Software defined radio decoding of DCF77: time and frequency dissemination with a sound card

Jean-Michel Friedt[1] Clément Eustache[2], Émile Carry[1], Enrico Rubiola[1]

[1]FEMTO-ST Institute, Time & Frequency, CNRS-UBFC

ENSMM, 26 rue de l'Épitaphe, 25000 Besançon, France

[2]Master PICS, Université de Franche Comté, Besançon, France

Key Points:

- Software defined radio decoding of atomic-clock controlled very low frequency signal
- Use of the phase modulation spectrum spreading for high resolution time of fligh measurement
- Ionosophere altitude variation measurement using a sound-card based setup

Corresponding author: Jean-Michel Friedt, `jmfriedt@femto-st.fr`

Abstract

We investigate time and frequency dissemination using Software Defined Radio processing of signals acquired from a Low Frequency emitter using a sound card. We use the resulting propagation time measurements for investigating some ionosphere physics and its interaction with cosmic ray flux. Rather than using the amplitude of the transmitted signal as classically considered, we here focus on a precise time of flight measurement by demodulating the spectrum spreading phase modulation added to the DCF77 amplitude modulation.

1 Introduction

Time and frequency dissemination has been an issue whenever a society aims at synchronizing activities (banking system, transports, power grid regulation) over a spatial range. Currently, Global Navigation Satellite Systems (GNSS) and the Global Position System (GPS) in particular, are amongst the reference time and frequency dissemination solutions exhibiting utmost stability, with accuracies ranging sub-100 ns when synchronizing a clock to the 1 PPS (Pulse Per Second) output of a GNSS receiver. However, these low power transmissions are prone to jamming and spoofing, so that alternative solutions are desirable. Low frequency (LF) solutions have been implemented well before the advent of GNSS [*Watt et al.*, 1972], and some emitters are still active, including the 77.5 kHz German DCF77 emitter located in Mainflingen (50°0'56"N, 9°00'39"E). This 50 kW emitter is powerful enough for its signal to be recovered over Western Europe [*Bauch et al.*, 2009; *Piester et al.*, 2011; *Engeler*, 2012], and the reader beyond this reach willing to decode its signal can collect records from websdr sites including `http://websdr.ewi.utwente.nl:8901/`. Most importantly for the physicist, the atomic clock-locked signal, with the reference signal provided by the German metrology laboratory PTB, interacts with the ionosphere, hence providing the means of probing ionosphere interactions with its environment [*Baker and Lanzerotti*, 2016], namely daily and seasonal cosmic ray flux fluctuations, and most significantly solar ionizing radiations. Such investigations have been classical since the 1960s [*Blackband*, 1964], but the proliferation of computers with huge computational power fitted with sound cards [*Schulte et al.*, 2012; *Carlà*, 2016] sampling at least at 192 ksamples/s allows for any curious experimenter to implement such a receiver at basically no cost since all demodulation schemes are implemented as software, the ultimate implementation of Software Defined Radio (SDR) principle in which the only hardware part is analog to digital conversion of the electromagnetic signal reaching the antenna [*Kamp*; *Dolea et al.*, 2013]. Indeed, the current trend to shift from analog to digital signal processing, especially in the context of time and frequency metrology [*Uchino and Mochizuki*, 2004; *Mochizuki et al.*, 2007; *Gotoh et al.*, 2011; *Huang et al.*, 2016; *Sherman and Jördens*, 2016], meets the requirements of improved stability, flexibility and reconfigurability [*Mindell*, 2011] provided by SDR, which has become practical lately with the advent of radiofrequency high resolution analog to digital converters.

Ionosphere property fluctuations are linked to the cosmic ray flux variations. The upper layers of the atmosphere are exposed to a flux of particles generated by the galactic environment on the one hand, and the Sun on the other hand. The orientations of the Earth with respect to this particle flux defines the ionosphere properties. Besançon, France (47°N, 6°E), is located about 370 km from Mainflingen, so that a direct time of flight of an electromagnetic wave lasts 1.2 ms. Assuming the same electromagnetic wave bounces over the D-layer of the ionosphere located [*Blackband*, 1964; *Davies*, 1990; *Johler*, 1962] at an altitude of about 50 km, the additional time delay is 100 μs (Fig. 1). Furthermore, assuming the ionosphere altitude varies from 50 to 90 km from day to night ionization conditions – whether the Sun illuminates or not the upper atmosphere – an additional delay of 95 μs is expected: all these numbers result from basic geometric considerations of straight paths between the emitter, the receiver and the reflector plane. Hence, investigating the ionosphere physics requires timing with sub-10 μs accuracy if these effects are to be observed.

Accurate timing requires some bandwidth spreading [*Raupach and Grosche*, 2014] since time resolution is given as the inverse of the bandwidth of the incoming signal. Such a requirement seems opposite to that of frequency dissemination which requires narrowband signals. This dual need was originally met in the case of DCF77 with an amplitude modulation once every second of an atomic-clock locked carrier, yielding timing accuracy in the hundreds of microseconds due to the poor resolution of amplitude variation detection. In the late 1980s an additional spread spectrum phase modulation scheme was added allowing for much better timing accuracy [*Hetzel*, 1988]. Despite very few commercial receivers using this additional mode – DCF77 receivers are fitted in most radio-controlled clocks including low-cost weather stations – we will see that the tremendous timing accuracy gain, over ten fold to reach sub-10 μs accuracy, will allow us to address

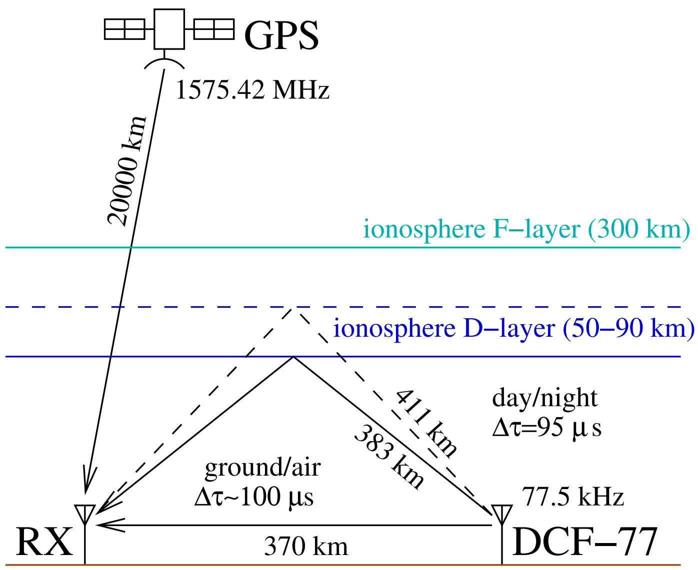

Figure 1. Schematic of the LF signal propagation between Mainflingen (DCF77, Germany) and Besançon (RX, France), 370 km geodetic distance. The time of flight difference between the ground and air wave bouncing off the ionosphere is 100 μs, and the ionosphere D-layer altitude variation between day and night induces another 95 μs delay in this geometric approximation. Ionospheric delay on the microwave GPS carrier is considered negligible in this application.

some of the ionosphere physics by processing the signal recorded by a personal computer sound card.

Based on these general considerations on long range wireless time transfer and the ability to probe ionospheric boundary conditions thanks to the high stability timing signal, the outline of the paper is as follows. First, we will describe the hardware setup for receiving the radiofrequency signal using a common personal computer sound card: the hardware is limited to a bare minimum antenna impedance matching circuit, which nevetheless requires some investigation considering the very short antenna dimensions with respect to the wavelength, its very high impedance and the need to buffer the signal before feeding the sound card. All demodulation and timing analysis are performed through software processing: implementation of the algorithms is developed in appendix A while the third section of the main text focuses on a description of the algorithm applied to extract first a stable phase and then a fine timing signal from the cross-correlation of the received signal phase with the known pseudo-random number sequence. Based on this analysis, the next section provides some measurement results demonstrating the timing accuracy is sufficient to observe daily ionosphere altitude variations between daytime and night time, as well as seasonal behaviour differences. Additionnaly, frequency stability measurements allowing for local oscillator temperature induced drift are demonstrated. The core reason for emitting timing signal over very low frequency signals being the long range synchronization of quartz-controlled clocks, we develop the analysis needed to tune such a control loop. Finally the last section is devoted to a comparison of the vertical and horizontal components of the electric field, representative of the two propagation paths through the ionosphere and over ground of the very-low frequency signal. Throughout this investigation, the GPS 1-PPS signal is used as a reference with respect to which the DCF77 timing signal is compared: a stereo sound card records simultaneously the two signals, hence rejecting the sound card clock impact on the measurement.

2 Hardware setup

SDR aims at limiting the hardware setup to an antenna connected to an analog to digital converter. Most radiofrequency applications require however an additional mixing step with a local oscillator since most analog digital converters (ADC) do not exhibit the sampling rate – typically a few MHz – needed to sample radiofrequency signals: shifting the signal under investigation from its carrier frequency to baseband, close to 0 Hz, also allows for filtering strong interference sources and prevents saturating the sampling stage with a wanted signal below the ADC resolution. VLF (Very Low Frequency, 3–30 kHz) and LF (30–300 kHz) allow implementing true SDR receivers: meeting Nyquist criteria of a sampling rate at least twice the targeted signal frequency range, recording DCF77 only requires an ADC with at least 150 ksamples/s sampling rate, a requirement met by most current sound cards sampling at 192 ksamples/s. Alternatively, we have successfully used a Terrestrial Digitial Video Broadcast (DVB-T) receiver fitted with a Realtek RTL2832U analog to digital converter sending data on a USB bus, as implemented with the Osmosdr GNURadio source, after removing the radiofrequency frontend: in such a configuration, the in-phase (I) and quadrature (Q) inputs of the RTL2832U are respectively connected to the DCF77 and GPS 1-PPS outputs, the latter as reference. GPS 1-PPS is defined as a 1-Hz digital pulse whose rising edge matches, to within a few tens to a few hundred nanoseconds depending on receiver technology and performance, the second of the time disseminated by the GPS satellite constellation. In both cases, whether using the sound card or the RTL2832U frontend, using dual-channel streams guarantees that the interleaved DCF77 and GPS measurements have been synchronously acquired. Since the electromagnetic signal emitted by DCF77 is vertically polarized, the coil antenna is, in this first setup, oriented horizontally with its normal pointing towards the emitter.

