Proceedings of the 2003 IEEE International Frequency Control Symposium and PDA Exhibition
Jointly with the 17th European Frequency and Time Forum

GPS DISCIPLINED OFFSET-FREQUENCY QUARTZ OSCILLATOR

Nigel C. Helsby
Time & Frequency Solutions Ltd, Eastways, Witham, Essex, England

***Abstract* - The Global Positioning System can be used to provide a precise time and frequency reference to control local oscillators. The local oscillator is necessary to provide holdover of time and frequency in the absence of GPS and to average short-term jitter in the GPS signal recovered by the receiver. This paper gives an overview of the hardware design and control strategy for converting the output of a free-running ovened quartz oscillator (OCXO) to the desired standard frequency output by reference to GPS. The method enables the use of a very low aging fifth overtone quartz oscillator, to which final plating of the crystal to trim 'on frequency' is omitted. Sufficient frequency control could not normally be applied to this very stable oscillator by the conventional method of varactor tuning. A compact implementation is achieved with very low spurious frequency content and low phase-noise of the standard frequency output. Performance data will be presented.**
***Keywords* – GPS, Disciplined, 5th Overtone, Aging**

I. INTRODUCTION

There are many applications for frequency standards which demand high accuracy and which are readily satisfied by the use of a chosen grade of quartz oscillator, usually within an oven, that is phase-locked to GPS signals. These oscillators are frequency trimmed by the use of varactor-diode tuning which allows a small amount of frequency control either side of the nominal value. The removal of GPS signal degradation by selective availability (SA) in April 2000 gave a performance boost to these instruments, allowing those using relatively modest oscillators to achieve the same or better performance than their predecessors. However some applications demand a high level of performance during the time that GPS is not available, for whatever reason. During this period the oscillator is held at the last best control value and free-runs until the return of GPS allows new corrections to be calculated. Some manufacturers have used adaptive temperature and ageing compensation during the holdover period [1] to improve performance but we have not found this a good substitute for the use of the best available oscillator that the budget permits. Prior to the offset frequency design we had used Brandywine Communications 3rd overtone 10MHz SC cut quartz oscillators in a miniature frequency and time receiver. No compensation system can take care of frequency jumps that many quartz oscillators exhibit [2] and which we recorded during the testing of the new instrument.

Taking the problem of aging into consideration we found that Brandywine Communications were able to make even higher grade low aging precision ovened quartz oscillators by employing 5th overtone SC cut quartz blanks to which the final plating stage was omitted. We can have any reasonable frequency around 5MHz but must accept a frequency tolerance of nearly ±0.1% which equates to ±4kHz. A narrow band high resolution synthesiser has been described [3] which allows the generation of frequencies around 10MHz with extremely small step size and high spectral purity. However, the task is greatly simplified if the originating standard frequency has a suitably chosen frequency offset which is then removed by subtraction of the offset using a DDS (Direct Digital Synthesiser) to apply this frequency correction. Several advantages ensue: the quartz oscillator can be chosen as 5th overtone (since we no longer intend to apply varactor control to 'pull' the frequency), the oscillator is always in the optimum free-running condition, and the frequency offset can be chosen to ease the problem of subsequent control. The control of oscillators via precision voltages applied to varactor diodes can be degraded by oven current variation causing small shifts in the control voltage which equate to unwanted frequency adjustments especially during the critical periods of reference loss.

Fig 1 shows the architecture of the GPS4A.

