

Clicklock

Weak Signal LF Propagation Measurement Tool

M. Greenman ZL1BPU
April 2006, Rev 3 June 2011

Clicklock is a graphical program for the PC which, in association with an LF/MF receiver and antenna, allows reception of very weak but stable signals, which can be studied over long distances and long time frames, using carrier synchronous integration techniques.

Clicklock uses a signal processing technique developed by Peter Martinez G3PLX. It enables a conventional communications receiver to be calibrated automatically to receive LF/MF signals with such precision that the carrier phase and power can be confidently monitored for days on end¹.

The technique demonstrates greater sensitivity than can be achieved using other techniques, and since more information is available about the signal (carrier phase and power), signals previously considered unusable can be detected and tracked to provide useful propagation indicators.

Measuring Phase

LF signal propagation cannot be assessed using HF techniques such as Doppler frequency measurement or radar-type ranging. The propagation mechanism is different, and the Doppler shift, which results from changes in reflective height in HF propagation, is far too small to detect as a frequency change since Doppler shift is proportional to frequency. The differences can however be detected as changes in carrier PHASE.

This technique is not practical on HF since the carrier phase changes are continuous and very marked, and so no meaningful phase comparisons can be made. It is for this reason that HF phase-shift keyed modes invariably use differential phase shift, and have limited performance where the incidental (path induced) phase shift is high. The upper limit of practical carrier phase measurement is about 500kHz.

Because of bandwidth limitations and the need for stable phase for carrier phase LF/MF measurements, conventional HF modulation techniques are not practical. However, two things can still be usefully measured: signal strength, and signal carrier phase. If the phase of the transmitted signal is known, or at least constant, changes in the time of flight of the signal can be measured directly by measuring the phase difference from source to destination against a common reference. There is no better, or universally available, common timing reference than the GPS system.

LF propagation is a combination of ground wave, which travels long distances, and D-layer or E-layer ionospheric propagation, where the ionosphere and earth act in the nature of a waveguide, since the height of the ionosphere is of the same order as the signal wavelength (measurable in kilometres).

The received phase of a ground wave component will be essentially constant if the transmitting and receiving sites are fixed, while the received sky-wave component will have a carrier phase which depends on the reflection height and the propagation mechanism (D-layer, E-layer) involved. The received signal will be a combination of these signals.

¹ See notes by Peter G3PLX in Appendix

As mentioned, at LF or MF, a conventional FFT-type signal detection tool such as ARGO may be able to detect a signal, but unless it has some very low speed modulation (such as QRSS Morse) or other unique signal feature, it is extremely difficult to identify the signal – to determine whether the signal is a real signal, or some artefact of the receiver. For example, Figure 1 shows a carrier near 20kHz (the receiver was tuned to 18.4kHz USB, and ARGO shows the audio tone to be about 1600Hz).

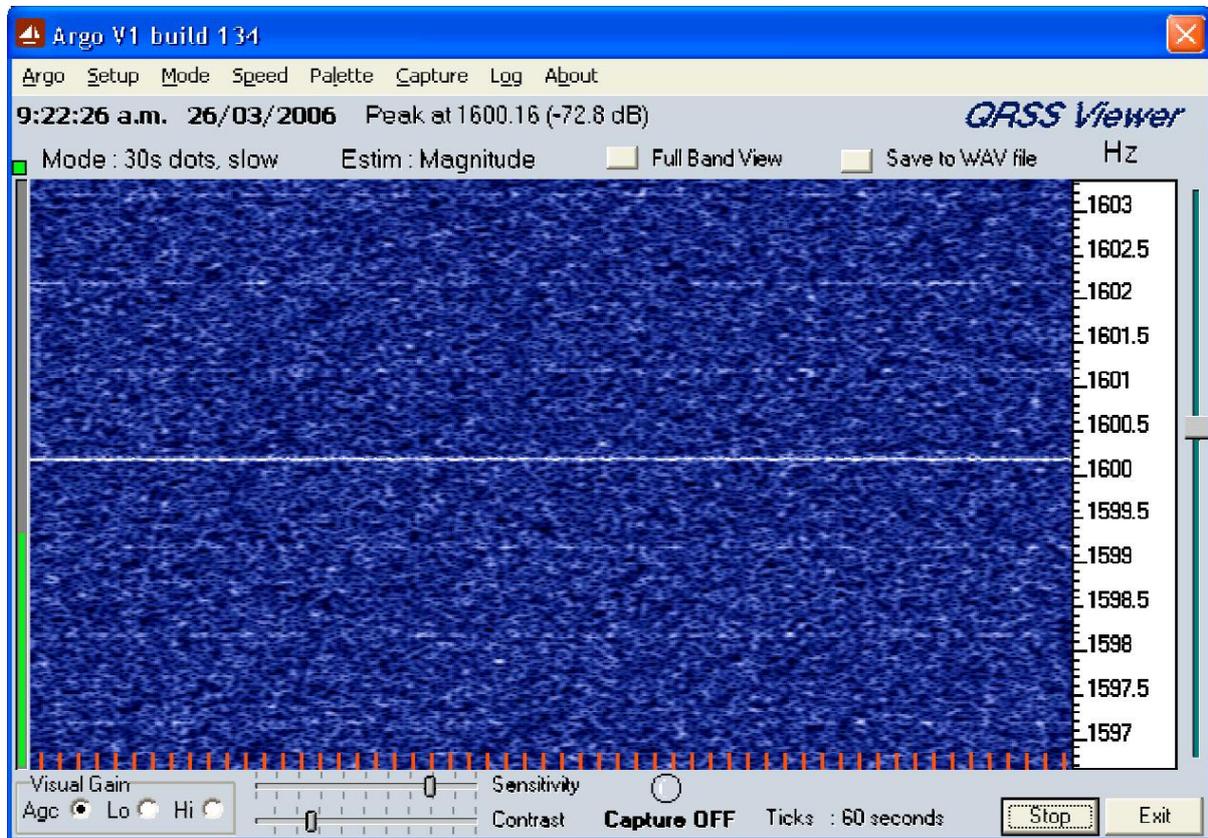


Fig. 1. ARGO trace of signal near 20kHz (at 1600.2Hz)

The bright line across the middle of the trace is about 6dB stronger than the background, and is clearly visible. There are other fine lines 1Hz apart caused by 1Hz pulses applied at the receiver antenna from a GPS unit (this picture was recorded simultaneously with the next picture). But is the frequency of the main signal exactly 20.000kHz? We can't tell, as there is no way to accurately calibrate out the receiver and sound card sampling rate errors, and certainly no phase information is available with this system. Even amplitude information is difficult to discern, and the picture is peppered with noise.

