the basis of the Mark two design. Generating the accurate 10MHz signal is a subject I'm also closely involved with, but that will have to wait for another article.

Another possibility for the next generation is using the AD9852 DDS. This chip can have a clock of up to 300MHz and an internal PLL which is programmable between 4 and 2. It also has a 48 bit accumulator so frequencies can be set to a resolution of 2^{-48} of the clock. At 108MHz clock frequency, followed by multiplication to 47GHz, this still allows a final frequency resolution of 0.16mHz (no, that is **not** Megahertz, it's **millihertz** !!!)

More afterthoughts ...

I'm not sure GPS LOCKING is necessarily the best route for high frequencies - with the long time constants needed, when used in /P or situations when the temp changes the freq can wander around quite a lot whilst still being "locked". I'd rather maintain a very good 10MHz source in a separate box with its own batteries that can be carried around and calibrated when necessary.

Locking is good idea, though, for beacons and home stations where a reference can be maintained always switched on - but still with battery backup for coping with power cuts. I much prefer using GPS to gate a frequency counter that then measures the reference oscillator freq over a long period - 1000 or 10000 seconds (17 minutes or 2.8 hours) for calibration only. Displaying 5 or 10MHz on a scope triggered at 1 pulse per second is fea-

sible for getting very accurate calibration, provided you do it in a darkened room.
For the best of both worlds, make a locked source for home use only (the MSF one is probably good enough), then check the portable unit against this using a scope which will only take a few minutes to set both traces to appear stationary against each other.

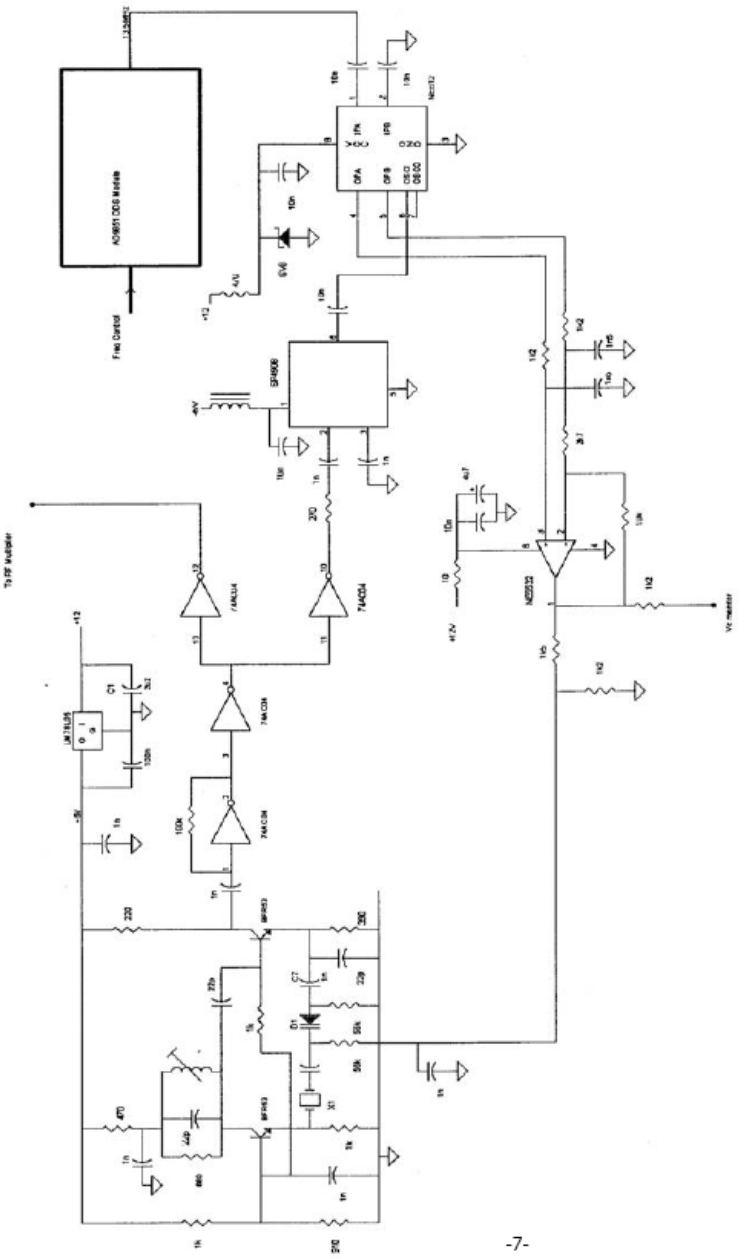

Fig 3 Circuit of DDS Locked Microwave Source Oscillator

A simple way of phase locking microwave local oscillators

Andy Talbot, G4JNT

In the article on 'Locked Oscillator Sources for Microwave Use' published in the last Microwave Newsletter, I made the throwaway comment, in the first paragraph, that phase locking local oscillators with nice round frequencies of integer MHz was simpler than producing an arbitrary value for beacons, etc. Well, after saying that I had to come up with something and here are a few notes on a simple breadboard method of phase locking these crystal oscillators to a reference source. The circuit is presented as a starting point - there will be some experimentation needed for each individual case.

For the popular microwave bands the LOs are usually generated from overtone crystal oscillators followed by multipliers, with the following table showing the usual LO frequencies for the narrowband segments along with the associated crystal frequency. The final columns show the highest frequency that is a sub multiple of both this and a 10MHz reference, the highest common factor or HCF, and the associated division. This HCF can become the comparison frequency in a phase locked loop, and is the highest frequency that is possible here.
Two things become obvious:
All the comparison frequencies can be derived from a 10MHz reference by making use of simple logic divider chips to give the divide by R function (all could be derived from 2MHz in fact) All the comparison frequencies are over 200kHz, so phase locked loops can be made with wide loop bandwidths. For those marked with a *, an even higher comparison frequency is possible, but the values stated keep the frequencies within a narrower band for a common design.
So now the only difficulty is providing the divide by N from the crystal frequency. An off the shelf synthesiser chip such as the MC145170, or those from other manufacturers, would make an easy job of this but there is an even simpler solution providing you are prepared to do a bit more adjustment and optimisation.

Anyone who has studied the 'brick' range of microwave sources will have seen how a high Q cavity oscillator is locked to a reference oscillator in the

Band	IF	LO	RF mult	Crystal	Division N	HCF(kHz) for 10MHz F_{ref}	Div R
1296	144	1152	12	96.00000	96	1000.00*	10
1296	28	1268	12	105.66667	317	333.333	30
2320	144	2176	24	90.666667	136	666.667	15
3400	144	3256	36	90.444444	407	222.222	45
5760	144	5616	54	104.000000	104	1000.00*	10
10368	144	10224	96	106.500000	213	500.000	20
10368	144	10224	108	94.666667	142	666.666	15
24048	144	23904	240	99.600000	249	400.000	25
24048	432	23616	240	98.400000	246	400.000*	25
47088	144	46944	432	108.66667	163	666.666	15
47088	432	46656	432	108.000000	108	1000.00*	10

100MHz region by a sampling phase detector. This device combines the functions of frequency multiplier and phase in one network. In the bricks, a snap varactor diode is hit with about 200mW of reference signal and so generates sharp sub-nanosecond pulses at this rate. These pulses are applied across a pair of microwave diodes, forming one input to a single balanced mixer whose other RF ports is fed with a portion of the cavity oscillator signal picked off via a small probe. This sampling mixer approach to PLL design makes for considerable simplicity as it inherently removes any need for a high frequency divide by N circuit - but has two major drawbacks. Firstly, the VCO can lock to ANY harmonic of the reference, and in the brick designs this is prevented by restricting the electrical tuning range of the VCO to less than half of the comparison frequency's Sometimes a pull in range of only 5-10MHz (at the fundamental L-Band frequency) can be observed.

