The Chirps Project: a new way to study HF propagation

In Radio Communication May 1998 I described how, by using DSP-based narrow spectrum analyser techniques, useful information about propagation modes can be obtained by studying the Doppler shift imparted to the signals reflected from moving surfaces such as meteor trails, ionospheric layers or even aircraft in flight. This technique provides a method by which anyone with an SSB receiver tuned to a suitable unmodulated carrier can display "dopplergrams" which show the presence of the reflecting surfaces, so long as they are moving.

Dopplergrams need movement.
The need for movement is an essential part of this technique: it is not possible, for example, to detect the presence of two separate reflecting layers if there is no relative movement between them. Further, since the amount of Doppler shift is a function of the angle-of-incidence, there is virtually no Doppler shift for low-angle reflection from the ionosphere, so this technique is not useful for long distance HF paths. For example, a dopplergram of an unmodulated carrier from a point nearly halfway round the world shows no sign of the presence of the two separate long and short path signals. This article describes a completely new technique for studying HF propagation which overcomes this problem and opens up many new avenues for exploration of the HF bands. Like the use of existing unmodulated carrier for dopplergrams, this technique borrows existing HF signals to make the details of the propagation visible.

I realised that for long distance paths on the higher HF bands, what was needed was a method of displaying the time-domain structure of a signal, not the frequency-domain structure as in the dopplergram technique. A few moments study of the theory showed that this would mean using some kind of wideband signal to probe the path rather than the "narrow" unmodulated carrier used in the dopplergram technique.

Pulses instead of carriers
A pulse is the obvious choice for the probe signal. If I could arrange for a friendly distant transmitter to emit a regular pulse signal and feed the received signal to an oscilloscope, the presence of two received pulses would indicate directly the presence of two separate paths, whether they were moving or not. More study of the theory showed that if I was using a standard SSB receiver with a bandwidth of 3 kHz, I might be able to just distinguish pulses that were separated by 0.33mS (the reciprocal of 3 kHz). This is just enough to separate reflections from, for example, the E and F layers of a near-vertical-incidence path but not enough to resolve the structure of propagation anomalies with smaller path differences than this. However, if I am just interested in studying HF propagation to understand how it affects the signals I normally receive in a 3 kHz bandwidth, I won't be missing anything if that's the bandwidth I use for my experiments.

I could certainly have tried pulses if I could have persuaded another amateur at a suitable distance to transmit such a signal. However, what I was hoping was that there might be some suitable HF signal already in common use which could be "borrowed" for this purpose, in the same way that unmodulated carriers which are plentiful on the HF bands were borrowed for the dopplergram technique. I tuned the bands looking for pulse-type signals that might be hijacked for the job, but it seemed that there was nothing suitable. Even the "Woodpecker" was nowhere to be found!

The phantom VFO-swisher unmasked
This idea would have remained just as a pipe-dream, but I stumbled on the breakthrough by chance soon after moving to a new location a few years ago. I had left the HF receiver on 28 MHz but came back into the shack late in the evening after the band had closed. I noticed a periodic sound like someone swishing a VFO across the band very fast and realised that this must be local to me and decided to investigate. By tuning to different frequencies and noting the precise times at which I heard the signal, I deduced that this was a carrier that was sweeping in one continuous scan from 2 to 30 MHz at a rate of 100 kHz/sec, repeating every 5 minutes.

Then I remembered that I had read about this somewhere before. This is a swept-frequency ionospheric sounder, sometimes known by the trade name "Chirpsounder". Radio researchers and professional HF operators use these for plotting the height and cut-off frequency of the ionosphere. The transmitter runs a few tens of watts into a broadband omnidirectional antenna and the return signal is received separately. A narrow-bandwidth receiver is swept up the band in step with the transmitter so that the direct signal is always zero-beat in the receiver. Think of it as a direct conversion receiver using the same swept oscillator that drives the transmitter. A signal reflected from the ionosphere will have originated from a slightly earlier point in the transmitter sweep because of the propagation delay and so give rise to a beat-note in the receiver. The beat-note frequency is a function of the propagation delay and the sweeprate, and can be in the range 0-500Hz for propagation delays of 0-5mS, which correspond to ionosphere heights from zero to 750km. A suitable spectrum analyser connected to the receiver output thus displays an "ionogram", a chart of the ionospheric reflection height over the band between 2 and 30 MHz.