In the next picture (Figure 2), the Clicklock receiver is operating locked to 20kHz EXACTLY, and is GPS synchronous. As the signal rises out of the noise, it becomes very clear that this signal is exactly tuned to the receiver, as it has reasonably constant phase. In other words, it is likely to be GPS synchronous. The signal started to appear at about midnight local time, and was still clearly there at 11AM. Each dot represents the result of integration of the signal for 100 seconds

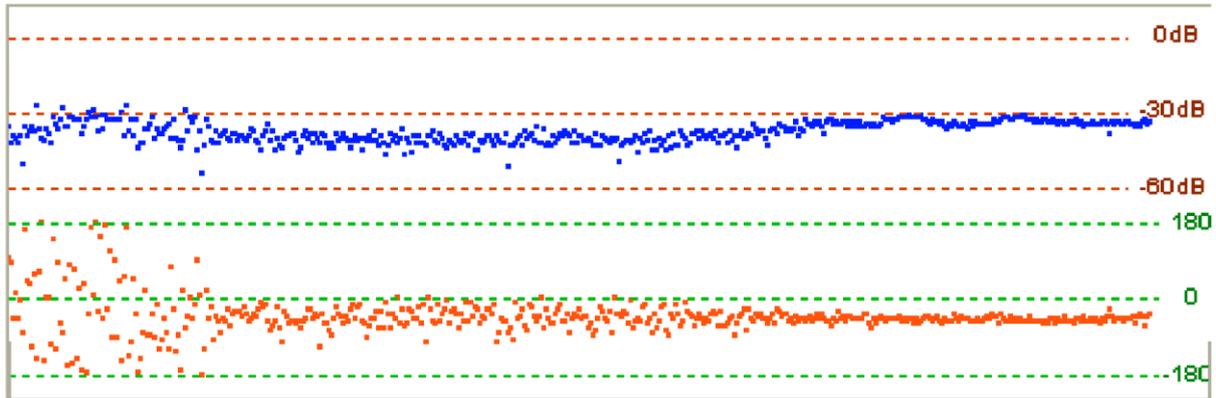


Fig. 2. Clicklock Power (blue) and phase (red) plots at 20kHz.

It is probably obvious that the **Clicklock** software rejects all other non-coherent signals, and all the noise, so the signal trace is very clean and sharp, even though weak. The signal being measured in this example was from a low power 20kHz experimental scientific station at the South Pole! Even though the range is about 8000km, reception is apparently surface wave (as the phase is constant). We can rule out it being from a local source as the signal has diurnal ionospheric effects.

In summary, on VLF, LF and MF, instead of conventional HF Doppler or radar techniques, using **Clicklock** two things can be usefully measured -signal strength, and signal phase; and in addition non-coherent (non-GPS related) signals are rejected. The system is more sensitive than any other known technique. Integration for 100 seconds (as in the example) is like having a receiver with a bandwidth of only 0.01Hz, and as importantly, the equivalent stability.

How does Clicklock work?

Communications receivers, even very good ones, have some uncertainties as a result of free-running (unreferenced) local oscillator phase and frequency. This is true even of the best receivers. There's no easy way around this: even if the frequency is spot-on, there's no way to meaningfully measure the received signal phase because the oscillators in the receiver do not have their phase referenced to any world-wide standard. Nobody wants to hack into a receiver to modify the oscillators, even if it was practical to do so, so how can we get around this?

In addition, there can be phase variations that creep into a receiving system as a result of tuning variations in the antenna system, thermal changes in cables, different propagation times in the receiver on different frequencies, with changes of frequency, use of different filters, and at different signal strengths. What we really need is a means of calibrating out (and compensating for) these receiver phase and frequency differences.

This is what **Clicklock** does. **Clicklock** is a PC sound card based system, and uses two versions of a universally available reference, the 1pps (1Hz) pulse from a GPS receiver, to determine and compensate for the receiver uncertainty. The 1pps pulses are generated by the GPS receiver from information received from the satellites, and have very high precision. Equally importantly, they are available around the world, and already form the reference for many existing transmissions.

In the **Clicklock** setup, the 1pps pulse from the GPS receiver is attenuated and fed directly to the PC sound card (the *PPS Click*), where it is recovered to generate the main software reference. The software uses this pulse to discover the exact location of a second 1pps event, the *RF Click*.

This second version of the 1pps pulse is fed into the receiver antenna, and is exposed to time delays and phase shifts through the antenna, cables and receiver. RF harmonics of the 1pps pulse are present at every Hz throughout the VLF/LF/MF spectrum, and those which fall within the receiver bandwidth cause an audible click to be heard. The relationship between these harmonics can be

analysed by the software. Without going into complex maths, the phase relationship between all the click harmonics in the receiver bandwidth will only be the same if the receiver is exactly on frequency², and at this point their absolute phase (i.e. delay) is also known by comparison with the audio click. If the frequency is slightly off one way, the phase of the different harmonics will generally increase across the receiver bandwidth, or decrease if the frequency is off the other way.

The software uses Digital Signal Processing techniques (DSP) to analyse the phase of the harmonics in order to tune the receiver to the exact frequency, in the correct phase. It can't do that directly, since there is no communication with the receiver, so another DSP technique, the 'zero IF receiver' is used. A local oscillator in the software (called a Numerically Controlled Oscillator or NCO) mixes the received audio (sound card digital samples) down to zero frequency. This oscillator operates at the designated centre frequency (called the lock frequency), say 1700Hz. The phase of the zero frequency result is analysed and the software oscillator moved slightly in frequency and phase until all the harmonics line up in frequency and phase at zero Hz. In other words the whole receiving system becomes phase locked by phase locking the NCO via the clicks. Hence the name **Clicklock**. When the receiving system has locked, the receiver is ready to look for signals that are coherent with the receiving system (on the same locked frequency and with a constant phase angle).

With this technique the receiver can be tuned in GPS-locked 1Hz steps, even if the receiver itself can only make 100Hz (or even approximate Hz) steps. This is because the software also provides an offset feature (for example to allow reception at 33.3333... kHz). However, the receiver must also remain stable within the tracking ability of the control loop, or the lock frequency might step up or down 1Hz.

The receiver signals come in from the antenna along with the RF click, and the software uses DSP techniques to determine the phase and amplitude of the signal, and to ignore the click products during signal reception (rather like a noise blanker). The signals can be extremely weak, well under the noise level, but by adding the samples together over a long time period (integration) the signal can be enhanced and the noise cancelled. The system operates best with unmodulated carriers, although it does work well with some interrupted carriers, such as LF time signals.

This property of essentially limitless enhancement only applies when the received signal is in exact carrier synchronism with the receiver. The amount of enhancement possible is limited only by the stability of the signal. The signal samples continue to add, and the noise that comes with it tends to cancel out, since noise is random and has no phase coherency. In addition, interfering signals and signals on adjacent frequencies, being non-coherent with the receiving process, are also cancelled out.

Coherent Reception

"Coherent" communication systems have been claimed in the past, but in reality they were at best time synchronized additions at bit level, NOT carrier phase and frequency synchronous. Repeated data can only be usefully added at a bit level to a limited extent, and after a time the signal phase wanders and the samples no longer add. Indeed, the long term sum of non synchronous signals is random, just like noise.