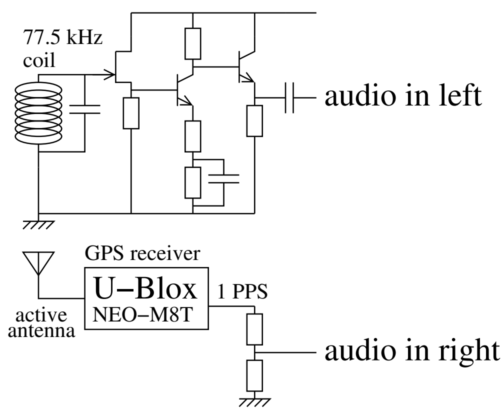

Figure 2. Schematic of the experimental setup: the coil receiving the DCF77 signal is tuned with a parallel capacitance for the antiresonance to match the targeted carrier frequency of 77.5 kHz. The high impedance output of the antenna feeds a FET transistor, for example BF245, before an amplifier and follower circuits based on NPN transistors, for example 2N2222, match the input impedance of one of the audio channels. The other audio channel is fed with an attenuated copy of the GPS receiver 1-PPS output as generated by a U-Blox Neo-M8T receiver.

Thus, we are able to connect an antenna straight to the sound card or RTL2832U input for further processing (Fig. 2): frequency shifting from LF band to baseband, followed by amplitude and phase demodulation. The only difficulty in setting up the antenna is the very long wavelength of the signal, meaning that the antenna is necessarily small [*ARRL*, 1997] with respect to the wavelength. Indeed, the 77.5 kHz of DCF77 has a wavelength of 3.8 km, so that a meter-long antenna will be considered infinitely small with respect to wavelength. It has been shown that such a small antenna necessarily exhibits high quality factor, a property usually frowned upon when designing an antenna aimed at operating over a wide frequency range, but here suitable since the antenna acts as a narrowband filter excluding strong nearby interferences, including switching power supplies and cathodic screens, and produces a strong voltage at the coil output. Furthermore, such a sub-wavelength antenna exhibits a much larger impedance at anti-resonance than the sound card input: an impedance matching circuit feeding a high impedance input with the antenna coil current (FET transistor grid) and generating a low impedance output is needed between the coil and sound card. Our circuit follows the inspiration from `www.qsl.net/dl4yhf/dcf77_osc/index.html` (accessed 2017), with the antenna scavenged from the DCF77 receiver circuit sold by Conrad (product reference: 641138). It is worth noticing that strategies for designing such very small antennas differ significantly from resonant antenna design: while in the latter case the impedance is close enough to 50 Ω for the reflection scattering coefficient (S_{11}) to be representative of the efficiency of the antenna at a given wavelength, small antennas operating in an anti-resonant mode exhibit very high impedance, well above 10 kΩ (Fig. 3). Under such circumstances, measuring S_{11} will not allow for tuning the antenna operating frequency: either a conversion to admittance (real part) exhibits a maximum at anti-resonance where a maximum voltage is generated by a given current induced by a magnetic flux flowing through the loop antenna, or a transmission measurement in which a function generator induces, in a forced regime, a voltage at the output of the tuned antenna in a transmission mode measurement, will allow for tuning the capacitance connected in parallel to the inductor formed by the coil antenna to operate at the wanted frequency.

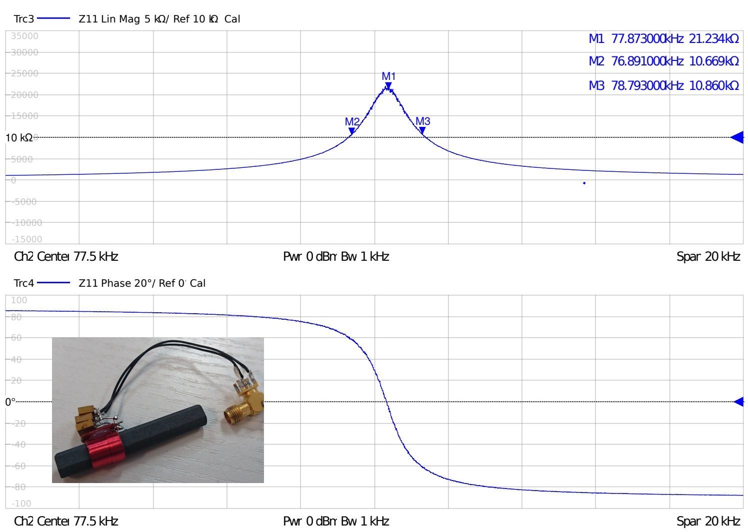

Figure 3. Coil antenna acting as an inductor, tuned to the operating frequency with a capacitor connected in parallel. Notice the maximum of the impedance at the operating frequency, as required to generate as high a voltage as possible for a given current induced by the magnetic flux flowing through the coil antenna. Both charts exhibit the impedance of the antenna as a function of frequency (linear scale), with the top figure displaying the magnitude and the bottom one the phase, in a frequency range of 77.5 ± 10 kHz.

The initial prototyping steps have been performed using GNURadio, a software framework designed to help SDR enthusiasts prototype digital signal processing functionalities yet provide real time signal processing and visualization, as opposed to post-processing using Matlab or its opensource implementation, GNU/Octave. Phase detection and automated analysis over long durations will be performed with the latter software. Transposing from radiofrequency band to baseband is such a common SDR processing task that it is implemented as an optimized processing block in GNURadio (Fig. 4): the Xlating FIR Filter. The time t dependent signal $s(t)$ received at the antenna exhibits a signal of interest modulated close to a carrier f_c, while recovering the property of the signal requires getting rid of the carrier: demodulating requires reproducing a local copy of f_c so that

$$s(t) \times \exp(j2\pi f_c t) \tag{1}$$

shifts the incoming signal to baseband (here $j^2 = -1$). Once the signal is shifted to baseband, the whole bandwidth, given by the initial sampling rate f_s, is no longer needed since the signal is band-limited: decimating, i.e. taking one in every N samples, reduces the bandwidth by a factor of N, easing processing steps since the datarate has been reduced. However, decimating brings all signals in the initial frequency band of $[-f_s/2; +f_s/2]$ to the new frequency band $[-f_s/(2N); +f_s/(2N)]$ by aliasing: low-pass filtering the frequency transposed signal prior to decimation is needed to get rid of these aliasing images, hence the inclusion of the Finite Impulse Response (FIR) filter in the GNURadio processing block. We now have a signal at baseband whose information content, lying in the amplitude and phase, must be decoded: such task will be performed solely by software processing.

3 Frequency lock

Amplitude demodulation is a crude processing step exhibiting the poorest noise rejection capability, but easiest to implement: the baseband signal is rectified and low-pass filtered. The local oscillator copy f_c only needs to be accurate enough for the signal to lie within the low-pass filter bandpass range. The narrower the low-pass filter the better the noise rejection, but also the longer the time response of the filter and hence the poorer the timing capability. We observe experimentally that a low-pass filter with 30 to 50 Hz bandwidth allows for observing the amplitude modulated pulses encoding time transfer, yielding time resolutions in the tens of millisecond. Practical amplitude pulse edge detec-