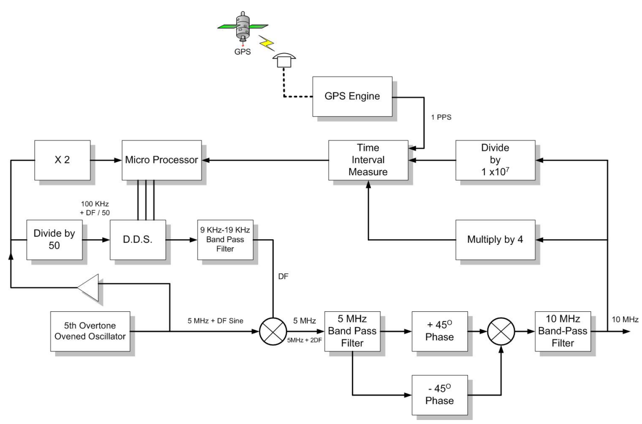

The frequency of the sine wave output of the ovened quartz oscillator is 5MHz plus an offset we have called DF. 5MHz+DF feeds a double balanced mixer signal input. The other input port of the mixer receives the filtered sine wave output of the DDS adjusted to the frequency DF. The mixer outputs an upper and a lower sideband with about 35dB carrier (5MHz+DF) rejection. The wanted lower side-band is selected by a 5MHz crystal filter; the upper sideband 5MHz+2DF and any remaining carrier 5MHz+DF are rejected. The 5MHz sine wave output is doubled to 10MHz, this time in a schottky diode ring double balanced mixer, and filtered again by a 10MHz crystal filter. The filters are implemented using simple low-cost quartz crystals, with damping to reduce microphony. This is the simplest method of single sideband frequency generation, with the addition of a frequency doubler at the output.

In order that the filtering of the unwanted frequencies is effective we have chosen an initial frequency offset with a nominal value of 14kHz. The crystal manufacturer, using the manufacturing process designed to minimise aging, can give us a frequency tolerance of about ±4kHz in the 5.014MHz nominal output frequency of the special 5th overtone SC cut unit.

The frequency out of the DDS is determined by its clock frequency *Fclk* and a 32 bit number *N* written to its registers. The value of *N* is added to an accumulator at each clock update, and the resulting ramp feeds a sine look-up table followed by a DAC that generates discrete steps at each update, following the sine wave form. The stepped waveform contains harmonics but they are relatively easy to filter using an active band-pass filter stage. The DDS output frequency is

$$F = N \times Fclk/2^{32}$$

The resolution of adjustment in terms of one step of the integer *N* is dF/dN which is $Fclk/2^{32}$. Clearly for the smallest frequency step we need to use a low clock frequency, but the lower the clock frequency, the harder it becomes to filter the clock components in the DDS output. In addition, we found that there is a high output from the synthesiser at the clock frequency minus the output frequency which also prevents us setting the clock frequency too low. A good compromise is a clock at about 100kHz, obtained by dividing the nominal 5.014MHz oscillator output by 50. Our approximate resolution in frequency out of the DDS is therefore

$$dF/dN = 100{,}280/2^{32} = 2.33E\text{-}5\ Hz$$

As far as our output is concerned this frequency is subtracted from the output frequency of the quartz oscillator. The minimum frequency step of the frequency corrector is therefore 2.33E-5 / 5E6 which is 4.6E-12.

As discussed, the conventional method of oscillator adjustment is to employ a 16-bit DAC that sets the tuning voltage of varactors. The adjustment range is typically set to ±1.5E-7 for a high grade SC cut oscillator, leading to a minimum frequency step of 3E-7/65535 which is also 4.6E-12. However, the DDS is able to control over a much larger frequency range with the same resolution and remove all of the quartz calibration error, as well as remove the deliberate nominal offset of 14kHz.

One of the interesting results of this method of control is that, because the oscillator is free-running, and we know the frequency of the correcting offset for a given code into the DDS, it is easy to deduce the exact behaviour of an oscillator that is under closed loop control and how it would have behaved (say during temperature testing) if the control was not applied. The same is true of a standard oscillator if the voltage to frequency characteristic has been logged, but, because the oscillator is not free running, the effect of adjustments on its characteristics are difficult or impossible to separate.

For the closed loop frequency control we have incorporated a small GPS receiver that outputs a reference 1pps for comparison with the divided down output of our 10MHz output. The receiver is allowed to find its antenna position for 24 hours, with coordinate values being saved in accumulators and then divided by the number of samples to produce a position average.

