MEASUREMENT OF FAST SETTLING MICROWAVE SYNTHESIZERS

Richard J Fawley MEng. (Hons) CEng. - Principal Engineer
Alex Scarbro BEng. (Hons) - Principal Engineer
Teledyne Defence Ltd. Airedale House, Acorn Park, Shipley, UK, BD17 7SW

Abstract

Microwave synthesisers employed in electronic-warfare applications are required to possess performance attributes that place extreme demands on the test and measurement methods utilized. Specifically, the task of measuring frequency-switching speed over a large bandwidth with sufficient fidelity is non-trivial. This paper presents a method of real-time analysis of frequency settling time using automated test routines. Issues related to achieving the required measurement fidelity are discussed, along with a summary of the theoretical performance bounds for frequency estimation.

Existing Measurement Methods

A tried and tested method of determining the frequency settling time of a synthesizer is to use a simple delay line discriminator. This method involves splitting the signal generated by the DUT in to two paths. One arm is delayed by 90° relative to the other arm and together they drive the LO and RF ports of a double balanced mixer acting as a phase detector. The IF port of the mixer is connected to a low pass filter to remove any high order products before being fed to the oscilloscope input.

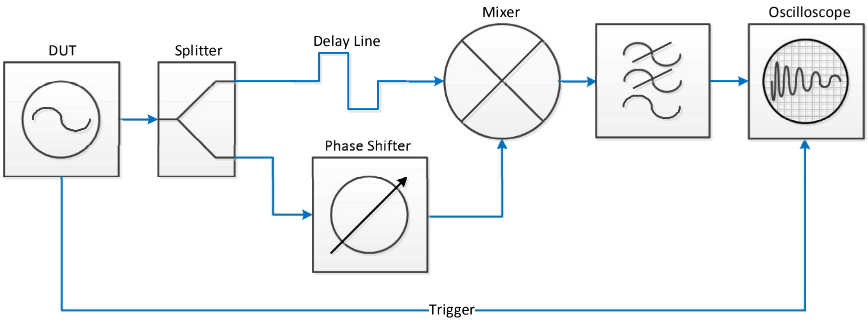

Figure 1 Delay line discriminator

By carefully adjusting the phase shifter in one arm, near-perfect quadrature between the mixers ports can be achieved at each frequency of interest. Quadrature of these signals is observed as a zero volt mean signal on the oscilloscope. A signal used to initiate frequency switching of the DUT is also connected to the oscilloscope to provide a trigger source. The DUT is then toggled between frequencies. Only when the waveform decays from a high frequency chirp to a zero volt DC level, can the synthesizer be declared as frequency and phase locked.

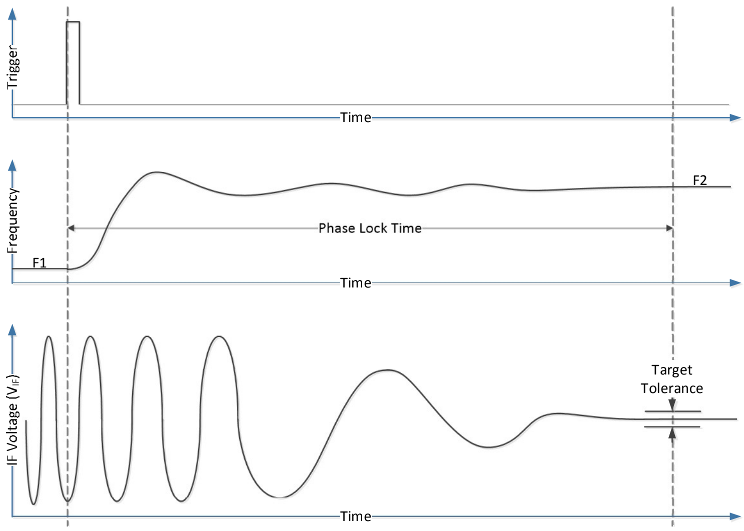

Figure 2 Discriminator output aligned with frequency and trigger data

This technique is relatively simple, gives good qualitative results, and only requires a few RF components & an oscilloscope with modest analogue bandwidth. There are however a number of drawbacks with this technique:

1. The phase must be carefully adjusted for each target frequency to achieve quadrature. This is not a significant problem, if you only wish to test the extremes of frequency, but compliance testing may deem every channel must be tested. A variant of this technique, *the phase detector method,* uses a second synthesizer locked to the same reference as the DUT and relies upon the adjustable phase offset capability of the second source to achieve quadrature. This is eminently practical to automate, but still time-consuming.

