The Quest for Microwave Frequency Stability
Or
How Far Should I Tune for You?

Tom Williams, WA1MBA

Situation

Whether on a hilltop or at home, knowing the right frequency half the battle of making a difficult contact on the microwave bands. Below 1296 most of us are sure just where we are, and warm-up time is quite small, usually just a few seconds. But on 10 GHz and especially on the EHF bands, we tune around several if not dozens of kHz to find the other signal. Usually, finding a very weak signal is a process that includes sweeping the antenna and sweeping the frequency. Reducing or eliminating one of the variables will significantly decrease the time to make a contact so that we can make more of them in the limited time that we are on a hilltop. Sometimes, not knowing the frequency can mean not making a contact at all. I can recall several times when having someone spot the frequency of the other station made an otherwise impossible contact happen.

Why do we need to do this on the higher bands and not on the lower ones? The answer is simply because frequency instability increases with frequency. An error of 1 kHz at 1296 is quite tolerable. Often we will move 2 kHz at 1296 and not even notice it because we hear the signal in our 2.7 kHz passband before retuning. But a 2 kHz error at 1296 equates to an error of 18 kHz at 10 GHz, 37 kHz at 24 GHz, and a whopping 372 kHz at 241 GHz. The few of us who work at 47 GHz and higher carefully plan for frequency accuracy and stability.

Some Definitions

Stability is the measure of change in frequency over the course of time. There are other, more specific stability measurements taken against external conditions, listed here in generally decreasing importance to hams. They are change in frequency with respect to change in:
1) temperature, 2) operating voltage, 3) electronic load, 4) pressure, 5) gravity (and acceleration), and 6) magnetic field.

Units of time-related stability are generally expressed in drift, where the expected change is a statistical measure of relative frequency change over different time scales. Sometimes we mix temporal (time) and temperature stability with the term “warm-up”. The very short-term time scales of less than one second are usually captured in phase noise measurements. The short time scales are from one second to tens of minutes. The mid-time scales are on the order of a day to weeks. The long time scales are usually months to years.

Time is a fact of existence. Its definition has kept philosophers, theologians and physicists puzzled for centuries. It is often described as the linear progression of events, from the past to the present. All definitions of time presuppose the experience of time passing. In engineering we generally agree that we all know what time is, and then attempt to measure it accurately and provide standards. The standards used by our governments rely on atomic vibrations.

Unfortunately, these standards are susceptible to atomic motions, and so the very best clocks contain highly specialized systems to keep the atoms being measured in a controlled state. The very best oscillators have a likely error or drift in the order of one second over thousands of years. Portable hydrogen maser clocks can be made with drift rates and accuracy on this order, but they cost hundreds of thousands of dollars.

Accuracy is the exactness with which a measurement is made or a frequency is generated. Precision is the exactness with which a measurement is reported, even though it might not be correct or accurate. For instance, a cavity "frequency meter" may read our oscillator at 10.22 GHz. The instrument has no more precision than this, and in fact the last place is a visual interpretation between two marks of 10.2 and 10.25 GHz. This is certainly not enough precision to determine that the oscillator is on frequency as a 10,224 MHz LO, where we need to know within less than 20 kHz what the frequency is in order to make contacts. Conversely, consider a highly precise meter that is not accurate. A nice 11-digit frequency meter may read 10,224,041,329 Hz, telling us that the oscillator is off by some 41 kHz. This is plenty of precision for our purposes –who cares about the last three places? But, if the frequency standard inside this nice frequency meter is off by only 5 parts per million, then this readout could be measuring an actual frequency somewhere between 10,224,092,449 Hz and 10,223,990,208 Hz. We might be fooled by an error of 50 kHz and be unable to find the other station without a heck of a lot of hunting. The point here is: to know that we are "on frequency" we need our oscillators to be accurate. To adjust our oscillators, we must have metrology that is sufficiently precise and also sufficiently accurate.

Frequency stability is afforded by physical systems that 1) have sufficiently narrow oscillation modes (low enough phase noise) and 2) we can sufficiently minimize the external forces that cause the frequency of oscillation to move. All electronic oscillators have an electro-physical (e.g. cavity or quartz crystal) or electronic (e.g. LC circuit) resonator and a feedback amplifier. For the most part, temperature disturbances will strongly affect the resonator and weakly affect the amplifier. Electrical disturbances will strongly affect the amplifier, but often strongly affect the resonator as well. We try to reduce the effects of temperature, voltage and load to a point where stability is good, and ignore the even smaller effects from pressure, magnetism and gravity or acceleration.

Figure 1. Here is a nice readout (there is a "1" off scale to the left) showing a 10 GHz CW signal "on frequency" to 100Hz – if the time-base standard for the frequency counter is accurate!

Portable Stable Sources

The portable and stable sources of frequencies that we (hams) use are based on either quartz or rubidium resonance. There are other, quite expensive sources, including cesium and hydrogen atomic standards, and the sapphire resonator.

Quartz crystals can be cut to provide stable resonators. Unfortunately, they are quite sensitive to temperature. The crystal can be cut in a manner that minimizes the effect of temperature within a range. Furthermore, we can control the temperature and also can compensate for changes in temperature. Controlling is accomplished with a heater that will maintain a constant temperature, significantly above ambient so that if the unit is exposed to the sun on a hot day it will not exceed its operating temperature. This is easily (albeit somewhat crudely) accomplished with a positive temperature coefficient thermistor, or PTC. This device can be soldered directly to a crystal can. The can and heater are enclosed in a heat insulator, such as polystyrene foam to make an oven. These oscillators are called OCXOs for Oven Controlled Crystal Oscillators. Such heated crystals can take up to 20 minutes to stabilize after power-on and will repeat the last warmed-up frequency to within 1,000 (poor performance) to 10 (excellent performance) parts per billion. This equates to not knowing your frequency at 10 GHz to within a range of 10 kHz to 100 Hz. Most of us experience frequency error in this range just after the warm-up period of ovenized LOs. Over a long term of operation, however, the crystal will age and drift some distance from its youthful frequency. Drifts over various time ranges are summarized in a table below.

Temperature compensation is usually accomplished by adding temperature sensors to both the resonator and critical circuit components. Control parameters are determined once the unit is constructed, and a feedback circuit uses these parameters to adjust circuit reactance in order to keep the frequency nearly correct. TCXOs are usually not as good as OCXOs, but because of miniaturization of electronic circuits, they are much less bulky than ovens and fit nicely into small equipment. The best stability crystal oscillators are DOXOs or Double Ovenized Crystal Oscillators. Often, high quality 10 MHz standards will be single or double ovenized and temperature compensated. Typical DOXO standards will slightly overshoot their final frequency when the oven reaches temperature, and then gradually close in. Drifts of non-ovenized TCXOs can be quite large as described below, unacceptable for EHF work, but still quite acceptable for many other communication applications. DOXOs and/or temperature compensated OCXOs, as are found in many 10 MHz standards have significantly better performance than a TCXO, and can provide acceptable stability, accuracy and drift for EHF work.

The rubidium standard is not nearly as affordable as a home-made OCXO or a 10 MHz standard DOXO. New rubidium units sell for about $1,500. I have heard rumors of them selling surplus for as low as $250, and have purchased them for a little under $500 each in small quantity. They are constructed in several form factors, and are available in 12, 15 and 24 Volt versions, positive and negative ground. Mine are probably the most common at 24 Volts negative ground. The unit is about as heavy and large as a microwave PLL Brick, and needs a small heat sink. It dissipates about 30 watts – slightly more during the typical 6-minute warm-up.

Figure 2. This photo shows a rubidium standard on the right, and for size comparison a “brick” oscillator on the left.

The rubidium standard is a very high quality frequency standard in a small package. It has stability and repeatability way beyond the capability of most crystal oscillators. There have been some recent developments in high-end DOXOs that approach some of the parameters of the rubidium, but they still fall short by a factor of more than 10 in most cases. There are some adjustments on the rubidium source, but they seem to move the frequency only a dozen or so millihertz at 10 Megahertz.