The wide tolerance on the frequency output of the ovened oscillator requires an initialising strategy for a new instrument. To allow automation of this process an amplitude detector is fitted at the output of the frequency doubler. The 5MHz and 10MHz crystal filters each have 3dB bandwidths which are about 100Hz so there is likely to be no output from the corrector for any new oscillator which has the nominal 14kHz offset – it may be 4kHz outside the narrow pass-band. The amplitude detector allows the controller to step through a range of trial frequencies for a new oscillator, searching for output in a first stage of adjustment. The second stage is to find a more exact value of the 32-bit DDS value *N* to correct the exact offset to zero by measurement of our 1pps against 1pps from the GPS receiver. This value is logged into EEPROM.

For continuous frequency correction the time interval between the GPS receiver 1pps and 1pps derived from the 10MHz output is measured, averaged and integrated. The continuous measurement and adjustment strategy results in final corrections being applied which are the smallest frequency step that can be made (4.6E-12). A second order phase-locked loop is implemented with variable time constants as used in our “Intelligent Frequency Standard” described some years ago [4] when the reference was the carrier frequency of an LF transmitter. The same principles apply.

There are some key characteristics for an instrument of this type which we will evaluate and compare with results taken from a similar instrument fitted with a more standard OCXO. Frequency aging and temperature stability will be investigated so that we can predict holdover performance. We also study phase noise, time domain stability (measured frequency stability over different time intervals), and spectral purity of the 10MHz output.

II. Aging

We took a newly delivered oscillator and plotted the frequency controller adjustments against time over the first few days of operation.

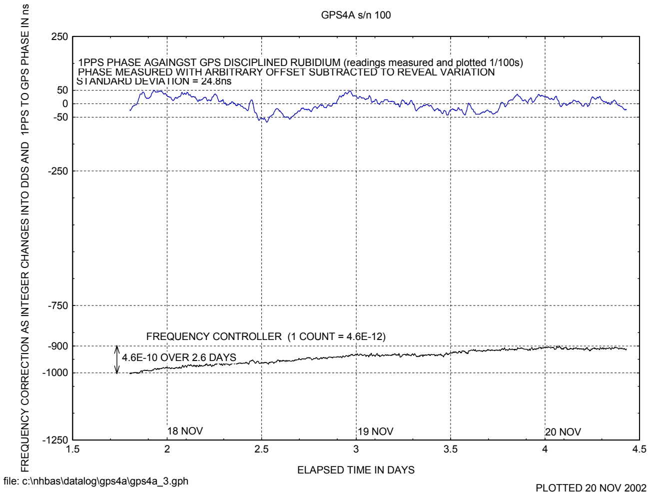

At the start daily aging is evident but averages less than 2E-10, not bad for a brand new oscillator, and it shows considerable improvement after 3 days.

The same unit was allowed to run for about 30 days, the usual period of pre-aging for precision oscillators, and then from the beginning of January 2003 it was logged over a 90-day period using Timetrace GPS Common View calibration of the laboratory monitoring system.

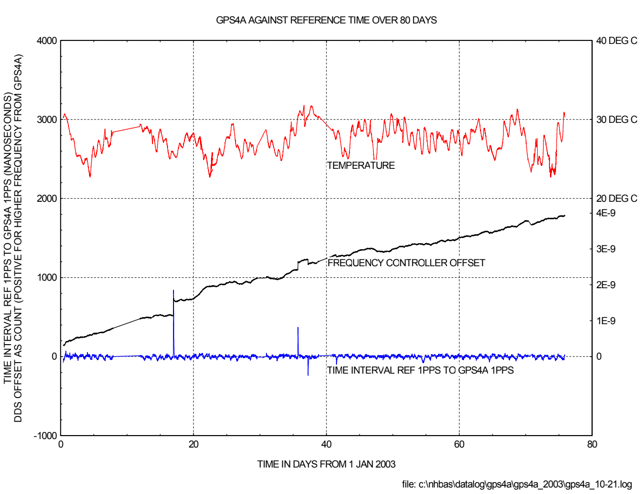

The frequency controller correction follows a long-term trend that is quite linear and amounts to a total increase of 4E-9 over 80 days, averaging 5E-11 per day. The oscillator has correspondingly aged lower in frequency by this amount. This is an unselected quartz unit taken at random from the small batch produced.