2. It is very difficult to declare "lock" to within a defined limit of frequency error. This can be a problem if the synthesizer is close to the target specification and the customer requires proof. The only way to determine the frequency error would be to demodulate the potentially low frequency FM superimposed on the DUT signal and place limits on that measurement.

The Proposed Approach

The proposed approach implements direct demodulation using a digital storage oscilloscope to sample the DUT output signal and provide a direct measurement of frequency settling.

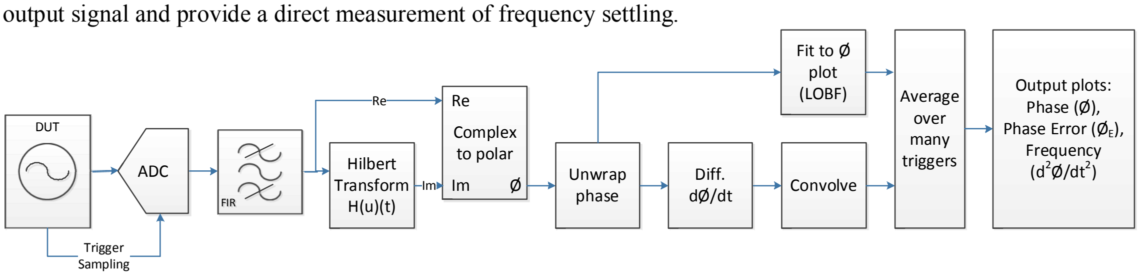

Figure 3 Direct digital demodulation

The sample rate must be at least twice the synthesizer tuning bandwidth as required by Nyquist sampling theorem. In practice the required sample rate should be much greater to ensure any low order high-level harmonics do not alias close to the carrier. Ideally a bandpass filter with equal group delay across the extremes of frequency should be installed between the DUT and the input of the oscilloscope. This group delay is subtracted from the measurement results. It should be noted that we couldn't use a sampling oscilloscope, as the signals we will measure are not periodic. Thankfully modern high-speed oscilloscopes can have analogue real-time bandwidths in excess of 3GHz and sample rates up to and beyond 40GS/s. The cost of such an instrument is typically less than $40k retail.

One of the advantages of this direct sampling method is that the oscilloscope does not need to be phase locked to the DUT or its reference oscillator. In fact, the averaging algorithms require that some degree of sample dithering take place, so the noise sources in the measurement system must remain uncorrelated for the most accurate results.

As with the analogue methods, the oscilloscope is configured to trigger from a frequency-switching signal. Some oscilloscopes have a dedicated trigger channel, which (on an oscilloscope with interleaved channel sampling) may give the greatest available real-time bandwidth to the single channel used to sample the DUT output.

The DUT is toggled between F1 and F2 with a period greater than the expected lock time and the oscilloscope captures the signal in the time domain for analysis.

As we cannot distinguish between amplitude or phase variation for a real signal, we need to extract the phase information by generating a complex version of the signal of interest by creating a quadrature component.

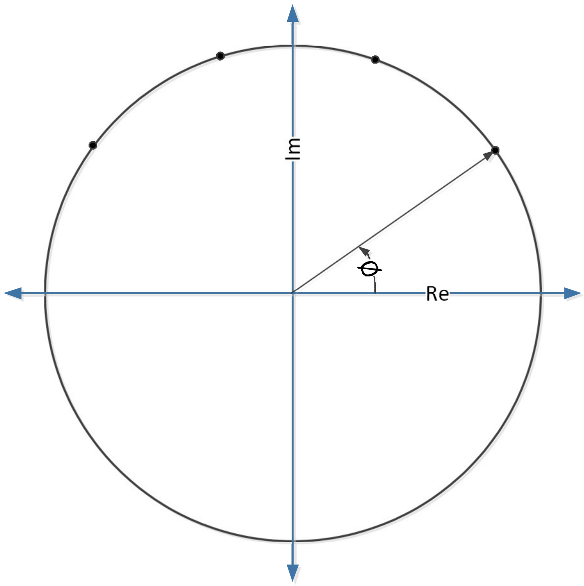

Figure 4 Phase from Re. and Im. signal components

This is achieved with the help of the Hilbert transform. This transform is essentially an all-pass filter with a 90° phase shift over its pass band. If we plot our samples Re=x(t) and the Hilbert transform of our samples Im=h(x)(t) we get a phasor rotating around the origin of our complex plane at a rate proportional to the frequency of the DUT.

