

Clicklock

LF Propagation Measurement Tool

M. Greenman ZL1BPU
April 2006

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

The technique demonstrates greater sensitivity than 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 signals can't 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 PHASE.

This technique is not practical on HF since the 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 phase shift is high. The upper limit of practical measurement is about 500kHz.

Instead, two things can be usefully measured - signal strength, and signal phase. If the phase of the transmitted signal is known, or at least constant, the time of flight of the signal can be measured directly by measuring the phase difference from source to destination. To do this, a common phase reference is needed.

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 something 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 phase which depends on the reflection height and the mechanism (D-layer, E-layer) involved. The received signal will be a combination of these signals.

As mentioned, at LF, 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 determine whether the signal is a real signal, or some artifact of the receiver. For example, Fig. 1 shows a carrier near

¹ See notes by G3PLX in Appendix

20kHz (receiver was tuned to 18.4kHz USB, ARGO shows the audio tone to be about 1600Hz).

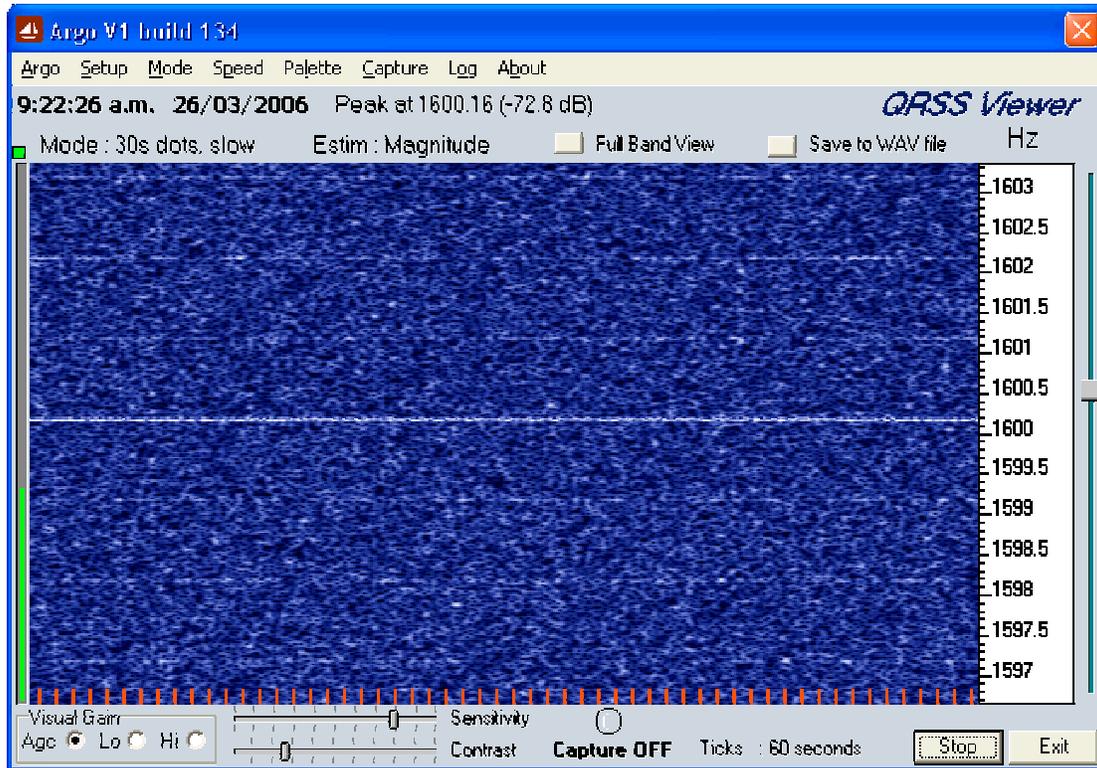


Fig. 1. ARGO trace of signal near 20kHz (at 1600.2Hz)

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 in the receiver 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. Even amplitude information is difficult to discern, and the picture is peppered with noise.

In the next picture (Fig. 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 coherent with the receiver, and 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.