Further research revealed that the signal I was hearing was about 50km away near Preston, Lancashire, and once I had recognised the sound, I realised that there were many more of them around the HF bands. Andy Talbot G4JNT sent me Fig 1 which shows a "waterfall" spectrum display of part of the HF band received in the south of England. Some marine coast stations can be seen on the right sending data and Morse code. The diagonal line on the left is a sweep sounder. A helpful research scientist provided a list of about 30 sweep sounders world-wide but indicated that it was incomplete. At this point it became clear that here was the missing wideband signal for my pipe-dream system and the Chirps Project was born.

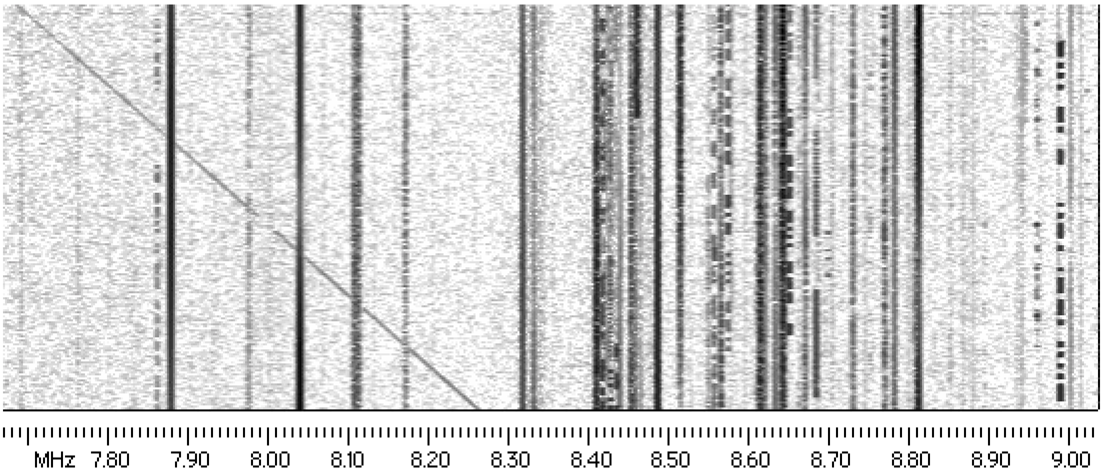

Fig 1. Spectrogram showing sweep sounder passing 8 MHz. The vertical axis represents about 5 seconds.

The magic chirp filter

A project to build a sweeping HF receiver was not something to be taken lightly, and I wondered what could be done with an ordinary SSB radio. If the receiver was left on a quiet frequency in the USB mode, it would produce a brief upward-sweeping tone when the sounder went past. I decided to build a tuneable audio filter that could be swept in step with the chirping tone. This was actually rather easier to do than it sounds and had some almost magical properties which transformed the pipe-dream into a working system.

To explain how this was achieved, consider first a circuit consisting of an audio delay-line. It's going to be a DSP delay-line eventually, but let's start by thinking of it as an analogue delay-line. Taps along this delay-line each connect to a potentiometer and the signals from all the potentiometers are summed together. To start with we turn all the pots to the top. In DSP-speak we say that the tap coefficients are set to maximum. The output will be proportional to the running-average of the input signal over time and that makes it a kind of low-pass filter. The cut-off point occurs when the input frequency is such that one whole cycle of the input signal just fills the delay-line and at that point the output sums to zero and there is a null in the output response. Fig 2 shows a block diagram of such a filter. Notice that if we feed a very short pulse into this filter, the output (the impulse response) is a square pulse with the same duration as the length as the delayline. Filters like this are easy to implement digitally and are sometimes called Finite Impulse Response (FIR) filters.