By taking the angle of the phasor, unwrapping the phase, and plotting it versus time we can observe a linearly increasing phase. This is indicative of the DUT being stable in frequency. We calculate two lines of best fit through data; One during the period where the DUT frequency is stable at F1 (before we switch frequency) and the other some time after we switch frequency during the period where the DUT is stable at F2.

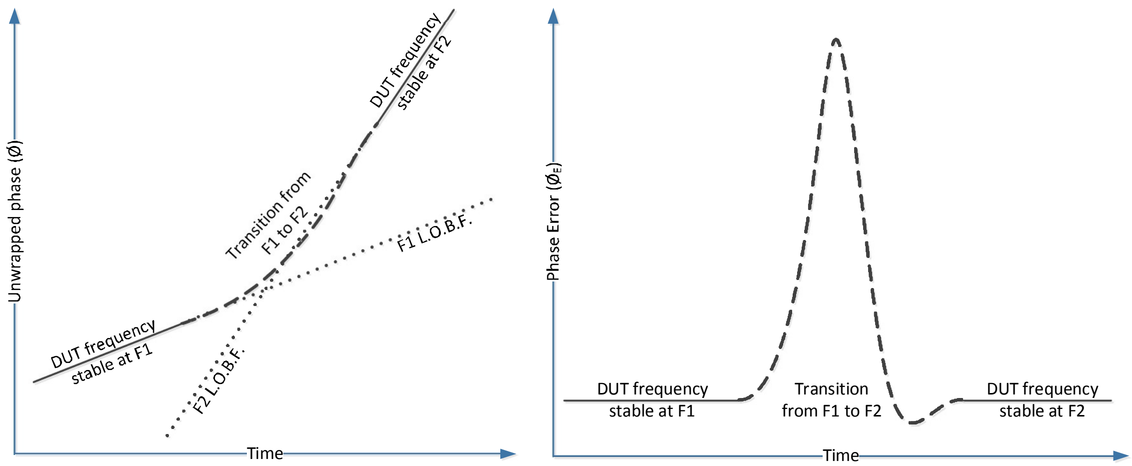

Figure 5 Unwrapped phase and phase error vs. time

This is done to automatically estimate an average of the start and end frequencies. These lines of best fit are then projected through the rest of the phase data. The least delta observed between each phase data point and the two lines of best fit is determined to be the phase error at that instant in time.

If we calculate the phase delta between consecutive data points and divide by the sample period, we get the phase change per unit time i.e. frequency. In theory we could now plot this frequency vs. time and our measurement would be complete, but unfortunately the data set contains various noise components.

These include, but are not limited to:

- The signal to noise ratio of the measurement chain. The gain of the oscilloscope channel and the DUT signal amplitude must be adjusted to maximise the swing of the sampled signal in relation to the vertical limits of the channel. This maximises the dynamic range of the oscilloscope such that the quantization noise will then become the limiting factor to SNR. A typical 8-bit oscilloscope will have a signal to noise ratio of around 40dB.
- Spurious signals generated by both in the DUT and the oscilloscope and harmonics which alias down into the band of interest.
- The sampling aperture jitter in the oscilloscope.
- The phase noise of the DUT signal.