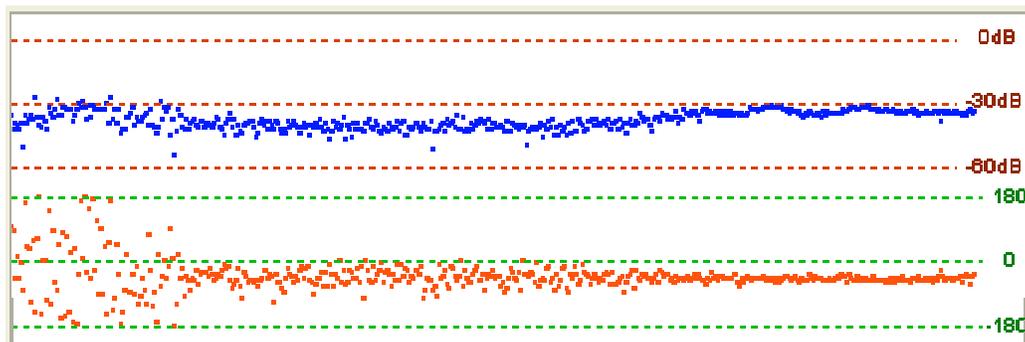


Fig. 2. **Clicklock** Power (blue) and phase (red) plots at 20kHz.

It is 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 is from a low power experimental scientific station at the South Pole! Even though the range is about 8000km, reception is apparently surface wave (as the phase is constant).

In summary, on LF, instead of conventional HF Doppler or radar techniques, using **Clicklock** two things can be usefully measured - signal strength, and signal phase, and non-coherent signals are rejected. If the phase of the transmitted signal is known, the time of flight of the signal, or at least changes in the time of flight, can be measured directly by measuring the phase difference from source to destination. To do this, a common phase reference is needed.

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 of even 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.

In addition, there can be phase variations that creep into a receiving system as a result of tuning variations in the antenna system, and different propagation times in the receiver on different frequencies, with 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.

Clicklock is a PC sound card based system, and uses two versions of a universally available reference, the 1pps (1Hz) pulses 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.

The 1pps pulse from the GPS receiver is fed directly to the PC sound card (the PPS Click), and forms the main software reference. The software uses this pulse to discover the exact location of a second event, the RF click.

This second version of the 1pps pulse is fed into the receiver antenna, where RF harmonics of the 1pps pulse every Hz throughout the LF spectrum are present, 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 harmonics in the receiver bandwidth will only be the same if the receiver is exactly on frequency, and at this point their absolute phase is also known by comparison with the audio click. If the frequency is slightly off one way, the phase of the harmonics will 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 down to zero frequency. This oscillator operates at the designated centre frequency (called the lock frequency). 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. In other words the whole receiving system becomes phase locked by phase locking the NCO. When the receiving system has locked, the receiver is ready to look for signals are coherent with the receiving system (on the same locked frequency).

The receiver signals come in 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. The

signal can be extremely weak, well under the noise, but by adding the samples together over a long time period the signal can be enhanced and the noise cancelled.

This property of essentially limitless enhancement only applies when the received signal is in exact carrier synchronism with the receiver. 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, are non-coherent with the receiver and are cancelled out.

Coherent Reception

Coherent communication systems have been claimed in the past, but in reality they were at best time synchronized at bit level, NOT carrier phase and frequency synchronous. Repeated data can only be integrated to a limited extent, since 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 intervals of time, and 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 fractions of a milliHz - and the stability to match.

The sensitivity offered is better than using a Spectrogram (such as Argo) - possibly by a factor of 20dB or more! This means that signals at least 40dB weaker than can be heard by ear can be detected reliably. It makes around the world propagation measurements possible, and allows measurements at up to 1000km range to be made with Amateur sized antennae and low power transmitters.

With really long integration times 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 coherent? (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).

Software

The software discussed here is **Clicklock** Version 2 (3 April 2006) by Con Wassilieff ZL2AFP. This software uses the G3PLX Clicklock algorithm, and provides a series of tools allowing analysis of GPS coherent signals. The main output is a pair of graphs plotting signal level (in arbitrary dB) and signal phase, measured in degrees. These two graphs (shown in Fig. 2) share a common timebase, 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 a filter.