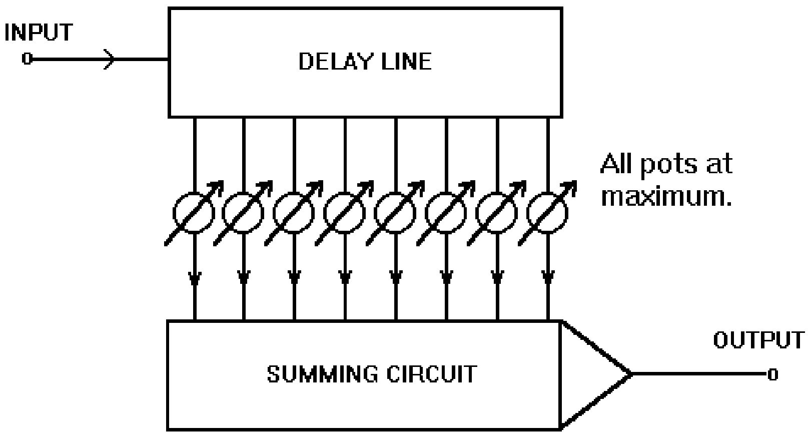

Fig 2. Simple low-pass FIR filter with delayline with a potentiometer on each stage and a summing circuit to form the output.

To make a bandpass filter centred on a frequency of, say, 1000Hz, two such filters can be combined in a circuit such as Fig 3, in which the input signal is first converted from 1000Hz down to DC in a pair of mixers fed with 1000Hz oscillators which are 90 degrees out of phase. That is, one oscillator is a sinewave and the other is a cosine waveform. The outputs of the two mixers are each passed through a low-pass FIR filter and then through two more mixers which convert the signal back up to 1 kHz. The final output is the sum of both channels. This is reminiscent of a third-method SSB system, in which the image signals of the first pair of mixers cancel out in the second pair. In this case the same frequency is used for both pairs of mixers so it's just a filter with no frequency translation, but one that can be tuned around the audio band simply by tuning the oscillator.

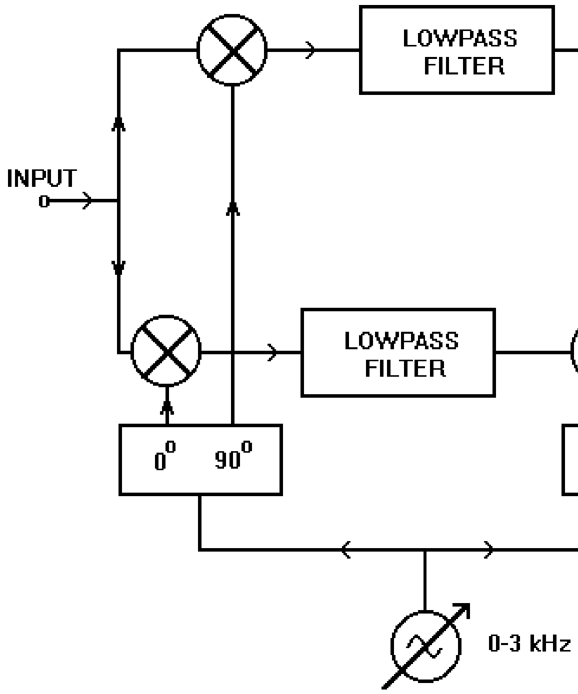

Fig 3. Tuneable bandpass filter made from two of the filters of Fig 2, built into an image-cancelling down-up converter.

Since we know that the expected signal sweeps at 100kHz/sec and the SSB receiver is 3 kHz wide, we can say that it will take 30mS for the received chirp to sweep through the passband, so that's how long we make the delay-line. This means that the low-pass filters will have a cut-off frequency of 33Hz (when one cycle of sinewave just fills the delayline), and thus the tuneable audio filter will pass frequencies up to 33Hz each side of the centre. When we know that the expected chirp has just started to arrive we can start the local oscillator sweeping from zero to 3 kHz at the 100kHz/sec rate. 30mS later when the received chirp just fills the delaylines, the filter output will be at it's maximum. We have received 30mS-worth of the 3kHz-wide sweep-sounder signal in a matched filter.

At first sight it might appear that this filter is going to give us an impulse response which is of 30mS duration, and we are really hoping for something better than this so that we can resolve closely-spaced paths: recall that the fundamental theory says we should be able to resolve 0.33mS with a 3 kHz bandwidth. But consider now what happens if the chirp signal arrives 0.33mS later that we expected. It's sweep will be 0.33mS behind, or 33Hz lower, than the oscillator sweep, and the outputs of the mixers will be at 33Hz rather than DC, right on the cut-off frequency of the lowpass filters. The same applies if the received chirp was 0.33mS early. So a chirp early or late by 0.33mS will be cut off by the filter. This is another way of saying that the filter has a time-domain response of +/-0.33mS, *but only to upward-sweeping chirps*. In other words, what we have made is a filter that selectively passes a chirping input signal of 30mS duration, with a passband of +/-33Hz, but gives an output pulse that is only +/- 0.33mS wide. We will be able to resolve chirps separated by 0.33mS after all.