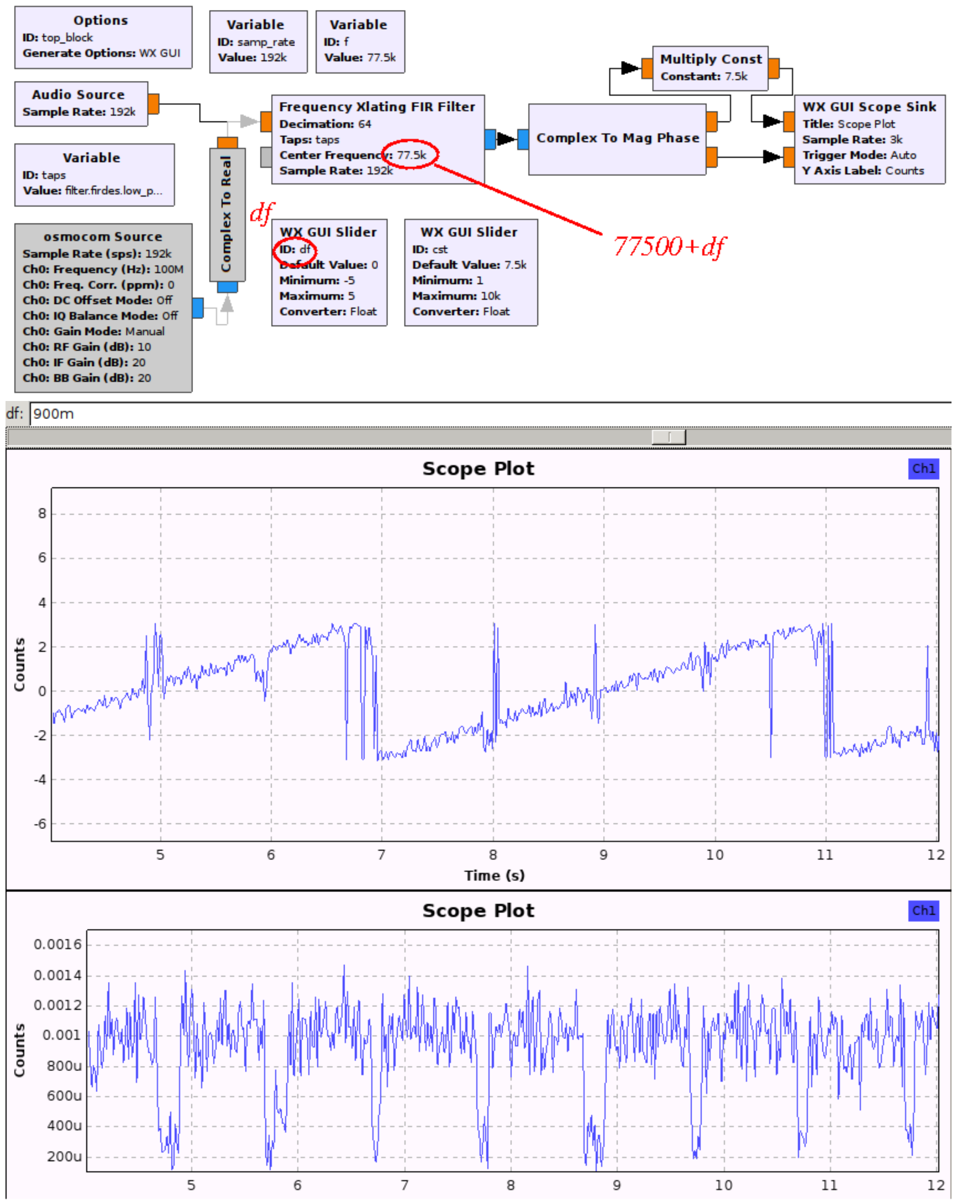

Figure 4. Top: GNURadio flowchart for realtime display of the decoded signal. The signal is sampled from a sound card (Audio Source), translated to baseband using the Xlating FIR Filter with a manually tunable frequency offset df with respect to the nominal 77500 Hz carrier frequency, and the phase and magnitude are displayed following low-pass filtering. Bottom: amplitude (bottom) and phase (top), with a slight frequency offset still visible as a linear phase drift over time. The amplitude modulated timing pulses are visible as signal drops every second on the bottom graph, with pulse width indicating the bit value.

tion shows that sub-millisecond time transfer is achieved using this strategy. Such a time resolution is insufficient to detect ionosphere variations, which were demonstrated previously to induce variations in the tens of microseconds range.

Pseudo-random phase modulation was introduced to spread the spectrum and improve timing resolution. The core aspect of this modulation scheme, also used in GNSS timing strategies with more complex implementations, is that the pseudo-random sequence is known, so that by cross-correlating a local copy of the code over the phase of the signal transposed to baseband, a sharp cross correlation peak occurs when the two copies of the code are synchronized: the cross correlation peak width is given by the inverse of the bitrate, and the noise rejection capability of the cross correlation is given by the number of bits in the code. Indeed, the pseudo-random sequence exhibits a 0-mean value, so that noise is averaged by cross-correlating with the code, and only the appropriate sequence of phase values coherently accumulates energy in the cross-correlation peak. Depending on the signal to noise ratio, cross-correlation peak fitting provides an additional timing accuracy gain equal to the signal to noise ratio. The pseudo random code generator implemented in DCF77 is known: the 9-th degree polynomial function x^9+x^5+1, whose implementation in C language is given at `https://en.wikipedia.org/wiki/DCF77#Phase_modulation`, feeds a linear feedback shift register generating a 511-bit long sequence with no repeating pattern over this duration, hence the spectrum spreading capability. The implementation informs us that the phase of the signal is updated every 120 periods of the DCF77 carrier, or at a rate of 77.5 kHz/120 = 646 Hz. Hence the expected timing accuracy is in the 1.5 ms range with the improvement brought by the cross-correlation peak fitting, which yields an observed sub-10 μs accuracy.

The challenge of phase demodulation lies in reproducing a local copy of the unmodulated carrier in order to allow for phase variation detection. Indeed, frequency f being the derivate of the phase φ in the expression

$$s(t) = \cos(2\pi ft + \varphi(t)) \tag{2}$$

the phase $\Phi = 2\pi ft + \varphi$ can be considered as split between a component linearly time varying with time $2\pi ft$ and a random component including the signal to be demodulated φ. Recovering φ hence requires an accurate estimate of f so that the mixing with the carrier yielding $f - f_c$ cancels and only φ is left in the expression of Φ. While amplitude demodulation only requires that $f - f_c$ lies within the low-pass filter bandwidth, and no feedback control is usually implemented on f which is nominally close to f_c in amplitude demodulation, phase demodulation requires f to track f_c to compensate for environmental fluctuations and oscillator aging of f, the local copy of f_c. A coarse approach is to bring the radiofrequency signal close to baseband by multiplying with the nominal value of $f_c = 77500$ Hz in our case, and then take the Fourier transform of the resulting complex signal. The abscissa δf of the maximum of the Fourier transform provides the frequency offset between f and its nominal value: the resulting signal is hence again multiplied by $\exp(2\pi\,\delta f\,t)$ for the baseband to be centered on 0 Hz. Such a strategy is only as accurate as one Discrete Fourier transform bin, which is the decimated sampling rate divided by the number of samples of the Fourier transform. An improved frequency offset estimation scheme is to perform a linear fit on the resulting phase, and compensate for any residual frequency offset by subtracting the linear trend. The latter processing step has been implemented but does not significantly improve our phase cross-correlation computation capability. The general algorithm used to measure accurately the time of flight of the DCF77 signal with respect to the reference GPS 1-PPS signal is summarized in Fig. 5. The practical implementation of these algorithm steps are given in appendix A.

However, this carrier frequency tracking solution already provides one result on frequency transfer: the sound card local oscillator will be affected by local environmental variations, most significantly temperature variations, readily observed with respect to the reference atomic clock signal received from DCF77. Plotting the frequency correction as a function of time – all records are timestamped with respect to Coordinated Universal time (UTC) – and comparing with the temperature history in Besançon (as provided by the local airport METAR logs provided at `https://www.wunderground.com/history/airport/LFSA/`), a clear correlation is observed (Fig. 6), as expected from the poor thermal insulation of the laboratory in which this experiment is taking place. The frequency offset of 0.8 Hz at 77500 Hz indicates a 10 ppm offset, with a temperature dependence of ±0.5 Hz for temperature variations of ±10°C, or a 0.6 ppm/K temperature dependence, a reasonable value for a quartz oscillator operating close but below its turnover temperature.

We have completed the frequency transfer investigation. However, observing ionosphere altitude fluctuation requires solving the time transfer issue, which is addressed in the next section.

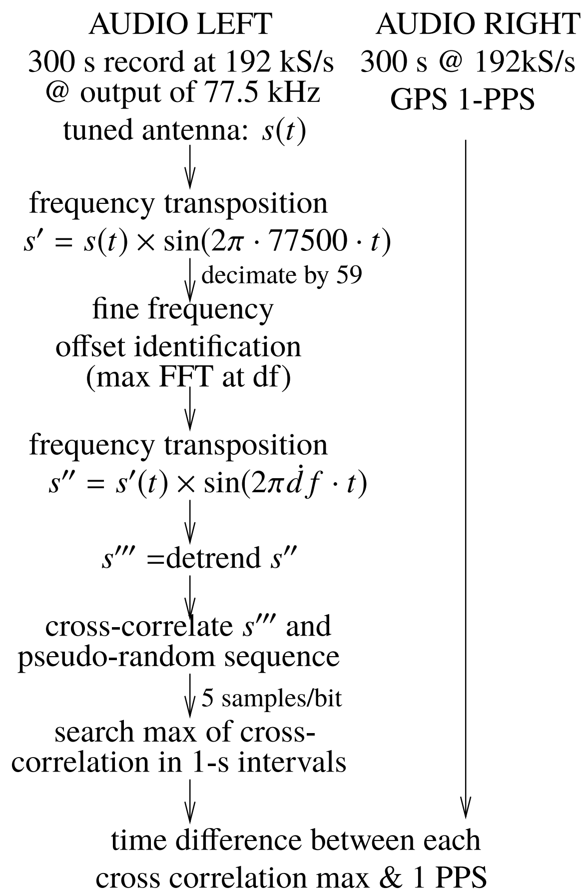

Figure 5. Processing steps applied to 60-second long measurement sequences of the DCF77 signal and the GPS-1PPS signal. Both signals are sampled with the stereo channels of a sound card clocked by the same reference signal, which is thus rejected when comparing one signal with respect to the other. Detrending involves identifying the linear trend on the dataset and removing this linear drift component. The decimation factor of 59 was selected for the decimated sampling rate of 192/59 kHz to closely match a small integer number of samples in the duration of one bit, namely 5 samples/bit as explained in the text.