If however the carrier frequency and phase are analysed in a synchronous manner, it is possible to add samples (integrate them) for very long periods of time, to have the signal build up and the noise cancel out. With a coherent system such as **Clicklock**, integration over HOURS is possible. This is like having a receiver with a bandwidth of small fractions of a Hz, and the stability to match.

The sensitivity offered by **Clicklock** is better than using a Spectrogram (such as Argo) - perhaps by more than 20dB. In addition, signals at least 40dB weaker than can be heard by ear can be detected

² To be really accurate, the harmonics will have a constant mean value of phase, since the group delay of the pulses through the receiver (especially in crystal filters) causes variations of phase, but these can be ignored here.

reliably. The technique makes around the world LF propagation measurements possible, and allows LF and ML measurements at up to 1000km range to be made with Amateur-sized antennae and low power transmitters.

With the really long integration times required for the highest sensitivity, it isn't possible to send meaningful data, as the bandwidth is so low, but plenty can be INFERRED from the reception -

- Is the signal there? (Does it coincide with an advertised schedule?)
- What frequency is it on? (Station A may be on one frequency, and station B another.)
- What happens to propagation over the course of the day -can you work out where the sky wave propagation time is? From this you can know something about where the signal is from.
- Does the signal vary over time? (If not, it might be a local spurious signal!)
- What happens to the signal phase? Is it leading (reflective height lowering)? Does it jump suddenly (sky wave or ground wave predominates)? Is the phase reasonably constant (ground wave)?
- Is the signal GPS synchronous? (Either a standard frequency station or a specially GPS synchronized signal). This may help identify the signal.
- Is the signal phase the same as yesterday at this time of day?

The system works best with signals that are transmitted on exact round frequencies (integer Hz). These are of course also easier to generate using GPS as a reference. However, the software has an additional fractional offset that can be added in order to operate with non-integer signals, but they must of course still exhibit very high stability – of the order of 1 part in 10^{10} or better (a typical crystal oscillator is about of 1 part in 10^5 , a TCXO 2.5 parts in 10^6).³

³ Even a Rubidium reference is of limited use. A GPS referenced signal should have the same received phase at the same time each day, which won't be the case with even the most stable independent reference. Fortunately GPS locked transmitters are not too difficult to make.

Software

The software discussed here is **Clicklock** Version 3 (31 August 2007) by Con Wassilieff ZL2AFP. This software uses the G3PLX Clicklock algorithm, and provides a series of tools allowing analysis of GPS synchronous signals. The main output is a pair of graphs plotting signal level (in arbitrary dB) and signal phase, measured in degrees. These two graphs (red and blue traces shown in Figure 3) share a common time base, so you can assess the signal phase and amplitude together, just as you would with a filter analysis tool. This information is extremely interesting, as after all, the ionospheric mechanism is not dissimilar to that of a filter.

There are three areas of interest in the software window. See Figure 3 (below).

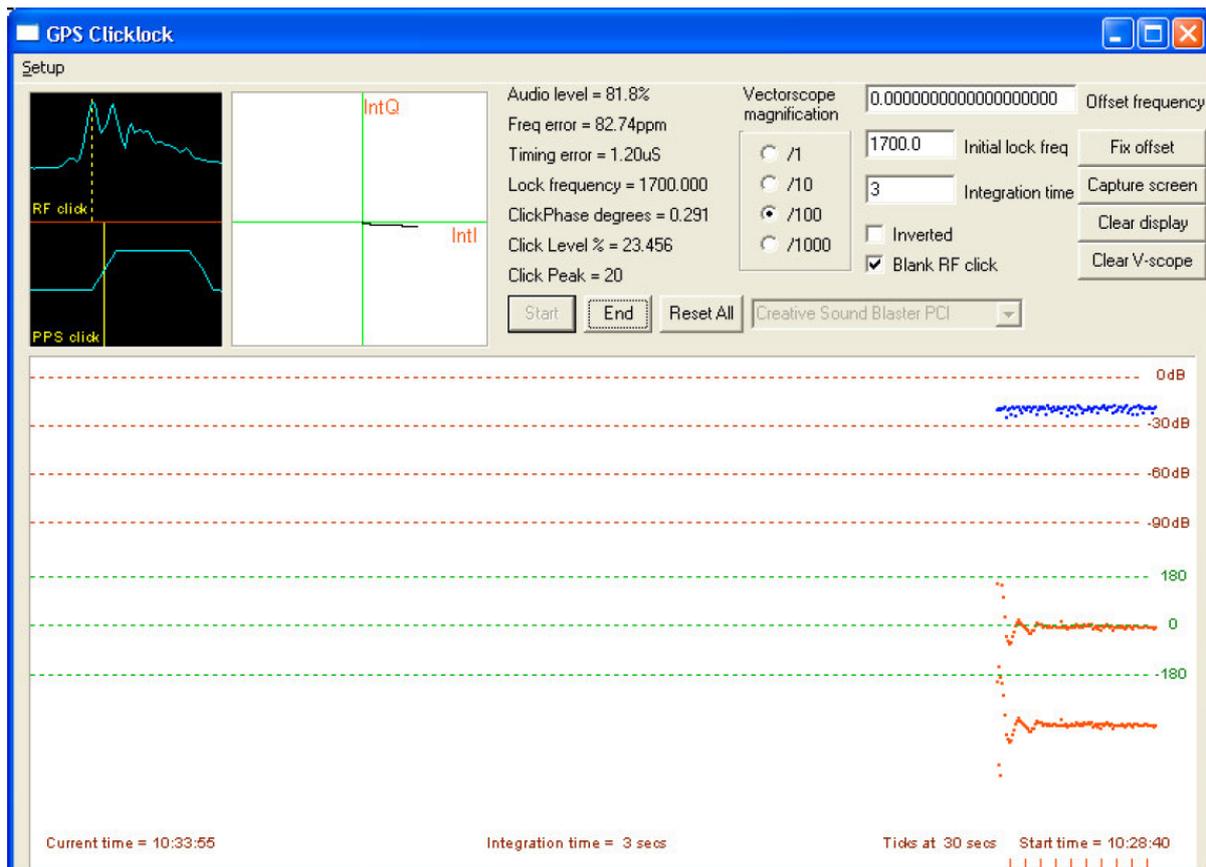


Fig. 3. The Clicklock Screen

Upper Pane

The upper third of the screen shows an oscilloscope display (left), an integrating vectorscope or phasor diagram (next), and the control panel (right). The lower trace of the oscilloscope shows the audio pulse (the **PPS Click**) fed into the sound card from the GPS receiver. The width of the pulse (in this case) is about 1ms. The software finds and synchronizes to the leading edge of the pulse, indicated by the yellow line through the rising edge.

The upper trace shows the receiver audio level, and the ragged **RF Click** is clearly seen. The user places the mouse on the first peak and presses the left button to indicate to the software where the RF peak reference point is. This point is then indicated by the dotted line.