Secondly, the output from the phase detector is at a very low level - typically a few tens of millivolts per radian rather than the 1.6V/radian of normal logic. In effect, the complete voltage range that would have been possible for the drive power is having to be shared over every one of the comb frequencies. However, a low noise op- amp can easily provide DC gains of the hundred or so required here, and differential amp doing this job inside the bricks is clearly visible when the side cover is removed.
So, lets try this idea at lower frequencies to lock a VCXO - the very restricted tuning range of crystal oscillators means that drawback 1 is not an issue. Look at the circuit of Figure 1, a two chip R divider (programmable for any value from 1 to 256) generates the reference. This is applied to an impulse generator using the propagation delay inherent in three high speed logic gates plus an additional capacitor to generate a series of negative going impulses of a few nanoseconds in width. If you look at the output at this point on a spectrum analyser, the spectrum will show the classic sin(x)/x shape with a null corresponding to the pulse width. The extra 100pF capacitor in the delay can be adjusted to ensure this null does not fall at the wanted crystal frequency - it can get a bit unpleasantly close with 74AC series gates in this position.

These impulses are at a level sufficient to directly drive a diode ring mixer, so all that is now necessary is to apply a portion of the VCXO to the other mixer port, amplify and filter the IF output and feed back to the VCXO for a complete PLL. As in all the microwave sources I've discussed so far, the PLL needs to have as wide a bandwidth as possible to remove VCXO close in phase noise and jitter, so there is not much effort that has to go into filtering, apart from removing the fundamental comparison frequency component. This last point is significant, don't go too low with the reference frequency as filtering it out will be more difficult, as well as the fact that the voltage swing available from the diode ring will be even less. The pair of 74HC161 devices forming the R divider could be replaced by appropriate sections of a 74HC390 chip, or an 'HC90 with feedback in some cases, but this solution gives a generalised divider allowing factors for R divisions not in the table, such as those with 7, 11, or 27 in them that couldn't be obtained from a simple configuration. It works be preloading in the number set on the links, then counting up to 256 where it overflows and loads in the preset value again repeating the process. So, the wire links have to be configured to load in a value of (256 - R). With HC CMOS, a logic '1' must not

be generated by leaving an input as an open circuit, as could be done in the old days of TTL (and even then it was unethical!), so each input must be tied to either +5V or ground.

Results:

My breadboard was not exactly as shown in Figure 1. Instead, I started off with a 1MHz signal rather than 10MHz and divided by 3 in a single chip to give 333.333kHz which was used to lock the 94.666MHz crystal in my 'JVL based 10GHz system. Figure 1 has been built, and just committed to PCB, but not yet into any finalised system. The first version for 10GHz worked fine, except for the fact that the VCXO drive level to the mixer was rather critical - this has to be kept at about 1mW for the mixer to remain linear, whilst maximising the DC output level from the IF.

If drive level is increased too much, the mixer saturates and, surprisingly, IF output falls off in this usage. I'm not too sure why this should happen, but is probably due to the short pulse forming the pseudo-LO, and its very low means value. Real experimenters may like to try replacing the packaged DBM with a pair of diodes plus transformer/ balun for a single balanced design as in the bricks. This may give more sensitivity and less dependence on RF drive level, but I couldn't be bothered with this level of fine tuning once the DBM had proved itself.

Note that the mixer also has to be one where both sides of the IF port are accessible. Most devices like the SRA-1 and SBL-1 offer this, but higher frequency mixers sometimes ground one side of the IF so in this event the op-amp will need to cope with a negative input voltage. The DC gain in the OP-amp circuitry is what was required for this 333kHz reference - other higher comparison frequencies will allow proportionately lower gains. The exact circuit configuration is not a completely true differential amplifier as this is not essential, it also has to transfer the 5V reference through to the output. As true differential operation is not necessary, to change the gain it is only really necessary to alter the single resistor shown as 330k in Figure 1.

An easy way to set up the necessary DC gain is to look at the op-amp output on a scope with both reference and VCXO signal applied, but with the input to the VCXO tuning diode clamped to the 5V centre value. Then by manually tuning the VCXO through its range a beat at the difference frequency will be seen, this will be centred on exactly 5 Volts and its peak to peak value must be wide enough to be able to tune the VCXO at all ranges of temperature and drift …. a value of 2 - 5 volts peak-to-peak will probably suffice for most VCXO designs. The op amp gain may need adjusting to get to this figure, and the value also depends to some extent on the RF level to the mixer.

With no reference applied, the VCXO needs to be adjusted to the wanted frequency with 5V on its tuning line …. try to get as close as possible as the tuning range is not very wide once temperature and crystal ageing are taken into account. My breadboard locks up near instantly when connecting the reference, provided the crystal heater has been pre-warmed! If a crystal heater is used, the initial cold start frequency will probably be outside the PLL lock range; in my case it takes about 10 seconds for warm-up before the PLL will lock from a cold start.

Now, the only drawback to getting within 0.1Hz on my 10GHz system (and yes, I can get 10-11 on the reference) is the accuracy of the 144MHz IF. The IC202 LO is derived from a DDS (more about this another day) controlled from a bog-standard packaged 100MHz crystal oscillator and rotary up down counter plus LCD display to 10Hz reso-

lution and 100Hz readout. The DDS could have been driven from the 10MHz reference multiplied up, but there is still the inherent accuracy of the carrier crystal inside the IC202, so I can only get to within tens of Hz of the wanted frequency at 144MHz without doing some pre-calibration of the IF, and even then it will drift a few more Hz over the day. When I designed the digital IC202 LO I hadn't anticipated wanting a better accuracy for microwaves than this, but some recent tests between G8ACE and G3NNS make this assumption invalid now.

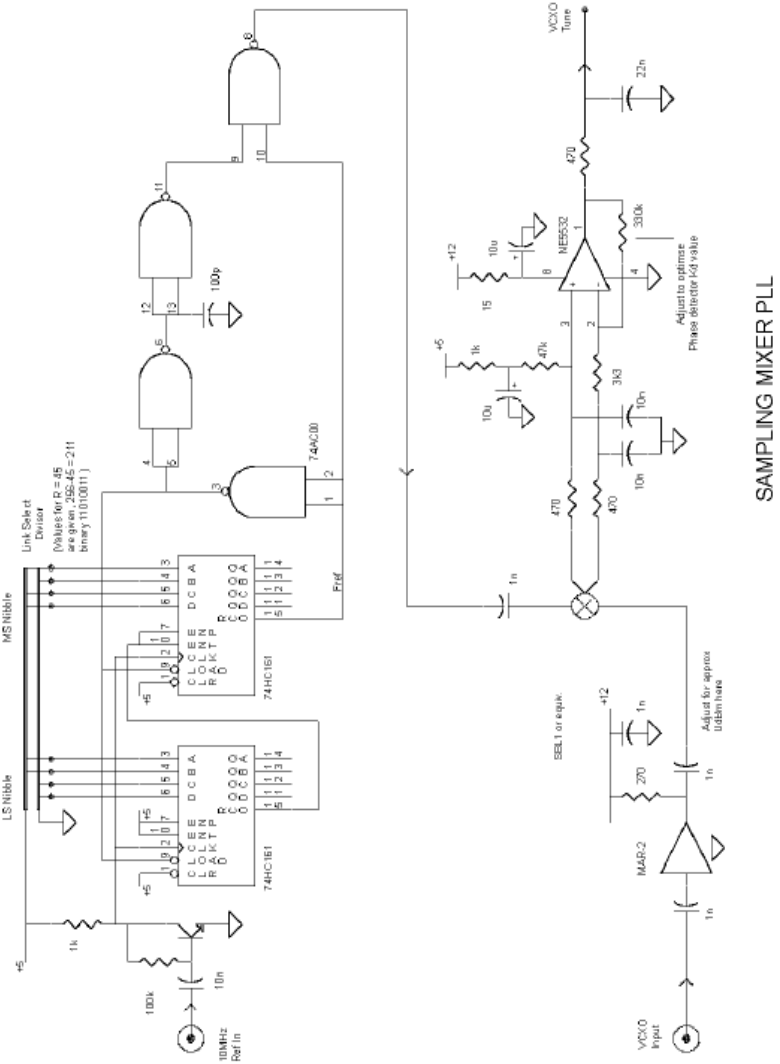
SAMPLING MIXER PLL
G4JNT 2004.04.07
MS Nibble
LS Nibble
Link Select Divisor
(Value for R = 45 are given 256-45 = 211 binary 11010011)
74HC161
74AC00
Fref
10MHz Ref In
VCXO Input
VCXO Tune
NE5532
MAR-2
Adjust to optimise Phase detector Kd value
Adjust for approx 1dBm here
SBL1 or equiv.
100p
100k
10n
1k
10u
47k
3k3
330k
470
22n
1n
270
+5
+12

DDS Module Controller Software

Andy Talbot, G4JNT

Overview

The DDS module is supplied with a PIC microcontroller that contains code to allow the AD9852 to be controlled via an RS232 serial link from an ASCII terminal such as a PC running Hyperterm software. A further input to the PIC can be used for an external trigger, so the updating of the DDS output can be synchronised to an external even such as a UTC pulse from a GPS receiver.