The time accuracy of the unit's 1pps output throughout the period is generally well inside ±50ns and often inside ±20ns for periods of days. There were two or three periods of lost of power during the logging requiring a complete re-start.

The exceptions are 3 notable 'spikes' in the 1pps time interval from the reference 1pps. These are caused by small frequency jumps in the oscillator; the largest frequency jump is -8.7E-10 at day 17, requiring the frequency correction to increase by the same amount. At day 35.7 a jump of -3.8E-10 occurs, most of which is recovered by a small jump back at day 37.2. These events cannot be predicted but are fortunately rare, only 3 being captured in data logged every 100s for 4 months. A technique that can help eliminate production oscillators prone to this type of instability has been reported [2]. The authors find that there is a correlation between the susceptibility to frequency jumps and Allan variance measured over 45minutes which can help the manufacturer to eliminate unreliable units.

III The effect of temperature

To measure the effect of temperature we placed the unit in the centre of a 9ft^3 freezer chamber with a heater installed to allow generation of both elevated and reduced temperatures under computer control. The air was gently stirred with a small fan with no direct flow of air over the unit under test. The disciplining was disabled at the start of the test so that the output frequency was corrected for the major frequency offset but not for any subsequent variations.

The easiest way to view the response is to apply a step temperature change and view the resulting frequency changes. To see the resulting frequency changes without interaction with the applied temperature step it is necessary to allow at least six hours at each temperature. A linear frequency change with time, due to aging throughout the 20 hour experiment, has been subtracted in the graph below. The temperature is controlled to about 0.1deg C during this test and can be ramped up at 2 deg C per minute or down at 1 deg C per minute.

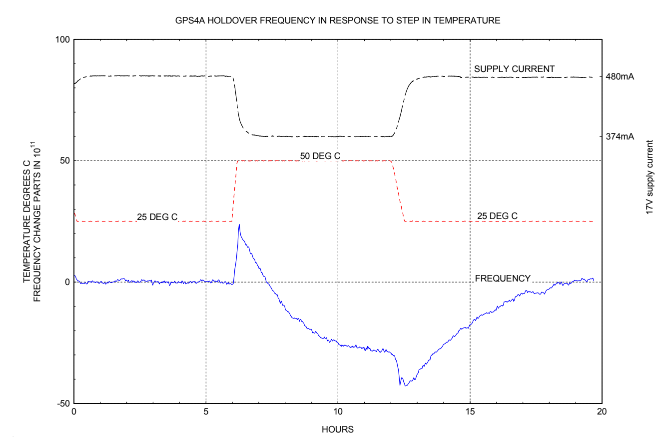

It is clear that the open-loop frequency variation with temperature is small, repeatable but complex in its characteristic. As a check on the effect of temperature on the unit the total supply current was measured. We can assume that the changes seen are due to the change in the oven current.

The initial frequency change is positive for a positive temperature step and lasts about 15 minutes, about the same duration as the time constant of the oven current in response to the temperature step. Subsequently the change in frequency reverses and appears to settle with a time constant approaching 2 hours. The frequency change settles at –2.8E-10 for the 25 degree temperature step. When the environment is stepped back from +50°C to +25°C similar changes take place in reverse. The initial frequency changes can be explained in terms of the temperature gradient that is introduced within the package as the oven responds to the input temperature step. We do not have an explanation for the subsequent 2 hour time constant of the remaining changes in frequency. In summary, the characteristic is encouraging because of its repeatability and because at about 1E-11 per degree C the frequency change is small compared with that seen in most ovened quartz oscillators of this size.

IV. PHASE NOISE

For many applications low phase noise of the 10MHz output is desirable. The graph below shows the phase noise measured against a good quality 10MHz SC cut oscillator. The noise is higher than expected from a conventional SC cut crystal oscillator (a low phase-noise disciplined oscillator is shown on the same graph and against the same reference for comparison), but close to carrier it is certainly quieter than a typical Rubidium oscillator. The frequency of the OCXO in the unit under test is 5MHz + 14.823kHz. As well as the wanted 5MHz the frequency corrector generates output at 5MHz plus twice the offset which is doubled again to 59.29kHz in the 10MHz output. It is visible but reasonably respectable at –140dBc. The largest peak directly attributable to the DDS is at 29.6kHz and is –116dBc. Whilst this is higher than we would like it does not compare badly with power-supply noise seen in the frequency outputs of some conventional instruments we have looked at.

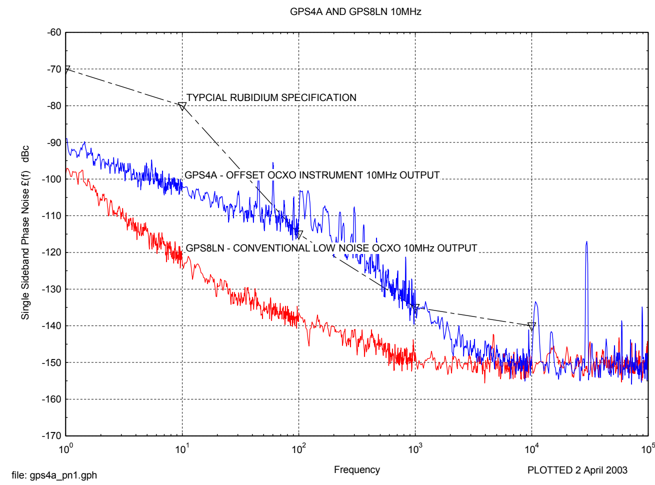

V TIME DOMAIN STABILITY

The time domain frequency stability was plotted as Allan Variance so that we could examine the stability of the disciplined oscillator over intervals of 1 second to several days. This graph will allow us to evaluate the effectiveness of the closed loop control and whether this control is damaging the stability of the oscillator especially over short to medium terms in the region 1 to 1000s. A trade off is required between too tight control which transmits jitter in the GPS 1pps onto the output frequency and too long a time constant which allows temperature changes to degrade the performance. The results obtained from a standard, non-offset oscillator are shown on the same graph for comparison. The curves are almost identical with just a suspicion that the short term stability of the offset oscillator is not quite as good. In fact we are not able to discriminate very well at values of τ that are less than 100s because our measurement resolution is limited to about 75ps as obtained from the SR620 Universal Time Interval Counter in use with the 1pps output of our laboratory reference.

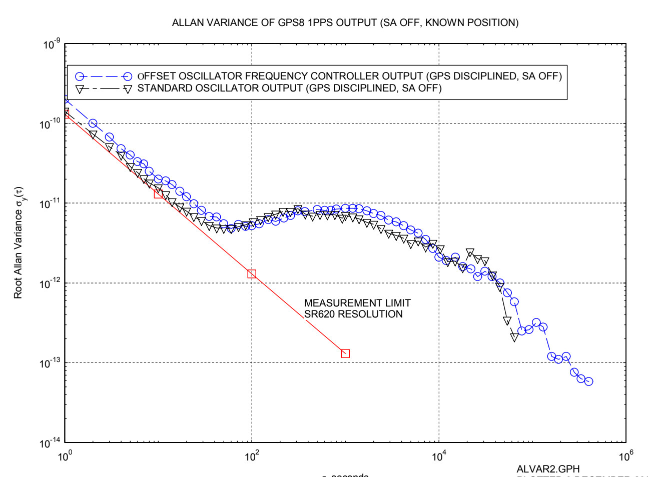

VI SPECTRAL PURITY

The 10MHz sine wave output is set to drive 1Vrms into 50 ohms (+13dBm). The output is examined on a spectrum analyser to check for spurious signals and harmonic levels. We particularly want to scrutinise the spectrum for unwanted signals from the frequency corrector. The scan below shows the result for 0 to 20MHz with the 10MHz output shown at –7dBm at the centre of the scan. A 20dB attenuator was fitted prior to the analyser to avoid overloading its input.