The integrating vectorscope (developed by Peter G3PLX) may be a little difficult to understand, but builds up a plot based on adding measurements of in-phase (I) and quadrature-phase (Q) measurements of the received signal over time. The I and Q information is separately integrated

(summed over time), and continues to accumulate until the plot of Int(I) vs Int(Q) runs off the edge of the plot, the vectorscope sensitivity is changed, or the vectorscope is cleared (the 'Clear Vectorscope' button, bottom right of the control panel section).

'Vectorscope Magnification' radio buttons, in the centre of the control panel, allow the sensitivity of the vectorscope to be set. In the setting '1', the current accumulated Int(I) and Int(Q) values are plotted. In the other modes, the values are divided by 10, 100 or 1000 before plotting, so the graph can accommodate longer periods of integration. Because the vector length is a measurement of the signal power over time, the vectorscope overflows more quickly on strong signals.

If the signal being measured is off frequency, a circle or spiral will be described on the vectorscope. The radius of the curve is dependent on frequency difference as well as signal strength, while the period of rotation is dependent on the frequency difference. If the plot describes a circle in ten seconds, then the signal is 1/10 or 0.1Hz off frequency.

Lower Pane

The lower two thirds of the screen contains the power (signal level) and phase plots. Signal strength is plotted in blue, against a 90dB scale. The scale is arbitrary, but 0dB represents a reasonably strong signal. Typically signals down to about -20dB on this scale can be heard by ear, and down to -30dB can be seen on ARGO. The plots move left one sample at a time, and one dot is plotted per sample. The time scale for this display is determined by the integration time (one dot per integration period), and since the plots are about 450 pixels wide, the visible duration of the plot is about 450 x the integration time. In Figure 3, the integration time is 3 seconds (the default value), and so the full plot duration is about 22½ minutes.

The phase plot is shown twice. This is so that when the phase wraps around at +180° to -180°, there is continuity of display. Without this feature it can be difficult to decide what the phase is doing if it is near 180°, or changes rapidly (as is often the case at sunrise/sunset).

At the very bottom of the lower pane are items that are stored when the plot is saved. Tick marks are placed every 10 samples, i.e. with a period of 10x the integration time. The text information includes the current time (time when plot is saved), the integration time (time between samples), the tick interval, and the start time of the plot (not the time at the start of the screen!)

Control Panel

Numerous items of information and various controls are presented in the top right control panel. Refer again to Figure 3. The first three buttons, bottom left, start the program, quit the program, and restart the program with all parameters reset. That's the easy bit! The Reset All button restarts the locking process, and must be used when the receiver frequency or bandwidth is changed, and if the Initial Lock Frequency or Offset Frequency (see later) are changed.

The 'Blank RF Click' tick-box should be left ticked. It enables the removal of samples containing the RF click from the phase and signal strength calculations. The 'Inverted' tick box allows LSB operation of the receiver⁴. Staying with USB is recommended to begin with (leave 'Inverted' unticked), as the software operation is confusing enough without adding further complication!

The sound card selection drop-down will be greyed-out if you have just one sound card, but if you have two or more appropriate cards, you should be able to select the one that is to be used. The menu item Setup / Input calls up the sound card recording applet to allow the input to be selected and the level to be set.

⁴ It inverts the direction of the NCO control loop. If you're on the wrong sideband, the software will not lock.

In the middle of the upper pane are a series of measurement parameters. 'Audio level' is the receiver audio level, and a level near 100% indicates risk of overload. The level is best kept between 50 and 80%.

'Freq Error' is the overall system frequency correction made by **Clicklock**. The example (Figure 3) shows a very small error, 82.74ppm, due solely to the PC sound card sampling error, since the receiver used at the time was operating locked to a GPS Disciplined Frequency standard.

'Timing Error' refers to the instantaneous measured difference between the most recent RF click and the long term averaged time reference known to the software.

'Lock Frequency' shows the actual frequency of the NCO used to correct for receiver frequency errors. In the example, there is no difference from the Initial Lock Frequency, although it hunts up and down ± 1 from time to time, and moves considerably while acquiring lock. With your receiver there could be several Hz difference from 1700Hz.

'Click Phase' indicates the phase of the most recent measured click relative to the long term phase reference derived from the PPS pulse. It should average close to zero.

'Click Level' is an indication of the level of the RF click in the receiver audio. It should be adjusted to be between 5% and 20%. Any less, and lock will be uncertain. Any more than this, and the receiver is at risk of overload, and of course the AGC (if used) will act to reduce sensitivity. To some extent the level may need to be adjusted depending on the level of the DX signal being measured. Make sure you have clicked on the exact peak of the RF Click.

'Click Peak' is the number of samples delay between the PPS audio click and the RF click, and must be set by the user by clicking on the peak in the oscilloscope window. The value will be reasonably constant for each receiver, IF filter and antenna, and the click point must be set whenever the 'Reset' button is pressed.

Towards the right of the controls section are several data entry boxes. The bottom one is 'Integration Time', which essentially controls the sensitivity of the system and inversely the graph speed. It can be any value (in seconds) upward from 1. For weak signals values from 10 to 100 seconds or more would be appropriate. This integration time interval also defines the speed of the graphs in the lower part of the window.

'Offset Frequency' is used (like an RIT control) to adjust the receiver to fractions of 1Hz, in order to receive coherent but non-Hz spaced signals. The value is in Hz, and is additive to the receiver frequency (in USB mode), so for example to receive a LORAN line on 100,010.01001001Hz (100kHz / GRI of 9990), set the receiver to 98,410Hz (received frequency – 1600Hz Initial Lock Frequency, see next) and the 'Offset Frequency' to 010.01001001Hz

For convenience, the 'Fix Offset' button to the right will reset the necessary parts of the program for the new offset, without requiring complete reset. You must use it when you change the offset.

'Initial Lock Frequency' is the starting frequency for the NCO phase-locked loop, and must be within 0.5Hz of the intended signal (otherwise the lock may end up 1Hz off). Because of limitations in the internal filters, this value should be between 1500 and 1900Hz, and must be integer. In the previous example, the receiver could be tuned to 98400Hz, the 'Initial Lock Frequency' to 16010Hz and the 'Offset Frequency' as before. Change this value to suit your receiver IF passband and to add small offsets required if your receiver will only move in 100Hz steps.

If the Initial Frequency is changed, or the receiver frequency or filter is changed, the 'Reset All' button must be used subsequently, to ensure correct operation.

'Clear Display' and 'Clear V-scope' buttons are also provided for convenience. The Display will not clear until the next sample is ready, and this may be some time if a long integration time is set. The 'Capture Screen' button saves the graph and parameters in the current directory in .BMP format.

Menu

The Setup menu drop-down has three item choices. 'Save I,Q to File' stores the values for each integration period to file along with a time stamp, in .CSV format. This is very handy for post analysis. 'Save screen to file' saves the main graph and legend text in BMP format.