4 Timing analysis

Having shifted the frequency to a baseband centered on 0, the phase $\Phi = \varphi(t)$ only exhibits variations introduced by the phase modulation scheme. Reproducing this sequence locally, and resampling so that an appropriate number of phase values match the duration of each sampled bit, a cross-correlation of both signal yields sharp cross-correlation peaks once every second (Fig. 7). The GNU/Octave listing given in appendix A exhibits the core processing steps and illustrates a typical processing chain implementing as software the most common components found in a typical radiofrequency receiver, including

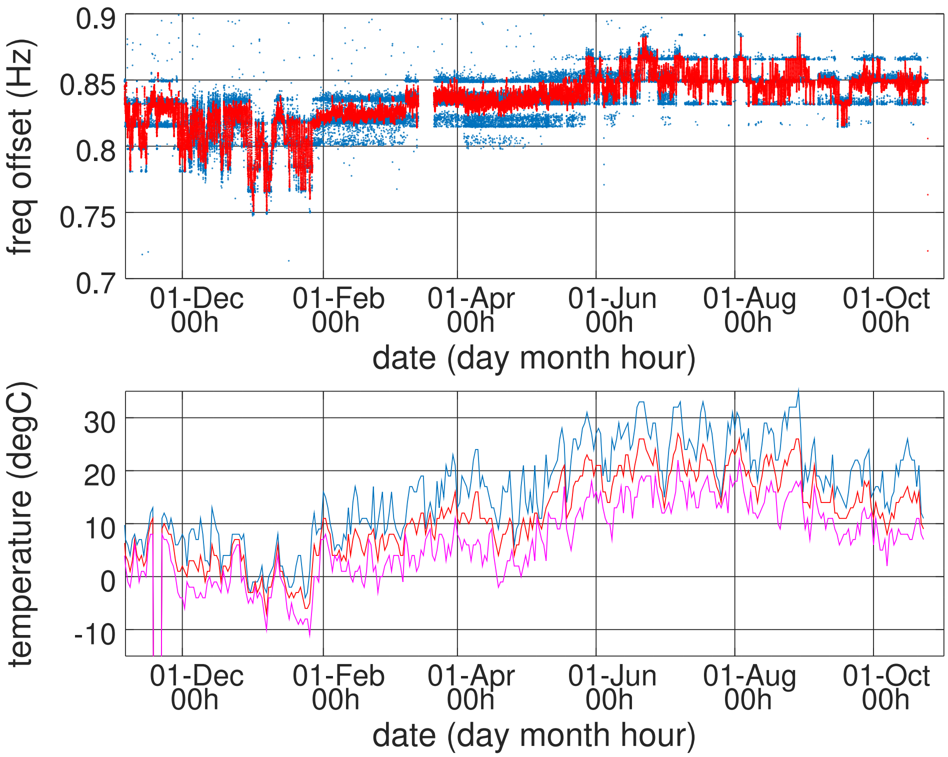

Figure 6. Top: frequency offset between the nominal received frequency of 77500 Hz, and the correction brought to bring the LF signal to a baseband centered on 0 Hz. The red curve is a sliding average over 10 samples (50 minute integration time) of the raw data shown in blue (each blue dot is the result of processing 60 second acquisitions). Bottom: history of the daily average of the temperature of Besançon airfield in Thise (METAR logs of LFSA callsign). The red curve is the mean daily temperature value recorded at the airfield, blue is the maximum and magenta is the minimum tempature recorded during each day.

mixing, low-pass filtering, remote oscillator frequency tracking by the local oscillator (i.e. demodulation), and signal decoding. A first coarse frequency offset between the received signal and the local oscillator is estimated from the position of the Fourier transform maximum. From this offset, a local oscillator signal is generated using a time signal synthesized with steps equal to the inverse of the sampling rate, and a dot product simulates the multiplier component that would be used otherwise for frequency transposition. Having removed the coarse frequency offset, a linear fit on the phase removes the residual linear trend of the phase, also known as frequency offset (since the derivate of the phase is the instantaneous frequency). A pre-computed pseudo-random code sequence is loaded and re-sampled at the same rate as the data recorded by the sound card. Following all these steps, the cross-correlation between the pseudo random sequence and the phase whose linear drift has been removed must exhibit a sharp peak once every second, when both patterns match.

The cross-correlation between the detrended phase and the pseudo-random sequence is computed, having previously removed the mean value of each signal to prevent a triangular baseline variation due to the integral over a constant offset: the cross-correlation exhibits maxima every time the pseudo-random pattern is met in the phase of the recorded signal, as seen on Fig. 7 (b) and (d). The improvement in the timing accuracy is emphasized by comparing the amplitude modulation (Fig. 7 (a) and (c)) indicating the beginning of each second, with the phase cross-correlation peak (Fig. 7 (d)): amplitude modulation being prone to link budget fluctuations and not being locked on the carrier during the demodulation which only consists of a rectifying and low-pass filtering, a narrowband low-pass filter induces bit spreading and degrades the timing resolution. Nevertheless, the two possible widths on the amplitude modulation encoding the one and zero values (short and long pulse) are well observed (Fig. 7 (e)). On the other hand, the spectrum spreading introduced by the phase modulation narrows the cross-correlation peak, allowing for much better timing analysis (Fig. 7 (f)). The time resolution gain on the phase cross-correlation measurement is visible by observing the width of the cross-correlation peak rising edge with respect to the amplitude pulse rising edges, both signals being synchronized on the falling edge.

Estimating the accuracy of this decoding step requires a local copy of a timing signal assumed to be a reference. We have compared the DCF77 cross-correlation peak timing with the 1 PPS of GPS receivers designed for timing application: U-Blox (Switzerland) provides low-cost ($<$ 90 euros) GPS receivers with the timing option of the 1 PPS output. The sound card recording DCF77 is hence configured in stereo mode, with the second channel recording the GPS 1 PPS output.

Comparing the time of arrival of DCF77 and GPS, the latter assumed to be negligibly affected by ionosphere delay in this configuration (sub-100 ns [*Giffard*, 1999]), yields a chart of time evolution exhibited in Fig. 8. The records are performed once every 5 minutes, timestamped with the computer time set to UTC, with 1 minute long records requiring 4 to 5 minute processing on the low performance DELL Latitude E6500 (Intel Core2 Duo CPU, 2.53 GHz, 4 GB RAM) laptop used here. As expected from the litera-

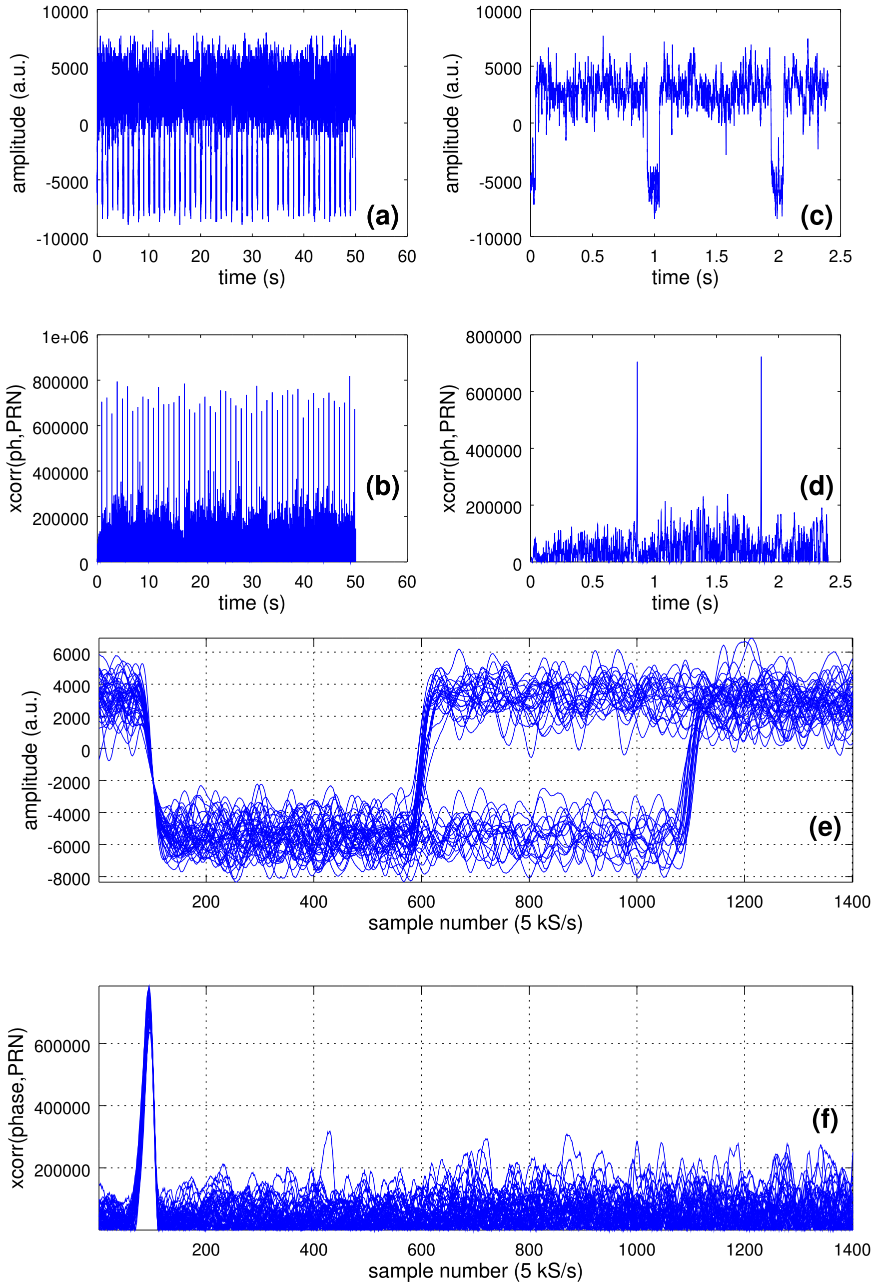

Figure 7. (a) and (c): amplitude demodulation, exhibiting dips every second (a) representative of timing marks (c). (b) and (d): phase cross correlation, again with cross-correlation peaks repeating every second (b) for a precise time transfer (d). (c) and (d) are zooms on 2.5 s-long parts of the (a) and (b) records. (e) and (f): comparison of the AM v.s PM cross-correlation timing accuracy by displaying a stack of 20 consecutive pulses. The Y-axis labeled "xcorr(ph,PRN)" indicates that the magnitude of the cross correlation between the phase samples and the Pseudo Random Number (PRN) sequence encoding the DCF77 phase is displayed.