Setting up the serial link

Set your terminal programme to 19200 baud, 8 data bits, no parity and 1 stop bit (19200 N81). Turn off local echo and do not enable CR-LF translation on receive. Connections between the PC and the DDS module are defined in Table 1:

Connect the serial link and switch on the DDS module. After about half a second delay, a display similar to that in Table 2 will appear.

All the registers in the AD9852 chip are loaded with the values stored in EEPROM, either the default initial values, or any that have been subsequently changed by the G command (see below). As supplied the default will result in an RF output at exactly 0.25 times the clock frequency, with no PLL multiplier in use. All other

Controller Software

The AD9852 DDS chip has got quite a comprehensive set of capabilities which are controlled by writing appropriate values to its working registers then triggering (updating) the device. The controller software has two main functional capabilities.

For simple frequency generation and phase shifting requirements, a straightforward command structure for quickly updating CW frequency or phase is implemented, with separate commands to update the DDS with these changes either immediately, or on an external pulse edge. The option of writing the current frequency to non-volatile storage (internal EEPROM) for immediate start up next time the module is turned on is also possible. The other mode of operation allows any of the internal registers to be written individually, giving full access to the chip's functionality. All register contents altered using these commands are stored in EEPROM and loaded in the next time the module is turned on. A final option is available to allow users to store a string of up to

Connection	9 Way Dtype Pin	PIC Connection
TXD	3	Port B0 (via resistor)
RXD	2	Port B3
Gnd	5	Gnd

Table 1: RS232 interface connections

9852 DDS Controller G4JNT

Qxxxxxxxxxxxx[cr] Pxxxx[cr] U W R V K R

0 = 00 00
1 = 00 00
2 = 00 00 00 00 00 00
3 = 00 00 00 00 00 00
4 = 40 00 00 00 00 00
5 = 00 00 00 00
6 = 00 00 00
7 = 10 64 01 20
8 = 00 00
9 = 00 00
A = 00
B = 00 00

K 240000000

Table 2

15 characters in EEPROM for reading back on the serial link; they perform no action on the DDS chip itself. This can be used, for example, to store the clock frequency so any DDS driver software can read this value back, and so be used with several different modules, each with their own clocks.

Commands

Immediate programming

Pxxxx[cr] Sets the phase of the RF carrier output. The format must be exactly as shown with xxxx replaced by hexadecimal ASCII characters. e.g. P8000[cr]. The new phase is programmed into the AD9852 (although not into EEPROM) but does not take effect immediately. The characters are not echoed back to the terminal, so when typing in by hand this has to be done blind. If the command is recognised, the controller responds with a single P followed by [cr][lf].This command controls the contents of DDS register 0 **Qxxxxxxxxxxxx[cr]** Sets a new frequency of the single carrier output. The format must be exactly as shown with xxxxxxxxxxxx replaced by hexadecimal ASCII characters. e.g. Q0280000000AF[cr]. The new frequency is programmed into the AD9852 (although not into EEPROM) but does not take effect immediately. If the command is recognised, the controller responds with a single Q[cr][lf]. This command controls the contents of DDS register 2

U This command updates the DDS chip with the new P or Q values set in the above commands. When complete, the controller responds by sending Z[cr][lf] No carriage return is needed after any single letter commands.

T Triggers the controller to wait for a positive edge on the external timing input (Port B2) before updating

the values from the P or Q commands. 'Z'[cr][lf] is returned when the update is done. Note that while waiting for the positive timing edge, the controller will be deaf to any further serial commands and may have appeared to 'hang'. This situation has to be checked for in any driver software by looking for the acknowledgement Z before issuing any further commands.

W Writes the current frequency to EEPROM. This command will only be accepted immediately after a P command has initially been issued. If accepted, the controller responds with a 'Z'

Register Programming

The individual AD9852 registers can be updated one at a time and the values are always stored to EEPROM for immediate start up. The AD9852's registers have different lengths depending on their function, from one to six bytes in length. There is plenty of scope for incorrect operation and unexpected results, particularly when programming the control register. Read the data sheet carefully before changing registers!

The G command is used to update an individual register: functionality. All register contents altered using these commands are stored in EEPROM and loaded in the next time the module is turned on.

A final option is available to allow users to store a string of up to 15 characters in EEPROM for reading back on the serial link; they perform no action on the DDS chip itself. This can be used, for example, to store the clock frequency so any DDS driver software can read this value back, and so be used with several different modules, each with their own clocks.

Commands

Immediate programming

Pxxxx[cr] Sets the phase of the RF carrier output. The format must be exactly as shown with xxxx replaced by hexadecimal ASCII characters. e.g. P8000[cr]. The new phase is programmed into the AD9852 (although not into EEPROM) but does not take effect immediately.
The characters are not echoed back to the terminal, so when typing in by hand this has to be done blind. If the command is recognised, the controller responds with a single P followed by [cr][lf].This command controls the contents of DDS register 0

Qxxxxxxxxxxxx[cr] Sets a new frequency of the single carrier output. The format must be exactly as shown with xxxxxxxxxxxx replaced by hexadecimal ASCII characters. e.g. Q0280000000AF[cr]. The new frequency is programmed into the AD9852 (although not into EEPROM) but does not take effect immediately.
If the command is recognised, the controller responds with a single Q[cr][lf]. This command controls the contents of DDS register 2

U This command updates the DDS chip with the new P or Q values set in the above commands. When complete, the controller responds by sending Z[cr][lf] No carriage return is needed after any single letter commands

T Triggers the controller to wait for a positive edge on the external timing input (Port B2) before updating the values from the P or Q commands. 'Z'[cr][lf] is returned when the update is done. Note that while waiting for the positive timing edge, the controller will be deaf to any further serial commands and may have appeared to 'hang'. This situation has to be checked for in any driver software by looking for the acknowledgement Z before issuing any further commands.

W Writes the current frequency to EEPROM. This command will only be accepted immediately after a P

command has initially been issued. If accepted, the controller responds with a 'Z'

Register Programming

The individual AD9852 registers can be updated one at a time and the values are always stored to EEPROM for immediate start up. The AD9852's registers have different lengths depending on their function, from one to six bytes in length. There is plenty of scope for incorrect operation and unexpected results, particularly when programming the control register. Read the data sheet carefully before changing registers!

The G command is used to update an individual register: Type G and the controller responds with 'Reg No' No [cr] is needed after the G

Enter a value from 0 to 9, A or B. There is no need to enter a carriage return. The controller responds with 'x bytes of Reg. Data' , where x is a number from 1 to 6. Enter the data in hexadecimal (with no [cr]), and as soon as the final character of the requisite number is entered, the controller will respond with 'Z'. (Note, one byte of data requires two characters, 6 bytes requires 12 characters).

The new value is immediately written to both the AD9852 register and the appropriate EEPROM register The V command is used to dump the entire EEPROM contents, in the format shown in the central portion of Table 2.

User Data

These two commands allow up to 15 characters of user data to be stored and read back from EEPROM.