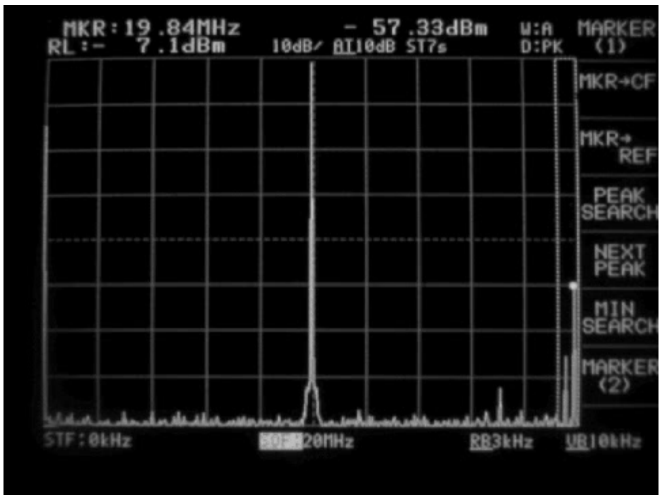

The spectrum analyser has a dynamic range of nearly 80dB. There are no unwanted frequencies below 10MHz clearly visible. Above 20MHz there is a peak at 17.04MHz which appears to be an artefact of the analyser – it is seen when we view other totally unrelated good quality 10MHz signal sources on the same instrument. The peak at 19.6MHz is induced by the output of the general purpose synthesiser implemented on the same printed circuit card that generates a logic level square-wave. Finally, at the end of the scan the second harmonic output is visible at –50dBc.

VII PHYSICAL CHARACTERISTICS

The unit is constructed on a single printed circuit card. The unit is constructed to accept DC power input and generates four output signals as well as providing two serial interfaces. The outputs are the modulated IRIG B timecode, 1pps, 10MHz sine wave and a synthesised frequency from a PLL synthesiser that is locked to the disciplined 10MHz.

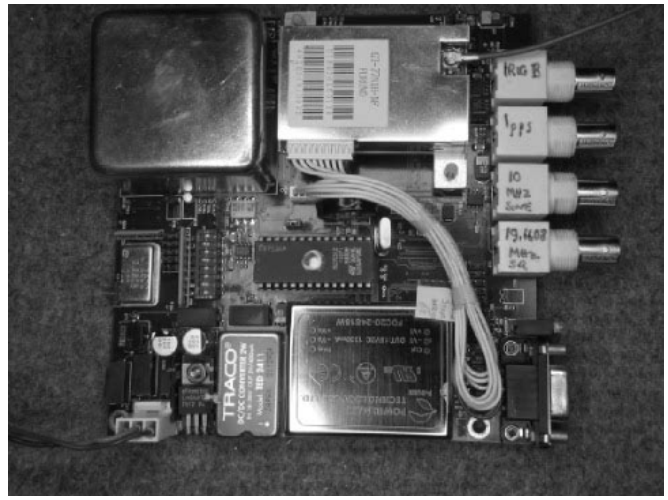

The card fits into a case which is about 5½" square and 1¾" high.

VIII CONCLUSIONS

The control system was successfully implemented within a small package size already developed for a non-offset version of the instrument. The quartz oscillator aging was low and consistent with the exception of 3 small frequency jumps noted over a 90 day period. Temperature stability of the unit during loss of GPS reference is exceptionally good. Phase noise at the 10MHz output is higher than is seen from non-offset quartz oscillators but lower than that seen from Rubidium oscillators; the frequency spectrum of the 10MHz output shows that the signal processing has suppressed unwanted frequencies. Work on lowering the aging and temperature effects is continuing.

ACKNOWLEDGMENT

Brandywine Communications built and supplied the special oscillators and the unit for which the above results are reported.

REFERENCES

[1] Bruce M. Penrod "Adaptive Temperature Compensation of GPS Disciplined Quartz and Rubidium Oscillators," *1996 IEEE International Frequency Control Symposium*.

[2] Charles A. Adams, John A. Kusters, Catherine M. Sousa "Improved Quality Factor for Quartz Oscillators" *1996 IEEE International Frequency Control Symposium*

[3] Richard Karlquist "A Narrow Band High Resolution Synthesiser Using a Direct Digital Synthesiser followed by repeated dividing and mixing" *1995 IEEE International Frequency Control Symposium*

[4] Nigel Helsby "Frequency Standards" *British Patent No 2,167,254 Published 1 June 1989*