Finally, if we consider only the instant of time when the expected chirp signal has just filled the delayline (which is the only time when we need to calculate the output), and we realise that the mixers can be implemented digitally by multiplications, we can dispense with the first pair of mixers in Fig 3 by replacing the constant tap coefficients in the filters by values pre-computed from the sine and cosine values of the sweeping local oscillator. The sequence of tap coefficients is no longer constant along the delayline, but varies in a sinusoidal manner along one filter and in the complementary cosinusoidal manner along the other filter, with the frequency varying along the sequence. When we do the summing of the taps now, it's a sum of products and one of the terms in these products is the sequence of pre-computed oscillator signals, spread out in sequence along the delayline rather than spread out in time from the oscillator. To further simplify the chirp filter, we can note that if we are only interested in measuring the amplitude of the received chirp and we don't need to reproduce the original signal at the filter output, we can also dispense with the second pair of mixers, and just combine the outputs of the low-pass filters together and achieve the same result. This means that we can now also throw away the oscillator since all the outputs we needed from it have been pre-calculated and "hard-wired" into the filter tap coefficients.

The finished chirp filter

This gives the final form of the chirp filter in Fig 4, with the audio input feeding into a single 30mS delay line with two sets of "chirped" tap coefficients alongside, each with a sum-of-products function, and a combining arrangement at the far end which is simply done in DSP by a sum-of-squares function. In practice the chirp filter coefficient sets are shaped slightly in order to eliminate the sidelobes outside the first nulls at +/- 0.33mS (+/-33Hz) but this has the effect of filling-in the first nulls so that they become the 6dB points of the finished chirp filter. We end up with a chirp filter with a 66Hz bandwidth and a 0.66mS wide chirp response, both figures now measured between the 6dB points.

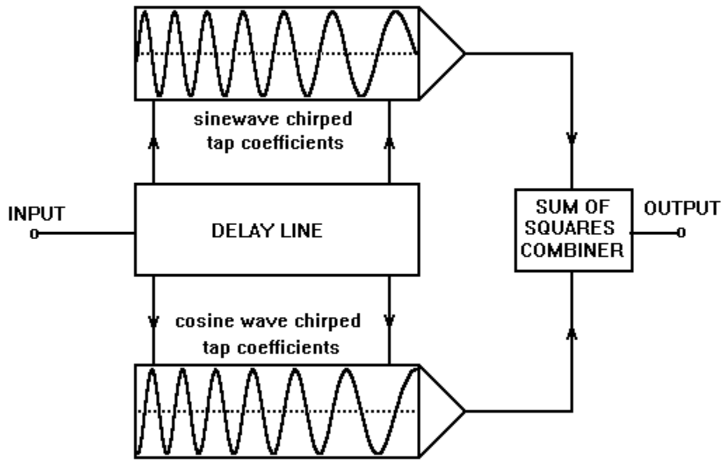

Fig 4. Final chirp filter consisting of one delay-line, two sets of chirped sine/cosine tap coefficients, two sum-of-products circuits and a sum-of-squares combiner.

To use the chirp filter in this form, we feed the incoming audio continuously through the delayline, and perform the DSP calculation whenever we wish to see the output response. If we are expecting a single chirp we need only do this calculation once when we expect it to have just filled the delayline, but if we want to listen continuously for possible chirps, we can do the calculation as often as we wish, typically at 8000 times per sec, the sample-rate of the DSP system. We have effectively created a continuous-time chirp filter: a strange device for which there is no equivalent in the analogue world. It's a filter that passes only upward sweeping tones and rejects everything else. Not only that, but when we input a 30mS-long chirp, it delivers a 0.66mS-long output pulse. We can just distinguish two chirps which are separated by 0.33mS even though they overlap each other for most of the time that they are sweeping through the passband. It is effectively a device for transforming chirps into pulses. It's even 17dB better than if I had persuaded someone to transmit pulses for me, because there is 45 times more energy in a 30mS chirp than in a 0.66mS pulse.