The K command allows data to be entered. Type K and the controller responds with 'Enter < 15 chars. of user data' Enter the characters required, followed by a [cr] to terminate. Any ASCII character is accepted, and all letters are converted to upper case. If an attempt to enter more than 15 characters is made, the controller responds with 'Overflow' and the first 15 entered are accepted. In the interests of conformity, if this data is used for clock frequency it is suggested this takes the form of , for example, '240000000.00Hz' so the readout is meaningful and can be easily read and interpreted by driver software.

The R command reads back the user data Issue R (with no carriage return), and the controller responds with the stored string

Frequency markers and references

Versatile Microwave Marker Generator

The circuit shown was designed by Kerry Banks, K6IZW, and previously published in the San Bernardino Microwave Society newsletter, which we gratefully acknowledge as the source of information. It is a very simple comb generator that produces low level marker signals every 10MHz, from VHF through microwave frequencies. It could be useful for experimenting and when testing equipment, providing reference points at regular intervals and at a rate more frequent than the normal references which are often based on a starter frequency of 1152MHz.

The 10MHz TCXO in the original design was a surplus Qualcomm item. These devices were available at low cost from Chuck Houghton, WB6IGB by emailing him at:
clhough@pacbell.net
The TCXO drives the circuit, producing 1 nanosecond pulses at 10MHz intervals. If the output is amplified in a simple external broadband MMIC amplifier, you could observe a "comb" of signals with a spectrum analyser. Pass this spectrum through a filter and observe the output on the analyser to assess the filter's passband.

In the circuit shown please note the following:

- The Qualcomm TCXO requires +12V but may be substituted with a lower stability TTL Clock Oscillator if the supply voltage is run off +5V.
- The logic IC 74AC00 must be used as other CMOS logic families are too slow.
- * the .001uF capacitor on pin 1`4 of the 74AC00 to ground must have the shortest leads possible.
- The output can be connected to coax or a small piece of wire as an "antenna".
- Changing the output capacitor to a small value such as 1pF reduces low frequency harmonics.

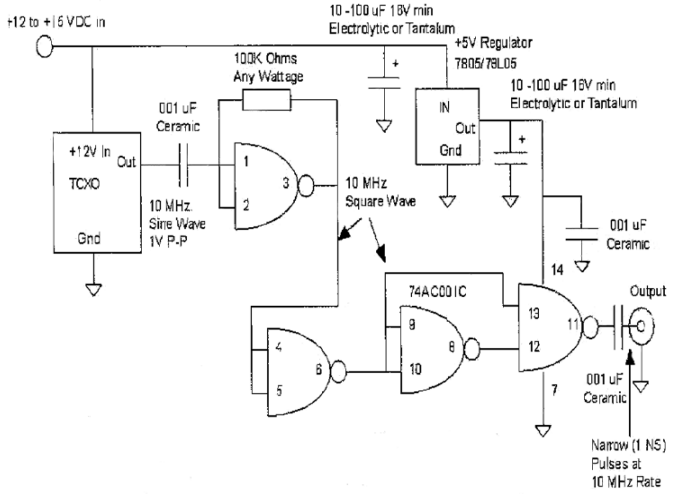

Practical Implementation of the K6IZW Microwave Frequency Marker

Allan, G8LSD

I have made the marker as described Kerry Banke, K6IZW. I had no joy at 10GHz with an 8.5 MHz oscillator but at 24 MHz harmonics are just audible. The good news was when I fed it from an Adret 5104 synthesiser (90 to 120MHz range) I found very pure signals and at very low power. For construction, I used a strip of fibre-glass PCB as a ground plane between the pins. The components were SMDs soldered directly to the pins. The cross connections were made outside and over the top of the pins. For output, I tried a short wire, and this was successful. Then I used a short length of thin semi-rigid soldered to the ground plane, the centre conductor being connected to the output pin by a 2pF SMD capacitor. The output now goes to a transition and a small horn. If the Adret is accurate, my offset on 10GHz is about 5kHz. The short spikes from the 74AC00 are clearly visible on a scope.

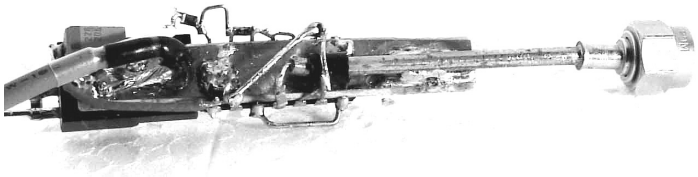

Above: the pcb marker Assembly

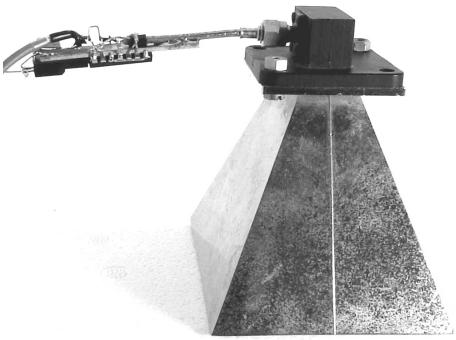

Right: The marker mounted on small 10GHz horn for test purposes

Probably the Simplest GPS Disciplined Oscillator Possible!

Andy Talbot, G4JNT

GPSDOs

Using the GPS Satellite system offers the advantage of very accurate timing and by extension, frequency control. The long term error is to all intents and purposes zero, with time and frequency accuracy being comparable to the international standard.

The traditional route is to use a relatively low cost GPS receiver module which outputs a 1 Pulse per second signal (1 PPS) aligned to UTC.

Basic GPS modules such as the Garmin GPS25 and Motorola Oncore have been around for some years and are available at low cost. It is possible to phase lock a divided down crystal oscillator to this 1 PPS signal and transfer its long term stability to, say, a 10MHz reference which is subsequently used for deriving any LO and beacon frequencies. The subsequent PLL system is usually described as a GPS Disciplined Oscillator rather than locked, since it is not, strictly speaking, actually 'locked' to the GPS system at all; just controlled by it via the 1 PPS generated by software in the receiver module.

The first GPSDO to appear in the amateur press was by Brooks Sheera W5OJM, described in QST July 1998. A lower stability, simpler version, suitable only for low data rate signalling on the LF bands, was published by myself in Radio Communication October 2002. Both of these, in different ways, demonstrates the problem with using the 1 PPS signal. On all these receiver modules the 1 PPS signal can have up to 1us variation from pulse to pulse, and this varies randomly. Later modules reduce to a few 100ns, but it is still there. Consequently, for a frequency standard with a short term stability measured over a few tens of seconds, this 1 PPS jitter has to be averaged out over many hundreds or thousands of seconds - so giving very long lock up times and loop tracking constant. Now, as the PLL has a time constant of many tens of minutes or hours, the voltage controlled crystal oscillator has to be stable over this loop time constant - particularly if it is to be multiplied up to GHz where a short term wander of a few Hz (parts in 10^-10) is noticeable. So a good quality oscillator has to be used here - typically an ovened high spec standard in its own right.

This was the approach taken by W5OJM with a microcontroller based digital PLL and loop time constant of hours. My design went the other way, and accepted a poor short term stability for LF use only, where the phase wander over a few seconds was inherently averaged by the LF signalling interval. Many manufacturers now offer off-the-shelf GPSDO modules with varying specifications between these , ranging in price from a few hundred pounds, to thousands.

The Jupiter-T Solution

Which brings me onto a new GPS module that makes a homebrew solution very much easier. The Jupiter-T module made by Navman (originally Conexant) has an output at 10kHz 'locked' to GPS time. Initially I was sceptical, thinking it probably only consisted of 10000 pulses per second, which could have been no better than the 1 PPS signal itself in the short term. However, after making extensive measurements, I came to the conclusion that it really was quite a respectable

signal. In particular, I could not detect any discrete sidebands at sub Hz frequencies.