ture, the ionosphere is unstable during winter time, with fluctuations in the hundreds of microsecond range. More interestingly, spring time brings ionosphere stabilization, with a

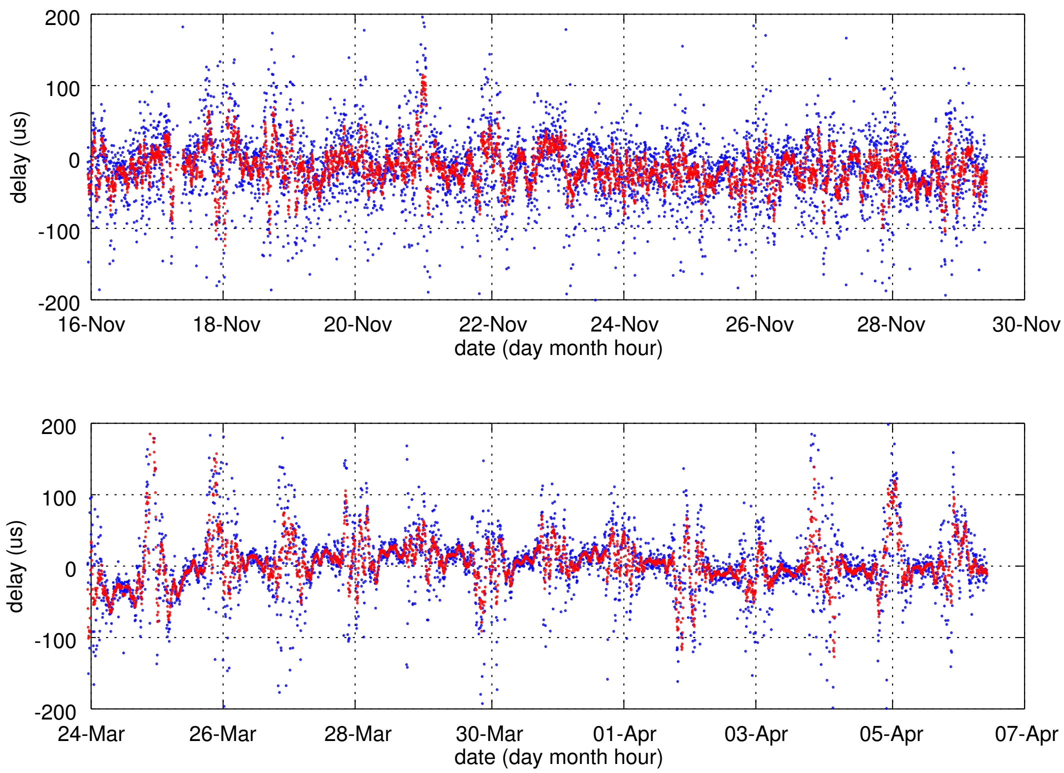

Figure 8. Comparison of the time difference between DCF77 and GPS 1 PPS in November (top), as the ionosphere is not stable during winter, and April (bottom), as the ionosphere stabilizes during daytime in spring and summer. The red dots represent data resulting from a sliding average over ten samples of the raw data shown as blue dots, which are themselves measurements integrated over 1 minute intervals (average of 60 DCF77 timing estimates with respect to GPS 1 PPS). All chart abscissa refer to time in UTC, with the date refering to the 0:00 hour of each day.

clear observation of the ionosphere delay stabilization during day time, as the upper layers of the atmosphere are exposed to solar ionizing radiation particles, and loss of stabilization during night. The stabilization matches the sunrise and sunset dates (Fig. 9).

Amongst the fascinating consequences of monitoring the LF propagation duration over a long duration is the hint of some interaction between the upper Earth crust – the lithosphere – and ionosphere as observed during earthquakes. [*Kumar and Kumar*, 2007; *Molchanov et al.*, 1998; *Chakrabarti et al.*, 2005; *Hayakawa et al.*, 1997]. The carrier frequency considered here seems to be too high to allow for the observation of cosmic particle fluctuation as observed from NOAA's geostationary GOES satellites. Such effects – Sudden Ionospheric Disturbances (SID) monitoring – is classically performed [*Dolea*

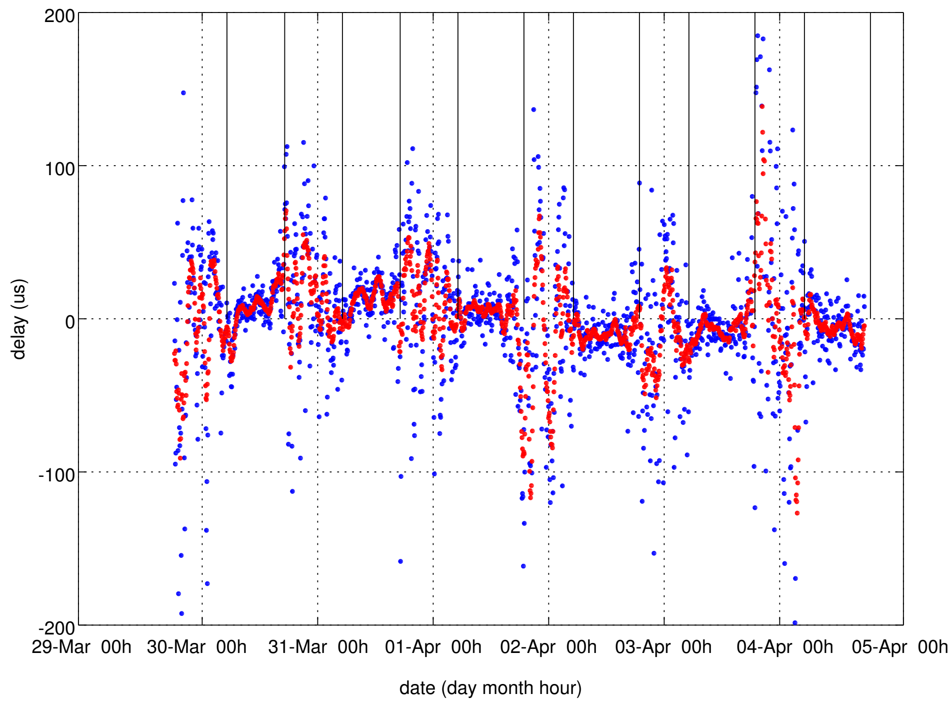

Figure 9. Short term analysis of the DCF77 timing delay with respect to sunrise and sunset times as calculated by the USNO application available at `http://aa.usno.navy.mil/data/docs/RS_OneYear.php`: the ionospheric delay stabilization when sun rises (vertical lines, alternatively sunset and sunrise time) is clearly visible in this chart. All chart abscissa refer to time in UTC, with the date refering to the 0:00 hour of each day.

et al., 2013] by observing the *amplitude* variation of the LF signal rather than its time of flight as considered here.

5 Timing accuracy

A detailed estimate of the accuracy of the time transfer needs to consider the evolution of the offset between GPS 1-PPS and DCF77 (Fig. 10) with integration time. Furthermore, let us remember that the rationale for maintaining VLF timing broadcast systems such as DCF77 (similar to WWVB in the North America or JJY in Japan) is the long term synchronization of quartz-controlled clocks whose excellent short term stability is given by the resonator but long term stability is poor due to aging, temperature dependence and offset with the nominal frequency with respect to the primary standards: despite the daily fluctuations of several tens to hundred of microseconds, the long term mean value exhibits no visible drift (Fig. 10, top) despite varying environmental conditions including space weather (Fig. 10, bottom). Controlling the quartz oscillator with a signal extracted from the VLF timing measurements to generate a stable composite signal exhibiting the best stability of both systems requires assessing the time constant of the feedback loop. Such a measurement is classically performed through the Allan deviation analysis of both clocks: the integration time at which the curves intersect defines the feedback loop time constant, as illustrated in Fig. 11.

DCF77 measurements are computed every 5 minutes following an integration of 60 pulse timings with respect to GPS 1-PPS. The timing accuracy is hence given by averaging the time offsets normalized to this measurement duration: as an example, a 50 μs uncertainty over a 5 minute measurement interval yields a relative accuracy of about $50 \cdot 10^{-6}/(5 \times 60) \simeq 2 \times 10^{-7}$. This result is indeed the first value in the Allan deviation plot exhibited in Fig. 11, in which the $1/\tau$ slope with τ the integration time is observed, indicating the lack of long term drift and stable time transfer with improved accuracy as integration time increases. Such a trend contrasts with that of a quartz tuning fork controlled oscillator, which exhibits better short term stability owing to the high quality factor of the quartz tuning fork, but drifts over long terms to exhibit long term instability greater than those of the VLF signal. The intersection of the two curves, around 1000-2000 s, defines the feedback loop constant to control the quartz tuning for with the VLF signal. The proposed setup is hence well suited for a digitally controlled quartz oscillator locked on the phase information provided by DCF77: we are aware of a single commercial product implementing such a functionality, namely by Meinberg (Bad Pyrmont, Germany).