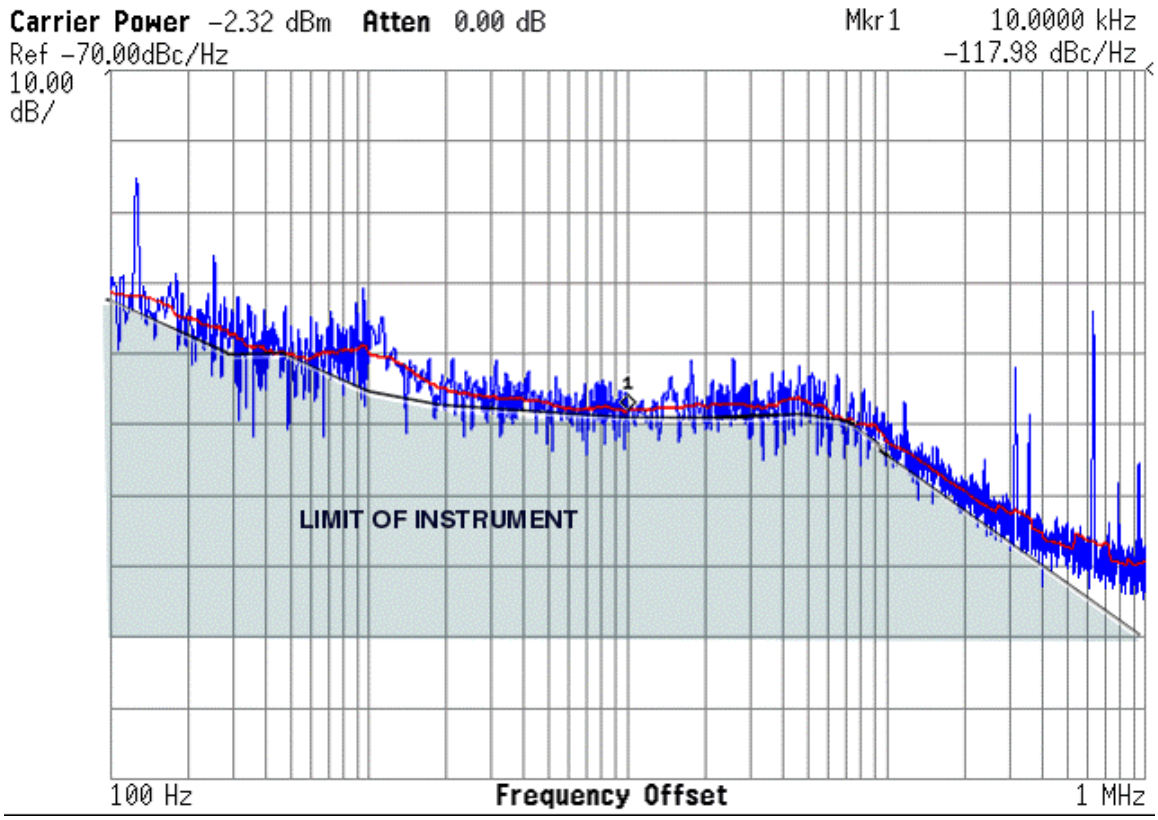

Figure 3. This plot shows the phase noise of one of the rubidium standards in the author's possession. The top of the plot is at –70 dBc. The average drops from –100 dBc to –110 dBc quickly, then plateaus at about –115 dBc up to about 50 kHz where it continues its descent. Although respectable for an oscillator of such stability, ones with superior phase noise are on the market. One notable problem is the set of spikes in the 200 to 500 kHz range. These are likely do to the internal phase lock loops that are synthesizing 10 MHz from the microwave frequencies that are atomic resonances of the rubidium. The peak at 120 Hz is due to the power supply. Unfortunately, much of the plot is near the limits of the instrument, an Agilent E4440A with phase noise option.

The table below shows approximate frequency error (1 standard deviation) translated to 47 GHz. Some designs will be better and some worse than shown here. The values shown are at +/- one standard deviation. This means that statistically, 65% of the oscillators will be this good or better. To get equivalent values at 10 GHz, just divide these by 5. To get equivalent at 145 GHz, multiply by these values by 5.

Time Frame & Temp Range	XO	TCXO	OCXO	DOXO	Rubidium
Accuracy, 20 minutes after startup @ 25 C	20 kHz to 50 kHz	20 kHz to 50 kHz	5 kHz to 10 kHz	500 Hz	2.5 to 10 Hz
Accuracy , 1 hour after startup @ 25C	5 kHz to 20 kHz	5 kHz to 20 kHz	1 kHz to 5 kHz	50 Hz	2.5 to 10 Hz
Drift per day	25 Hz	25 Hz	25 Hz	5 Hz	.05 to 0.5 Hz
Drift per year	25 kHz	25 kHz	5 kHz	750 Hz	0.5 to 1 Hz

Drift, 10 years	150 kHz	150 kHz	5 kHz	5 kHz	5 to 50 Hz
15C to 35 C	200 kHz	1 kHz	250 Hz	5 Hz	0.5 Hz
0 to 70 C	750 kHz	2.5 kHz	500 Hz	5 Hz	2 Hz

The DDS and the PLL

Fairly stable and reasonably accurate single ovenized crystal oscillators can be constructed or purchased. If we need to be on frequency within ½ of a voice bandwidth (about 1.5 kHz) then they are not very suitable for 47 GHz and higher bands, and can even be far enough off to cause some problems at 10 GHz. If we are willing to tune 5 to 10 kHz, or wait an hour for warm-up, then well constructed versions of such oscillators may be OK for 47 GHz use, but not for 145 GHz. Fortunately, it is possible to synthesize a frequency from a standard, such as 10 MHz. This way, a small standard, such as an inexpensive DOXO affordable at flea market prices around $40 or a rubidium standard for around $300 can be used to assure both accuracy and stability sufficient for dead-on applications up to and beyond 145 GHz.

Two methods are generally in use for frequency synthesis. One is Direct Digital Synthesis (DDS) and the other is the fractional Phase Lock Loop (PLL). The DDS is easier to describe, so I will do that first.

The DDS

A memory chip is loaded with the value for the sine function at hundreds or thousands of points along the waveform. A digital to analog converter is connected to the output of this memory (the memory data is delivered to the D/A converter). Then, a clever address circuit produces an address sequence for the memory, each successive address being some increment over the last one (and it wraps around at the end). The address circuit is given the standard as its clock. To change the output frequency, one simply changes the increment. In order to synthesize useful frequencies, the clock must be run at 200 to 300 MHz, so the external standard is multiplied up to this range before being applied to the address circuit.

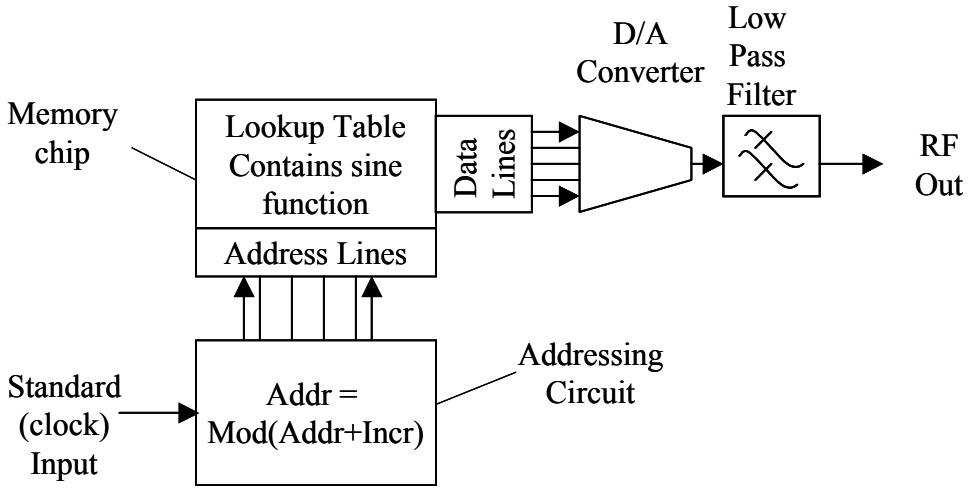

Figure 4. The DDS directly synthesizes a frequency numerically and converts it to analog.

The DDS has excellent frequency agility. In fact, it is so agile that it is often used as a digital VFO in modern transceivers. Fortunately, in those transceivers, one can operate the DDS at a low enough frequency and use sufficient filtering to significantly reduce the spurs that a DDS naturally generates. In order to generate frequencies useful to microwavers, a DDS output in the 60 to 120 MHz region is desired. If a single frequency or a few nearby ones in this range is desired for a particular application, it would be advisable to construct a narrow bandpass filter to reduce the effects of any output spurs on subsequent circuits. This will reduce out-of-band emissions and possible adverse mixing products on receive. Furthermore, there are some degenerate frequencies that produce significant nearby spurs. By changing one or two Hz these cases can be avoided.

Very high quality modern DDS chips are available at low cost and have been employed in some amateur designs for VFOs. I have taken one of these, the Analog Devices 9852, and with careful isolation and shielding, examined spurs. A plot below shows some example spectra.

The phase noise of a DDS generally follows that of the clock. This particular chip allows for the input of a 10 MHz clock and then will internally multiply it by an integer factor of up to 20 for the clock it uses. Unfortunately, the internal multiplier introduces more phase noise than a good external clock, so I used a high quality external clock at 240 MHz, itself locked to a 10 MHz standard. Multiplication of the standard to 240 MHz increased phase noise within the locking bandwidth of the multiplier, and so the resulting phase noise of the DDS is significant.