This suggested a simple GPSDO solution. By taking a simple, low cost 5MHz voltage controlled TCXO (VTCXO) module and dividing down to 10kHz, this can be phase locked to the output from the Jupiter in an analogue PLL with a time constant of a few tens of seconds.

The circuit diagram **Figure1** shows how simple this can be. Obviously, without the ability to be able to tell if the GPS receiver has locked up by reading the NMEA or binary data it sends from its communication port, there is no way of knowing if the system is functioning properly.

The Jupiter module does output its 1 PPS signal and a nominal 10kHz immediately after switch on, but the timing of these is way off and the initial case of no GPS lock can be inferred from the large frequency error. In fact this is so large that the PLL is out of lock range and the resulting frequency is sweeping so wildly that it is obvious. When the module does lock up to GPS after a few minutes, the frequency and phase of the 1 PPS jump immediately and abruptly to their correct values, with the PLL taking a short time after this to stabilise. Although not shown in the diagram, an LED connected to the phase pulse output of the 4046 chip will slowly change brightness over a few seconds during the GPS lockup, and then much more slowly as the PLL locks, eventually settling to a stable half-maximum brightness

The R/C values forming the PLL filter are optimised to my particular VTCXO which has a sensitivity of 125 Hz/V at 5MHz, and a required tuning voltage in the 0 to 1.5 Volt region. As the bandwidth and tracking performance of the PLL depends on this filter, it is worth spending a bit of time optimising the values in this area.

Some Test Results

The 200_{th} harmonic of the 5MHz output, at 1GHz, was monitored on a UHF communications receiver in CW mode and the output at 1kHz monitored with Argo to show short to medium term frequency shifts. All the local oscillators in the receiver were locked to a Rubidium frequency standard that has been calibrated to an accuracy of a few parts in 10^{-11}. At 1GHz, a frequency shift of 1Hz corresponds to 10^{-9} frequency error.

Figure 2 shows the plot after the system has locked up and been running for about 30 minutes.

It can be seen that the frequency is maintained usually within a couple of Hz, and randomly wanders over a mean period of something like 20 - 30 seconds - this being a function of the PLL bandwidth. The breadboard which produced these results was lying open on the bench, and susceptible to perturbations when I touched it - it is quite possible this trace would be cleaner still if the unit was packaged in a screened metal box.

Figure 3 shows the effect of disconnecting and then reconnecting the GPS antenna. Presumably, the quite fast variation during the period of no GPS signal is the GPS receiver going through its search routine to find the satellites. The faint line that remains fixed at 1kHz exactly is caused by leakage from the Rubidium source controlling the communications Rx.

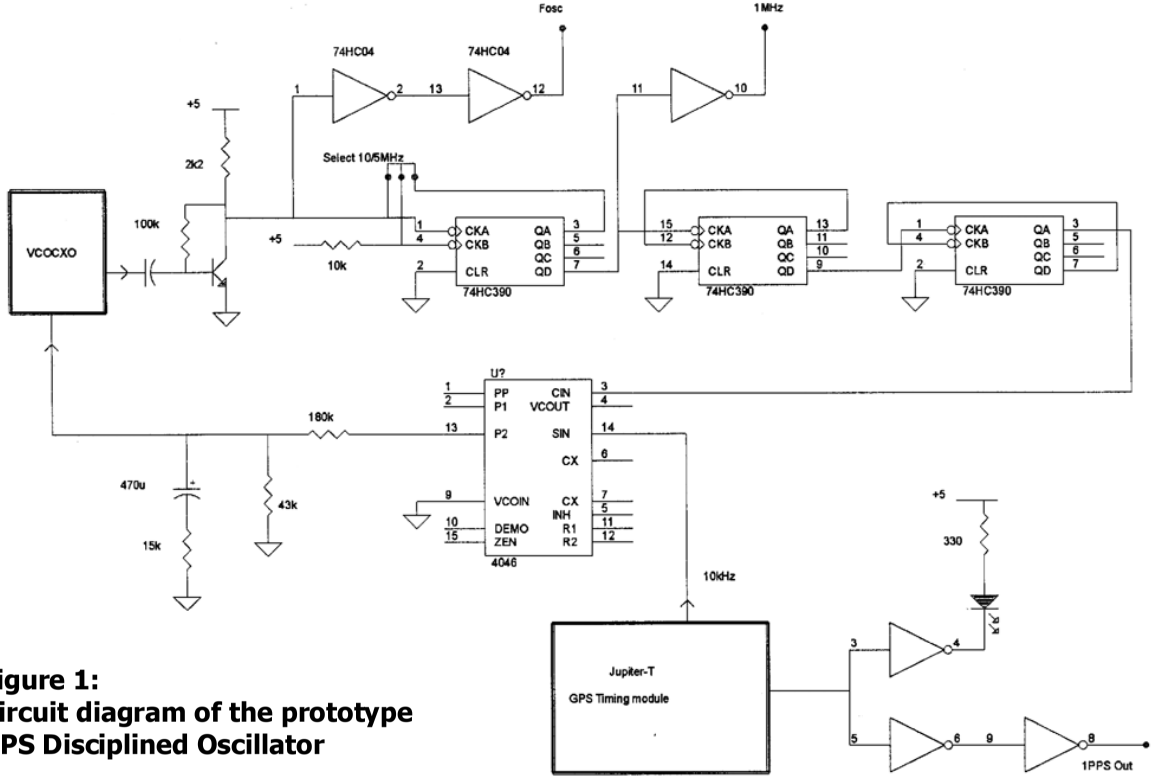

Figure 1:
Circuit diagram of the prototype GPS Disciplined Oscillator

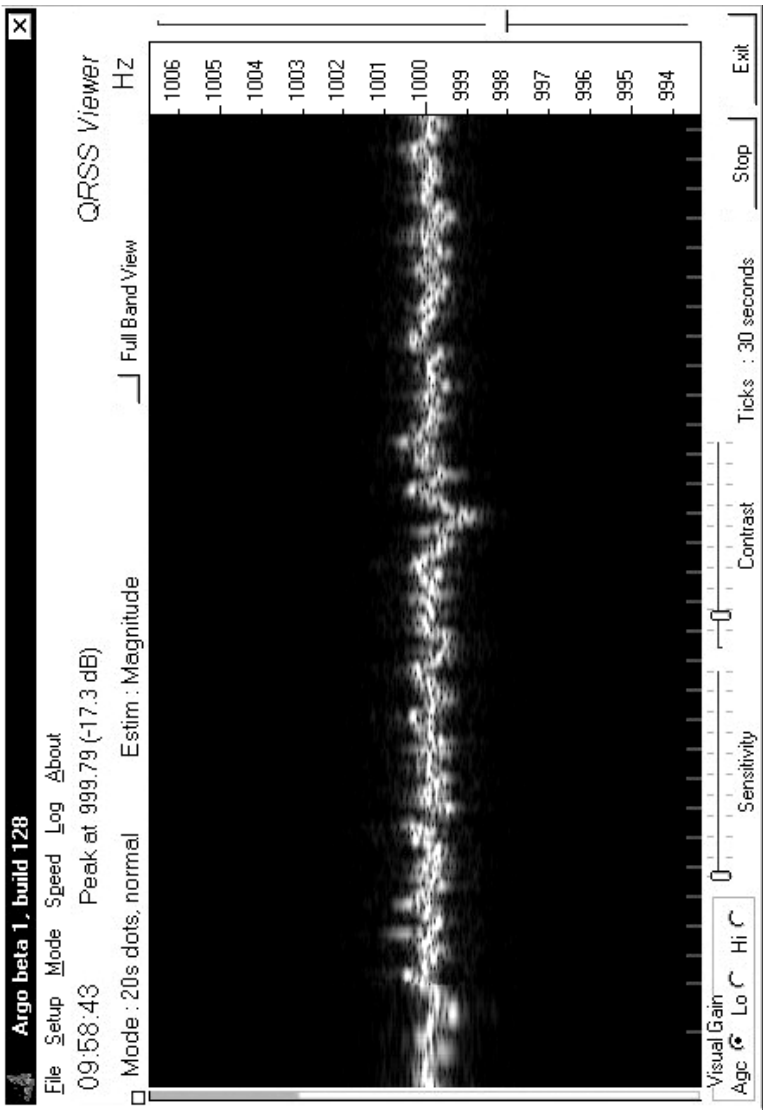

Figure 2 Frequency tracking performance when locked, VCXO output multiplied to 1GHz

