SDR and VLF time signals

Poul-Henning Kamp
The FreeBSD Project
phk@FreeBSD.org
http://phk.freebsd.dk/loran-c

Abstract

VLF signals have been used to transmit precise time and frequency signals since roughly the second world war and even in these GPS days, DCF77 and WWVB signals are probably still responsible for getting more people out of bed in time than any other time and frequency distribution method.

This paper outlines some experiments on receiving a decoding these signals using a software defined radio, and the results obtained. For illustrations and data please see the above URL.

1. Interesting VLF time signals

There are basically three classes of interesting time signals in the VLF spectrum: dedicated CW time signals, Loran-C and other frequency stabilised CW stations.

1.1. Dedicated CW time signals

Amongst the dedicated CW time signals, DCF77 and WWVB are undoubtedly the most well known, but in addition to these two, there are at least a handful more around the world, all located in the 40-80 kHz spectrum.

Typically these stations employ a binary AM modulation of a once-per-minute "timegram" which can be trivially demodulated by a small microprocessor or even discrete logic. In addition to providing time of day and calendar information, other data may be present, for instance UT1 corrections, daylight savings or leap second change warnings. The carrier frequency of these signals are typically derived from a steered Caesium standard and therefore often have specifications in the 10^{-12} range.

1.2. LORAN-C

LORAN-C is a hyperbolic radio-navigation system which transmits highly precise synchronised signals from a mesh of transmitters throughout most of the northern hemisphere. For navigation LORAN-C has been marginalised by GPS, but since the signal has entirely different properties than GPS there is still a strong resistance to closing the system down, the aviation community in particular wants to retain it as an independent backup system to GPS.

On very desirable feature of the LORAN-C signal is that it is virtually impossible to jam it. Loran-C transmitters operate from 200kW to 1MW power and the 100kHz signal goes through practically anything, including a fair amount of normal building materials. GPS on the other hand is a very weak microwave signal which it only takes a battery-operated transceiver of a few watts to jam over a large area.

1.3. Frequency stabilised CW stations

Many commercial long-wave broadcast stations, and for that matter many other transmitters in low end of the spectrum, use atomic standards to stabilising the carrier frequency, and these can therefore be used as frequency standards as well.

2. Hardware

Because the interesting frequency band for VLF time-signals is more or less from DC and up to 500 kHz and because it is desirable for time-signals to be able to calibrate the receiver delay from antenna to detection I decided to go for direct reception without a heterodyne mixer stage.

This more or less means that the receiver has only one hardware component in addition to the antenna: the A/D converter.

Finding a suitable A/D converter proved tricker than I had expected. Ideally I would want a 16 bits with a sample frequency in the 1-2 MHz range, and preferably a PCI card which I could plug directly into my PC computer, but nothing even close like it were available. Eventually I settled on a 20 MHz, 4 channel, 12bit PCI card instead, despite this being a fair bit more expensive.

The sample clock for the A/D converter is derived from one of the various frequency sources I have in my lab. The A/D card has a built in divider to which a standard 5 or 10 MHz signal can be connected, this has worked without a hitch.

I built a simple 1' diameter loop-antenna and added a one stage amplifier based on the AD797 op-amp for driving the signal through coax from my attic to my basement lab.

3. Device driver software

I've spent ten years of my life on the FreeBSD UNIX operating system so obviously my software would run under FreeBSD. The goal was to do only low-level work in the device driver for the A/D card and do the rest from various application programs.

Since my main interest was LORAN-C signals, I gave the device driver the ability to perform GRI/FRI level averaging of the signal since that would have to be a real time function.

This is done with a simple exponential averaging into a circular buffer, for instance to receive the 7499 LORAN-C signal, a buffer containing 74990 microseconds worth of samples would be configured.

On my Athlon 700 MHz test machine, the CPU gets saturated tracking about 10 signals which is plenty for my purposes.

4. Sampling rate

I had initially expected that a sampling rate of 1 MHz would be a good choice for receiving LORAN-C signals, but my initial experiments revealed that having 10 samples per 100kHz cycle left something to be desired, in particular it was hard to find the top of the individual half periods precisely enough to reliably detect and lock onto the 3rd zero crossing.

Obviously, the more samples the better, but since this had a proportional effect on the available CPU power, a too high sample frequency was not desirable either. Further experimentation showed that 1.25MHz gave sufficient improvement to get me reliable locking.

5. LORAN-C reception

For LORAN-C signals can be described by two convolutions of a periodic unit impulse, first the basic LORAN-C pulse and then the master or slave FRI-coding of the pulse train.

Viewing the LORAN-C signal in this way allows us to order the reception stages in the order requiring the least necessary CPU power: First the FRI period averaging which is done in the device driver. Second the FRI-coding deconvolution which is a simple summing operation over the 8 or 10 relevant samples in the FRI period and finally a FIR bandpass filter to remove unwanted signals.