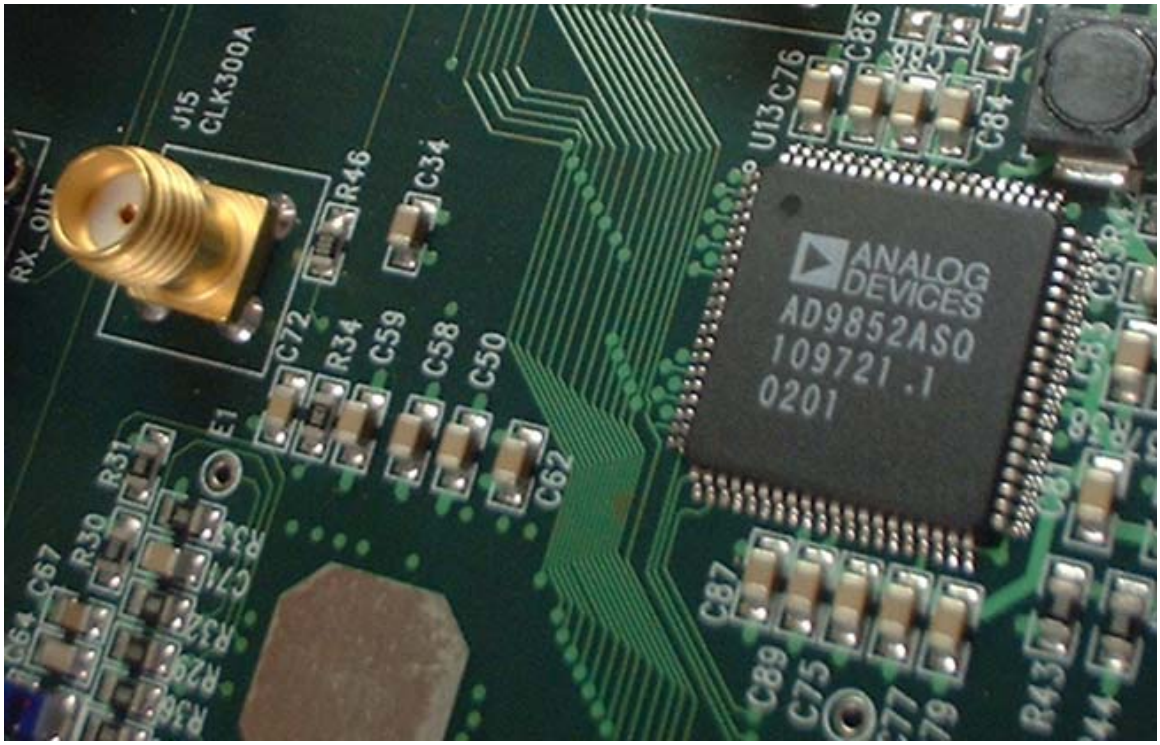

Figure 5. Here is a close-up photo of a surface mount version of the AD9852 used for the accompanying plots.

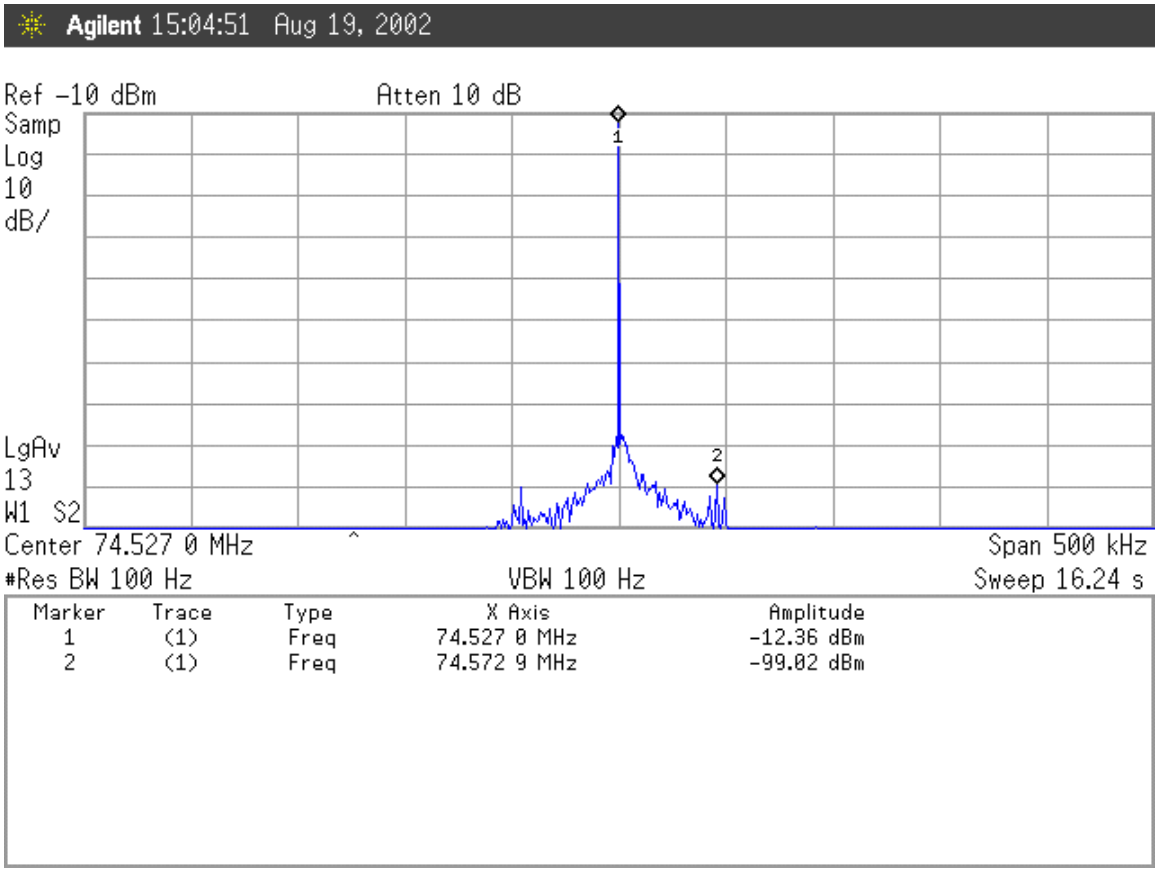

Figure 6. This plot shows the close spectrum around a 74.527027 MHz DDS. There are some spurs at -88 dBc about 46 kHz away from the primary signal. Otherwise, the near-in signals for the DDS are pretty clean. This might not be suitable for a high performance VHF receiver, but this spur is unlikely to be much concern for microwave use.

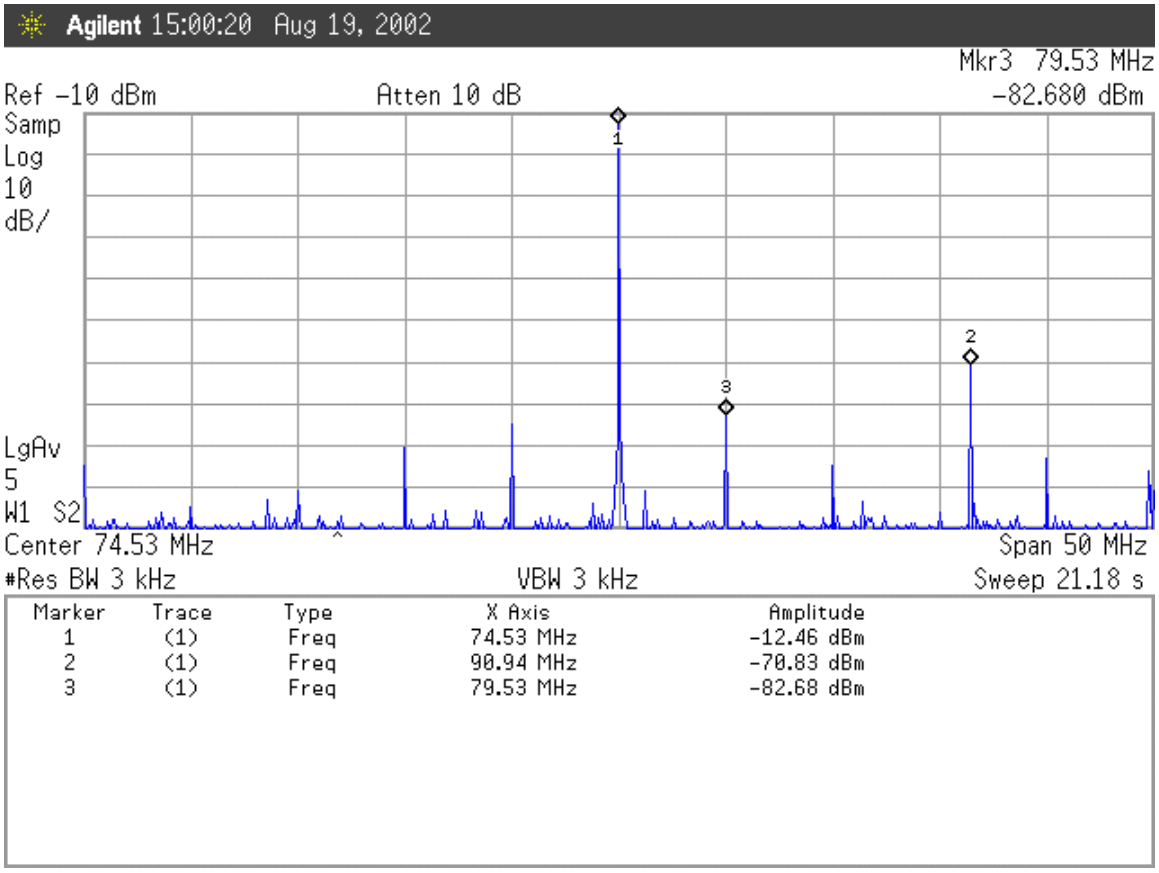

Figure 7. In this plot some of the wide band spurs from the DDS become apparent. Marker number two is only –58 dBc and number 3 is –70 dBc. If a narrow filter were used on the output of this DDS, it would provide a very accurate and usable locked source that can be frequency controlled quite precisely. These spurs, which can in the worse case be as strong as –53 dBc are difficult to predict, and so a good spectrum analyzer is needed to find a “sweet spot” close to the desired frequency that puts the spurs far enough away to filter out.