I wrote the DSP code to implement such a filter in the Motorola DSP56002EVM kit, the same one that I used for the dopplergram experiments and for other projects. In addition to the receiver audio, an accurate timing signal was fed into the DSP card, initially derived from a crystal oscillator but later from a GPS navigator module equipped with a 1 Hz output which is derived from the atomic clocks in the GPS satellites. The chirp filter output from the DSP card goes via a serial link to a computer which displays it on the screen in a raster scan with elapsed time along the X-axis, propagation delay along the Y-axis, and the signal level shown as the intensity of the displayed pixel, with a range of 64dB between black and white. This is very similar to the waterfall display used in the dopplergrams, but with the vertical axis showing time delay rather than Doppler shift. The Y-axis scan was configured like a delayed timebase, so that it displayed the signal received in a chosen 150mS slot in every 5 minute cycle. Each received chirp thus appeared as a blip in a vertical stripe painted on the screen once every 5 minutes, each stripe appearing slightly to the right of the previous one. A single path displays as a thin horizontal line over a period of hours, and multipath, for example, shows up as fuzziness or extra horizontal lines above the main trace.

The first results on screen

A typical trace of the local sounder, taken on 10 MHz in March 2000, is shown in Fig 5. The line at zero propagation delay is the direct groundwave. In the middle of the day when there is no skip zone, the vertical skywave reflection can be seen very strong at 2mS delay, with characteristic turned-up ends. Before and after this a fainter, more diffuse trace is visible at 4mS or more. This is the backscatter path, reflected from the ionosphere then scattered back from the ground beyond the skip zone. The varying backscatter delay shows the varying skip distance.

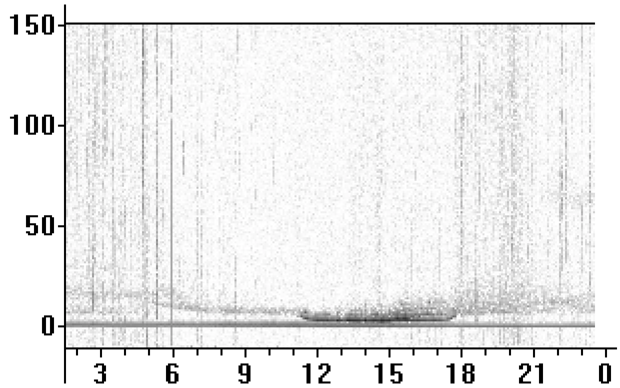

Fig 5. Waterfall of local UK sounder at 10 MHz showing groundwave, midday skywave and backscatter from the edge of the skip zone in the morning and evening.

This was a promising start. I went on to monitor other sounders and found that the chirp filter was far better at detecting chirps than I could by ear. Many signals showed multipath, scatter, and other interesting propagation effects. In particular on the higher bands many signals showed a weaker echo after the main signal, delayed by up to 150mS. This was clearly the long path, but it seemed to be present far more often than I had expected. Even the local sounder produced a "long path" at a delay of 139mS on some occasions, having travelled one complete circuit of the globe. We are normally only aware of a long path if it's open when the short path is not, but evidently both paths are often present together. Fig 6 shows a trace of an Australian sounder showing the well-known morning long path and afternoon short path, while Fig 7 shows a south-east USA sounder with a less well known longpath across the Pacific. Fig 8 shows a sounder in Cyprus received direct and briefly via an even longer path across the Indian and Pacific oceans and North America.

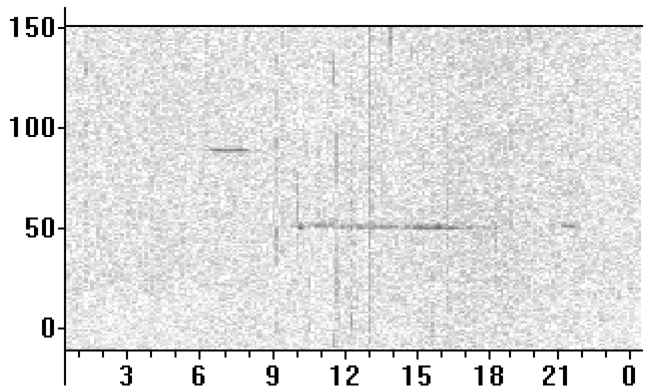

Fig 6. Waterfall of sounder in Northern Australia received in UK on 16MHz, showing morning longpath at 89mS and afternoon shortpath at 50 mS.

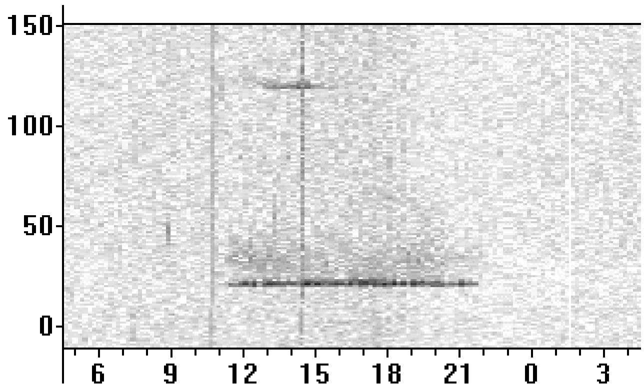

Fig 7. Waterfall of sounder in the Eastern USA received in UK on 21MHz, showing the direct path during the day at 21mS with simultaneous long path at 118mS in the afternoon.

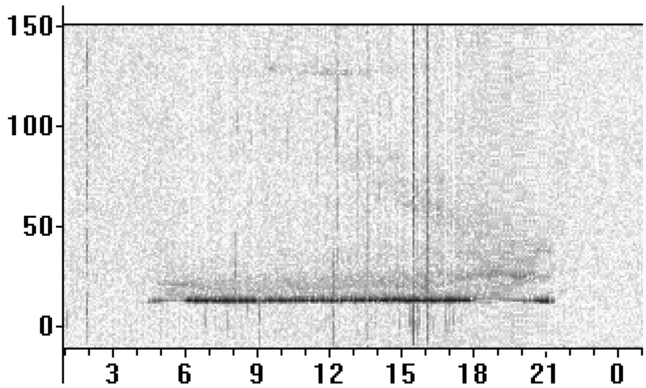

Fig 8. Waterfall of Cyprus sounder received in UK on 29MHz, with the direct path at 12mS some sidescatter paths, and brief traces of longpath at 127mS.

Identifying the sounders.
The timing of the sounders and the receiving system was stable enough to be able to follow the same one from day to day. I had a list of sweep sounder locations, but no way of relating the names in this list to the sounders I could hear. I was not able to find any published information which linked the location of the sounder to it's timing. So how did I know that the signals in Figs 6, 7, and 8 came from Australia, USA, and Cyprus?

This leads on to phase two of the project. I realised that I already had enough information to find the distance to those sounders which I could hear by both short and long paths. If my local sounder gave a direct signal at one instant and a round-the-world echo 139mS later, I could work out the propagation delay from other sounders by measuring the interval between the received longpath and shortpath chirps, subtracting that from 139mS and dividing the result by two. If I could do this from two places in the world, I could find the sounder.

ZL1BPU joins the project
To test this idea I sent the DSP software and the associated PC program to Murray Greenman ZL1BPU near Auckland. He has the same DSP kit and was keen to join the project. By this time I had established that the sweepsounder transmitters operate by starting their sweep at a specific number of seconds past the hour, and at 5 minute intervals thereafter, with a sweep rate of 100kHz/sec. To identify each chirp we heard, we noted the time at which we heard it and subtracted the received frequency divided by the sweep rate. This effectively gave the time at which it started sweeping from zero MHz. We called this the chirptime and it meant that we didn't need to co-ordinate the frequencies on which we were both listening. Because the sounders repeated every 5 minutes we didn't need to synchronise our listening periods either, just record the chirptime as a number of seconds between 0 and 300.

We found about 15 sounders that we could both hear and we could measure longpath-minus-shortpath delays on some of these and hence work out the distances from each of our respective locations to the unknown transmitter. We then plotted these as lines of constant distance on a map and found the points of intersection. Some of these lines crossed close to islands which we already knew were the locations of sounders. One of these was Cyprus and other crossing points indicated locations in Australia, USA and other parts of the world.

By this time I had calibrated the chirp measurement system so that we could not only get stable traces of the received chirps and read off the interval between longpath and shortpath, but we could also measure the actual received chirptime to a resolution limited only by the DSP sample-rate. This was possible because the GPS units we were using as frequency-standards not only gave us stable 1 Hz pulses but these pulses were accurately aligned to universal time: the GPS manufacturers could assure us that the leading edge of the 1Hz pulse from my GPS was within 1uS of the leading edge of the 1Hz pulse at ZL1BPU. We even devised a way to measure and compensate for the delay (about 1.2 mS) through the filters in our SSB receivers.

The time-difference method

We could now extend our ability to locate sounders to those which we could not both hear by long and short paths. We did this by measuring the chirptime at both stations, subtracting one from the other and then plotting a "line of constant time-difference" on the map. Another chirptime observation from a third station gave a second time-difference line and the unknown transmitter would lie at the intersection of the two. This is the same process as used in hyperbolic navigation systems such as Decca, Loran and GPS.

PY3CRX and KC7WW were recruited to the project and several more sounder locations were added to the collection. As an example, Fig 9 shows the position-lines for one set of observed chirptimes, 224.154 at ZL1BPU , 224.166 at KC7WW, and 224.168 at G3PLX,. The time-difference line for G3PLX-KC7WW=2mS is shown in red and that for KC7WW-ZL1BPU=12mS in green. Note that these lines cross in two places, not one - an annoying side-effect of the earth being round rather than flat! In this case we knew from the times that we could hear this sounder that it was probably at the Falkland Islands rather than the Philippines, but we confirmed it when PY3CRX reported a chirptime of 224.134 so that we could also plot the blue KC7WW-PY3CRX=32mS position line.

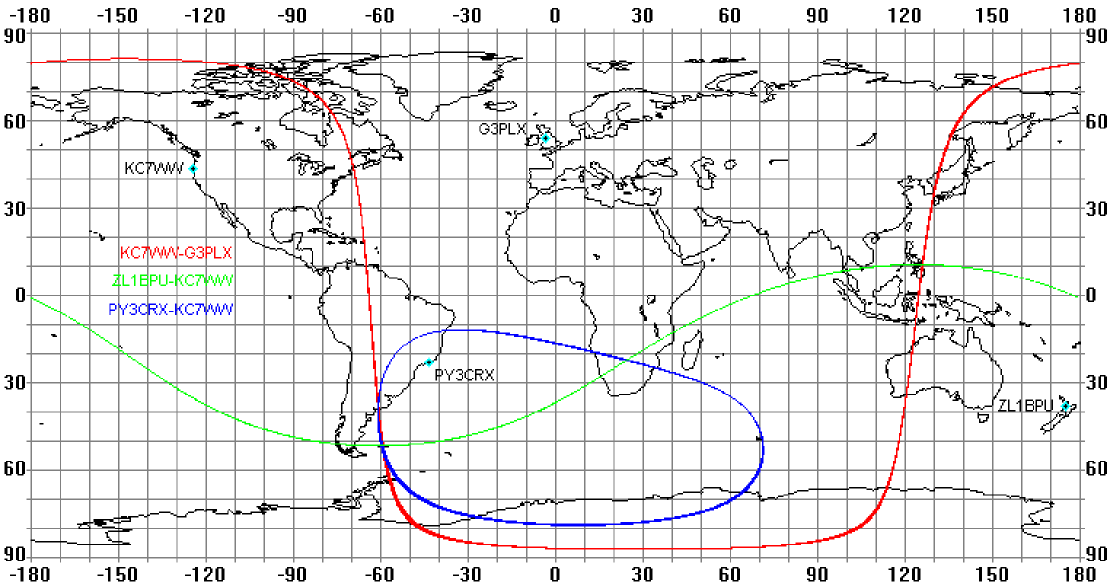

Fig 9. World map showing time-difference lines from KC7WW, PY3CRX, ZL1BPU and G3PLX crossing at the Falkland Islands.

The absolute time method

Once we had the location of a sounder we could calculate back to find the chirptime at the transmitter. It soon became apparent that many sounders started their sweeps exactly on a one-second mark and we noted that these were also the sounders which remained absolutely stable in their timing from day to day, while others drifted slightly to the extent that one would expect from a crystal-controlled clock. It seemed reasonable to suppose that the stable ones were using the 1 Hz pulse from GPS for their timing too, and we later verified this from manufacturer's information gleaned from the internet. This gave us a third method of measuring the distance to a GPS-locked sounder by simply using the decimal part of the received chirptime. Thus a stable sounder with a chirptime of XXX.001 was probably 1mS away (300km), one with a chirptime of XXX.0695 was exactly halfway round the globe, and if it was later than that it was longpath. We located a few more this way and also cross-checked some of the previous observations.