The graph is also saved as XXXXXX.BMP in the current directory when the 'Capture Screen' button is pressed., where XXXXXX represents the time when the file was saved.

The graph and all parameters (complete application window) can also be saved using ALT + Print Screen to save to the clipboard, and then the image can be pasted into a graphics program or other program.

Signals to Receive

You will of course need a stable SSB receiver with LF capability. You will also need a specialized LF antenna. One of the simplest to build and get working is the PA0RDT Mini-Whip⁵.

Obviously if no suitable Amateur signals are available (and it's not often that they are!), it is best to use commercial GPS synchronous signals. There are numerous suitable commercial signals that can be used as test transmissions, but these may not be on a frequency of interest or in the required direction.

Suitable signals need to be CW or at least pulsed CW (such as time code transmissions). PSK, MSK and FSK signals are generally not useful. The best useful examples are standard time and frequency stations, such as JJY40 (40kHz), RTZ (50kHz), many stations including MSF, JJY60 and WWVB on 60kHz, HGB (75kHz), DCF77 (77.5kHz) and BBC Radio 4 on 198kHz. Sidebands of LORAN stations can also be used, but these are generally not on integer frequencies, and you need to know how to mathematically predict where they are. Interestingly, both tones of the FSK transmission from California on 135925 and 135975Hz are GPS synchronous.

Transmitting GPS Synchronous Signals

It might surprise you to know that it is quite practical to generate your own GPS synchronous transmission, or have a friend do it for you. The simplest way to generate a synchronous signal is to use a temperature compensated crystal oscillator (TCXO) at some easy multiple of the intended frequency (for example $137.5\text{kHz} \times 64 = 8.8\text{MHz}$ $182.5\text{kHz} \times 20 = 3.650\text{MHz}$), and divide down to the operating frequency. To achieve GPS phase synchronism, use the GPS 1pps signal leading edge to reset the frequency divider every second, and if the oscillator is carefully adjusted, the divider will generate a GPS synchronous carrier with low 1Hz phase noise sidebands. A very stable oscillator (TCXO or OCXO) will give best results. You could for example use a 10MHz OCXO and divide by 72 in order to experiment on 138.888...kHz.

It is not enough to just use a Rubidium Standard or GPS disciplined reference (say 10MHz divided by 72 as suggested), because you cannot guarantee the carrier phase angle, and the listening stations will not be able to make phase comparisons with each other or with previous transmissions. If you use

⁵ See <http://www.radiopassioni.it/pdf/pa0rdt-Mini-Whip.PDF>

the suggested 1pps reset to the divider, you will generate the same phase angle every time – even if you do not use the equipment for a week.

The reset pulse must be very narrow – 30ns or so would be best. This can be achieved easily using a CMOS ‘race’ circuit:

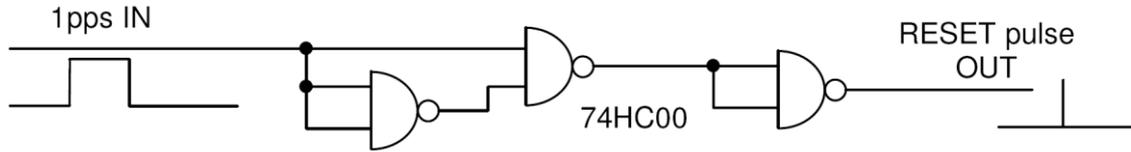


Fig. 4. A suitable 1pps synchronizing circuit

Another way generate a suitable GPS synchronous signal is to use a Direct Digital Synthesizer (DDS) which can have the phase synchronized by the GPS 1pps event. The ZL1BPU LF Exciter⁶ also has this ability. It includes a stable reference, milliHz frequency resolution, GPS synchronization, and the ability to set the exact transmitted phase angle at will, for signalling and ID purposes.

⁶ See www.qsl.net/zl1bpu/MICRO

Setting Up Clicklock

There are two phases to a successful setup, the radio and computer connections, and the software itself. This is tricky stuff, so read carefully and digest thoroughly! The description refers specifically to Clicklock3.

Equipment

You need a VLF/LF/MF receiver and antenna, and a reasonably fast computer with sound card operating at least Windows XP.

You need a source of GPS 1pps (1Hz pulses). The pulses should be at least 1V high, have a rise time better than 100us and ideally a pulse width between 10us and 2ms. A GPS Disciplined Reference such as the HP Z3801A will do nicely, but many GPS receiver 'timing engines' are suitable – UBLOX Antaris or LEA-5, SiRF SiRFStar 3 or newer, Rockwell/Navman Jupiter or Canadian Marconi Superstar for example.

You need to build a Click Interface Box (see next page).

Ideally you should inject the RF Clicks into the antenna using an adjustable auxiliary probe antenna, although it is much more convenient to inject directly at the Click Interface Box. If you do this, you will also probably need a variable or stepped attenuator (10dB per step will do) between the antenna and Click Interface Box, in order to prevent the antenna signal from overloading the click harmonics.

Connections

This description essentially describes how to connect up the Click Interface Box.

Start with a stereo audio cable with 2.5mm plugs at each end. Plug one end into the PC sound card LINE IN socket. Connect the other end of the cable to the Line In socket of the box. The box includes a line transformer to provide isolation.

Connect the GPS Receiver 1pps output (should be TTL or CMOS level, 0V with pulses to 3.3V or 5V) to the GPS 1pps input of the box. Set the sound card (via Adjust Audio Properties) to use the LINE input, and set the record level to about ¼.

The box also injects the 1pps GPS signal into the receiver antenna output. This is suitable for use with an active whip, or loop antenna with preamp, although purists will tell you to use a probe or link antenna near the receiving antenna in order to remove antenna phase errors. (This will be important if you use a tuned antenna such as an active loop).

The series diode and resistors in the box define the RF Click harmonic level, but be prepared to adjust the resistors to get the level correct.

Connect the active antenna to the From ANT socket on the box, and the To RX socket to the receiver antenna input. It can also be helpful to include a switched attenuator in the antenna preamp output to allow better adjustment of the click-to-signal ratio. With the RF Click injected at the box, I find it useful to reduce the antenna output while setting the RF Click level.

If using a loop antenna, fashion a loop of wire to drive from the GPS 1pps output via a 1k resistor, and couple the loop of wire into the antenna enough to achieve a reliable click.