6 Cross-polarization measurements

In a propagating beam model, as opposed to a waveguide model in which the Earth surface and ionosphere define conducting boundary conditions, the LF wave propagates along two paths, one along the Earth surface and the other one reflecting on the ionosphere. Since the emitter generates a vertically polarized wave and the receiver coil is horizontal for the magnetic flux to induce a current in each coil, the strongest wave component dominates the received signal, making the identification of the wave bouncing off the ionosphere challenging. Since the wave reaching the ionosphere interacts with an ionized medium with free charges in a magnetic field, polarization rotation occurs through the Faraday effect, which might provide a solution for separating the air wave from the

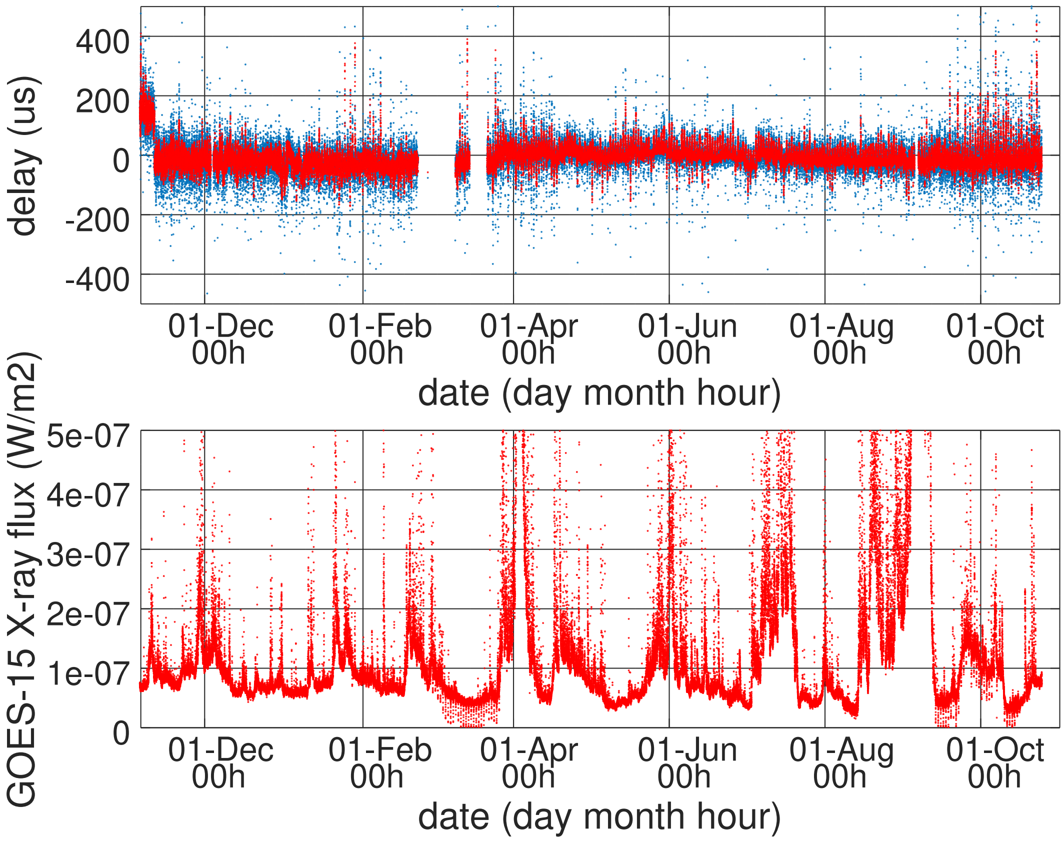

Figure 10. Long term investigation of the delay between the atomic clock-disciplined DCF77 and GPS 1-PPS (top), compared to the X-ray flux observed by NOAA's GOES geosynchronous satellite observations, as available from `ftp://ftp.swpc.noaa.gov/pub/lists/xray/` (5 minute interval records from the primary sensor). No correlation between the two quantities is visible, probably because 77.5 kHz is too high a frequency to detect such phenomena. The stabilization of the ionospheric propagation properties during spring and summer are well visible as the reduced fluctuation in the middle part of the top chart (spring and summer) with respect to the left and right (winter and autumn), with zooms in relevant regions provided in Fig. 8. The phase jump after the first week of measurement is associated with a slight change in the phase slope analysis for unwrapping, emphasizing the influence of the signal processing chain on the absolute phase evaluation. The algorithm was no longer modified after this initial change to ensure continuity of the measurements. All chart abscissa refer to time in UTC, with the date refering to the 0:00 hour of each day.

ground wave. By performing simultaneously two measurements, one with a horizontal coil (sensitive to the ground wave – no polarization rotation) and a with a second setup using a vertically oriented coil (insensitive to the ground wave), the air wave is separated and the time delay analyzed (Fig. 12). Since an electromagnetic wave propagates with the

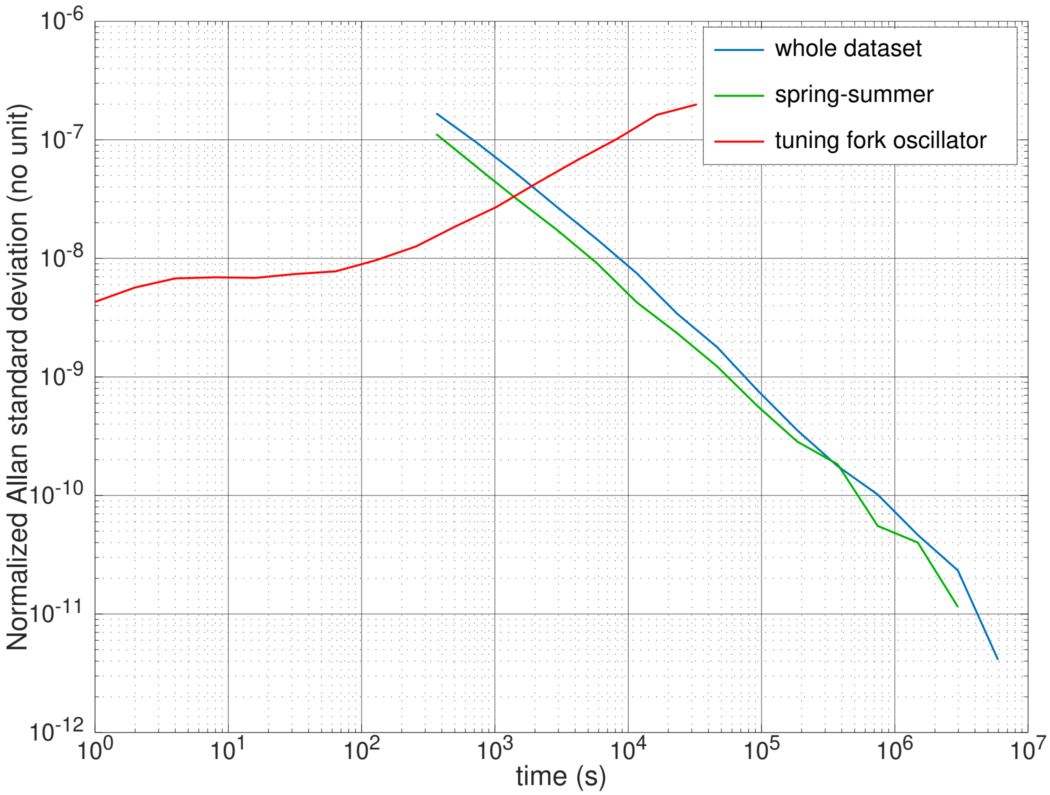

Figure 11. Allan deviation of the time offset between GPS 1-PPS and DCF77 (blue), and of a 32768 Hz tuning fork oscillator as classically found in wrist watches (red). The intersection of the two curves provides the time constant of the composite clock in which the DCF77 signal could be fed back to the tuning for oscillator to correct long for term drift of the latter. The green curve exhibits the Allan deviation of the spring and summer dataset, starting April 1st, when the ionosphere has stabilized during daytime, improving the time transfer stability.

wavevector $\vec{k}$, electric field $\vec{E}$ and magnetic field $\vec{B}$ normal to one another, the detected electric field is along the radius of the coil. Hence, an horizontal ferrite antenna with the plane containing the coil oriented vertically detects the vertical linearly polarized electric field, and a vertical ferrite antenna with the plane of the coil horizontal detects the linearly polarized horizontal electric field.