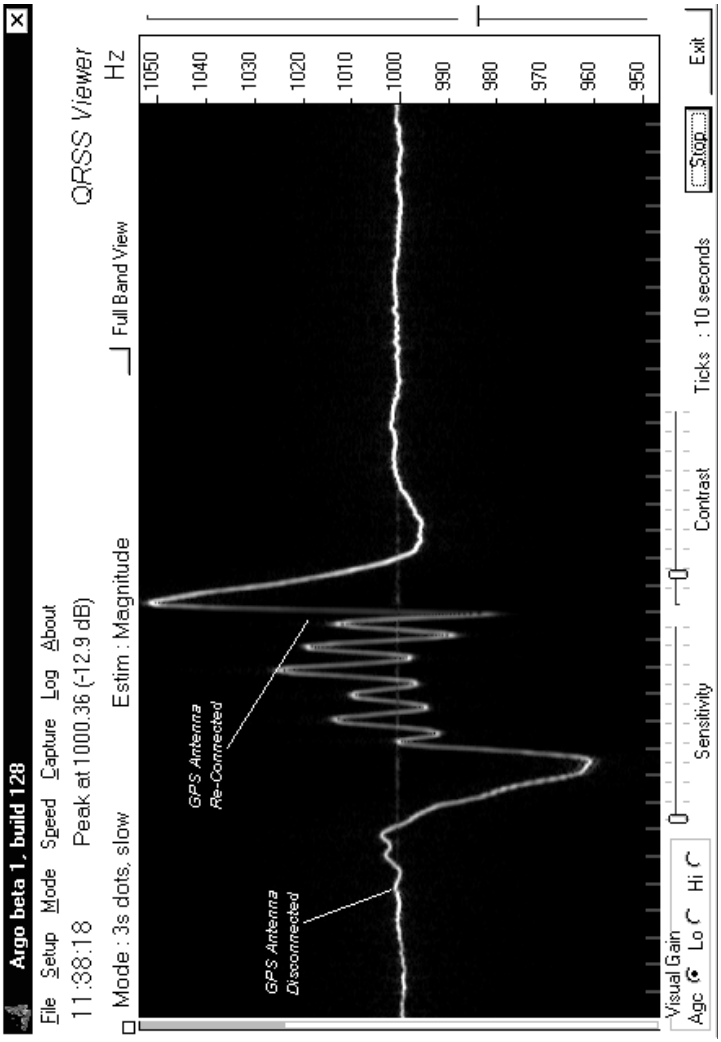

Figure 3 Frequency tracking with loss and re-acquisition of GPS signal—multiplied to 1GHz.

Jupiter-T

This design is based around one specific GPS receiver - others with a 10kHz output may be available, but I don't know of any. This is the Jupiter T GPS Timing Board.

Off Air Frequency Standard

David Johnstone, GM4EVS

Design notes

This design is based on the G4JNT Mk2 MSF Standard published in the RSGB Microwave Newsletter in April 1996. Changes have been made to the design of the MSF receiver and the 10MHz master oscillator.

MSF Receiver

Good receiver performance is one of the key components of a reliable off-air standard. This is especially true for users located 200 miles or more from the transmitter at Rugby.

Some have used directional frame aerials for MSF reception. Although cumbersome, these can provide an excellent signal. I was fortunate to have an 8in long x 0.5in diameter ferrite rod available.

A winding of enamelled copper wire, 6.5in in length, gave an inductance of some 10mH. This was resonated at 60kHz with a combination of fixed and variable capacitors. This included an allowance of some 180pF for the 6 foot length of RG174 coax used connect the rod to the receiver.

A 12 inch length of 7/8 in OD plastic tubing was used to house the rod, together with a 250pF compression trimmer. The cork from a wine bottle was cut in half to provide a 'plug' for each end if the 12 in tube. One of the plugs was drilled to accommodate the RG174 cable.

The receiver itself was built on double-sided PCB, with the top-side acting as a ground-plane. This PCB was housed in a 4.5 x 2.5 x 1in aluminium die cast box. This receiver module was mounted within the main case used for the off-air standard. Good shielding ensures that the receiver amplifies only the off-air 60kHz signal, and not the nominal 60kHz signal generated by division of the 10MHz master oscillator. The receiver circuit followed the ideas proposed by G3LYP to introduce additional LC tuned circuits to improve selectivity. Indeed, those of you with the last 32 years worth of RadCom in the loft, will know that the April 1970 issue provided the design for a Droitwich-locked (then 200kHz) standard by GM3TFY. The receiver used 4 cascaded NPN bipolar LC tuned amplifiers and employed an external antenna.

In addition to the LC filtering, the main FET RF amplifier was changed from an MPF102 to a J310.

Biasing of the FET amp was adjusted to give a drain current of some 8mA. The LC tuned circuits were under-coupled to compensate for the greater gain of the J310. The net result was about the same gain but better selectivity and IMD performance.

During the design, a series LC circuit was included to allow notching-out of a chosen unwanted LW signal e.g., Radio 4 on 198kHz. Use of this was found to be unnecessary in practice.

The biasing of the bipolar amplifier was changed to improve its temperature stability. The emitter voltage was made large relative to the nominal 0.6 volts V_{BE} of the transistor. So when the V_{BE} varies with temperature, the biasing of the transistor remains largely unaltered. Also, some emitter degeneration was used to control the gain of the stage.

The emitter follower design added a 220R series resistor in the base feed to pre-empt any tendency towards oscillation.

The aggregate performance of the

revised receiver was confirmed using the ARRL Radio Designer CAD package. In practice, at a distance of 300 miles, a clean 0.8v p-p output was obtained from the receiver.

With this, the standard remains locked 24 hours per day despite propagation changes at dawn and dusk. As yet, the 10MHz output from the standard has not been multiplied up to check the jitter at 10GHz.

10 MHz Oscillator

Some pulling of the 10MHz oscillator was reported by G4JNT when using the 10MHz output from the Mk2 standard. The oscillator design was revisited. Inspection of the master oscillator circuits for the TS930, TS940, TS850, TS570 and FT1000MP showed them to be all nominally the same. A Colpitts design was used employing a bipolar device having a high gain-bandwidth product. This choice ensures a low base-collector capacitance (C_{BC}) minimising the effect of temperature-driven C_{BC} changes on the frequency of oscillation.

A Colpitts oscillator was built using a 2N5179. This was lightly coupled to a bipolar emitter follower using a 2N2222. You may be inclined to choose an FET follower because of its high input impedance. But overall, the bipolar emitter follower is better-behaved. The prototype oscillator was breadboarded using a small piece of single-sided PCB, face-up, cut into small islands with a hacksaw. The prototype worked so well that it was used for production! The oscillator unit was housed in its own die cast box within the main unit.

While free-running on the bench, the 10MHz oscillator moved about 3Hz in the first 30 minutes from a cold-start. Thereafter, a variation of less that 1Hz per hour was observed during the day with the shack temperature remaining around 64-66F. The measurements were made using a Racal 9916 frequency counter with an OCXO reference that had been on for several weeks.

As the oscillator was operated at low power levels, the 10MHz output from the emitter follower was only 0.4 volts p-p. This level must be raised to 5 volts p-p to drive the 74HC390-based divider chain. This can be achieved cheaply and effectively by wiring a CD4046 as an amplifier – a technique used by Brooks Shera in his GPS-locked frequency standard (QST July 1998). All sections of the 74HC390 divider were used to give 10MHz, 1MHz and 100kHz outputs. Thereafter the G4JNT Mk2 design was followed. A 3MHz signal was filtered out from the 1 MHz square wave, and divided by 50 to get a nominal 60kHz signal for phase comparison with the off-air 60kHz signal. A CD4046 phase comparator was used followed by an LF353 as a level shifter, to provide the drive to the varicap diode in the 10MHz oscillator. In all, four separate PCBs were used within the unit – MSF RX, 10MHz oscillator, 10MHz divider and phase comparator / level shift. Apart from the MSF RX, the remaining 3 boards were each fitted with separate 78L05 voltage regulators and extensive decoupling capacitors. An external 13.8 volt supply is used to power the unit.