The easiest way to achieve all these complex connections is to build a little interface box, for example using the following circuit (Figure 5, next page):

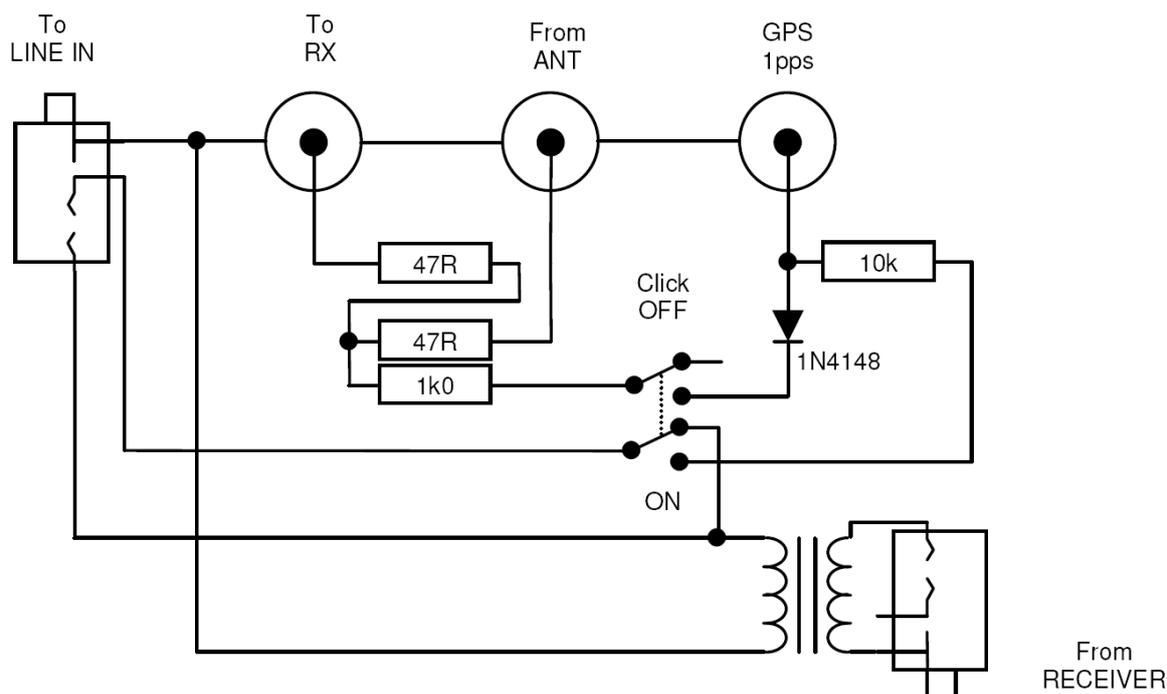


Fig. 5. A suitable click and signal injection interface box

This interface allows normal receiver and sound card operation when the switch is in the OFF position. Use panel mount BNC connectors for the antenna connections, a panel mount PHONO connector for the GPS 1pps, and 2.5mm panel mount Stereo sockets for the receiver and sound card audio connections. A small plastic box will suffice to house the components.

Software Setup

Start with a signal on a known accurate frequency, for example JJY40 on 40kHz. At your first attempt choose a stronger signal, as these are easier to tune and give faster results. Make sure the GPS receiver is running and has a reliable fix (preferably at least 7 satellites in view 99% of the time). Check the 1pps output to ensure that the rise time is good (a few ns) and the pulse width is suitable (1ms positive going is ideal).

Having made the connections as described above, tune the receiver on USB 1.7kHz LOW, so the signal (if audible) makes a 1.7kHz tone in the audio output – for example, tune to 38.3kHz for JJY40 or 58.3kHz for WWVB. Switch on the clicks at the interface box. Set the receiver bandwidth as wide as possible and the AGC off (if you have a choice, that is!).

There is no advantage to using a receiver bandwidth less than 1.5kHz, and a narrow filter may make it difficult to see a recognisable RF Click on the software oscilloscope.

Start the software **Clicklock3.exe**. On the Control Panel, press the 'Start' button. Check that the **PPS click** oscilloscope view soon shows a reasonably clear pulse, perhaps about 2/3 of the height up to the red line. It should take a few seconds to settle with the yellow marker on the leading edge. If the level is too high or low, lock will be poor. The 'Audio Level' (shown at top of centre of display) should be between 50 – 80%. If necessary adjust the sound card record volume to suit. You may not see the falling edge of the pulse, if the GPS 1pps output has a pulse duration longer than about 5ms.

Press the 'Reset All' button, and you should observe the **PPS click** move slowly along horizontally, and settle with the centre of the up edge aligned with the yellow line.

Observe the **RF click** oscilloscope view. The click should be obvious, nicely rounded and stable, preferably not too ragged. If it is very ragged, use a wider receiver filter. You may need to attenuate the signal from the antenna to see it, especially at the upper end of the operating range (MF). Adjust the receiver audio level to give a suitable height of pulse. Audio adjustment is something of a compromise since you need to adjust the computer audio input level to suit the Audio Click. Hence it's best to use receiver gain to set the RF Click and leave the computer audio level where it is.

You don't want the RF Click to be strong enough to activate the AGC or overload the receiver. If you can't even see the RF Click (or hear it in the receiver audio) with the antenna preamp disconnected from interface box, you may need to lower the value of the 1k resistor in the box (this is unlikely). To clearly hear the click with the antenna and preamp connected, you may need to attenuate the preamp signal, as the preamp output level is often quite high. Up to 30dB may be needed. You should be able to hear the receiver audio and the click from the receiver clearly at the same time, and 'Click level' should be about 5 – 20%.

Don't confuse the click from the receiver audio with the thump that you hear in the sound card from the audio pulse in the other channel! It is possible that you have the left and right connections reversed.

A good way to check the connections is to temporarily plug stereo headphones into the Line In socket on the box. You should hear the RF Click and receiver noise in one ear and the Audio Click in the other.⁷

When the **RF click** is clear and stable, move the mouse so the cursor is over the first rising edge of the **RF click** pulse, and click the left mouse button. The Click Peak will be set and a dotted line added to indicate the reference edge. The Click Peak value will typically be between 5 and 25 (depends on the receiver).

By now the receiver should be locking (look for a stable value for 'Lock Frequency'). Once this is stable, the Vectorscope and Chart plots begin to make sense.

Set the Integration Time appropriately (try 3 sec for a strong signal, 100 sec for a weak one). If the Vectorscope now starts to show a straight line, you may have a good coherent signal being recorded.

Adjustment of levels is something of a compromise. The sound card level and balance can be used, but it is best to independently adjust the level of the three input signals – the audio 1pps pulse, the RF click, and the receiver audio. A switched RF attenuator with 10dB steps, fitted between the antenna preamp and the click injection point and receiver will be very handy.

For the really advanced user – yes, you CAN run multiple instances of **Clicklock** on the one computer at the same time, and you can also run a Spectrogram as well (illustrated by Figure 1 and Figure 2, which were recorded together). If you run **Clicklock** with ARGO, you must start Clicklock first, or it won't operate correctly.

If you run two instances of **Clicklock**, you can for example record two different sidebands from the same signal, or record two adjacent signals. You can for example pick out the first two 1Hz sidebands of the JJY40 signal on 40kHz, and even see the stable phase relationship between the sidebands on either side of the carrier.