It is not necessary to filter the entire FRI period with the FIR filter only the relevant window around the deconvolved pulse form, this reduces the work by a factor of about 100 divided by the number of transmitters tracked in this LORAN-C chain, and therefore filter kernels with more than 1000 points become affordable.

The next task is to find the 3rd zero crossing, a task which has been the subject of much research over the years because doing it without programmatic support is very very hard.

I have settled on doing a match with a theoretical template pulse and using a simple best fit selection.

The LORAN-C signal is actually very wide-band, and therefore it even a 80-120kHz band-pass FIR filter will distort the waveform slightly. Fortunately, this can be compensated trivially by also running the template impulse through he FIR filter so it becomes subject to the same distortion.

Once the 3rd zero crossing has been identified, an interpolation is performed using the surrounding 3 samples. This interpolation must account for the uneven amplitude of the half-period right before and after the zero-crossing.

At this point, all that is left to do is keep track of that 3rd zero crossing and outputting its position in the FRI window and convert this position to a phase offset between the sample clock and the received signal.

6. LORAN-C results

Despite the fact that I have insufficient amplification on my antenna signal and therefore only use about 8 bits of the A/D converters resolution, I have been able to successfully track the master signals of five different LORAN-C chains at the same time: 6731M, 7001M, 8000M, 7499M and 9007M.

The strongest of these, 7499M, which is only 213 km from my location show a phase noise relative to my old HP5061A caesium standard of approx ± 20nsec, and a modified allan variance of 2^{-11} @ 1sec.

The other signals have performance which is consistent with other empirical data on LORAN-C propagation.

7. LORAN-C future work

The next obvious step is to track the slave signals in addition to the master signals.

It could be interesting to develop an adaptive model for the propagation path for each individual transmitter tracked, but the necessary accumulation of raw data would obviously be a multi-year exercise.

Since some of the transmitters are "dual-rated" (for instance are both 7499M and 6731Z transmitted from the Sylt LORSTA) these signals are subject to the same propagation conditions and can be treated as one signal, thereby getting more samples and better S/N ratio.

It would be possible to decode the "EuroFix" modulation on the LORAN-C signal, but this would require a different receiver topology since the FRI averaging done in the device driver averages the EuroFix modulation out.

8. DCF77 Reception

The DCF77 signal is interesting because in addition to have a tightly controlled carrier frequency and a 1 bps AM modulation, the carrier is also phase modulate with a PRN sequence which should make it possible to detect the start of UTC seconds much more precisely than the AM modulation would allow.

When a buffer in the device driver is configured to contain an entire second worth of samples and use an averaging factor of 1, it is possible for a program to grab a snapshot of the buffer once per second and process it.

The best way to detect the phase encoding is to build an ideal signal and find the best correlation, this is unfortunately rather expensive computationally when the sample frequency is 1.25 MHz, but fortunately, once the correlation has been found the DCF77 carrier signal can be tracked instead, and this is computationally very cheap.

9. DCF77 Results

Detecting and decoding the AM modulation was trivial, and yielded a good improvement over the normal single-chip "Temic" receivers found in radio-controlled alarm clocks and similar devices.

Detecting the phase-modulation was also found to be reliable, and once detected, tracking the carrier phase was trivially possible.

I have not made quantitative measurements of this available yet, but the phase-modulation correlation was into the 1 microsecond territory.

10. Decoding other CW signals

Decoding other AM modulated CW signals were also achieved, both the "HBG" and 75 kHz and the "MSF" signal from Rugby.

An interesting idea is to try to derive the receivers location using the DCF77, MSF and HBG signals as hyperbolic beacons.

The French long wave station at 162kHz uses an interesting phase modulation to transmit a binary data stream, but I have yet to write software to decode this modulation, but it is clearly visible in a number of test runs I have made.

11. Phase tracking

Tracking the phase of multiple CW stations in general can be done with only a single 1 second period buffer in the device driver because none of the stations transmit on carrier frequencies which are not high multiples of one Hz.

In fact, since most stations transmit on integer multiple of one kHz, a buffer of one millisecond is generally sufficient.

Detecting the carrier phase is done by passing the buffer through a narrow IIR bandpass filter, mixing it with a locally generated sine wave to eliminate the carrier and calculating the phase from the averaged I and Q signals.

This is still a work in progress, but initial results show that for typical signals the phase can be tracked with very little computational effort since the kernel averaging does most of the heavy work.

12. Future work

The ultimate result of all this work, would be an frequency and time receiver which received and tracked all available signals, applied a time-scale algorithm to them to derive a local timescale used to steer for instance an OCXO or Rubidium frequency source. Time will show if I manage to pull that off before they pull the plug on the LORAN-C signals.

It should also be possible to receive the amateur band at 137 kHz with my receiver, but it would probably require more bits on the A/D converter before it was interesting.

Obviously decoding the various modulations on other signals in the VLF band could also be attempted, but right now this is not in my sphere of interest.