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

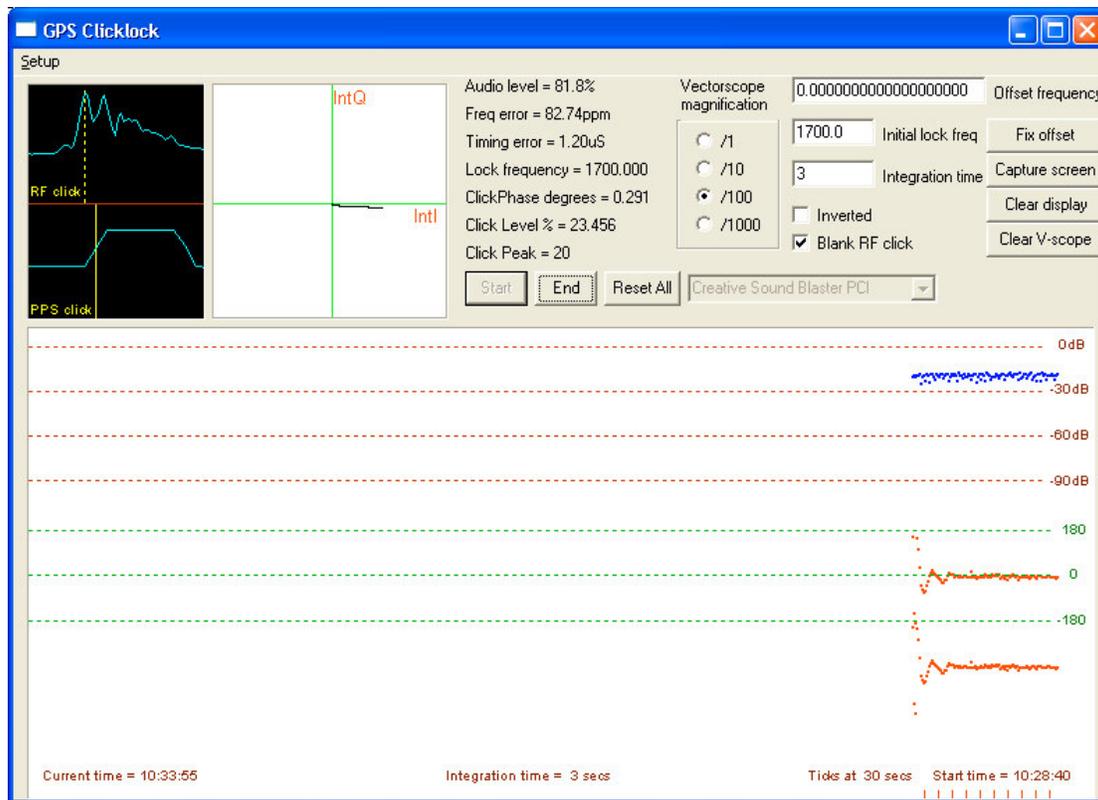


Fig. 3. The **Clicklock** Screen

Upper Pane

The upper third of the screen shows an oscilloscope display (left), a vectorscope (next), and the control panel (right). The lower trace of the oscilloscope shows the audio pulse 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 RF click is clearly seen. The user places the mouse on the 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 vectorscope is a little difficult to understand, but builds up a plot based on measurements if in-phase and quadrature phase measurements of the received signal over time. The information is 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 below).

'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, the vectorscope overflows more quickly on strong signals.

If the signal being measured is off frequency, a circle will be described on the vectorscope. The radius of the circle is dependent on frequency offset 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 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 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, and since the plots are about 450 pixels wide, the visible duration of the plot is about 450 x the integration time. In Fig. 3, the integration time is 3 seconds (the default value), and so the 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 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 information is the current time (time when plot is saved), the integration time, the tick interval, and the start time of the plot (*not* the time at the start of the screen!)

Control Panel

Numerous information and controls are presented in the top right control panel. Refer again to Fig. 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 is changed, 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 receiver. Staying with USB is recommended ('Inverted' unticked), as 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 appropriate cards, you should be able to select the one that is 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.

