

Results of a Software Defined Radio (SDR) implementation of Two Way Satellite Time and Frequency Transfer (TWSTFT) emitter and receiver system

J.-M. Friedt*, B. Chupin†, M. Lours†, É. Meyer*, O. Chiu†, F. Meyer*, W. Daniau*, J. Achkar†

*LNE-LTFB, FEMTO-ST T/F & Observatoire de Besançon, Besançon, France

†LNE-SYRTE, Observatoire de Paris - Université PSL, CNRS, Sorbonne Université, Paris, France

Emails: jmfriedt@femto-st.fr, Joseph.Achkar@obspm.fr

Abstract—Spectrum spreading of the carrier of an ultrastable clock using pseudo-random sequence generators implemented in a Field Programmable Gate Array (FPGA) is broadcast after bandpass filtering as a Binary Phase Shift Keying (BPSK) modulated signal towards a geostationary satellite using a commercial upconverter to the 14-GHz Ku band. The 70 MHz Intermediate Frequency (IF) numerically controlled oscillator, BPSK modulation and digital information carrying the timestamp of the second in the emitter are implemented as software in the FPGA gateway. The downlink and loopback signals are recorded by coherent dual channel SDR receivers and post-processed for fine time of flight analysis and TWSTFT link comparison. Results and lessons learnt during the development are described.

I. TWO WAY SATELLITE TIME AND FREQUENCY TRANSFER EMITTER AS SOFTWARE DEFINED RADIO

The challenge of implementing TWSTFT hardware and processing software lies in the ability to monitor time of flights with a resolution well below the inverse of the bandwidth available on the satellite transponder. While the classical time-frequency duality states in the general case that the time resolution of a measurement is the inverse of the signal bandwidth – also known as the radar range resolution – the assumption that a single correlation peak between the received signal and the local copy of the pseudo-random spectrum spreading sequence is a single sampling period wide allows for peak fitting and achieving better resolution than the inverse of the bandwidth. It is shown that the time of flight resolution improvement over the sampling period is equal to the signal to noise ratio of the received signal [1], so that high signal to noise ratio of a 2.5 Mcps/s BPSK signal recording (inverse of 5 MHz bandwidth would be 200 ns resolution) can achieve sub-picosecond time of flight measurement [2]. The requirement of a SDR implementation of TWSTFT emitter is hence to generate a spectrum-spreading pseudo-random sequence, either using Binary Phase Shift Keying (BPSK) or Quad Phase Shift Keying (QPSK) for redundancy [3], and on the receiver side to correlate the recorded signal with the emitted code after compensating for frequency offsets between local oscillators. Results on such a functional link has been

achieved [4] (Fig. 1) and the technical lessons learnt from this achievement are described here.




Fig. 1. Evolution of time differences (left) and corresponding time instabilities (right) of the two-way link OP-LTFB equipped with SATRE modems and SDR Tx/Rx platforms over a 3.5-week measurement. The “all SDR” measurement includes all SDR measurements at beginning and end of odd UTC hours, “ITU SDR” only uses 2-minute records once every odd UTC hour similar to SATRE schedules.

On both ends of the link, the IF signal output from the downconverter is recorded by an Ettus Research X310 SDR coherent dual-channel A and B receiver, with channel B recording the loopback signal transmitted locally for timing reference (Fig. 2). Post-processing involves identifying (correlating) the local copy of the pseudo random sequence in the

received and loopback signals, and oversampling by fitting the correlation peak in order to improve the timing resolution with respect to the sampling period by a factor equal to the signal to noise ratio. The targeted 200 ps standard deviation on the time of flight measurement requires a 1000-fold improvement over the 200 ns inverse of the channel bandwidth, hence determining the targeted signal to noise ratio on the downlink signal.

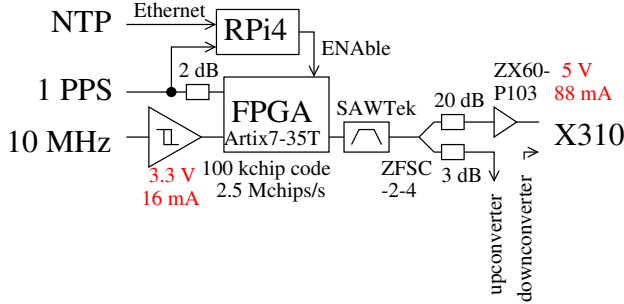


Fig. 2. Experimental setup of the SDR emitter: as described in the main text, the Raspberry Pi 4 (RPi4) actively monitors an NTP server for triggering the emission of the pseudo-random number (PRN) sequence by the FPGA on both ends of the link, with a reference channel connected to the loopback input of the dual coherent input X310 SDR receiver, the other input being connected to the VSAT downlink.