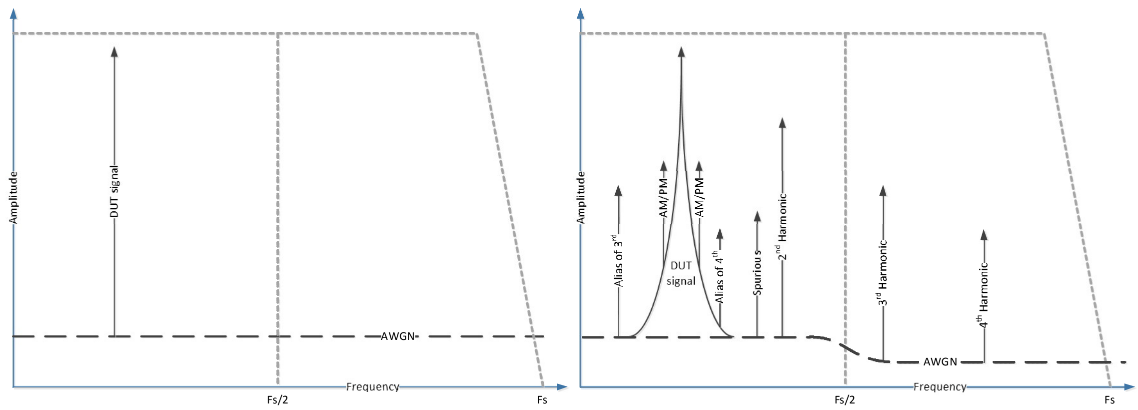

Figure 6 Ideal and real world sources of noise and interference

With this relatively simplistic test method, there is very little we can do about the systematic noise sources and so our only option is to average many processed results. This will artificially improve our SNR and ultimately increase the accuracy of our frequency estimation as long as the noise components are uncorrelated between measurement cycles.

We must however consider a few factors:

- The wider the tuning range of the DUT we are trying to measure, the higher the sampling rate of the oscilloscope needs to be.
- The slower the tuning time, the greater the number of samples required.
- The poorer the SNR of the measurement, the more averaging is required.

Therefore measuring a high frequency slow tuning oscillator in a low SNR environment means that we need to capture and store huge amounts of data. For example, if we set the oscilloscope to have a sweep time of 1ms at 40GS/s with 8-bit resolution, each record will be 40MByte long and take between 0.5s and 40s to transfer from the instrument to the connected PC (depending on the medium used). Therefore a 100-sweep average would require 4GByte of RAM just to store the sweeps and anywhere from 50s to 4000s to transfer the data.

There are optimizations that can be made (including down sampling the DUT signal), but it should be recognised that this particular implementation of direct digital demodulation is best tailored to fast switching synthesizers unless there is enough time and memory available to transfer and store the data sets.

Test Results

The following results show the frequency locking performance of a fast settling microwave synthesizer based on a PLL with some technical enhancements. This synthesizer is being tested over a portion of its full tuning range.

It is instructive to perform an FFT of the complete data set to understand the average level of noise and spurious products before further processing.

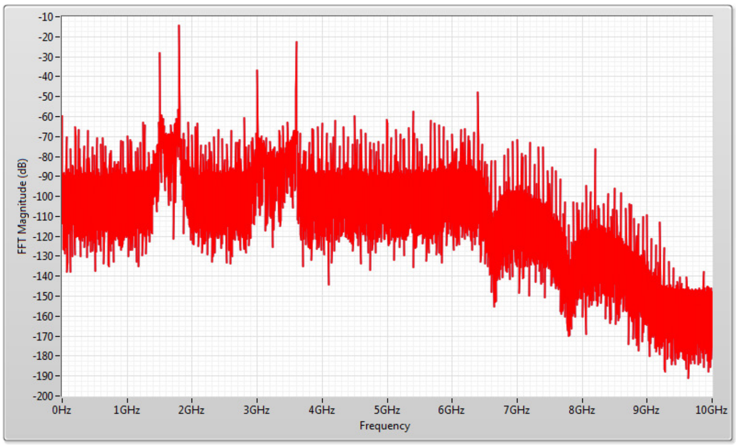

Figure 7 FFT of sampled data before filtering

Here we see our two expected signals at 1.5GHz and 1.8GHz. These are the defined F1 and F2 tones and their relative levels relate to the period of time spent at each frequency (i.e. synthesizer spends more time at F2 than F1 in this particular data set). Harmonic products of both F1 and F2 and interfering tones can also be seen extending out the full bandwidth of the oscilloscope.

If we just perform frequency estimation and apply a 100 point moving average, a significant amount of noise is observed and hence ambiguity in the measurement.

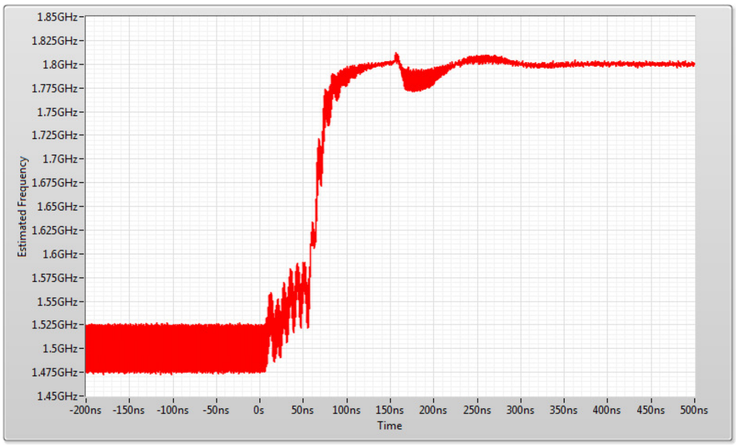

Figure 8 Unfiltered estimation of frequency as the DUT steps from F1 to F2

The primary cause of this noise is the out of band interfering signals. Next a 91-tap FIR bandpass filter was applied to the data with its passband set to the frequencies of F1 and F2 respectively. As can be seen, this reduces the highest-level nearby spurious tone to better than –40dBc.