The unit has performed well and credit must go to Andy, G4JNT, for the robustness of his original design. If you won't always have access to a high-accuracy TV signal, or shy away from the complexity of a digital GPS-based solution, then an MSF-locked standard may well fit your requirements.

Finally, I have some people to thank. These include Dave, G4AON, for his help and inspiration.

We started our MSF projects together, but he built the G4JNT Mk2 design in much less time than I took, and kept reminding me of this!

Also, I am most grateful to Peter, G3PHO, and Mike, G3LYP, for the information and help they provided to make this project a success.

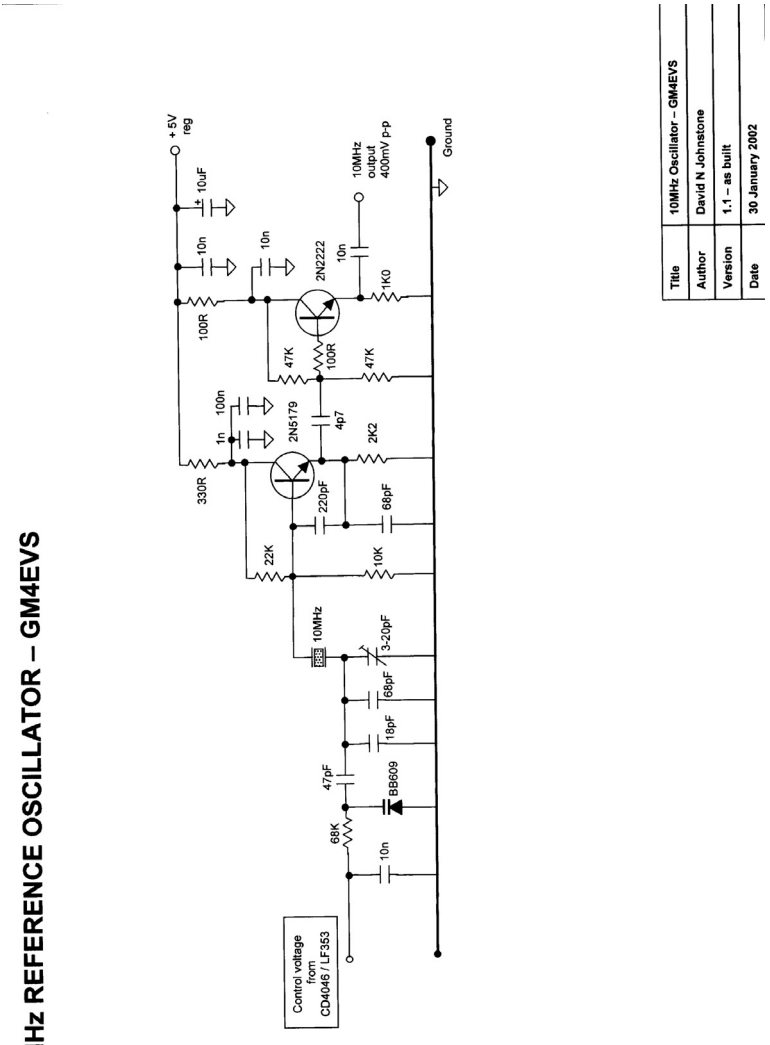

General

Crystal Heaters

Some useful observations by Doug Friend, VK4OE

I recently did some investigations on two Murata crystal heaters that I had available. One type I had bought from Charlie and Petra back in the mid 1980's which I believe were designed for operation at 40 degrees C, and the other type was obtained locally in Australia and intended for operation at 50 degrees C. My friend Rod VK4KZR and I had assumed that temperature was stable, somewhat independent of operating voltage, but some undesirable frequency variations led us to do our investigations. Shown below are the results of our tests. Clearly, the older 40C devices don't vary much with voltage, but 8V would seem to yield the actual specification temperature. But our interest was not much in these 40C heaters because that temperature is often the ambient temperature is Summer here, with internal equipment temperatures expected to be higher. The heaters intended for 50C very obviously should be used only at 5V. At that voltage there is reasonable repeatability between individual devices, which is not the case at higher supply voltages. I can vouch for the very good temperature performance of oscillators using these heaters at 5V. The only drawback is that the crystal should be specified/ ordered for operation at 50C. I have used up the two older 40C heaters in circuits where the crystal has a normal ambient temperature specification. I much prefer the 50C heaters.

MURATA Model	Unit No.	Temp. @ 12.6V	Temp. @ 8V	Temp. @ 5V
BM500N	1	43 ºC	NT	37 ºC
BM500N	2	43 ºC	NT	36 ºC
BG330N	1	58 ºC	55 ºC	51 ºC
BG330N	2	61 ºC	55 ºC	49.5 ºC
BG330N	3	58 ºC	56 ºC	50 ºC
BG330N	4	62 ºC	55 ºC	48 ºC
BG330N	5	63 ºC	55 ºC	49 ºC
BG330N	6	62 ºC	NT	50 ºC
BG330N	7	NT	57 ºC	50 ºC
BG330N	8	NT	58 ºC	50 ºC
BG330N	9	NT	55 ºC	48 ºC

Reducing Warm Up Drift in Local Oscillators

From and idea by David Hall, G8VZT

Murata crystal heaters are commonly used to reduced warm up time and drift in oscillators such as the DDK004. Unfortunately this is not always optimised for minimum drift as the crystal is usually fitted flush to the pc board and the Murata has to heat both crystal case and the surrounding ground plane on the board.

The warm up time can be significantly reduced (by as much as 50%) by simply insulating the base of the crystal can from the pc board. Fit a small PTFE (Teflon) shim or spacer under the crystal, between it and the board. A 1000pF ceramic plate capacitor maybe soldered from the crystal can to the ground plane to effectively ground the can to external RF influences. The ground plane around the crystal can also be removed to some extent in order to reduce the area affected by the Murata heater.

Further improvement can often be made by running the heater off a 9V or 10V regulated supply.

G4DDK004 Oscillator Upgrade

Notes by Kevin Murphy, ZL1UJG

I did some improvements to the last DDK004 oscillator that I built. Listed below are the modifications and the reasons for them. You will need the original circuit diagram and construction notes for reference:

MODIFICATION	REASON
1. Fit a 330nH axial inductor across the crystal (1nF series capacitor Cx must be fitted to prevent DC flowing between TR1 and TR2 emitters)	This gives a stable pulling range due to the xtal capacitance (~7pF) being cancelled
2. Ground TR3, TR4 and TR5 emitters directly. R11, R15, R19, C12, C17, C22, C24 are not fitted. The holes for TRs 3, 4 and 5 were carefully filed with a small round needle file so that the collector and base leads would fold down in little notches while the emitters were soldered directly to the ground plane. The PCB areas where the emitter Cs and Rs were are covered with copper tape and soldered.	This gives a direct path for the emitters and reduces potential problems with emitter bypassing.
3. The 1st striplines in TR4 and TR5 collectors are now directly grounded with copper tape. The transistors' outputs are fed via small chip caps (e.g. 33pF 0805 types). Tr4 has its collector volts fed via a 100nH 1206 size chip inductor. TR5 pcb stripline extended with a small wire. C21, C25 and C29 are not fitted. Where the first tuned striplines in TR4 and TR5 collector circuits are grounded, there are 2 cuts in the 1st striplines. The 1st cut is to isolate the collector from the tuned stripline and the second cut is to isolate the grounded end of the 1st stripline from the Power feed circuit. The 2 pcb cuts can be done in one motion with a scalpel and metal ruler.	This increases the Q of the striplines by grounding them directly.