Down the left side of the centre 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 (Fig. 3) shows a very small error, 82.74ppm, due solely to the PC sound card sampling error, since the receiver used was operating from a Rubidium standard. 'Timing Error' refers to the instantaneous measured difference between the most recent RF click and the long term averaged time known to the software.

'Lock Frequency' shows the actual frequency of the NCO used to correct for receiver frequency errors. In the example, with a Rubidium reference, 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. '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 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, and the receiver AGC is at risk of overload, and of course the AGC 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 and antenna, and must be set whenever the 'Reset' button is pressed.

In the middle of the centre pane are several data entry boxes. The bottom one is 'Integration Time', which essentially controls the sensitivity of the system and inversely the 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 bottom third 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 additive (in USB), 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, a 'Fix Offset' button will reset the necessary parts of the program for the new offset, without requiring complete reset.

'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.

If the Initial Frequency 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 '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. 'Save screen to file' saves the main graph and legend text in BMP format².

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 can also be saved using **ALT + Print Screen** to save to the clipboard, and then then image can be pasted into a graphics program or other program.

Signals to Receive

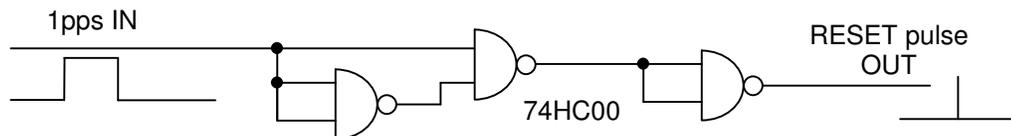
Obviously it is best to use signals that will teach you something about the propagation path you have an interest in, and so should originate at or near the planned transmitting site, on or near the frequency of interest, and of course be GPS coherent.

You will of course need an 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-Whip.

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 examples are standard time and frequency stations, such as JJY (40kHz), RTZ (50kHz), many stations including MSF 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 predict where they are. The FSK transmission from California on 135925 and 135975Hz is GPS coherent.

² I've not yet worked out where the files are saved!

It is quite practical to generate your own GPS coherent transmission, or have a friend do it for you. The simplest way to generate a coherent signal is to use a crystal oscillator 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. Use the GPS 1pps signal leading edge to reset the divider every second, and if the oscillator is carefully adjusted, the divider will generate a coherent carrier with low 1Hz phase noise sidebands. The reset pulse must be very narrow – 30ns or so would be best. This can be achieved easily using a CMOS ‘race’ circuit:



Another way generate a suitable signal is to use a Direct Digital Synthesizer (DDS) which can have the phase synchronized by the GPS 1pps event. The ZL1BPU LF Exciter³ has this ability. It includes milliHz frequency resolution, GPS synchronization, and the ability to set the transmitted phase angle at will.

Unfortunately it is not enough to just use a highly accurate reference (such as an Oven oscillator or even a Rubidium Standard) for the transmitter divider, as although the phase may stay fairly constant for minutes (OCXO) or hours (Rubidium), it will never be possible to predict the transmitted phase or infer anything from the received phase.

Setting Up *Clicklock*

There are two phases to a successful setup, the radio and computer connections, and the software itself. This is tricky stuff! The description refers specifically to **Clicklock2**.

Connections

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 receiver audio to the TIP of the other end of the cable, using the sleeve as ground. Preferably use a line transformer to provide isolation.

Connect the GPS Receiver 1pps output (should be TTL level, 0V with pulses to 5V) to the ring of the audio cable via a 10k resistor, using the sleeve as ground. Set the sound card RECORD applet to use the LINE input, and set the record level to about $\frac{1}{4}$.

Then inject the 1pps signal into the receiver antenna. If using an active whip, or loop antenna with preamp, inject into the coax feed to the receiver using a 50 Ohm tee. Use a series diode (anode to GPS receiver) and resistor of about 1k, and be prepared to adjust the resistor to get the level correct. It can also be helpful to include a switched attenuator in the preamp output to allow better adjustment of the click-to-signal ratio.