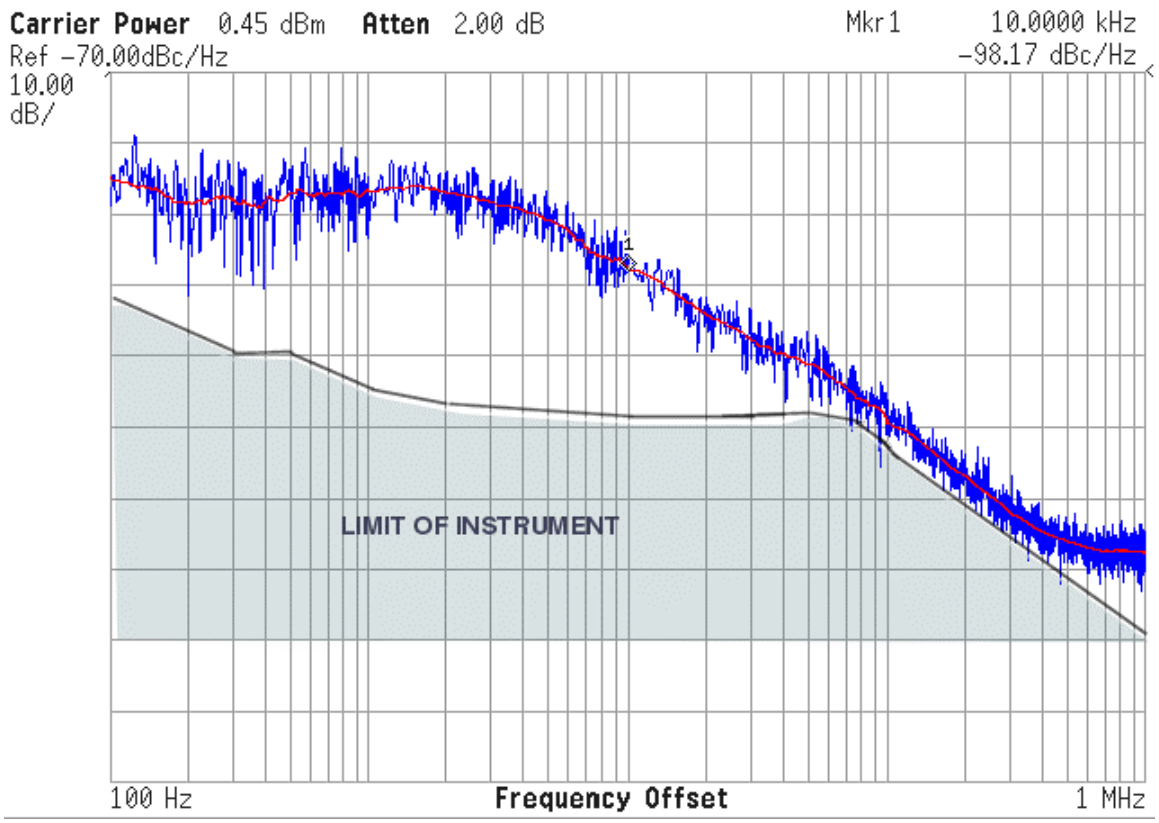
Carrier Power 0.45 dBm
Atten 2.00 dB
Mkr1 10.0000 kHz
Ref -70.00dBc/Hz
-98.17 dBc/Hz
10.00
dB/
LIMIT OF INSTRUMENT
100 Hz
Frequency Offset
1 MHz

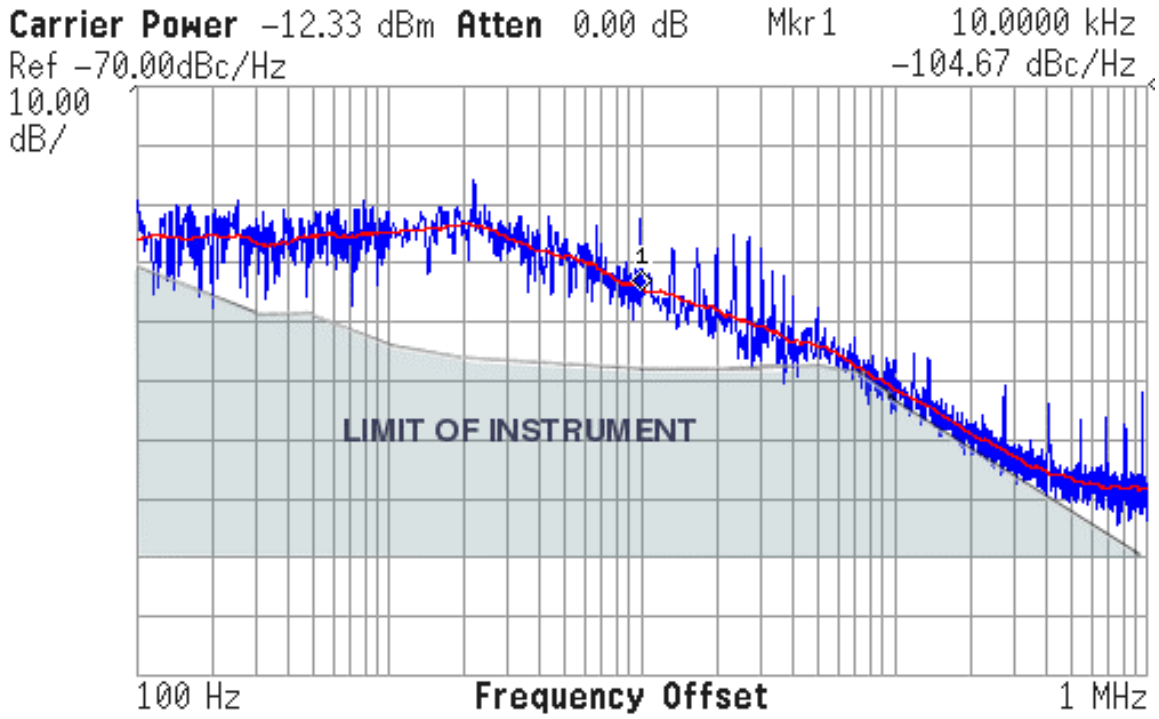

Figure 8. The top plot shows the phase noise of the 240 MHz locked clock that drives the DDS. The bottom one is phase noise of the DDS output at 74.527027 MHz. The top of each plot is -70 dBc. The average up to about 4 or 5 kHz is in the –95 to –100 range. It eventually drops off to a respectable –140 dBc. This is slightly worse than the original phase noise of the rubidium standard that is used to drive the DDS, but about 10 dB better than the 240 MHz clock up to about 20 kHz. Phase noise spikes around 10 to 40 kHz can be attributed to some digital circuits on the DDS board that are communicating with the controlling computer. In a careful design these would be reduced or eliminated.

How the PLL works

A phase (and frequency) locked loop two primary components. One is an oscillator with voltage control (a VCO) operating at the desired frequency. The other is a lock circuit that compares the VCO frequency to the standard, and produces a positive output when the difference goes one way and a negative when it goes the other. When the frequencies are identical, the comparator produces a voltage proportional to the phase difference. This locks the VCO to the standard in both frequency and phase.

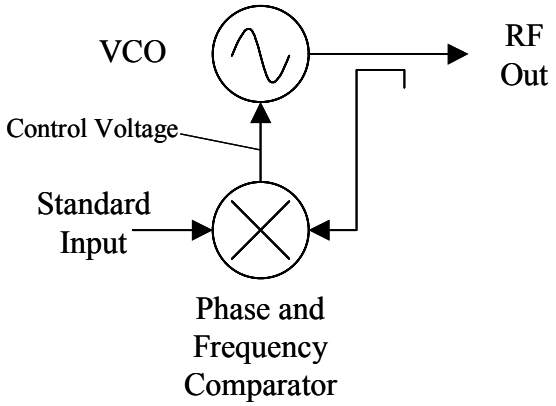

Figure 9. Simple Phase Lock Loop

Of course, a simple PLL will only lock a VCO to the frequency of the standard. To get an output at some different (e.g. higher) frequency than the standard, either the VCO frequency, or the standard frequency, or both must be divided or multiplied so that the comparison can be done between two sources at the same frequency. For instance, if the desired frequency were 110 MHz and the standard were 10 MHz, then the VCO frequency would have to be divided by 11, or the standard multiplied by 11 prior to comparison.