In order to reject systematic delay, the setup was rotated 90° half-way during the experiment to check that the delayed channels would switch as the horizontal and vertical antenna channels were exchanged. Such a result was indeed observed. The mean value of the delay between the two channels is 170 ±60 μs (Figs. 13 and 14), surprisingly close to the expected value deduced from a geometric raytracing model. However, the poor signal

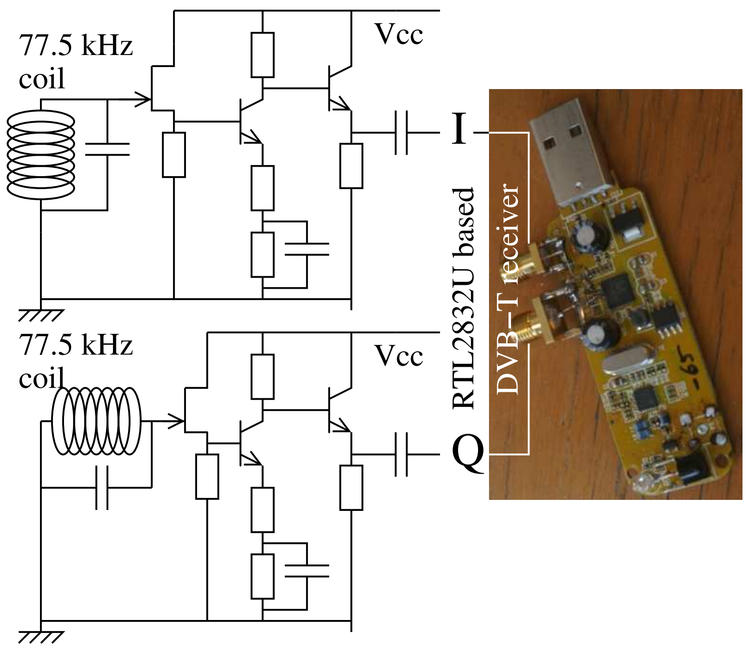

Figure 12. Crossed polarization measurement: two identical setups are connected to the I and Q inputs of a RTL2832U based DVB-T receiver.

to noise ratio of the vertically polarized antenna prevented identifying day/night fluctuations. Indeed, some negative delay was observed, as opposed to the predicted delay of the air wave with respect to the ground wave: such measurements were however excluded following a quantitative criterion of signal to noise ratio on the vertically polarized receiver. Fig. 13 illustrates this analysis: the selected criterion is inverse of the average of the two cross-correlation values located at the vertical arrows **c1** and **c2**. Since the cross-correlation peaks have been normalized, the inverse of the mean value of **c1** and **c2** provides an indicator of a signal to noise ratio, with measurements rejected if this criterion is below 15. Each curve set in Fig. 13 includes two traces: one for the horizontal polarization and one for the vertical polarization. Since the horizontal ferrite antenna (vertical electric field component) always exhibits excellent signal to noise ratio, all curves overlap on the left-most reference cross-correlation peak. Poor signal to noise ratio exhibited by the red and magenta curve yield strong dispersion on the position of the second cross-correlation peak, while acceptable signal to noise ratio following the proposed criterion yields to overlapping blue and magenta measurement correlation peaks (right, horizontal electric field component), allowing for precise time of flight difference measurement with respect to the reference cross correlation peak (left, vertical electric field component). The values to the right of the graph in Fig. 13 indicate the measured time of flight difference: the two cases of acceptable signal to noise ratio yield close results of 187 and 200 μs respectively, while the two cases of poor signal to noise ration yield excessively dispersed results, here 291 and 83 μs respectively.

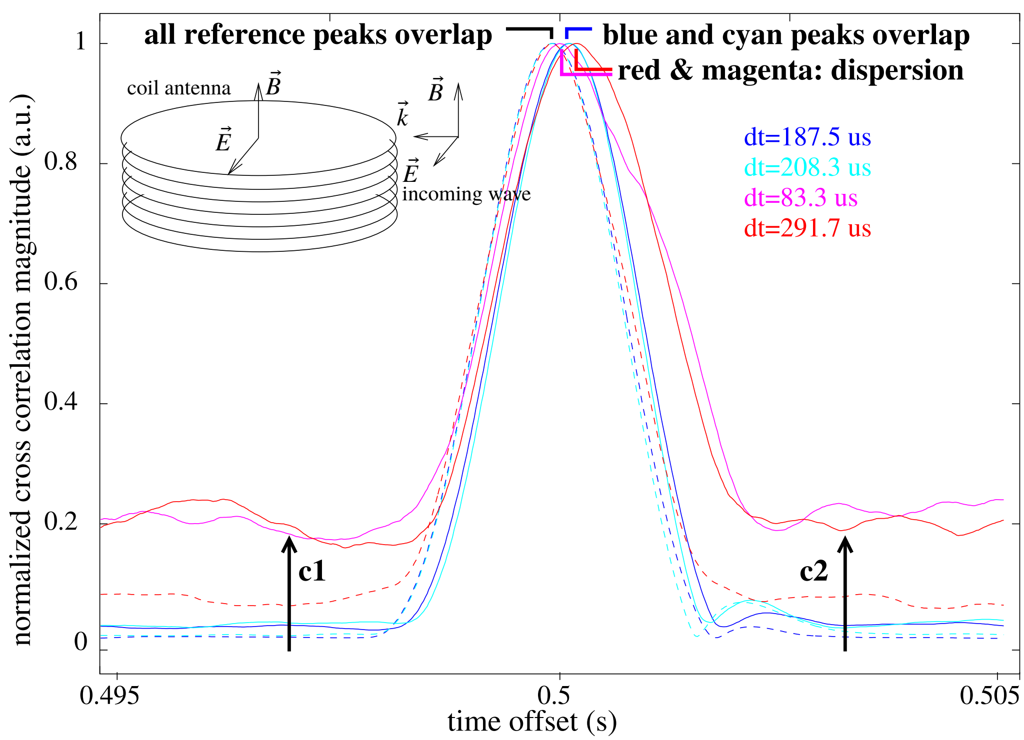

Figure 13. Cross-correlation curves for high signal to noise measurements (blue and cyan), and low signal to noise (red and magenta) ratios. Each dataset exhibits two curves, one for the reference (vertical polarization – dashed line) and one for the measurement (horizontal polarization – solid line). Signal to noise ratio (SNR) is defined by the normalized cross correlation peak maximum to the baseline (positions **c1** and **c2**) value. High SNR yields accurate time delay difference between the vertically (left-most cross-correlation peak) and horizontally polarized (right-most cross-correlation peak) waves. Inset: the antenna current is generated by the magnetic flux through the coil.

7 Conclusion

Software defined radio and digital signal processing are used to analyze a high stability time and frequency transfer signal emitted at very low frequency by the German DCF77 emitter. Since the propagation of this signal is dependent on ionospheric conditions and especially the altitude of the layer with the electron density whose plasma fre-

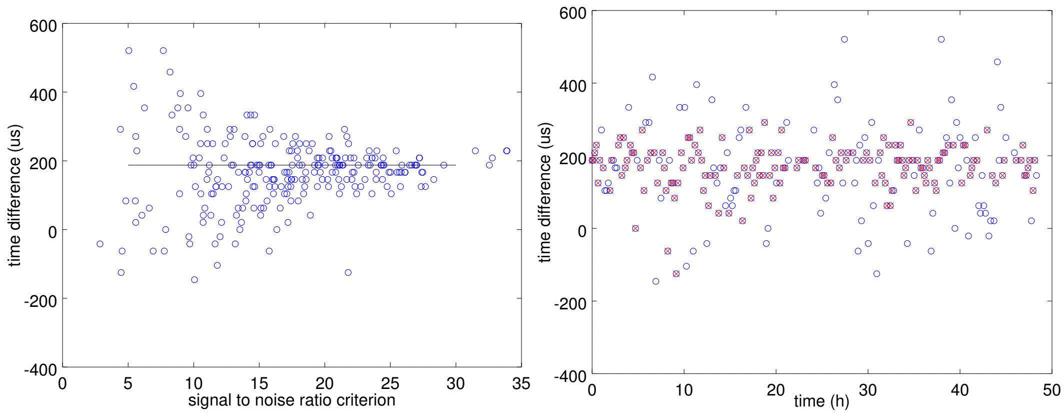

Figure 14. Crossed polarization measurement: the SNR criterion (left) was applied to reject erroneous measurements. Right: while all measurements over a two-day period exhibit significant dispersion, primarily due to the poor SNR of the vertically polarized antenna, selecting the data with a criterion above 15 yields a time delay between vertically and horizontally polarized signals of 170 ±60 μs or a median value of 190 μs. Blue circles are all the measurements, amongst which only the red crosses meet the criterion defined above and are considered in the delay calculation.

quency matches the radiofrequency wave frequency, the time of flight is representative of the ionosphere altitude variation. Daily and seasonal variations are readily observed, thanks to the improved timing capability of the pseudo random phase modulation added to the coarse amplitude modulation used for time transfer. The temperature dependence of the local oscillator of the receiver is also observed with this setup, which solely consists of an antenna, impedance matching circuit and personal computer sound card.

Such a basic setup is designed for dissemination and long term monitoring activity for its low cost and ease of assembly. The performance, allowing for 10 μs time of flight measurement, is suitable for observing daily ionospheric condition variations through timing analysis rather than the classical amplitude measurement. Daily variations of more than 100 μs are readily observed, as are the seasonal ionosphere stabilization during spring and summer and instability from the end of autumn to winter. From the authors laboratory location at a range from the emitter at which the ground wave and air wave exhibit comparable amplitude, the vertical (direct) and horizontal (reflected) components of the electric field exhibit a relative time delay consistent with the expected geometrical model of wave reflection on the ionosphere.

Acknowledgments

Andreas Bauch (PTB, Germany) prompted this investigation with his course on time transfer at the European Frequency and Time Seminar (`efts.eu`). Franck Lardet-Vieudrin (FEMTO-ST, France) provided support in designing and understanding the operating principle of the short antenna. François Vernotte (Besançon Observatory, France) provided the explanation on the conversion of time intervals to normalized quantities for Allan deviation analysis. Eric Meyer (Besançon Observatory, France) prompted the investigation on the cross-polarization time delay measurement. This work was partly supported by the Programmes d'Investissements d'Avenir (PIA) FirstTF and Oscillator IMP grants. All datasets are made available to readers at `http://jmfriedt.free.fr/dcf77`.

References

ARRL (1997), *ARRL Antenna Handbook, 18th Ed.*, chap. Small Loop Antennas, ARRL.

Baker, D. N., and L. J. Lanzerotti (2016), Resource letter SW1: Space weather, *American Journal of Physics*.