If using a passive 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, using the following circuit (next page):

³ See www.qsl.net/zl1bpu/MICRO

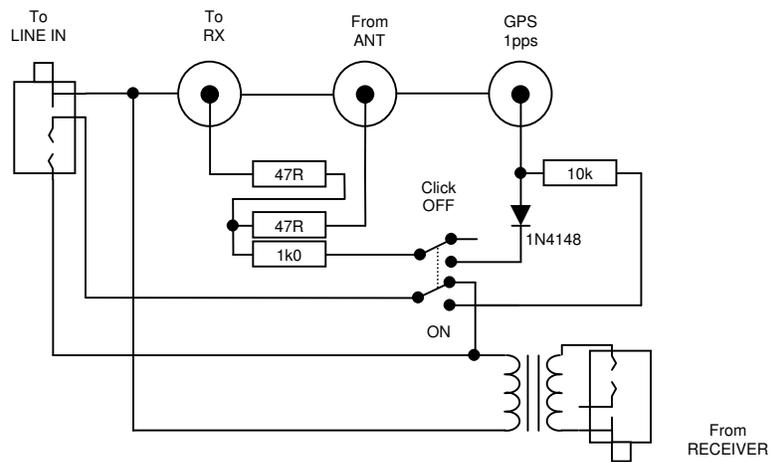


Fig. 4. Suitable Clicklock Interface

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 JJY on 40kHz. Make sure the GPS receiver is running and has a reliable fix (preferably at least 7 satellites in view 99% of the time). 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. Switch on the clicks. Set the receiver bandwidth as wide as possible and the AGC off (if you have a choice, that is!).

Start the software **Clicklock2.exe**. On the Control Panel, press the 'Start' button. Check that the PPS click oscilloscope view shows a reasonably clear pulse, perhaps about 2/3 of the height up to the red line. If the level is too high, lock will be poor. The 'Audio Level' should be between 50 – 80%. 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 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. You may need to attenuate the signal from the antenna to see it. You don't want the click to be strong enough to activate the AGC or overload the receiver. If you can't even see the click (or hear it in the receiver audio) with the antenna preamp disconnected from interface box, you may need to lower the 1k resistor (this is unlikely). To clearly hear the click with the 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 with the thump that you hear in the sound card from the audio pulse in the other channel!

When the RF click is clear and stable, move the mouse so the cursor is over the first rising edge of the 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 10 sec for a strong signal, 100 sec for a weak one). If the Vectorscope now shows 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 use – 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 Fig. 1 and Fig. 2, which were recorded together). 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 two 1Hz sidebands of the JJY signal on 40kHz, and even see the 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, and set the Initial Lock Frequency of one to 1701 and the other to 1702Hz. The 1702Hz sideband should be about 3dB power in level, but have the same phase. If you record at 1701 and 1699Hz, you will see the first sidebands on either side of the carrier, and should clearly see their relationship, although they rotate together with ionospheric effects.

If you receive two 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 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 are recordings of AMATEUR signals on 181400Hz. Some of the pictures were made with earlier versions of the software. 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, and an EIRP for 1W transmitter output of about 4uW! The recordings were made at a range of 500km.

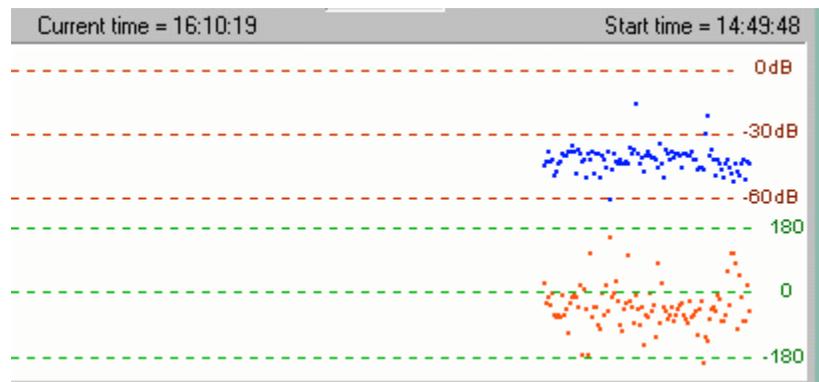


Fig. 5. 100mW on 181400Hz at 500km range

This first recording (Fig. 5) shows the signal to be very weak – but was from a transmitted signal of only 100mW, and recorded during the daytime at a noisy location! EIRP of 0.4uW received at a range of 500km! The transmitter was the ZL1BPU LF Exciter (running barefoot). 30 second integration was used.