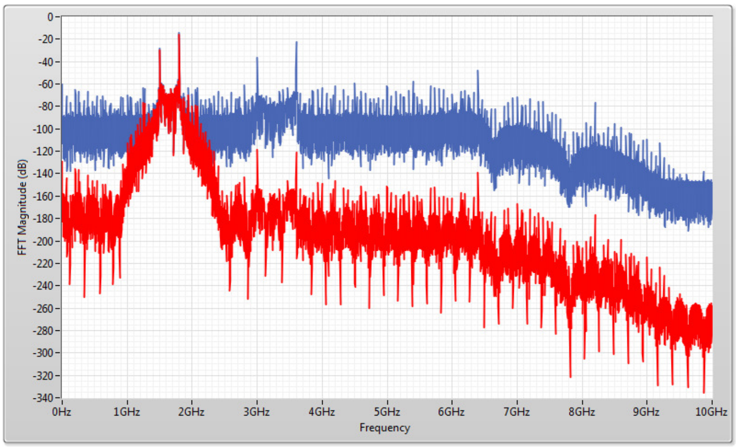

Figure 9 FFT of sampled data after filtering. Blue = Original, Red = Filtered

Now we can repeat the frequency estimation process after convolving the data over 100 points. Now there is a substantial reduction in noise.

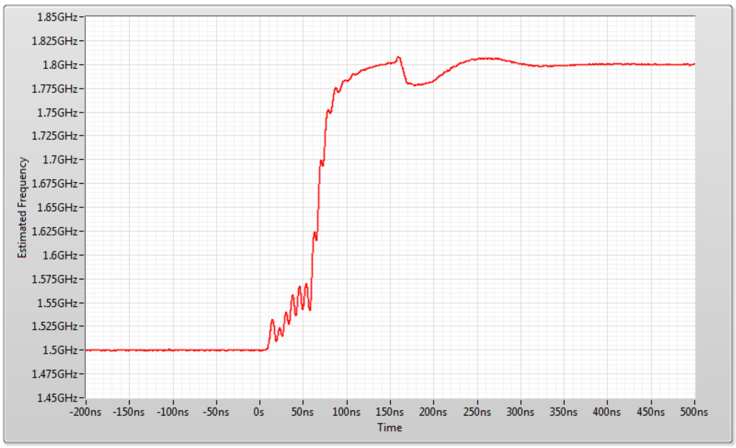

Figure 10 Filtered estimation of frequency as the DUT steps from F1 to F2

In fact, if we zoom in on a portion of the measurement where the frequency is stable we see a peak to peak variation in our estimation of approximately +/-500kHz. Averaging many sweeps will reduce any uncorrelated errors by 1/sqrt(N), but to achieve a resolution of +/- 100kHz would require 100 averages, +/-10kHz would require 10,000 averages, +/- 1kHz would require 1000,000 averages etc.