Bauch, A., P. Hetzel, and D. Piester (2009), Time and frequency dissemination with DCF77: From 1959 to 2009 and beyond, *PTB-Mitteilungen*, *119*(3), 3–26.

Blackband, W. (1964), *Propagation of Radio Waves at Frequencies Below 300 Kc*, on behalf of Advisory Group for Aeronautical Research and Development, North Atlantic Treaty Organization.

Carlà, M. (2016), Measure of 1/f noise using the sound card of a PC, *American Journal of Physics*, *84*(4), 311–316.

Chakrabarti, S. K., M. Saha, R. Khan, S. Mandal, K. Acharyya, and R. Saha (2005), Unusual sunset terminator behaviour of VLF signals at 17kHz during the earthquake episode of Dec., 2004, *Indian J. Radio and Space Phys*, *34*, 314–317.

Davies, K. (1990), *Ionospheric radio*, 31, IET.

Dolea, P., V. P. Dascal, O. Cristea, and T. Palade (2013), In-situ measurements regarding lf radio wave propagation using DCF77 time signal transmitter, in *Telecommunication in Modern Satellite, Cable and Broadcasting Services (TELSIKS), 2013 11th International Conference on*, vol. 2, pp. 449–452, IEEE.

Engeler, D. (2012), Performance analysis and receiver architectures of DCF77 radio-controlled clocks, *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, *59*(5), 869–884.

Giffard, R. (1999), Estimation of gps ionospheric delay using L1 code and carrier phase observables, in *31st Annual Precise Time and Time Interval (PTTI) Meeting*, pp. 405–416, http://www.dtic.mil/get-tr-doc/pdf?AD=ADA497270.

Gotoh, T., J. Amagai, T. Hobiger, M. Fujieda, and M. Aida (2011), Development of a GPU-based two-way time transfer modem, *IEEE Transactions on Instrumentation and Measurement*, *60*(7), 2495–2499.

Hayakawa, M., O. Molchanov, T. Ondoh, and E. Kawai (1997), On the precursory signature of kobe earthquake on VLF subionospheric signals, in *IEEE International Symposium on Electromagnetic Compatibility*, pp. 72–75, IEEE.

Hetzel, P. (1988), Time dissemination via the lf transmitter DCF77 using a pseudo-random phase-shift keying of the carrier, in *Proceedings of the 2nd European Frequency and Time Forum (EFTF)*, pp. 351–364.

Huang, Y.-J., M. Fujieda, H. Takiguchi, W.-H. Tseng, and H.-W. Tsao (2016), Stability improvement of an operational two-way satellite time and frequency transfer system, *Metrologia*, *53*(2), 881.

Johler, J. R. (1962), Propagation of the low-frequency radio signal, *Proceedings of the IRE*, *50*(4), 404–427.

Kamp, P.-H. (), A cheap SDR Loran-C frequency receiver, phk.freebsd.dk/AducLoran/AducLoran-0.3.pdf.

Kumar, S., and A. Kumar (2007), Diurnal variation of 19.8 kHz signal propagation over long path to suva, *The South Pacific Journal of Natural Science*, *11*, 67–69.

Mindell, D. A. (2011), *Digital Apollo: human and machine in spaceflight*, Mit Press, Cambridge, MA, USA.

Mochizuki, K., M. Uchino, and T. Morikawa (2007), Frequency-stability measurement system using high-speed adcs and digital signal processing, *IEEE Transactions on Instrumentation and Measurement*, *56*(5), 1887–1893.

Molchanov, O., M. Hayakawa, T. Oudoh, and E. Kawai (1998), Precursory effects in the subionospheric VLF signals for the Kobe earthquake, *Physics of the Earth and Planetary Interiors*, *105*(3), 239–248.

Piester, D., A. Bauch, J. Becker, and A. Hoppmann (2011), Time and frequency broadcast with DCF77, in *Proc. 43rd Annual Time and Time Interval (PTTI) Systems and Applications Meetings*, pp. 185–196.

Raupach, S. M., and G. Grosche (2014), Chirped frequency transfer: a tool for synchronization and time transfer, *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, *61*(6), 920–929.

Schulte, C. H., G. M. Müller, H. Horn, J. Hübner, and M. Oestreich (2012), Analyzing atomic noise with a consumer sound card, *American Journal of Physics*, *80*(3), 240–245.

Sherman, J. A., and R. Jördens (2016), Oscillator metrology with software defined radio, *Review of Scientific Instruments*, *87*(5), 054,711.

Uchino, M., and K. Mochizuki (2004), Frequency stability measuring technique using digital signal processing, *Electronics and Communications in Japan (Part I: Communications)*, *87*(1), 21–33.

Watt, A., R. Plush, I. Brown, and A. Morgan (1972), Worldwide VLF standard frequency and time signal broadcasting, *Precision Measurement and Calibration*, *5*, 297.

A: GNU/Octave implementation of the decoding sequence

```
x=read_complex_binary(d);
dcf=real(x);
gps=imag(x);
fs=192e3;
time=[0:length(x)-1]'/fs;          % fs = sampling rate
```

The file named `d`, created by GNURadio as a binary record with floating point format alternating the left and right audio channels, recording the DCF77 antenna output and GPS 1 PPS signal respectively, is read and the time index is created with steps given by the inverse of the sampling rate `fs`.

```
dcf=dcf.*exp(j*2*pi*(77500)*time);
lpf=firls(250,[0 720 790 fe/2]*2/fe,[1 1 0 0]);
dcf=filter(lpf,1,dcf);
x=dcf(1:59:end);
time=time(1:59:end);
```

The signal is transposed from radiofrequency band (77.5 kHz) to baseband by a multiplication with the local oscillator synthesized digitally as a sine wave with angular pulsation $2\pi \times 77500$ rad/s. The low-pass filter removes noise and unwanted parasitic components from the mixing step: indeed, the magnitude of the Fourier transform of the real signal `dcf77` is even, and the frequency transposition creates a spectral component at -77.5-77.5=-150 kHz which is aliased to 192-150=42 kHz, eliminated by the low-pass filter. Once the signal is brought to baseband, the whole bandwidth is no longer needed since the signal is only located a few kHz around baseband: excess samples are discarded by decimating by 59, and time is decimated similarly, equivalent to dividing the sampling rate by this same factor. The decimation factor of 59 was selected considering the known bit-rate of the signal emitted by DCF77, namely 120 periods of the 77500 Hz carrier, or 1.5484 ms. The decimation factor of 59 was selected to have a small integer number of samples during each bit: $59/192 = 0.3073$ ms which is $1.5484/0.3073 = 5.04$ close to 5 samples/bit. Such a selection will make the cross-correlation with a pseudo-random code re-sampled to the selected sampling rate easier to analyze.

```
[yf,xf]=max(abs(fft(x-mean(x)))); % coarse frequency offset identification
xf=xf-length(x)-1;
df=-xf/length(x)*fs               % index to frequency conversion
lo=exp(j*2*pi*df*time);  % transpose by xf (fe->fe+xf ou fDCF->fDCF-xf)
x=x.*lo;
```

Following the transposition from radiofrequency band to baseband by the nominal frequency offset, a fine tuning of the difference between the local oscillator frequency and remote oscillator frequency is identified as the frequency at which the Fourier transform is maximum. This Fourier transform index is converted to a frequency by remembering that a discrete Fourier transform over N samples spans from minus half of the sampling frequency to half of the sampling frequency, or a bin size of `fs`/N. Again the multiplication brings the signal exactly on the baseband 0-Hz frequency.

```
[u,v]=polyfit(time,xp,1);         % once coarse offset removed, linear fit on phase
x=x.*exp(-j*time*u(1)-j*u(2));    % linear phase shift = frequency offset
xp=angle(x);                      % phase modulation ...
```

Since we aim at demodulating a phase-modulation, any leftover phase drift must be removed. The frequency is the derivate of the phase, so that the previous step might have left a fine phase drift with a slope below the bin size of the Fourier transform: a linear polynomial fit gets rid of the fine linear drift, or residual frequency offset. These last fine-tuning steps must be repeated for each new record since the local oscillator frequency, clocking the sound card, fluctuates over time with environment (Fig. 6).

```
load lfsr.dat
np=192000/59*(120/77500);         % PRN chip length (120 periods of carrier)
oldP=0;
for k=1:length(lfsr)
  P=round(k*np);                  % resample
  if (lfsr(k)==1) longlfsr(oldP+1:P)=ones(P-oldP,1);
  else longlfsr(oldP+1:P)=zeros(P-oldP,1);
  endif
  oldP=P;
end
```

Having recovered a fine estimate of the received signal phase, we aim at extracting the pseudo-random phase sequence imprinted on the carrier. The bit-sequence generated by the polynomial was computed and stored in a `lfsr.dat` file as described in section 3, with a rate of 1 sample/state. The sampling rate resulting from the decimation was selected to have a number of samples of the phase close to an integer number of samples of the phase encoding: at 120 periods/phase state, the number `np` of samples is $192000/59 \times (120/77500) = 5.04$, close enough to 5 for the 512 sample long pseudo-random code to be easily re-sampled to match the current sampling rate: each bit is copied enough time for the sampling rates to match, resulting in the `longlfsr` vector.

```
yc=xcorr(xp-mean(xp),longlfsr-mean(longlfsr));
yc=yc(floor(length(yc)/2):end);  % cross correlation result
```

Finally, the cross-correlation between the phase `xp` and the pseudo-random sequence `longlfsr` is computed, having previously removed the mean value of each signal to prevent a triangular baseline variation due to the integral over a constant offset: the cross-correlation `yc` exhibits maxima