Locked 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 (eg 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 http://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.

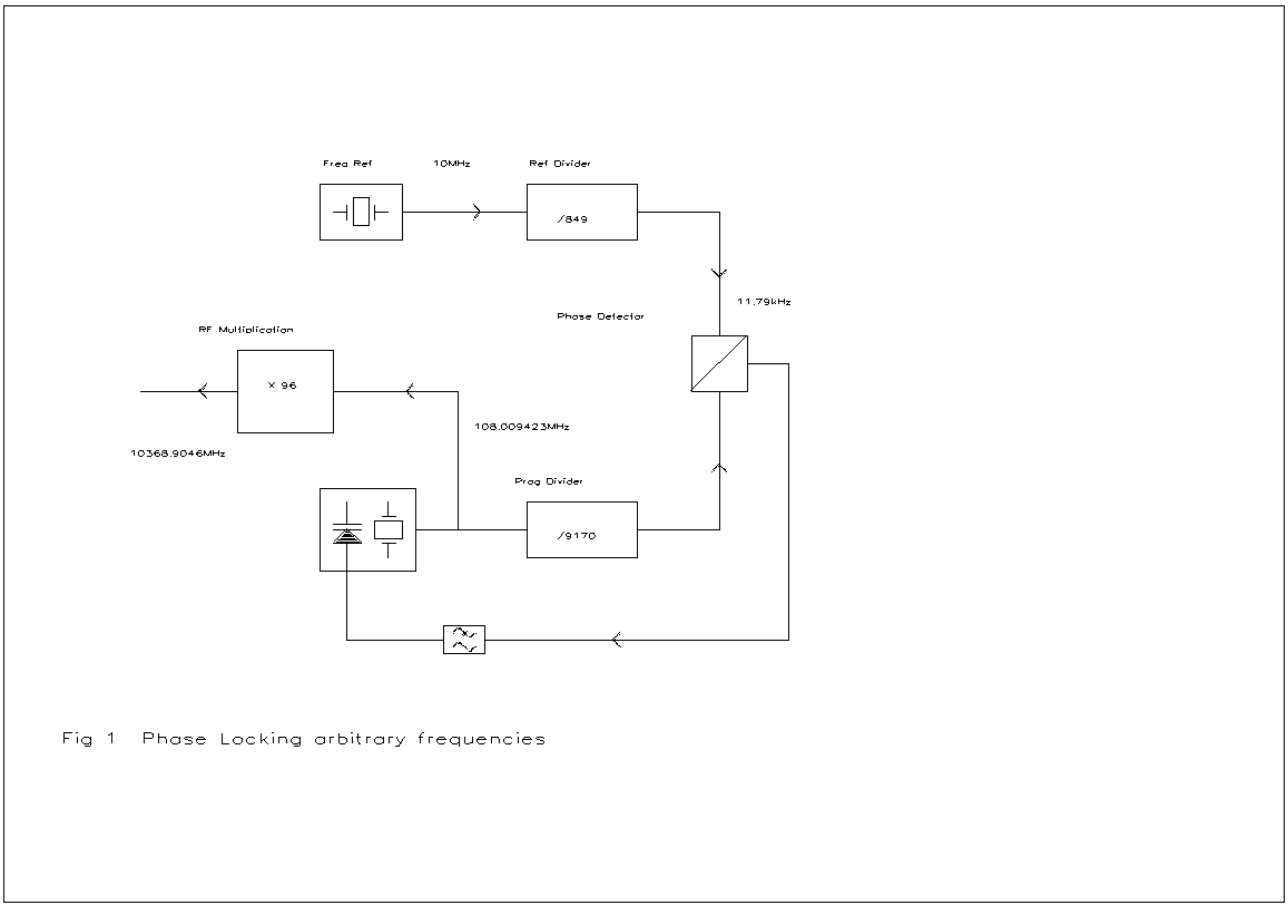

Fig 1 Phase Locking arbitrary frequencies

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. Microwavers 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 :

$$BW = Kv \,.\, Kd / N$$

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 :

$$BW_{(Max)} = 250 \times 1.6 / 9170 = 0.04\text{Hz or 1 cycle per 23 seconds}$$

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 unovenned 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 RadCom 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 pin compatible with the AD9850 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.

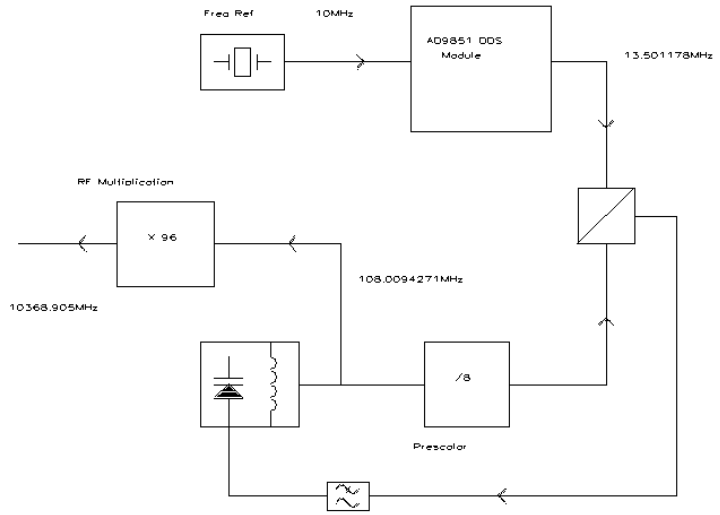

Fig 2 DDS Source for generating the PLL reference

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. Instead of using a prescalar to divide the VCXO down to a frequency that could be generated by the DDS, why ot 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 DDS device device is only operating at 100MHz the original AD9850 can be used in this adaptation. 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 unovenned crystal to be used over a wide frequency range, but with decent overtone devices can usually be achieved.