To record adjacent 1Hz sidebands, tune the receiver 1700Hz low, and start two incidences of the software, setting (for example) the Initial Lock Frequency of one to 1701 and the other to 1702Hz. The 1702Hz sideband should be about 3dB lower in level, but have the same phase. If you record at 1701

⁷ Headphones are not reliably L-R connected, or may be reversed on your head, so it's not possible to say which sound will be in each ear here. You may have to reverse the connections in your box to get the correct signals into the sound card.

and 1699Hz, you will see the first sidebands on either side of the carrier, and should clearly see their phase relationship, although they rotate together with ionospheric effects.

If you receive two stable signals close together on a Spectrogram, and wish to identify them, you can run a Clicklock session on each of them, and prove (a) whether they are either or both GPS synchronous, and (b) whether they are from the same source (will have almost identical fading and related phase patterns).

Picture Gallery

The following pictures include some recordings of AMATEUR signals on 181400Hz, and some were made with earlier versions of the software which lacked the double phase plot shown in Figure 3 and Figure. 7. For the Amateur transmissions the transmitting antenna was a modest top-loaded Marconi with an effective height of about 5m. This means a very low (but typical) overall antenna efficiency. The recordings of Amateur transmissions were made at a range of 500km.

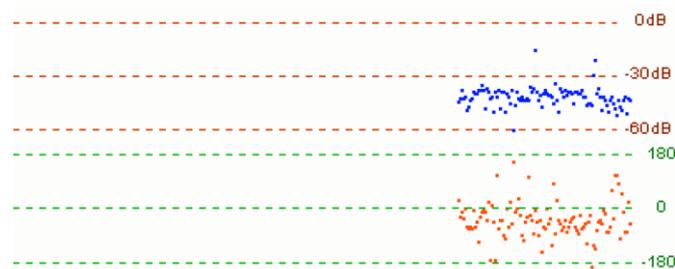


Fig. 6. 100mW from ZL1BPU on 181400Hz at 500km range

This first recording (Figure 6), of just the phase/power plots, shows the signal to be very weak – certainly not audible or detectable using ARGO - but was from a transmitted signal of only 100mW (perhaps 0.4uW EIRP), and recorded during the daytime at a noisy location! During the day, only ground wave propagation is possible, and the attenuation is high. The signal was received at a range of 500km! The transmitter was the ZL1BPU LF Exciter (running barefoot). 30 second integration was used at the receiver.

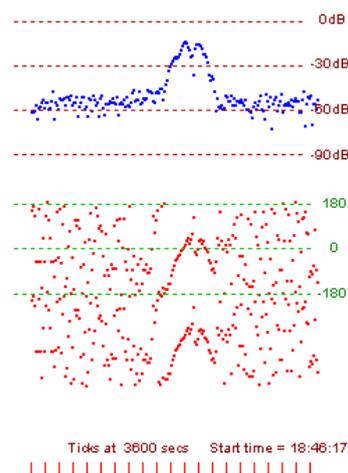


Fig. 7. DCF77 pops out of the noise at sunrise – 15,000km range!

Figure 7 is an example of the extreme sensitivity of the system with 360 second integration time. This is the German DCF77 transmission on 77.5kHz, received at 15,000km range. At the peak it is 30dB above the noise (but still barely detectable on ARGO at any speed).

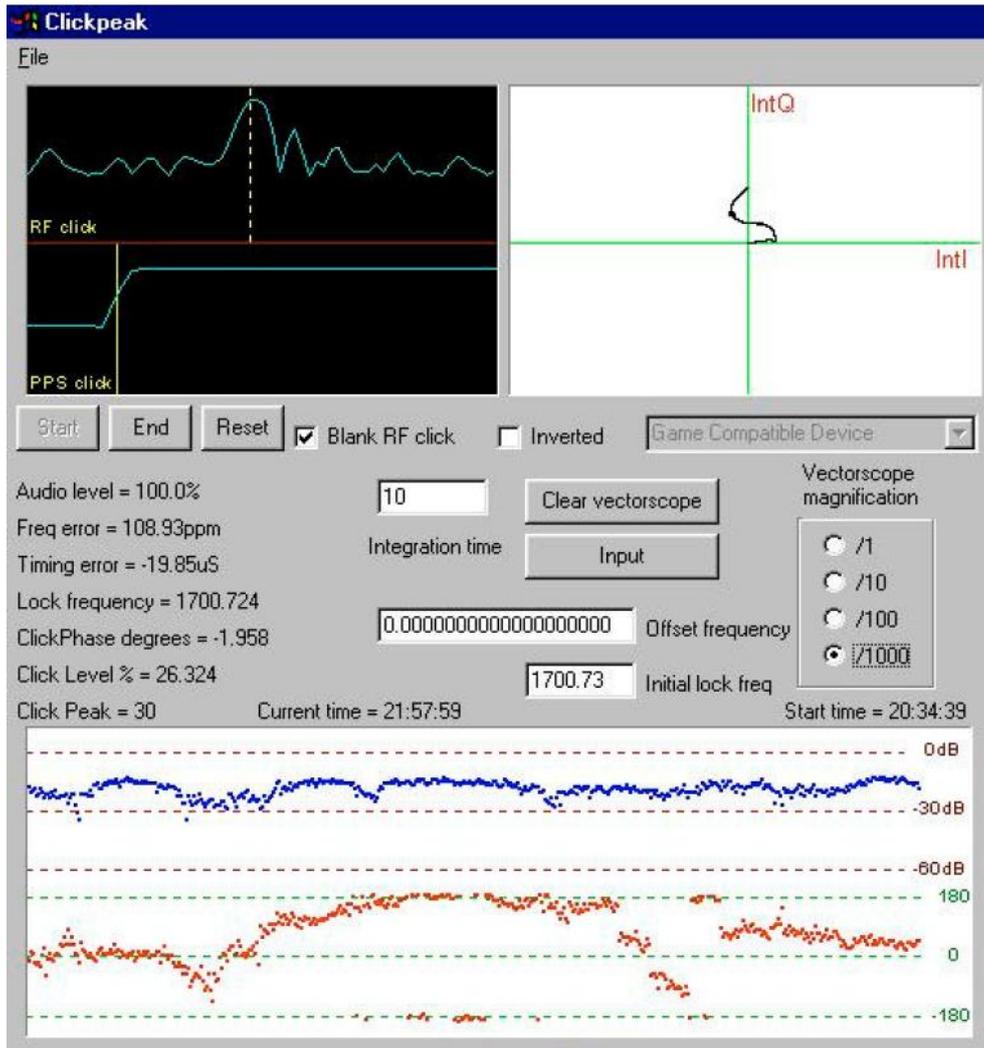


Fig. 8. 2W signal showing phase signalling

This next picture (Figure 8) was recorded from an old version of the software, but the main features are the same. It shows a much stronger Amateur signal (2W transmitted, EIRP about 8uW). The transmitter was the ZL1BPU LF Exciter with a power amplifier, adjusted for 2W output, range of reception 500km.