MODIFICATION	REASON
4. Collector resistors R12, R16 and R20 (associated with TR3, TR4 and TR5) are increased to 150 ohms. The collector's side of each R is fed to the base bias circuits.	This is called a self-regulating, saturated multiplier circuit. (The collector voltage reduces under drive and reduces the base bias). Values other than 150 ohms were not tried. Increased power may be obtained by reducing the resistor values to 120 or even 100 ohms.
5. A 3 terminal regulator (78LO8) is fitted on the pcb with associated tantalum caps of 10uF on the output and 1uF on the input. Note the 4 holes (each 1mm diam) drilled with clearance pads. (use on outboard 7808 1 amp regulator if a crystal heater is fitted and run it off 8 volts).	This reduces the effects of the operating point of TR3, TR4 and Tr5 changing due to supply variations and reduces chirp and other random variations.
6. The small 3 terminal regulator, IC1, already fitted to the pcb for TR1 and TR2 is removed. A new NPN transistor (e.g. BC547 or 2N2222) is fitted (collector to +8V, base to centre hole and emitter towards the Butler oscillator). A 1K resistor from the 8V to the base of the transistor and a 10uF tantalum from base to ground is used. Note: the middle hole must now have a clearance pad so that the base lead doesn't short out.	This reduces regulator noise feeding into TR1 and TR2 circuitry.
7. Fit low cost and more common Philips 1.4 – 5pF trimmers for the 1200MHz multiplier section by reversing the middle trimmer (like an interdigital filter).	Availability and the original trimmers were over coupled, causing the middle trimmer not to tune.
8. An extra 2.2pF NPO ceramic capacitor is fitted on TR3's base and the position of R9 is moved. Two holes need to drilled, with clearance pads.	To improve the multiplier performance (increased collector current)

Notes:

Where tantalum capacitors are mentioned in the original DDK004 documentation and in these modifications, it is very important that they are used as they have superior filtering characteristics. Even if you are not doing these modifications to an existing pcb, it is still recommended that an additional 10uF tantalum is fitted close to the output of IC1 to suppress noise.

The level of crystal harmonic products on the output (e.g. 1200MHz) is better than –45dBc. The board draws ~60mA. MRF901 transistors were used for Tr3, 4 and 5 multipliers and +7dBm (5mW) output was obtained. R8, 18 ohms, was changed to 100 ohms but it is not a necessary modification.

Overall, the modifications are made to improve the keying characteristics, stability and noise characteristics of the DDK005 oscillator/multiplier. I am upgrading four other units to this standard so that their performance is enhanced for transverter and beacon use.

Since the crystal is sensitive to temperature, it is recommended that a crystal heater is fitted. These are obtainable from Downeast Microwave un the USA or from VK5EME Minikits in Australia. Failing that, some sort of thermal jacket from foam or a polystyrene packing bead can be fitted over the crystal.

For even superior stability, an external crystal oscillator such as that available in kit form from John Hazell, G8ACE, may be fitted and fed into the first stage of the Butler Oscillator, after removing the existing crystal and choke. For matching reasons it is recommended that a J310 FET is fitted instead of TR1 and that R2 is changed to 220 ohms. TR2's emitter in this case can have the 22pF NPO capacitor replaced with a 1nF type.

A Distribution Amplifier for the G4JNT PS-DSO

Paul Marsh, M0EYT

A while back, Scatterpoint published an article by Andy, G4JNT, on making a simple GPS disciplined oscillator based on the cheap Jupiter GPS boards. Since I had a couple of these modules spare, I decided to make one. I used a good quality 10MHz oscillator from John, G8ACE, as the ovened source – then the simple circuit from Andy to lock the oscillator to GPS. The output was locked on 10MHz,and I was very pleased with the results. Then I looked at the 10MHz output on the spectrum analyser.

At this point, I noticed that the 10MHz output and its harmonics extended up over 1GHz – it does make a nice marker generator as is but, since I wanted to lock a receiver to this GPS-DO, I needed to clean the output up a bit. After some hours wasted trying to get good results with multi pole filters I'd bread boarded, I needed a new solution. I looked around the garage for inspiration and found an old box of 10Mbps Ethernet LAN cards that I'd kept. Whilst looking at the isolated BNC sockets on them, I noticed the transceiver modules, so decided to "Google" them and sure enough, one module was a multi-pole filter for the RX and TX paths. After hooking one from the board, I checked it on the spec-an with the results below:

10MHz output before filtering

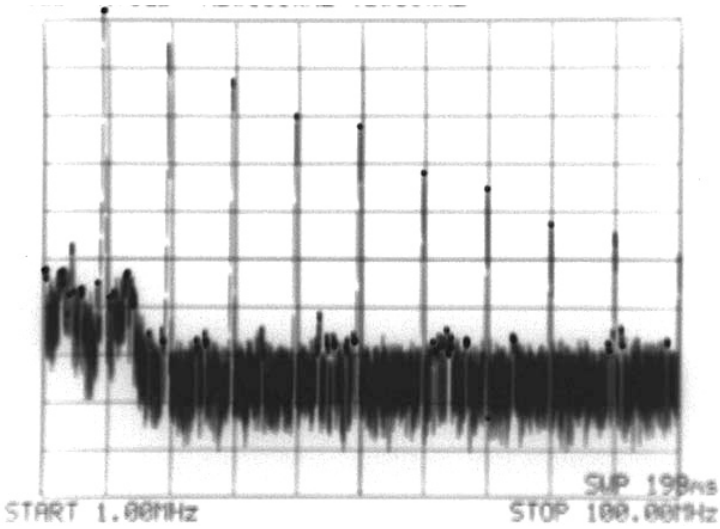

10MHz output after filtering

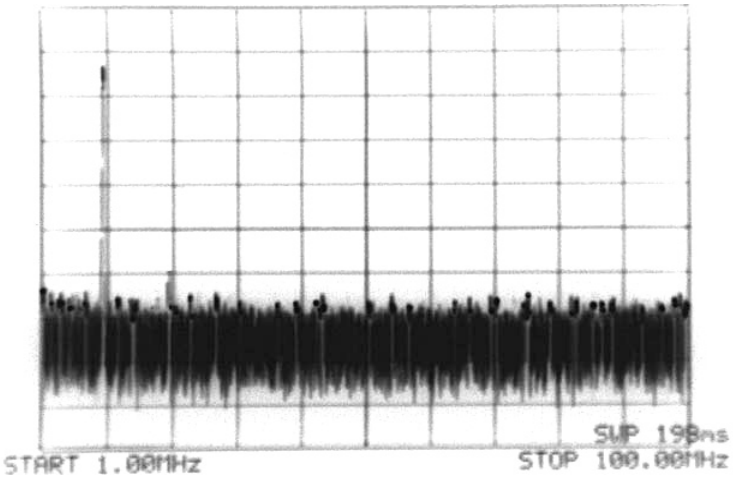

The output from the JNT GPS-DO is first buffered by a 74HC04 Hex Inverter, 2 gates in series, then the second gate driving the remaining 4 gates in parallel.

The 10MHz output from the 74HC04 of course contains harmonics. Each output is fed via a 27 ohm resistor, then into the filter network. The filters are marked 'YCL 20F001N' Each YCL 20F001N package contains one 7 pole TX filter and one 5 RX filter – you can either use the two filters to drive 2 separate outputs, or as I did, wire both filters in series to get further rejection of unwanted harmonics. The pin outs for the filter block are on the schematic diagram below:

The other good thing about these filter blocks are the isolated outputs – if you use the isolated BNC socket commonly found on the old LAN cards, you can get rid of problems caused by earth loops too.

I found once I had built the buffer/filter circuit, the outputs resembled a clean sine-wave on the oscilloscope. I'm sure the output from my old Wavetek synthesiser is cleaner now too. I did make a PCB for this buffer/filter; it is for surface mount DIL packages.

Schematics: