Short Term Stability Measurements of Several 10MHz Reference Sources

Andy Talbot G4JNT
November 2013

Introduction

I am fortunate in having an HP5061A Caesium Beam frequency standard that can generate a 5MHz reference signal accurate to something like a few parts in 10^{-13} when properly set up. It is difficult to be precise about the accuracy of Cs standards, because the Caesium resonance actually forms the definition of the second. However, as I have not calibrated out the Zeeman lines for local magnetic field, it is probably only within a few parts in 10^{-12} and *could* even be as bad as 10^{-11} for all I know ! Without access to at least two more Cs sources it is impossible to check However, the HP5061A is still a couple of orders of magnitude better than the majority of local sources we frequently use as references and better than most people can manage.

In order to see small short term frequency shifts of the order of a few parts in 10^{-10}, the 10MHz needs to be multiplied up so that short term instability appears as a frequency shift of a few Hz. I have several microwave synthesiser modules that can take in a suitable reference signal and generate a precise microwave frequency that is an exact known multiple of the reference input. A suitable common frequency for two such modules was chosen at 2.5GHz.

Test Setup

A reference 2500.000000 MHz signal was generated using an LTC6946 Integer-N synthesizer module *[1]* driven from the 5MHz output from the HP5061A. The 5MHz also goes to a frequency doubler to provide a master 10MHz reference used in the first plot.

An LMX2541 Fractional-N synthesizer *[2]* was driven at 10MHz by the various reference sources under test. This module was set to generate 2500.000900MHz when driven with 10MHz exactly, ie. Exactly 900Hz higher than the reference

The outputs from the two synthesizers were connected to the RF and LO ports of a Minicircuits ZEM-4300 0.3 to 4.3GHz mixer. The IF output nominally at 900Hz was taken to a PC soundcard line input where *SPECTRAN* audio analysis software was used with a resolution bandwidth of 0.18Hz (sometimes 0.37Hz) to examine the short term stability of the resulting 900Hz tone.

The results are shown on a waterfall display, with a typical vertical duration of three minutes. Each red tick on the vertical axis corresponds to one minute.

A Caveat : *This test is not very kind to fractional-N synthesisers, and shows then in their worst light.*

Fractional-N synthesisers when used to generate frequencies close to integer multiples of their references are not terribly clean. They can, and do, generate relatively high levels of sidebands at multiples of the fractional step frequency. 900Hz is very close to the 10MHz reference frequency of the PLL and is not a combination that would normally be preferred for a practical RF situation where a clean signal is desired. Here, however, a particularly clean signal is not needed; we only want a stable main one to measure small shifts.

The Fract-N unit used here is set for a step size of 10Hz, and on the plots that follow, 10Hz sidebands can often be seen. For the purposes of this test, these sidebands *in no way detract from measurement accuracy* – they just don't look very nice. In all cases, the only trace that matters is the central main one.

Also, at all times there appears a faint line at exactly 900Hz that bears no relationship to the test input and is there even with no test input signal to the mixer. It is arriving through the reference synthesiser alone (the LTC6946) , and cannot be explained; it just should not be there! There is no mechanism in the synthesiser that could generate it. However, it does provide a convenient marker to see and measure frequency errors against.

Ground loops, unshielded synthesizers modules and leakage are all present.

First Test, Audio Calibration

The soundcard clock itself is a source of error in the measurement of the audio frequency. A sampling rate of 48kHz was used, and it was already known there is no inherent error in the soundcard when this sampling rate is selected, other than that of the crystal oscillator itself which it probably out by several parts, or tens of parts, per million.

The LMX2541 was driven from the doubled 5MHz signal that supplied the reference synthesiser, so both RF signals 900Hz apart are now coherent and the generated audio should be exactly 900Hz. Figure 1 shows the resulting tone. The 10Hz sidebands are very evident, as are other spurious artefacts due to unshielded RF modules and interference and the use of a coherent reference, but the main carrier is clearly where it should be at 900Hz. Note that the indicated frequency is shown as exactly 900.00Hz. Since the resolution of the FFT is 0.18Hz it is perhaps somewhat fortuitous the indicated value is precisely this - but such things do happen, Murphy's Law does fail occasionally.

The sidebands at plus/minus 4.5Hz I can't account for – but with a completely unshielded test setup are probably some locally generated artefact, or are in the LMX2541 somewhere, or interference leakage, or a ground loop, or… or… or…

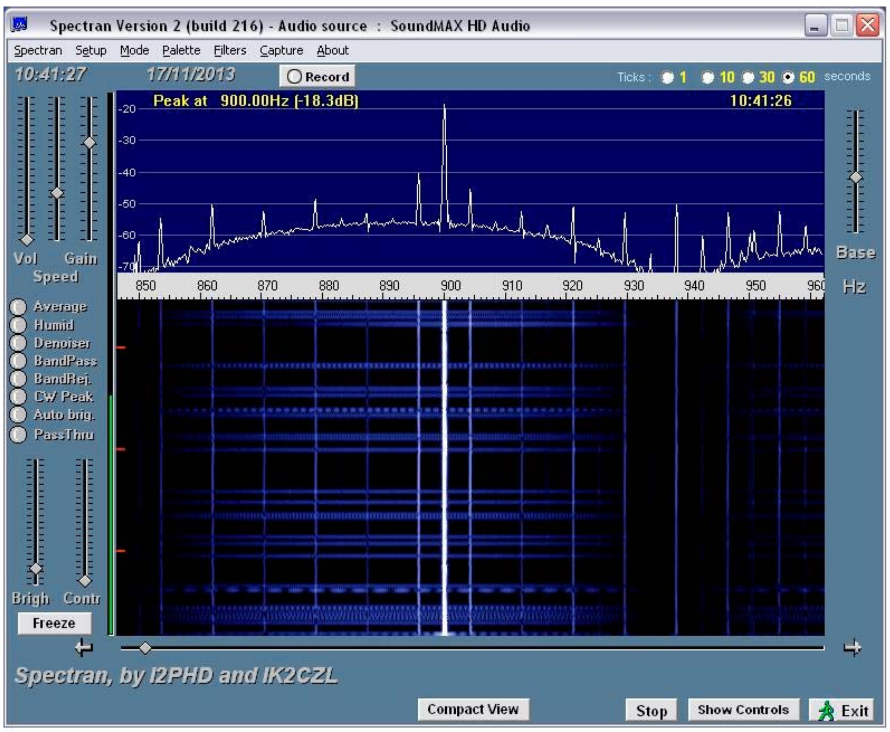

Figure 1 *Coherent Reference Calibration (a high level of 10Hz sidebands due to the Fractional-N Syhthsizer are clearly visible, as well as the 4.5Hz ones unaccounted for)*

Fast Warm up Oven Controlled Oscillator

This OCXO was taken from a defunct Rhode and Schwarz Synthesizer of mid 1980's vintage. It has an oven status output that shows lockup is typically achieved after about 5 minutes, and it then needs another 15 minutes to stabilise properly. This plot was generated after it had been on for 30 minutes.

The OCXO is running 31Hz low corresponding to $12.4*10^{-9}$. While this can be tweaked out with care, it is typical of the setting and retrace capability of such a single ovenned oscillator. More importantly, note how after just 30 minutes of warm up, the frequency is staying stable to within a Hz or two over several minutes. An OCXO of similar capability is used as the reference for the GB3SCK 24GHz beacon where short term frequency shifts of only a few Hz are seen over several minutes, but whose absolute frequency is only accurate to a few hundred hertz (at 24GHz)

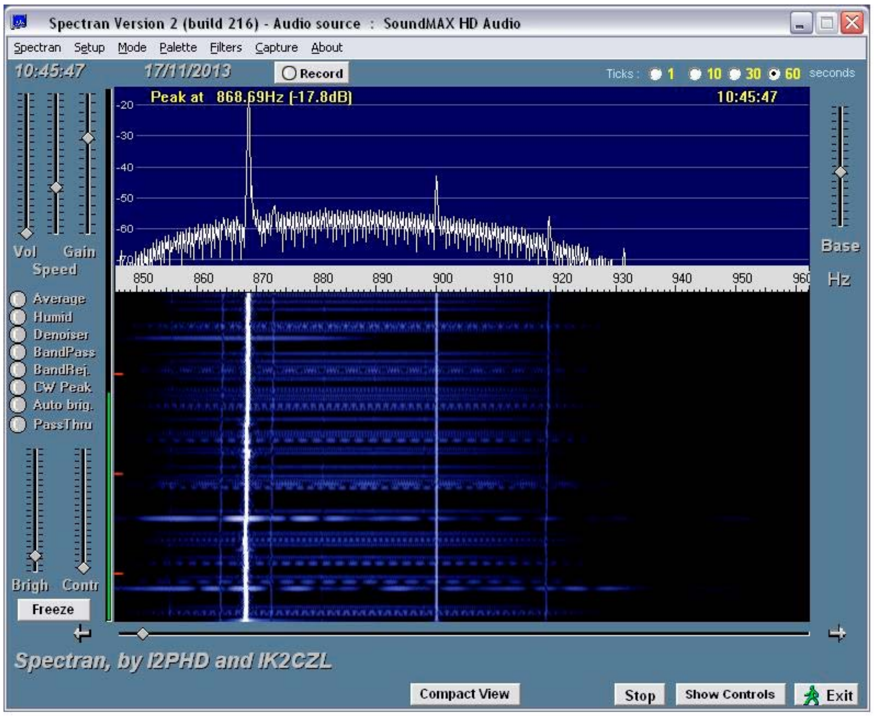

Figure 2 *Fast Warm-up OCXO*

Temex / SpectraTime LPFRS01 Rubidium Source

This unit is advertised as having fast warm up, and is quoted as being within specification of a few parts in 10^{-10} eight minutes after switch on. The first plot in Figure 3 is after 9 minutes: Figure 4 shows the plot after it has been running for 20 minutes.

Absolute frequency is in the region of 0.25Hz high, but with a frequency measurement resolution of 0.18Hz (FFT bin size) the exact value of the error could easily be 50% either side of this. At 2.5GHz this corresponds to one (or perhaps two) parts in 10^{-10} Typical of a roughly calibrated Rubidium source. On the first plot, started only one minute after the specified warm up time, a slight shift of perhaps 0.5Hz at the beginning (bottom of the screen) is visible, as well as a bit of instability a minute or two later. The second plot after 20 minutes shows the typical result, with no short term wobble visible at all.

The same 4.5Hz sidebands seen in the reference plot also appear to be present here.

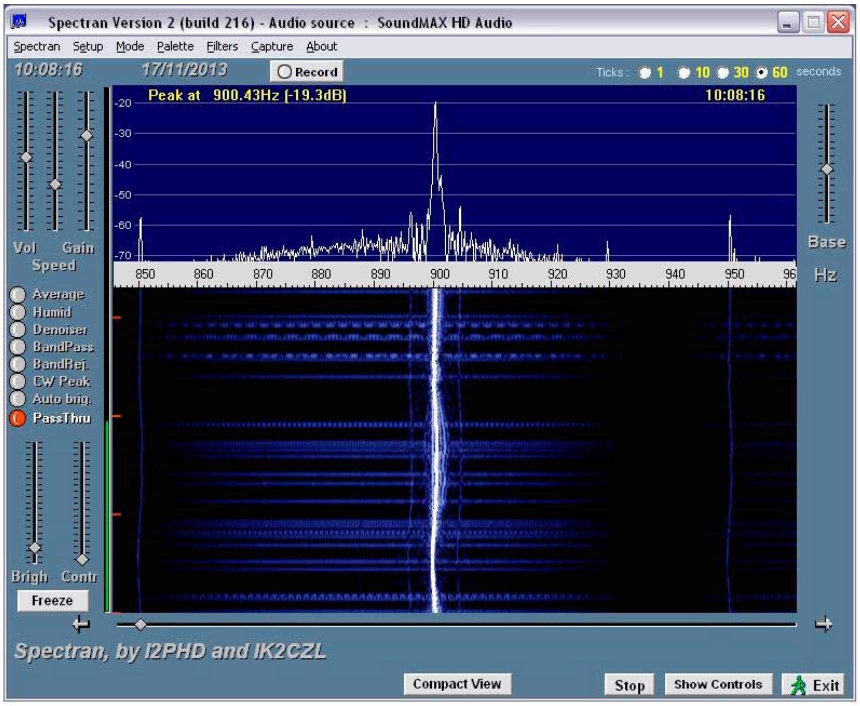

Figure 3 *LPFRS01 Rb Source, 9 minutes after Turn On*

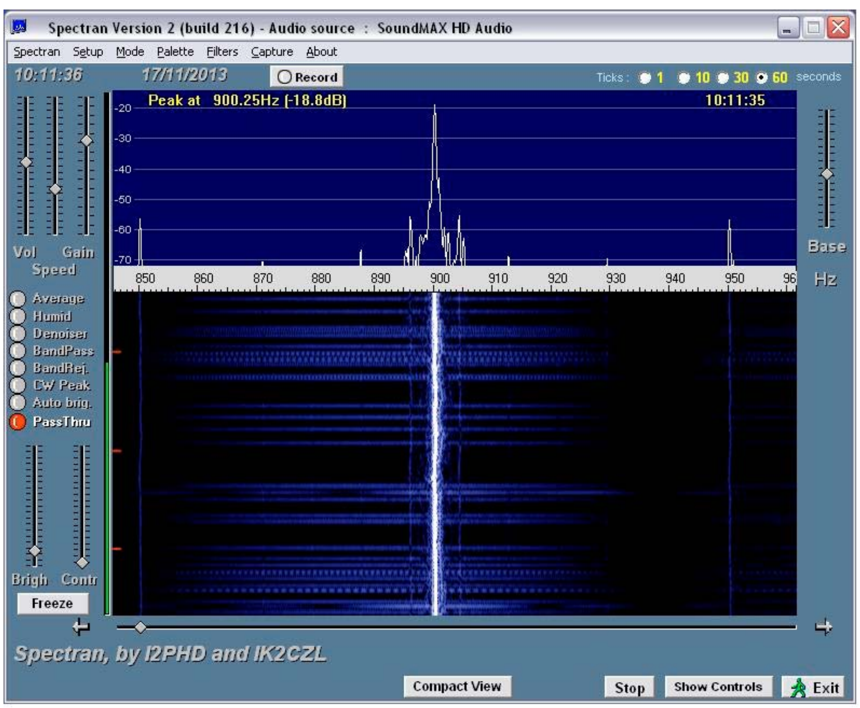

Figure 4 *LPFRS01 Rb Source after 20 minutes*

GPS Disciplined Oscillators

Three GPSDOs were tested:

1) A commercially produced unit made by Connor Winfield and aimed at the telecomms market where the primary requirement is timing and long term accuracy.

2) A “Cheap and Simple” GPSDO based around a Jupiter-T GPS receiver with 10kHz output which locks a small TCXO module used as the master oscillator. This is nearly identical to that shown in *[3]* The loop time constant is short – only tens of seconds - so it faithfully tracks the output of the GPS receiver itself, providing virtually no smoothing. All artefacts due to GPS receiver tracking and constellation changes will be visible

3) A VE2ZAZ GPSDO. This is a frequency locked loop, unlike the other two which are phase locked. The ‘ZAZ unit counts frequency using a conventional 16second gating period, then averages / sums N successive values, sending periodic corrections to the OCXO used as the reference oscillator. The integration time in use here is of the order of 6 minutes, Electrical design of the ZAZ GPSDO is not ideal, the 10MHz signal distribution is integrated onto the same PCB as the processor, making it prone to spurious pickup. When multiplied to 2.5GHz these spurii frequently appear in the output spectra

Figure 5 shows the **Connor Winfield unit**. Although the mean frequency is spot on – centred on 900Hz - fast variations of up to +/- 5Hz are common, peaking nearly 10hz at times. A mean value looks to be around 2 to 3Hz RMS. Figure 6 shows the same unit with the trace expanded (0.37Hz resolution, 1 minute top to bottom) and shows a particularly bad glitch of around 15Hz. The rate of change of frequency after any glitches, which is mostly set by the internal loop constants, seems to be in the order of 3Hz / second

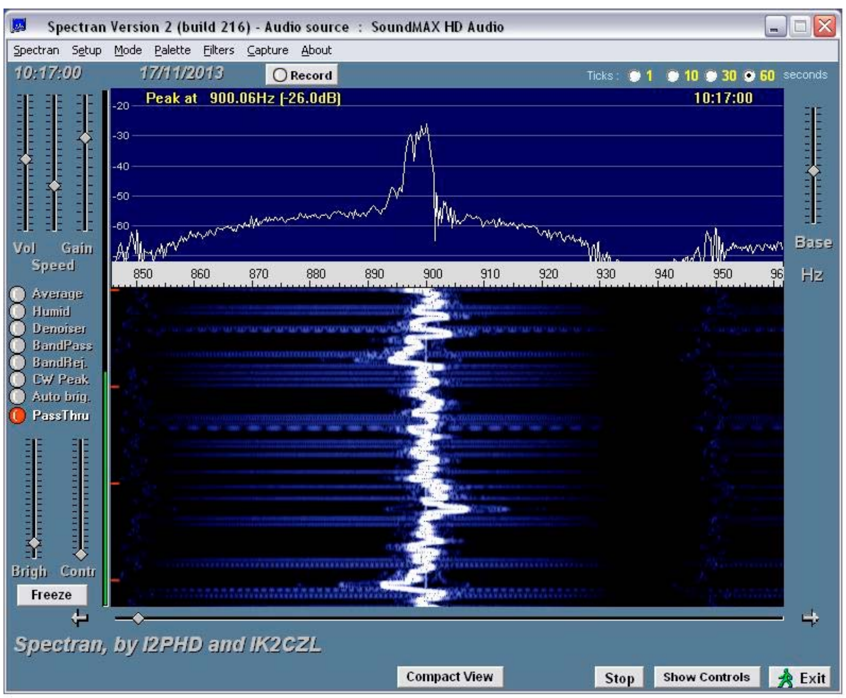

Figure 5 *Connor Winfield GPSDO*

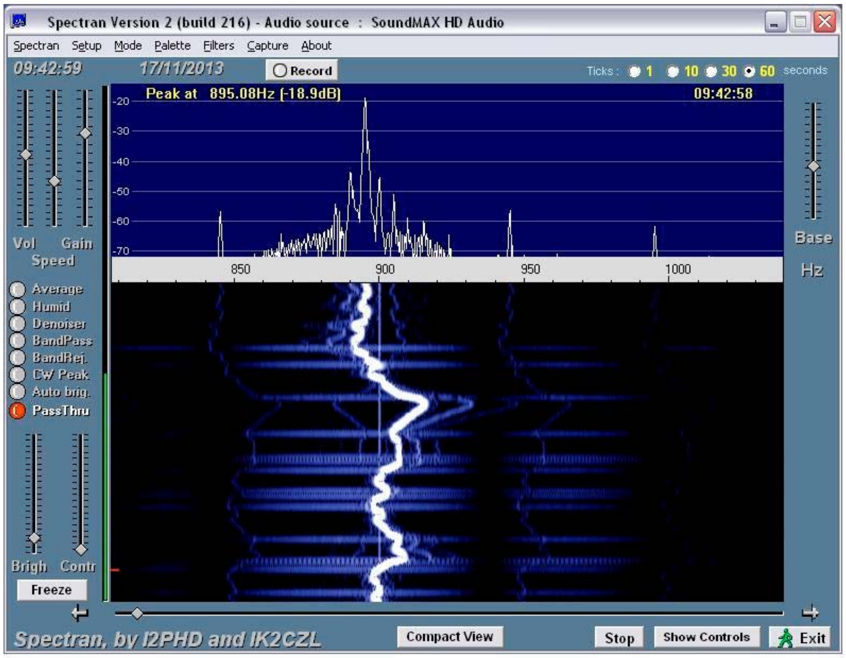

Figure 6 *Connor Winfield GPSDO, expanded trace showing a particulalrly bad GPS induced glitch*

Figure 7 shows the Cheap and Simple GPSDO based around a Jupiter-T GPS receiver, locked at 10kHz with a fast loop time constant and some noticeable glitches Figure 8 shows this on an expanded scale; slightly lower levels of glitches appear to be present.

The loop time constant is faster than that on the Connor Winfield unit, so rate of change of frequency is faster.

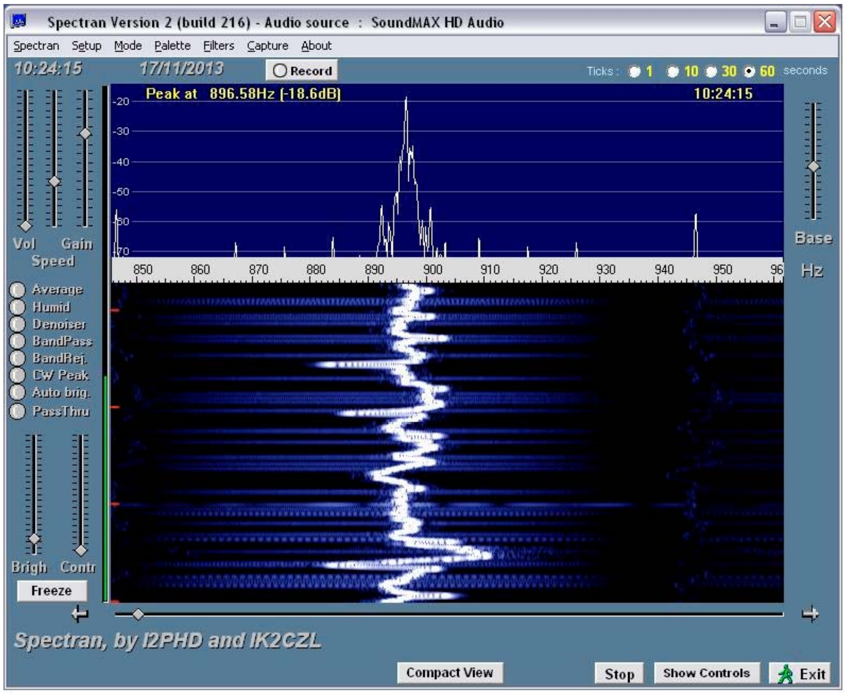

Figure 7 *Cheap and Simple GPSDO*

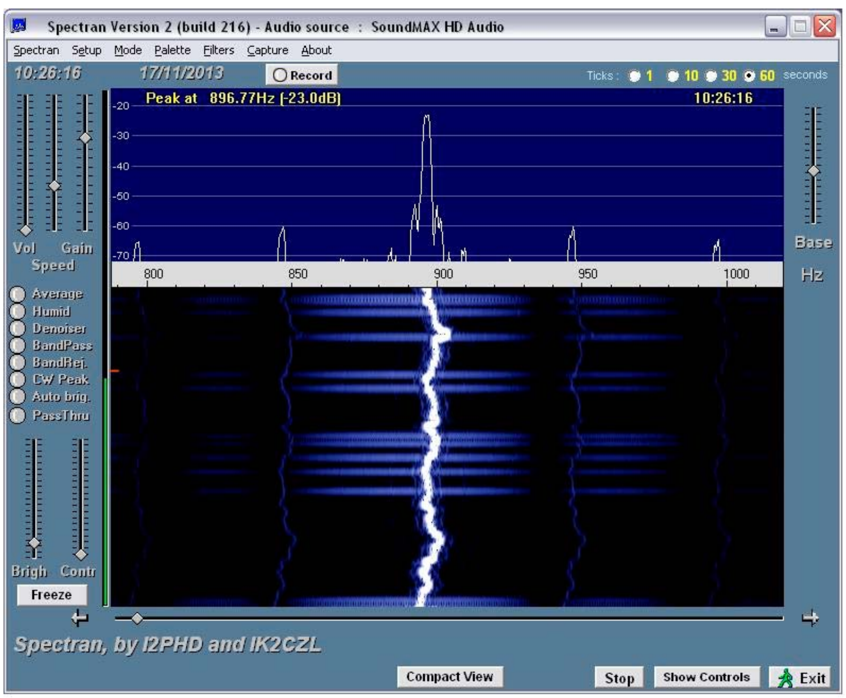

Figure 8 *Cheap and Simple GPSDO, expanded scale*

Figure 9 shows the VE2ZAZ unit. Unlike the previous two GPSDOs this is a frequency locked loop and can be expected to have a continuous; albeit slowly corrected, frequency error. A constant offset of -2.7Hz can be seen, corresponding to around 10^{-9} With an averaging period of 6 minutes, and considering the way the

VE2ZAZ issues update commands to correct frequency, this error will most likely only be corrected after a couple of hours, by which time external temperature may have changed things anyway.

It is possible I haven't optimised the averaging times and correction values within the GPSDO setup – there are a plethora of user settings to be changed, but my gut feeling is that with a basic single oven OCXO a performance much better than 10^{-9} in the medium term (hours, days) is the best that can probably be obtained

However, note the short term variation which is excellent, and comparable with that from the fast Warm up OCXO above (it uses a similar type of OCXO) and the Rubidium source.

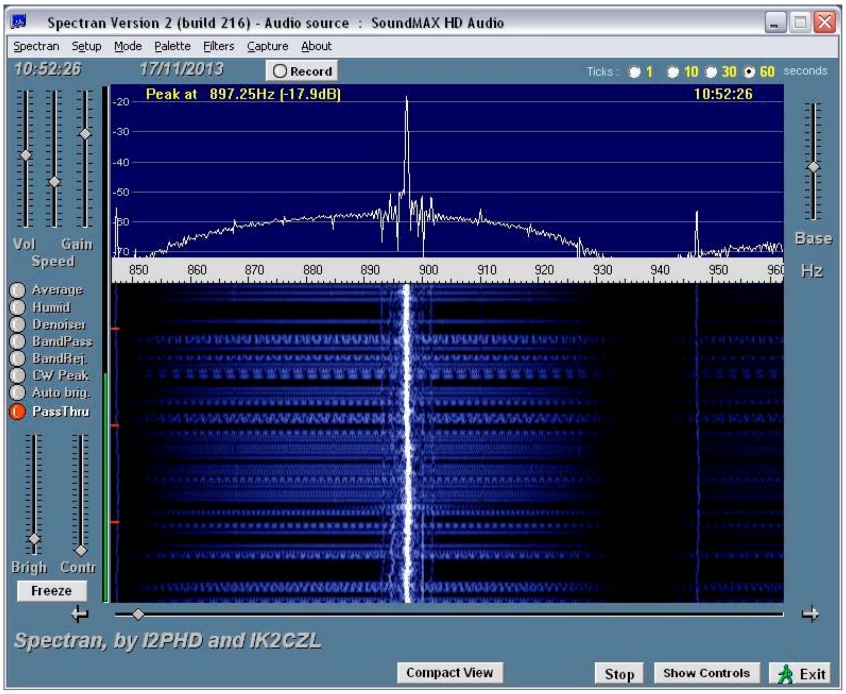

Figure 9 *VE2ZAZ GPS Disciplined Oscillator showing excellent short term frequency stability, but with a fixed offset error that 'will' eventually be corrected – very slowly, perhaps, by the frequency-counter-based correction loop.*

References

[1] LTC6946 Integer-N Synthesiser Module http://www.g4jnt.com/LTC6946_Synth_Module.pdf
[2] LMX2541 Fractional-N Synthesiser Module http://www.g4jnt.com/LMX2541_Synth_Module.pdf
[3] "Probably the Simplest GPSDO" http://www.g4jnt.com/SimpleGPSDO.pdf
"an even simpler one" http://www.g4jnt.com/EvenSimplerSimplestGpsdoPossible.htm