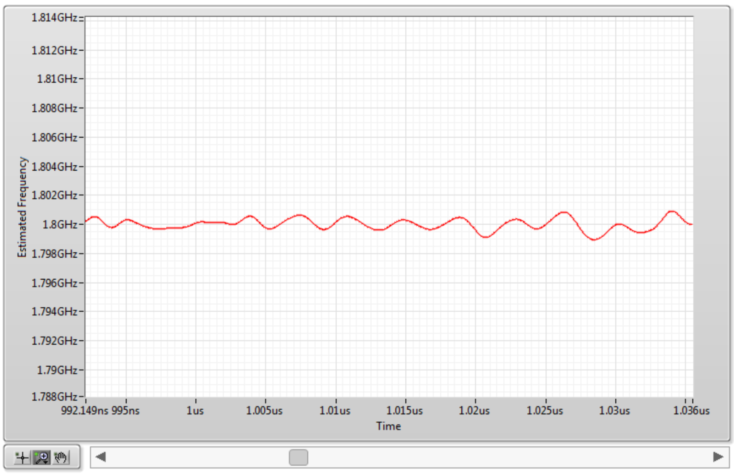

Figure 11 Zoomed in view of frequency estimation results

Theoretical Limits on Frequency Measurement Uncertainty

Having implemented the frequency estimation method as described above, it is instructive to examine the effectiveness of the measurement technique and compare the results with expected performance, theoretical or otherwise. As discussed, we are required to estimate the instantaneous frequency of the signal at points in time as the synthesiser switches between two frequencies using a finite number of noisy discrete time observations. Although the parameter of interest is by definition non-stationary, it can be assumed to be stationary over a short observation interval. If so, the likely performance of the estimation technique can be predicted. The Cramer Rao Lower Bound can describe the theoretical limit for the measurement uncertainty of any frequency estimation technique:

$$\mathrm{var}(\hat{f}) \geq \frac{12F_s^{\,2}}{(2\pi)^2 \times SNR \times N^3}\ (Hz)$$

Where, Fs is sample rate, N is number of samples and SNR is the signal to noise ratio. This relationship holds whether we have knowledge of the phase or amplitude of the signal or not. The variance of the frequency estimation decreases with either or both of the following:

- Improved SNR
- Increased number of samples for a given sample rate, i.e. longer observation intervals.

This relationship is valid when the signal has a single frequency component corrupted with additive white Gaussian noise. Although these conditions are unlikely to be present in a practical situation, the CRLB provides valuable insight into the best-case performance and is a benchmark most estimators are judged against.

The following figure illustrates how practical estimators perform compared to the CRLB. Typically, estimators depart from the theoretical best case at the extremes of SNR, however most efficient estimators do indeed reach the CRLB over a specific range of SNR.

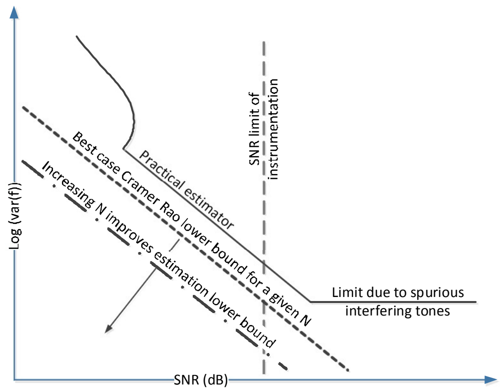

Figure 12 Practical estimator performance

It is then possible to calculate the CRLB for the ideal conditions assumed to be present for our test scenario, namely 100 points (5ns observation interval) and 40dB SNR (dominated by the noise floor of the oscilloscope). Using the equation provided earlier we can determine that the best-case measurement uncertainty is 110kHz RMS. Our measured results of approximately 350kHz RMS are clearly not optimal, but are of the order of the best-case result.

The Phase Differencing Algorithm

As an interesting aside, it must be noted that an Instantaneous Frequency Measurement (IFM) Receiver estimates the frequency of wideband radar emitters by performing a very similar phase differencing function using analogue and/or digital techniques. These receivers have been utilised for many decades and remain unparalleled in performance with respect to wide bandwidth and short pulse processing.

As can be seen from Figure 15, a high performance IFM receiver can estimate frequency to within 1MHz of the CRLB in a 50ns interval at 0dB SNR. This puts our results over a 5ns interval in to context and gives us additional confidence in the technique.