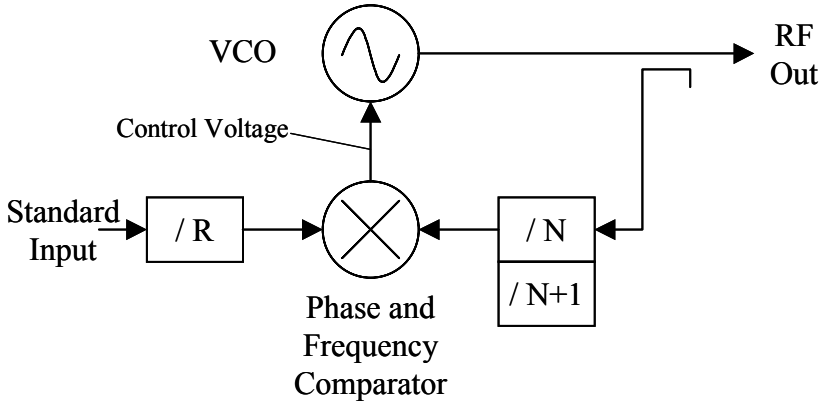

Figure 10. Fractional N Phase Lock Loop

The fractional N PLL consists of a conventional PLL and a fractional standard based lock generator. The fractional circuit divides the standard (called the "clock") by some number. The prescalar divides the VCO frequency by some other (usually larger) number before comparison. There are practical limits regarding the divisors, and they must be integers because the circuits are actually counting signal zero crossings. For example, we might have a 10 MHz standard, and a desirable VCO at 94.666666…MHz (when multiplied by 108 this can be used as a 10GHz LO of 10,224 MHz). By dividing the standard of 10 MHz by 150 we get 66.666…kHz and by dividing the VCO of 94.666…MHz by 1,420 we also get 66.666…kHz. These two frequencies are compared and a phase lock signal is generated to keep the VCO at the desired frequency (based on the standard).

A trick is employed to produce the /N number conveniently. The /N circuit is actually a prescalar, which is a circuit that divides by a fixed number, such as 32. Because we wish to divide by numbers that are not just multiples of 32, the "dual modulus prescalar" divides by both 32 and 33. Another circuit then tells the dual prescalar how many of each divisor to use in order to

accomplish the correct division. In our case 32 counts of 32 = 1024 and 12 counts of 33 = 396. Together, 1024 + 396 = 1420. That special circuit, the /R and the phase and frequency comparator are often integrated into one circuit called a "PLL Frequency Synthesizer with Dual Modulus Control" such as the Motorola MC145158. As you can see, the synthesizer circuit needs only to operate at some fraction of the final frequency, and the only component outside of the VCO that operates at a high frequency is the prescalar.

To keep the locking circuit stable, a comparison frequency must be chosen that is at most 1/50 of the standard (200 kHz for a 10 MHz standard), and to keep the locking fast enough to be effective the fractional frequency should be above about 50 kHz. These bounds put constraints on the frequencies that can be locked. At 24 GHz, for instance, there may be schemes wherein the closest frequency achievable for a 100 MHz region crystal VCO is 1 kHz away from the desired frequency. This will result in an error of 240 kHz at 24 GHz. If the situation can tolerate this, as might be with a tunable IF radio, then all is well. However, if the oscillator is to be used as a beacon, it might not.

PLL Terminology

These were adapted from [1] and other sources.

Free-running Frequency, also usually the VCO center frequency is the frequency at which the loop VCO operates when not locked to an input signal. Usually, the VCO control voltage is set to the midpoint of the control range and all circuit reactances are set to provide the greatest stability.

Lock Range is the range of frequencies over which the loop will remain in lock. Normally this is centered at the free running frequency.

Capture Range (also called Lock-in Range) is the range of frequencies over which the PLL will obtain a lock when it is unlocked.

Lock-up Time is the time required for an unlocked loop to lock. This depends on both the locking bandwidth and damping.

Phase Comparator Conversion Gain $K_{d,}$ is the relationship between the phase comparators output voltage to the phase difference.

VCO Conversion Gain $K_{0,}$ is the relation between the input control voltage and the frequency shift. It is measured in radians per second per volt (rad/sec/V).

Loop Gain $K_{V,}$ is the product of K_d and $K_{0,}$ and the low pass filter DC gain. K_V is evaluated at the appropriate control voltage level and at the appropriate center frequency. K_V has units of 1/sec.

Closed-Loop Gain (CLG) is the open loop gain K_V divided by $1+ K_{V.}$

Damping Factor is the standard damping constant of a second order feedback system. It refers to the ability of the loop to respond quickly to an input frequency step without excessive overshoot.

The Low Pass Filter

Of the many design elements of a PLL, the most important are around the low-pass filter. This element attenuates the undesirable high frequency components at the output of the phase

comparator, rejects sources of interference and therefore reduces locking bandwidth phase noise, slows the rate of frequency change resulting in quick recapture during a transient, and provides damping.

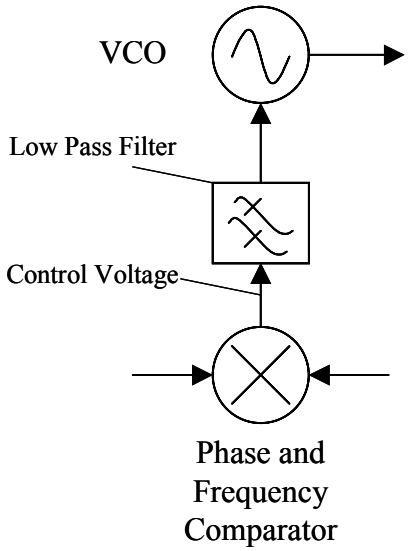

Figure 11. A low pass filter is incorporated in all PLLs between the comparator and the VCO.

When the low pass filter bandwidth is decreased (its time constant made larger) 1) the capture time (when a disturbance occurs) becomes longer, 2) the pull-in time increases, 3) the capture frequency range decreases, 4) the interference rejection properties improve, and 4) transients become under damped.

Under damping reduces stability, and if taken too far will result in a PLL that sweeps instead of locking. Under damping can be fixed by increasing the frequency of the filter (decreasing its time constant) or decreasing the gain of the loop.

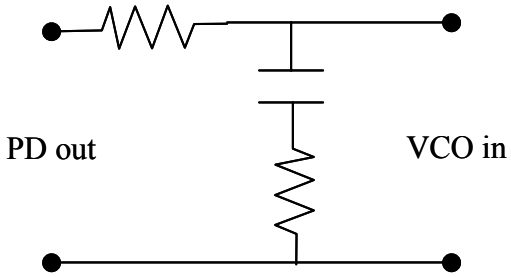

Figure 12. This type low pass filter is most often selected because it provides good stability.

A first-order "lag-lead" filter has the property that its frequency response is dependent on the damping. One has to be sure the keep the loop gain K_V low enough in combination with the filter to prevent under damping. To increase damping, either reduce the gain or decrease the time constant of the filter. As the value of the resistor in series with the capacitor gets lower, the circuit degenerates to a simple and effective first-order filter, but it may have less inherent stability than

one with a higher resistance value. If the resistance in the PD output leg gets high, thermal noise begins to appear in the phase noise of the VCO.

Most manufacturers of integrated PLL chips will give functions for optimizing the loop filter. Sometimes they are correct, and other times you may improve upon them by adjusting the parameters to suit your needs. Designers are asked to use caution when trying to apply second-order filters because of the likely poor stability. However, some second order filters are used successfully if designed with stability in mind.

PLL Noise

As in all PLLs, the locking bandwidth determines which oscillator becomes the dominant source of noise. Within the locking bandwidth, the phase noise of locking standard appears in the output. Also, any amplitude noise in the VCO control circuit becomes phase noise in the output. Therefore, care must be taken in the design of the locking output circuits and the voltage control inputs. Three common sources of noise in this circuit are the inherent noise in the control signal amplifier, excessively high source impedance of the loop bandwidth filter, and power supply noise both for the locking circuits and power for the VCO. Outside of the locking bandwidth, the output phase noise is that of the VCO. Additionally, any amplitude noise generated by the VCO or any subsequent amplifiers is apparent in the output.

A Low Noise PLL design

An article appeared in Nov 1999 QEX [3] “Stable, Low-Noise Crystal Oscillator for Microwave and Millimeter-Wave Transverters” by KD6OZH. This title sure turned me on! In late 2000 I designed a two-board set and made 15 of them. I made some modifications for convenience, although perhaps in so doing I may have compromised some circuit characteristics. In any case, I constructed 3 for myself, and made up kits for a dozen other amateurs. Some of them are in use today. I have one set in my 47 GHz radio. I suggest that you obtain and read that article if you are serious about locking a crystal oscillator for subsequent multiplication up to the EHF bands. (By the way, I have no more boards or kits.)

His design uses a straightforward phase comparator circuit and a PIC processor to load the values in on power-up. KD6OZH put significant effort into analysis and design of a voltage controlled crystal oscillator so that he could achieve very low noise –both phase noise and amplitude noise. No frequency agility here – to change frequencies one has to both re-program the PIC, and also order and wait for a crystal! Some fiddling with the oscillator was needed, depending on the crystal characteristics, much as KD6OZH had described in the article.

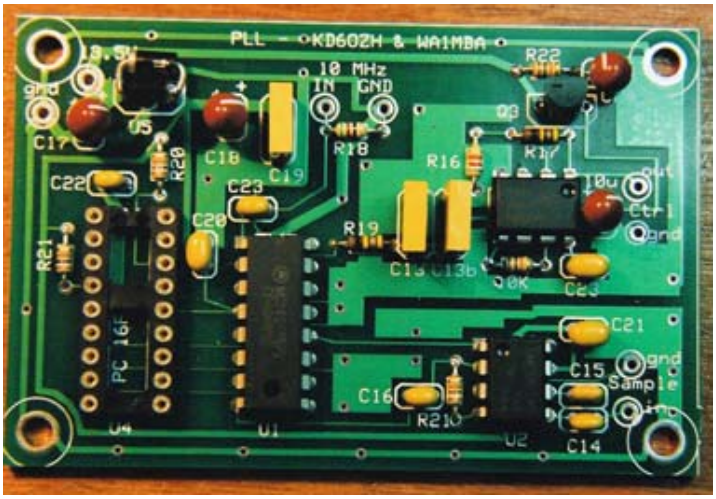

Figure 13. Here is an assembled PLL board. The PIC processor that holds the R, N and N+1 parameters is not plugged into its socket.

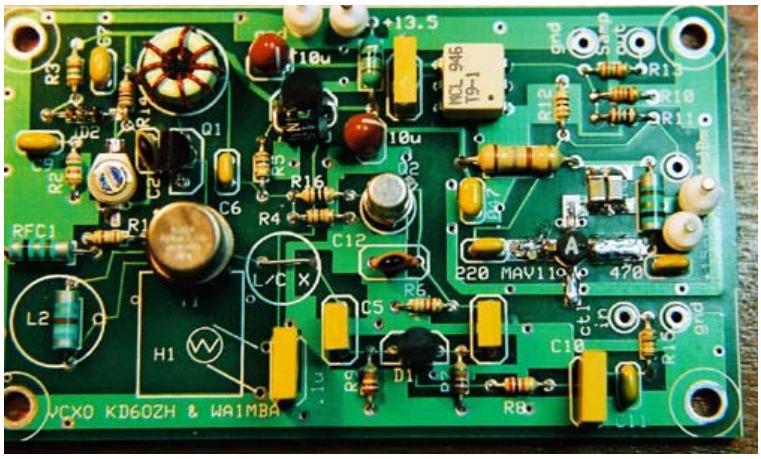

Figure 14. Here is the VCXO circuit board. This one has a crystal in a T0-5 can towards the left side of the board. Space was left for a standard HC9 can as well, and a PTC heater.

KD6OZH was right! Although the unit requires some careful adjustment of the reactances so that it becomes tunable across the expected control voltage range without loosing oscillation, and although, just like with any crystal oscillator, the natural frequency does drift with age so that it requires some retuning (like when you go out on the cumulative contest weekend), it has really low noise, and is super stable when locked.

Below are some plots of the performance of a 74.527027 MHz PLL VCXO implemented with KD6OZH's design. This frequency would be multiplied by 171 to produce an LO for 24,192 MHz with a 1296 IF. The exact frequency puts the IF off by an additional 43 kHz because the exact frequency could not be obtained by divisors that result in fractional frequencies within the stable region. The phase comparison (fractional or "/R") frequency is 67,567.567… Hz. I set up the instrumentation to examine this frequency, but quite honestly, could not find it.

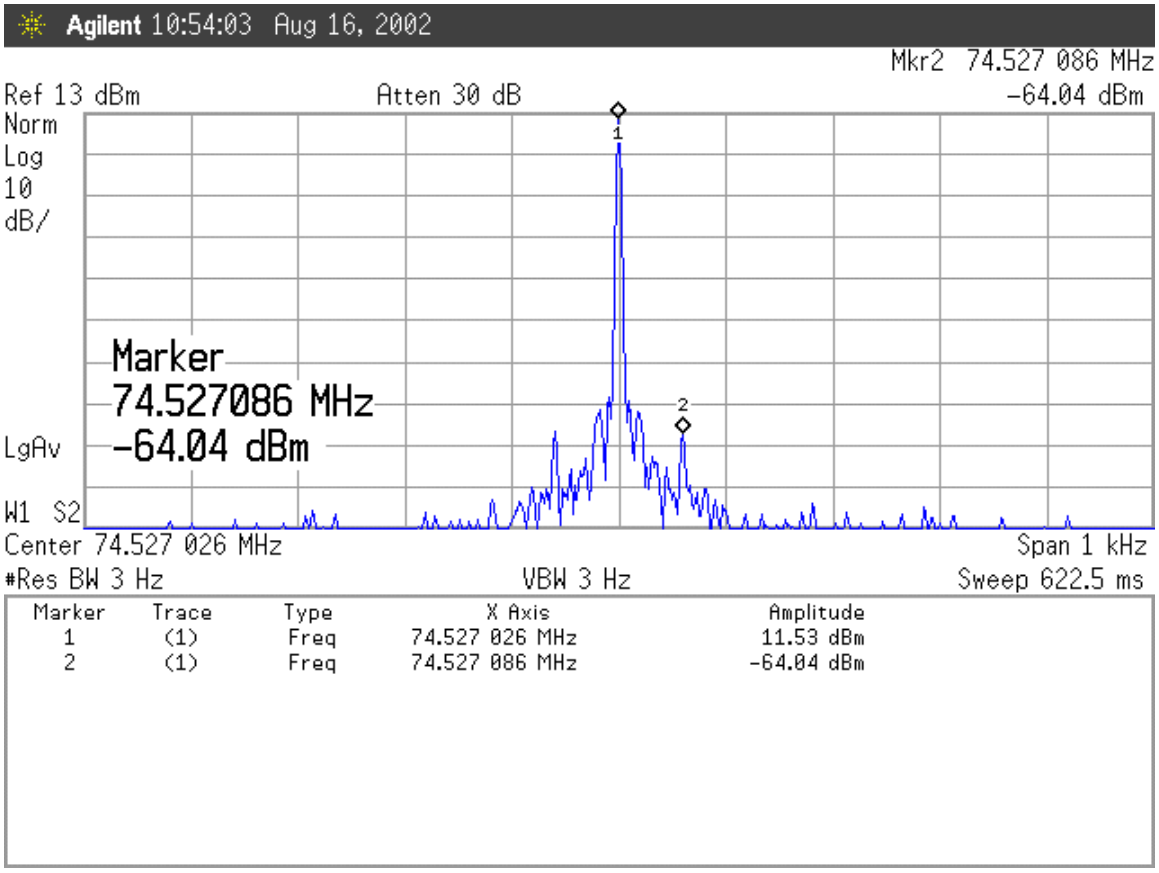

Figure 15. In this plot, the frequency span very close to the output of a fractional N PLL is shown. This is my implementation of the KD6OZH PLL described, where a crystal oscillator is voltage controlled as the VCO. Marker 2 shows some energy at –75.5 dBc that is 60 Hz away from the carrier. This is due to some hum in the laboratory power supply, and possibly some pickup in the power wiring. All of the noise in the bandwidth shown appears in the audio spectrum at 75 MHz. Components that are FM (such as phase noise) will be multiplied in offset, and increase in amplitude as this carrier is multiplied to become a microwave LO. Components that are AM, like this hum, will likely be spectrally preserved as it is multiplied. This very close to carrier noise is remarkably low. KD6OZH used a slightly better power filter for the VCXO than I did, and such filtering may reduce this residual noise.

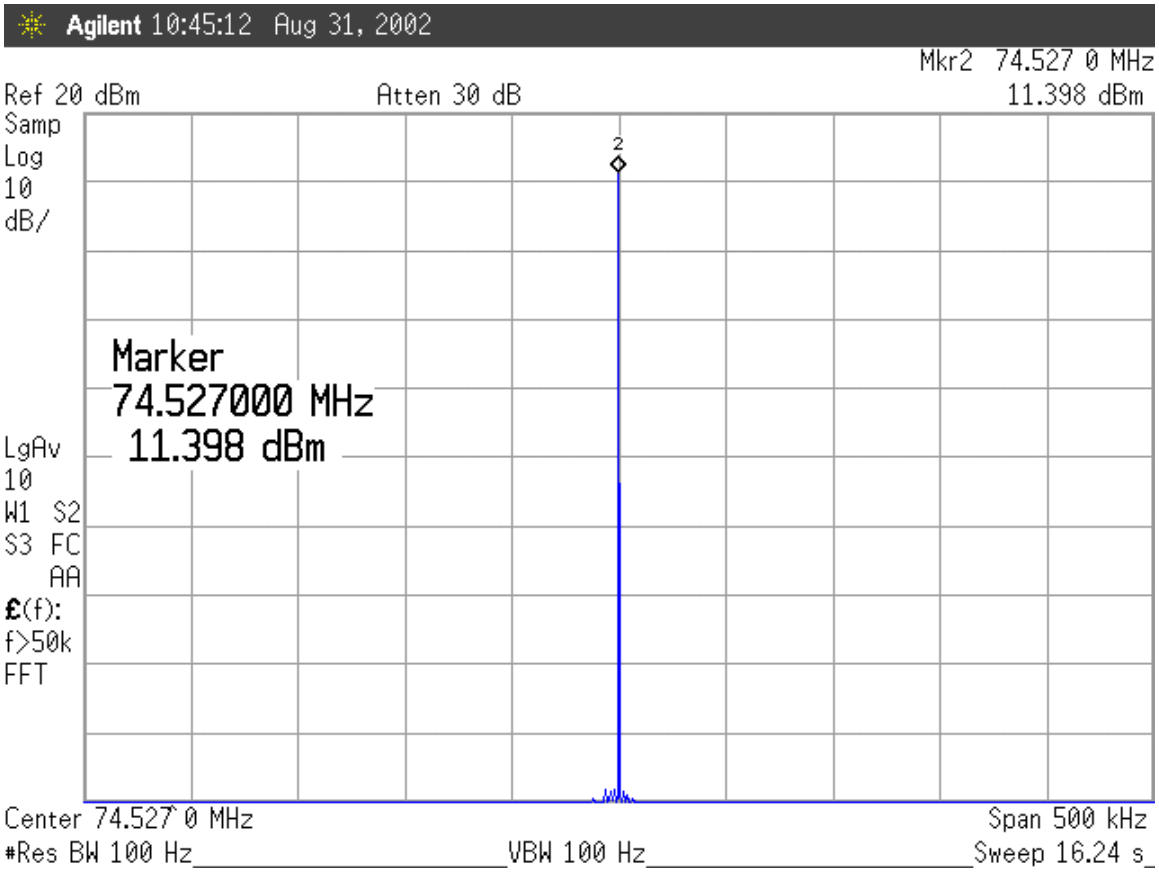

Figure 16. This plot, with similar scale to one of the DDS plots, shows the noise 500 KHz around the output of the PLL VCXO. The spectrum analyzer does not always show the exact frequency correctly. In this case it rounded off the last 27 Hz to 00. There is no apparent noise down to the limit of the instrument, except –90 dBc very close in to the carrier. Because the resolution bandwidth is set to 100 Hz, there is not enough resolution in this analysis to show the very close noise that is apparent in the previous figure.

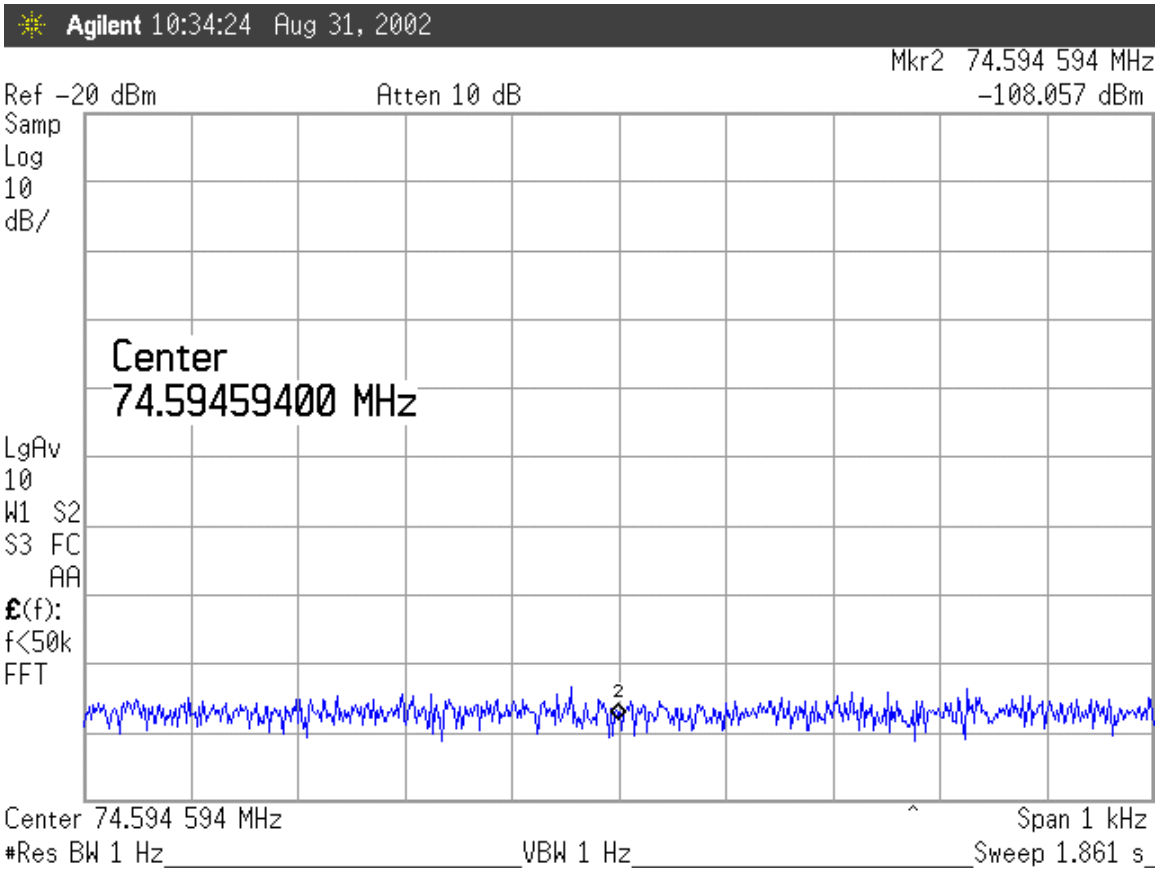

Figure 17. In this plot, by moving away from the carrier we can use the instrument to look down to the noise floor of –119.5 dBc where the /R frequency, offset from the carrier by 67.567 kHz to 74.594594…MHz might be examined. Although the instrument has another 20 dB of dynamic range before its (amplitude) noise floor (when no carrier is present), the phase noise of the instrument is approximately –119 dBc at this offset. So, we are probably looking at the instrument phase noise rather than any signals from the oscillator! Clearly, excellent low noise VCXO design and good practice of isolating the phase comparator from the VCXO can make a difference.

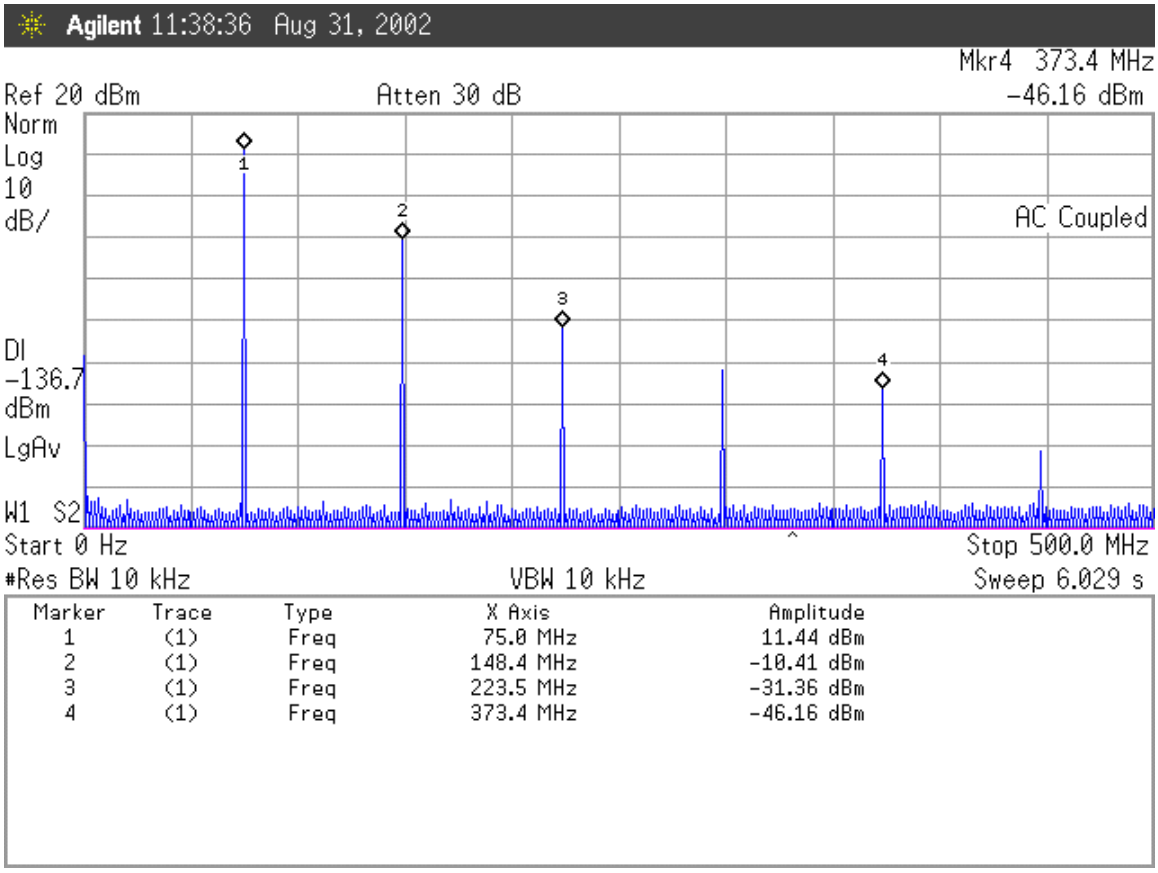

Figure 18. One compromise I may have made in my fiddling with the original design is the harmonic content of the output. Here we can see that the second harmonic is quite strong at only –21 dBc, and the third at another 21 dB down from the second. This may be no problem at all for applications that will use low multiples of this oscillator in the multiplier chain. It can cause havoc, however, with the PLL sampler, and so in some cases a simple low pass filter was required to keep the second harmonic out of the prescalar.

To be fair, DDS sources have outputs at aliases of the clock frequency. They are not harmonics, but are the mixing products of the clock and the output frequency. They also must be filtered, but can be done with a simple low pass filter with a knee at a little below one half of the clock frequency. Because there is no feedback path dependency in a DDS, this filtering is not required for operation of the DDS. Some designs, in order to get higher frequencies, will use these aliases to advantage, just as some PLL designs will use filter and use a harmonic for further signal development in subsequent circuits.

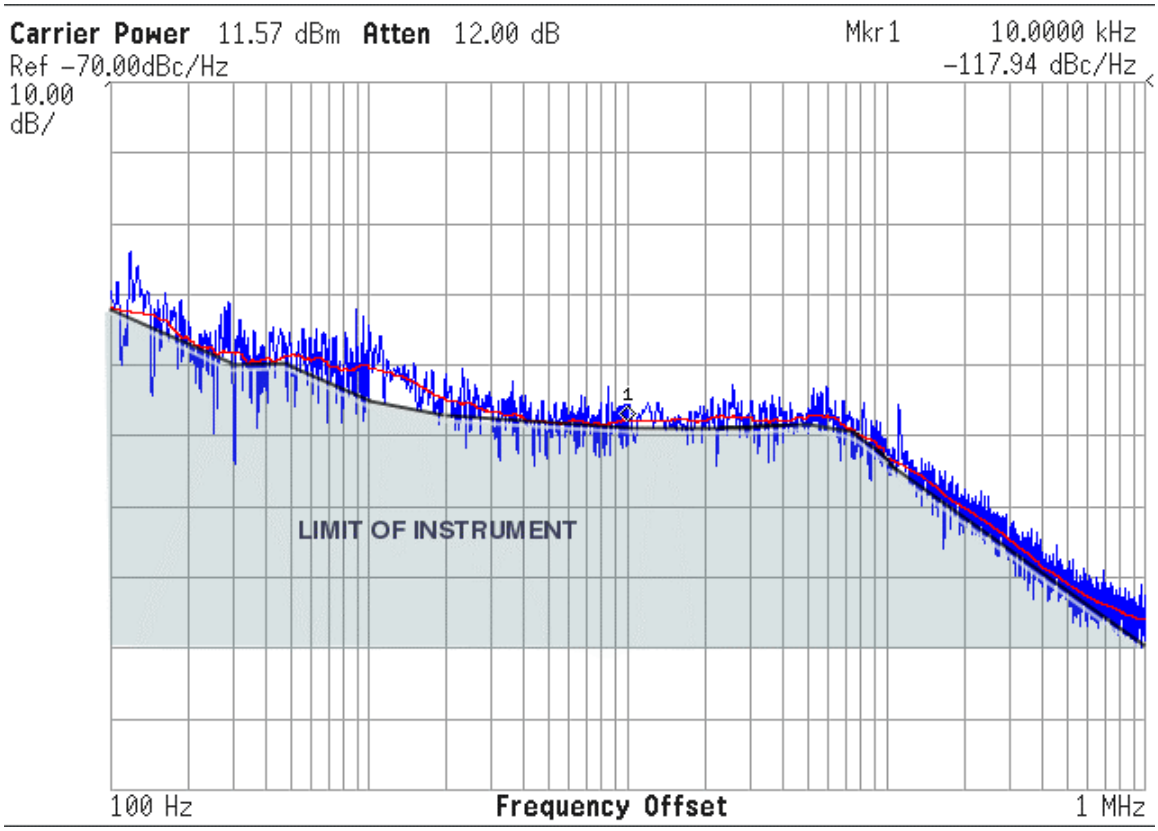

Figure 19. In this phase noise plot of the PLL, except for a very short span around 500 Hz to 2 kHz, the oscillator is performing at or better than the measurement instrument. KD6OZH reported lower phase noise than this plot indicates by as much as 20 dB. I had no way of measuring phase noise below the limits of this instrument.

There is no evidence in this example of the fractional signal (67.5 kHz) that often pollutes this type of synthesizer. Other fractional PLLs often produce a small but noticeable spike at some point in the phase noise spectrum. It may exist at a level below the ability of this instrument to measure.

Example idea for a locked EHF LO

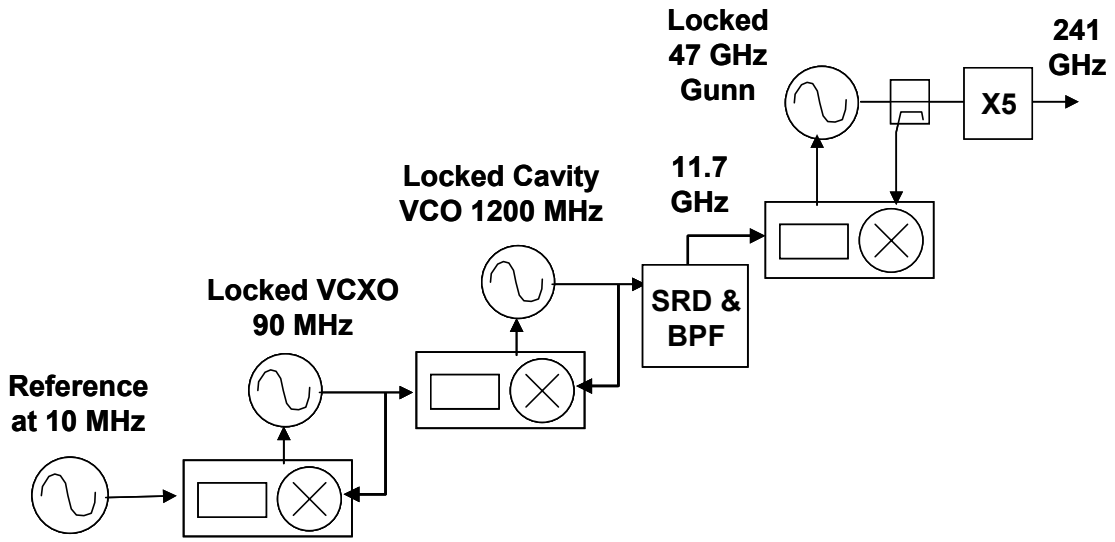

Figure 20. Here we look at a possible way to lock a signal at 241 GHz. The boxes each represent a complete phase comparator, data loading circuit and band pass filter. The reference standard at 10 MHz locks a VCXO, like the one described at about 90 MHz. This 90 MHz becomes the standard for a PLL "Brick" that is actually a locked cavity VCO running around 1200 MHz and driving a Step Recovery Diode multiplier. The output is 11.7 GHz, which can then be used to lock a Gunn oscillator at 47 GHz, which is multiplied by 5 to produce 241 GHz. Of course, many schemes are possible, and some may require fewer phase lock loops. One should consider means to produce LO for RX and TX frequencies with the fewest circuits. The scheme presented here has the possible benefit of also producing ample RF at 47 GHz.

References:

[1] Philips application note AN178.

[2] Motorola MC 145151 family PLL circuits product description.

{3} "Stable, Low-Noise Crystal Oscillator for Microwave and Millimeter-Wave Transverters", John Stephensen, KD6OZH, Nov 1999 QEX