Automating the search

I then went on to further refine the PC software so that it would detect and log chirptimes automatically and sort them into their correct periods, because by this time we could see that there were sounders with periods of 5, 7.5, 8, 10, 12, 15, and 30 minutes. Typically between 100 and 300 individual chirps per hour are detected during daylight at G3PLX, using a 6m vertical whip antenna. The software sorts out regularly occurring chirps to find the period and chirptime and usually logs about 30 - 50 different sounders over a 24 hour period. This automation made it possible to expand the network of observers to include EA2BAJ, DF7YC, and G4JNT. The lists of received chirptimes were saved to file and Emailed to me for analysis. As a result some 40 sounders have been located so far in all corners of the globe, but there are many more that have not yet been heard in three places. It became clear later that there are more sounders with sweep-rates other than 100 kHz/sec, although I have not so far written software to process these.

The way forward: all-band beacons?

There are several directions that the chirp project could now take. At the simplest level, anyone listening with just a conventional HF SSB receiver can hear the stronger sounders, and identify them with the help of a clock. You may like to try this yourself: the Cyprus sounder is easily audible by ear in the UK most days. It starts it's sweep (from zero MHz) at 250 secs past the hour and sweeps at 100kHz/sec, so if you tune to 21.0 MHz during the daytime, you may hear it at 250+210=460 seconds past the hour and every 5 minutes thereafter. The sounder near G3PLX will also be audible around the UK on the LF bands. It starts at 78 seconds, so if you tune to 3500 kHz you should hear a chirp at 78+35=113 secs past the hour and every 5 minutes thereafter. You may then find that you can hear many more once you know what they sound like. These could be useful additions to the existing HF amateur beacons, with the advantage that they are far more numerous and operate over the entire HF spectrum and not just on spot frequencies. Incidentally, none of these sounders actually starts sweeping at zero MHz: they are silent until they get to about 2 MHz, and sometimes they are also muted around distress and safety channels. More sounders are listed in Table 1. and there is a full list on the Chirps Project web site.

There is clearly scope for automatic beacon monitoring and It is possible that chirp-detecting software could be written to use a low-cost PC soundcard rather than the DSP card that we have used so far. Of course, all these ideas for using sweepsounders as beacons depends on the availability of accurate and up-to-date information about their locations and operating schedules. We may even be able to generate and maintain this information ourselves by using the triangulation methods I have described.

Another possible application suggested is to add gps-locked chirps at regular intervals to the emissions from existing amateur beacons. Not over the whole HF spectrum but just over a 3 kHz sweep. It would only need a 30mS chirp every 5 minutes to give propagation traces similar to those shown in this article. Although this would take up 3 kHz of the band, the duration of the chirp is only that of a single Morse-code dot at 40wpm.

The way forward: propagation research?

Even without this extra infra-structure and development work, this project has now opened up the possibility for using the signals from existing swept-frequency sounders for amateur HF propagation research in a variety of areas. It would be interesting to study long-path propagation more extensively now that we have a way to see it clearly. Other topics such as trans-equatorial propagation, long delayed echoes, non-great-circle paths, and aurora come to mind.

I would like to thank Marcus Ramos PY3CRX, Johan Forrer KC7WW, Andreas Gawron DF7YC, Eduardo Jacob EA2BAJ, Andy Talbot G4JNT, and in particular Murray Greenman ZL1BPU. The Chirps project would not have been possible without their help.

Further information about the Chirps project, including an up-to-date list of sounders and more examples of received chirp traces, together with links to related sources of information can be seen at the Chirps project website at www.qsl.net/zl1bpu/chirp/chirps.html

Table 1. A selection of sweep sounders logged in April 2000. You will hear a sounder with a chirptime C on a frequency F (MHz) at C+F/10 seconds past the hour, repeated at the appropriate interval.

Location	Period mins	Chirptime secs
S England	5	44
NW England	5	78
Italy	15	118
Spain	15	138
SE Australia	15	178

Ascension Is	5	222
Falkland Is	5	224
Cyprus	5	250
N Australia	7.5	306
Caribbean	15	342
Indonesia	15	390
New Zealand	15	500
N Scotland	15	662
Iceland	15	680