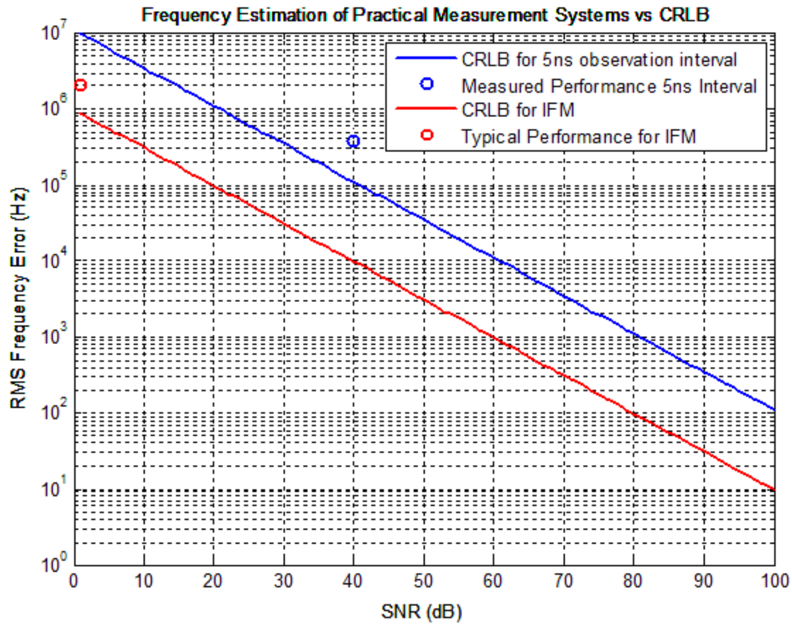

Figure 13 Measured results vs. the Cramer Rao Lower Bound for given intervals

Non-Ideal Factors Affecting Measurement Uncertainty

The noise present throughout the system and measurement chain is not necessarily white nor is it necessarily Gaussian. As indicated earlier, the broadband noise present in the system is likely to be dominated by the quantization and/or jitter noise of the broadband digitisation equipment, which may be correlated with the signal of interest. Furthermore, it is likely that there will be the requirement to utilise frequency selective filters, the resulting noise will be band-pass in nature, reducing the effectiveness of any time-domain averaging employed.

The signal of interest is likely to be corrupted by multiple coherent interferers, namely harmonics and spurious components generated by both the DUT and the measurement system itself. It is likely that certain interferers could be filtered out, however there is a potential for many spurious products to appear close the signal of interest, particularly high order components aliasing-down due the sampling process.

It can be shown that if we have a single component sinusoidal signal corrupted with a single sinusoidal interferer separated by Δf *Hz* with relative amplitude R, the instantaneous frequency will contain a time varying error component:

$$e(t) = R \cdot \Delta f \cdot \left[\frac{R + \cos(2\pi \cdot \Delta f \cdot t)}{1 + R^2 + 2R \cdot \cos(2\pi \cdot \Delta f \cdot t)} \right]$$

Which varies between the limits:

$$\pm \frac{R \cdot \Delta f}{1 \pm R}$$

For example, an interferer 100MHz away from the main signal of interest 60dBc down in power, the instantaneous frequency will have a cyclic error component, which is approximately sinusoidal with peak to peak value of +/-100kHz. It must be noted that these cyclic error components can only be reduced by filtering or averaged out if integer multiple cycles of the error component can be captured. A situation which in practice will be difficult to achieve when multiple agile interferers are present and are aliasing into the band of around the signal of interest, i.e. when Δf is small. A practical example of how the elimination of an interferer significantly improved the measurement performance was illustrated in Figure 10.

Multiple interferers will all contribute to the instantaneous frequency error and will add in RMS fashion if they are not correlated with each other.

Analytic Signal Generation

Creating the analytic signal from sampled data is potentially non-trivial. Typically, only an approximation of the true analytic signal can be derived. For a high fidelity frequency estimation process, deviation from the ideal can result in significant estimation errors. If the Hilbert transform is used, one must pre-process the real data with an appropriate time domain window. Failure to do this can result in an analytic signal with significant imbalance between the complex components. Just as imbalances in quadrature processing elements in receiver chains result in poor sideband suppression. Imbalances due to non-ideal analytic signal generation create spurius which will degrade the frequency estimation process. Alternative methods exist for generation of the analytic signal, including complex downconversion of the digitised data. These methods also have the potential to provide rejection of unwanted interferers.

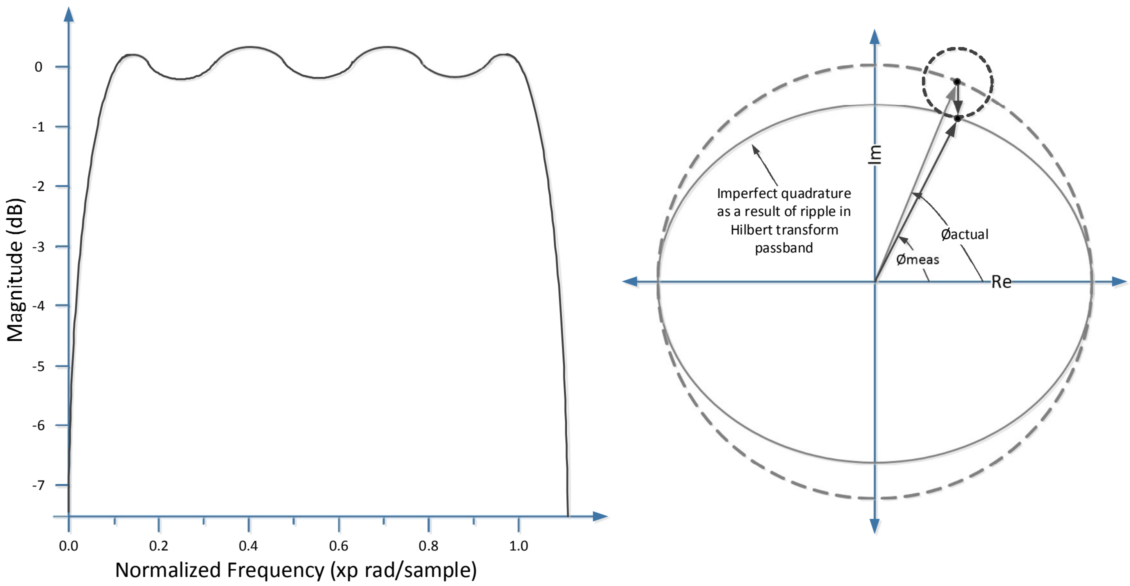

Figure 14 The Hilbert transform frequency response and its effect on complex frequency

Refinement Of The Technique Through Polynomial Fitting

In estimating the instantaneous frequency by calculating the phase difference between adjacent samples and smoothing, we assume that the phase of the signal can be approximated as a linear function of time. A sliding window with sufficient length is utilised to capture and process sufficient number of samples whilst tracking the agile signal. However, if the frequency law varies significantly within the window used, the estimate will be degraded.

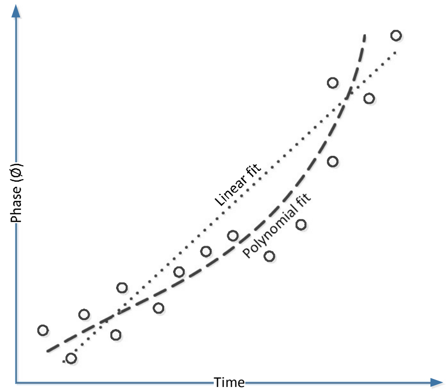

Figure 15 Linear vs. polynomial fit of unwrapped phase data

A refinement to this sub-optimal scheme recognises the fact that the signal of interest will likely have a phase law with higher order time varying phase components. The estimation technique can therefore be extended by modeling the signal phase law with a polynomial of appropriate order. A complex form of the signal can be described as:

$$x(n) = A(n)\,e^{j\left(a_0 + a_1 n + a_2 n^2 + ...a_p n^p\right)}$$

where a_0, a_1 ... a_p are coefficients of the polynomial of order p, assumed to correspond to the phase law of the signal. If the signal is assumed to follow a linear FM law, then $p=2$. The estimation process is simply a case of determining the coefficients of the polynomial, typically using least-squares techniques. The initial phase θ_0, of the analysis set can be determined from a_0, the instantaneous frequency $\delta\theta/\delta t$ and linear frequency rate $\delta^2\theta/\delta t^2$ from a_1 and a_2 respectively. An important benefit of calculating higher order phase derivatives such as FM rate is that they are a direct measure of the settled state of a frequency agile signal. A synthesiser can be deemed to be settled in frequency only when the magnitude of the a_2 coefficient corresponding to the FM rate is less than a prescribed value. Typically, initial phase θ_0 of the analysis window is not of particular interest for measurement of synthesiser frequency settling, However, having a robust method of ascertaining the absolute phase of a synthesiser at a particular instance in time is valuable when developing and measuring coherent synthesisers.

Conclusion

A reliable method of measuring the frequency settling time of fast tuning wideband synthesisers has been proposed and implemented. A fast tuning synthesiser was tested, demonstrating that the technique can operate in near real time. Factors affecting the fidelity and accuracy of the measurement process have been identified both theoretically and through practical experiment. In understanding the major contributors to measurement uncertainty, it is possible to appropriately configure various parameters associated with data conditioning and processing. In verifying the limitations of the achievable performance due to various real-world factors we are able to refine the technique further, to this end, enhancements to the technique are proposed.

References

Bruce G. Anderson, “Frequency Switching Time Measurement Using Digital Demodulation” IEEE Transactions on Instrumentation and Measurement, April 1990.

Boualem Boashash, “Algorithms for Instantaneous Frequency Estimation: A Comparative Study”, July 1990.