Because the received signal was (relatively) strong, it was recorded using 10 second integration. This transmission was during the evening. Some evidence of fading and phase shift following sunset is visible on the left of the graphs. About 2/3 of the way across the red graph is clear evidence of step phase changes, with no change in signal level. This was caused by deliberate advancement of the transmitter phase in four 90° steps, one every five minutes. This illustrates how it would be quite easy to send a phase-modulated identification pattern. Obviously the data rate (1 bit every 5 minutes) is too low to send much more!

The recording also shows how, even with perfect transmitter phase stability, the received phase can change quite markedly in a matter of 30 minutes (the length of the recording) due to ionospheric effects.⁸

⁸ Monitored during ground wave conditions, the ZL1BPU LF Exciter will record a dead straight line for days on end, and will also do so if power is removed and reapplied – perfect GPS synchronism. (It has to be said however that the phase of the Exciter signal is best not examined too closely with 1 second integration, as there is often a saw tooth phase pattern due to slight clock inaccuracy followed by regular GPS synchronization).

The vectorscope in Figure 8 was reset at the start of the recording, and shows a blob in the middle of an S-shaped curve. This blob was caused by the four-step phase change, as for a short period of time the signal would have been integrating in four different directions. The generally upward trend of the signal in the vectorscope is because on average across the duration of recording the phase was around 90°, and the S-shape because the recording started out at zero phase (you can see this at the left of the red trace on the graph in Figure 8), rose to 180°(curved back to the left on the Vectorscope), and after the intentional phase switching continued on at about 30°. All this you can see on the red graph, and is due to ionospheric changes.

As well as the obvious differences due to software version, this picture also shows two differences in the oscilloscope display when compared with Figure 3. First, there is no visible downward edge of the PPS click. This is because the Rockwell Jupiter GPS engine used for Figure 8 has a longer 1pps click duration, about 10ms, as opposed to 1ms in the CMC Superstar engine used for Figure 3. In the upper (RF click) trace, you see first that the click has much greater delay – this is a function of the receiver used. There is also only one peak to the pulse. Once again, this is because the downward edge of the pulse (which caused the second peak in Figure 3) is off the screen to the right. You should expect similar differences with your setup.

Figure 8 was captured from ZL2AFP's own computer running the WIN95 operating system and Clicklock1. Figure 3 was captured on a WINXP machine using Clicklock2. Clicklock3 looks much the same as Clicklock 2 – there are only internal improvements.

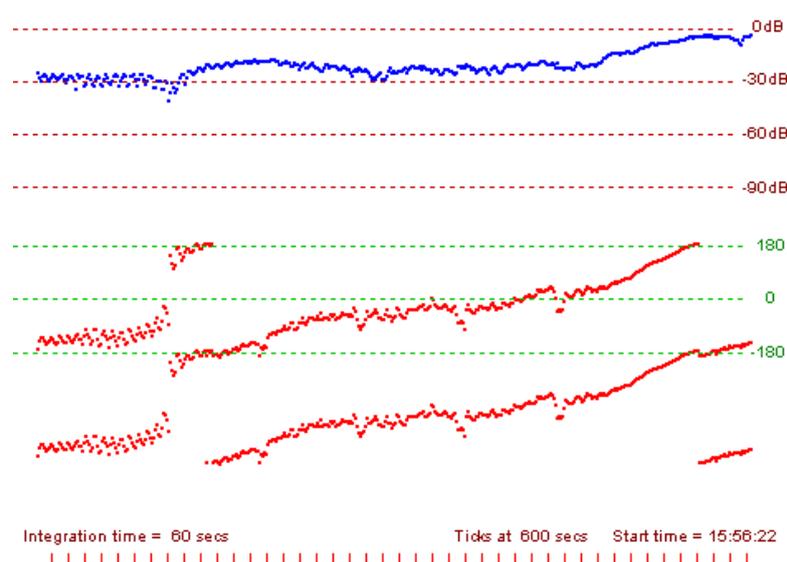


Fig. 9. Reception of WWVB at 10,000km range

Figure 9 is an overnight recording using 60 second integration at 60kHz. At least two signals are visible. The first (first quarter of display) is unknown, as there are several transmissions on 60kHz. However in the middle of the night WWVB becomes the strongest signal, and the rest of the graph shows the WWVB signal with increasingly strong reception and steadily increasing phase (implying shortening of the ionospheric path). Over the course of the night the phase advances about 180° (a reduction of 2.5km in the path length). The ticks every hour are a sure sign that this is WWVB – it is the only transmission on this frequency which uses a phase step to mark the hour.

Appendix

Comments by G3PLX

The Clicklock technique stems from my observation that if I connect the 1Hz pulse output of my GPS to my LF receiver, I can hear a repeating pair of clicks at 1 sec intervals which are the high-order harmonics of the fast edges of this pulse. Not only does this give me a precise time reference which I can use if I want to do precision off-air timing measurements, but if I feed the audio clicks to a phase-comparator, I can see the phase rotating slowly with time.

The rate of rotation is the 'beat-note' between my receiver frequency and the nearest whole Hz. This 1Hz, coming from the GPS module, is effectively derived from the most accurate source on the planet (or at least orbiting the planet). Not only is it stable in frequency but it can be used as a universal reference for phase measurement on any frequency. If I use that to lock the software, I can then demodulate incoming signals in such a way that ALL the residual frequency and phase drift of the receiver is cancelled completely. If I tune-in a signal which is also locked to GPS, and the propagation is stable, there will be no frequency error and no phase drift. At all. Ever.

Scott [VE7TIL] has [and Con ZL2AFP has also] implemented this so that the received signal is displayed on what we have christened an integrating vectorscope. This displays the signal phase and amplitude simultaneously on a circular display. With no signal, a dot appears in the centre of the display. If there is some noise present, the dot moves randomly. With the program set to receive on a specific GPS-locked frequency (which you can enter to as many decimal places of a Hz as you like), if there is a signal on that frequency, the dot will start moving off centre in a particular direction, this direction depending on the RF phase of the signal relative to the GPS reference. The weaker is the signal, the slower will the dot move, but, so long as the signal doesn't change its phase, there is NO LIMIT to how low in signal level you could go.

This is what we loosely refer to as 'coherent reception'.

For example, I could set the program to the frequency of a LORAN line from a local transmitter (<1000km), and the dot will move off towards the edge of the scope. The program is locked to GPS, the LORAN transmitter is stable in phase, so the direction of movement of the dot (the RF phase of the received signal) always stays the same. If I repeated the experiment on another day I would get exactly the same phase reading. I could have left the system on all day and detected a signal 24 times weaker than if I had left it running for an hour. If I move 1/4 wavelength closer to the transmitter I would see the phase change by 90 degrees, even if I switched-off the receiver and the computer completely during the move.

This opens up an awesome set of possibilities for really weak-signal reception. Many transmitting stations are already able to transmit phase-locked to GPS. This technique means we can explore all the possibilities for coherent transmission and reception, with just an LF receiver which is already stable to 1Hz, and a GPS module with a 1Hz output.