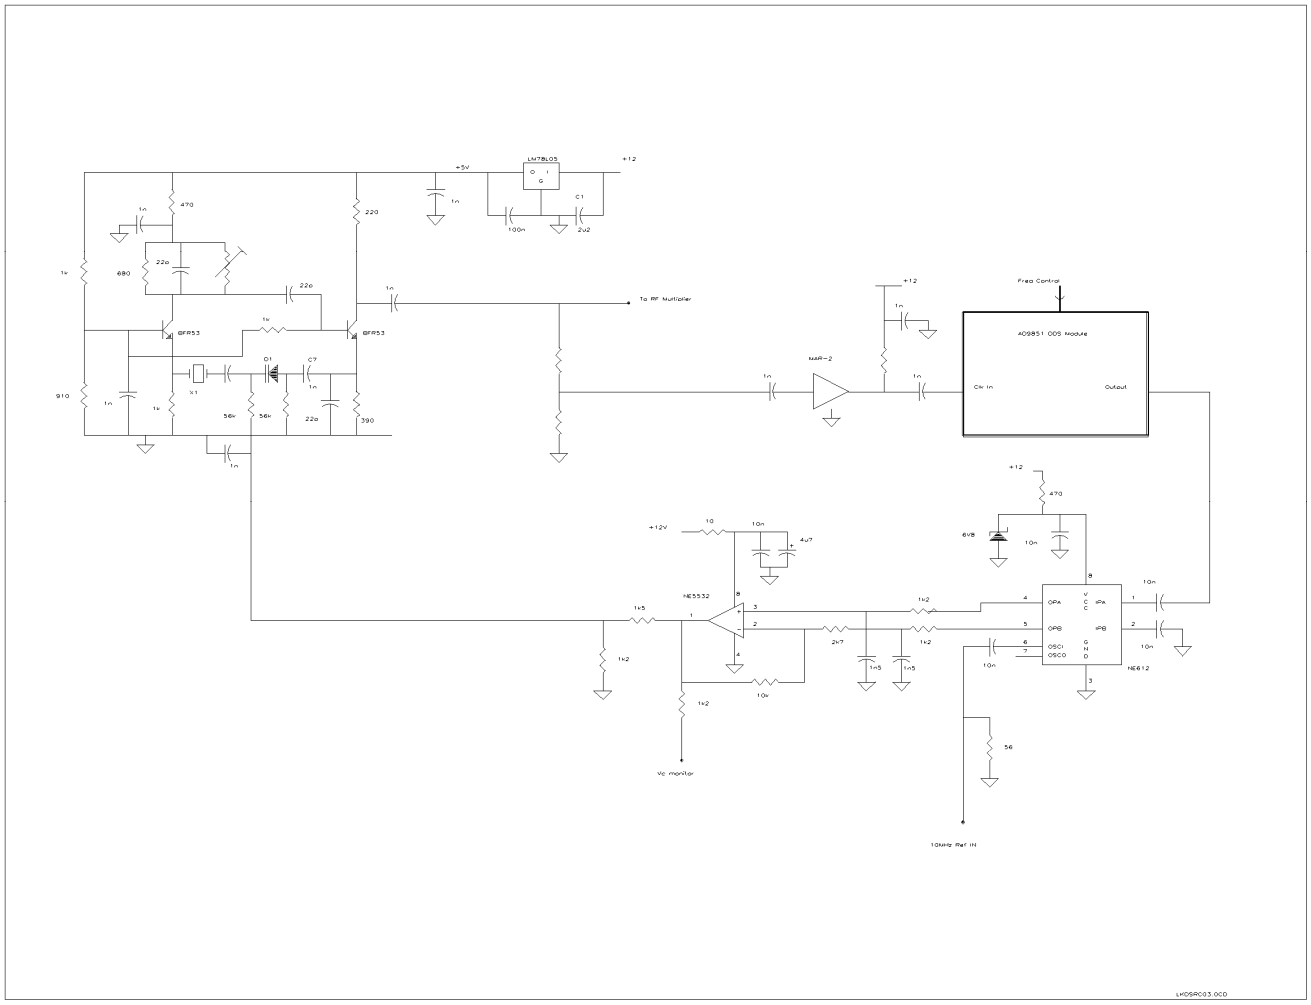

An incidental advantage of the loop locking up quickly it that is just about possible to generate frequency shift keying by reprogramming the DDS. Provided the frequency shift is small - no more than around 1 to 2kHz - the loop can track this change within a few milliseconds. If CW at around 12WPM is 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 approximately $F_{clock} / 2^{32}$ which for a 60MHz clock is 0.014Hz. Although calculations are complicated by the DDS working 'backwards' When multiplied up to 10GHz this step size translates to a little over 10.7Hz - probably accurate enough for most purposes!

Future versions will use the AD9852 DDS. This chip can have a clock of up to 300MHz and an internal PLL which is programmable between 4 and 20. 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 not Megahertz, its millihertz!!!)