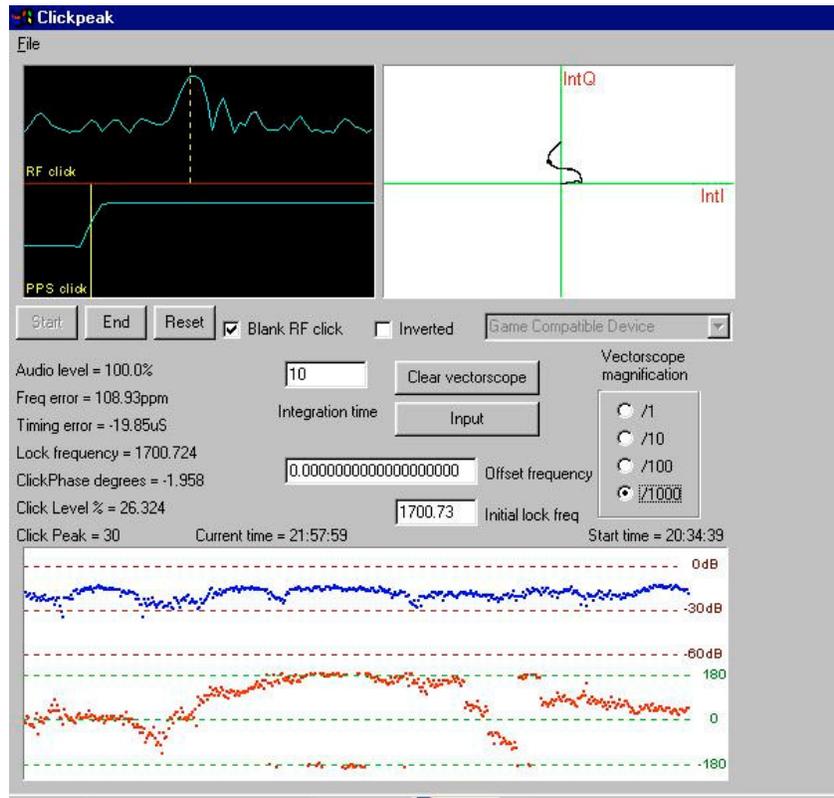


Fig. 6. 2W signal showing phase signalling

This next picture (Fig. 6) shows a much stronger signal (2W transmitted, EIRP about 8uW). The transmitter was the ZL1BPU LF Exciter with a power amplifier, adjusted for 2W output.

Because the received signal was strong, it was recorded using 10 second integration. This was recorded 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 graph is clear evidence of phase change, with no change in signal level. This was caused by deliberate advancement of the transmitter phase in 90° steps every five minutes. This indicates 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 shows how the received phase can change quite markedly in a matter of 30 minutes (the length of the recording). Recorded under ground wave conditions, the ZL1BPU LF Exciter will record a dead straight line for days on end – 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 a sawtooth phase pattern due to clock inaccuracy followed by regular GPS synchronization).

The Vectorscope in Fig. 6. was reset at the start of the recording, and shows a blob in the middle of the 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 different directions. The generally upward trend of the signal in the vectorscope is because on average the phase was around 90°, and the S-shape because the recording started out at zero phase (you can see this on the graph below), rose to 180° (curved back to the left on the Vectorscope), and after the intentional phase switching continued on at about 30°.

As well as the obvious differences due to software version, this picture also shows two differences in the oscilloscope display when compared with Fig. 3. First, there is no visible downward edge of the PPS click. This is because the Rockwell Jupiter GPS engine used has a longer 1pps click duration, about 10ms, as opposed to 1ms in the CMC Superstar engine used for Fig. 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 – and there is only one peak to the pulse. Once again, this is because the downward edge of the pulse (which caused the second peak in Fig 3) is off the screen to the right.

The different appearance of the software (figs. 5 and 6) is because this screenshot was captured from ZL2AFP's own computer running the WIN95 operating system and Clicklock1. Fig. 3 was captured on a WINXP machine using Clicklock2.

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 has [and Con has also – ed] 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 it's 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.