Great care was taken, when generating the IF signal from a digital system, to comply with emission requirements, by filtering unwanted out-of-band spectral components. A SAW bandpass filter (Sawtek 851547) centered on 70 MHz with better than 50 dB rejection removes any unwanted spectral component that might affect neighboring communication on the satellite transponder (Fig. 3).




Fig. 3. Signal generated by the FPGA at an IF frequency of 70 MHz and filtered by a Sawtek part number 851547 70 MHz bandpass filter with 2.5 MHz bandwidth (blue), demonstrating compliance with respect to the SATRE modem emitted spectrum (black) and lower sidelobes.

The drawback of recording and postprocessing is the size of the generated files: recording a complex (IQ) signal at 5 MS/s, with 16 bit/sample, on 2 channels (loopback and downlink) leads to $\Rightarrow 5 \cdot 10^6 \times 2 \times 2 \times 2 = 40 \text{ MB/s} = 13.2 \text{ GB}$ in 5.5 minutes or 144 GB/hour. Storing such files is no longer a technological challenge since consumer grade mass storage media will handle such files.

II. SURFACE ACOUSTIC WAVE DELAY LINE SIMULATOR

A major challenge when developing the processing algorithm is the slow motion of the satellite with a speed up to a few ns/s, well below the sampling period. While a Doppler frequency offset is readily simulated by offsetting the driving clock, such slow time delay variations are challenging to simulate in the digital domain. An high thermal sensitivity Surface Acoustic Wave (SAW) analog delay line operating at 70 MHz IF was designed and manufactured to act as a laboratory simulator of this slow drift of the time delay of the received signal with respect to the loopback signal. The low loss and large bandwidth (5 MHz at 70 MHz) is achieved by using the strongly coupled YXI/128° lithium niobate oxide (LNO) piezoelectric substrate with its nominal coefficient of velocity (frequency) with temperature of 70 ppm. Considering the Rayleigh wave acoustic velocity of 3979 m/s, then the $5.3 \mu\text{s}$ is achieved for a 21.3 mm-long delay line, patterned using a single lithography and aluminum patterning step on a 4-inch (100 mm) wafer (Fig. 4). The long delay line only induces minor losses – around 10 dB – thanks to the low-loss propagation and wide aperture preventing scattering of the Rayleigh wave (Fig. 5).



Fig. 4. SAW delay line for simulating the satellite motion by heating the piezoelectric substrate to induce acoustic velocity change and hence time of phase variations. The long delay line designed to induce sufficient delay variations when heating only induces minor losses since propagating a Rayleigh wave. The interdigitated electrodes visible on both ends of the propagation path free of electrodes, next to each SMA connector.

The long delay is selected to induce large enough temperature induced delay variations. Indeed, a 100 K temperature variation leads to a group delay variation of 37 ns/100 K considering the 70 ppm/K temperature sensitivity of YXI/128° LNO. Although this delay variation is only one sixth of the period at 5 MS/s, it is representative of half a period after interpolating the recorded signal by a factor of three (zero padding both ends of the Fourier transform when computing the correlation in the Fourier domain). Tripple transit echoes from re-emissions are visible in the time domain characteristics, but much shorter than the 40 ms (100 kchip long at 2.5 Mchips/s) code length we are using in our implementation of TWSTFT.

III. EMITTER SYNCHRONIZATION

Because the satellite is moving at a rate of $X \text{ ns/s}$, $X \in [-5 : +5]$, a timing mismatch between two speakers comparing their clocks using TWSTFT of 1 s would lead to an unacceptable error of up to $\pm 5 \text{ ns}$ between sessions. Hence, comparing signals broadcast by the two speakers within 40 ms is needed if sub-200 ps time delay fluctuations are to be




Fig. 5. Top: transfer function of the SAW delay line patterned on YXl/128° black lithium niobate with nominal 70 ppm/K temperature sensitivity – centered on 70 MHz and about 5 MHz wide – with 10 dB losses, as a function of substrate temperature. The ripples in the band pass are due to triple transit and re-emission since all interdigitated electrodes are split finger designs and the sides of the piezoelectric substrate were coated with acoustic wave absorbing resist to avoid mechanical reflections. Bottom: time domain response as inverse Fourier transform of the frequency domain characterization, exhibiting the main transfer delay of about 6 μ s and the subsequent odd triple-transit reflections leading to the ripples in the bandpass visible in the top chart.

achieved. The proposed scheme for synchronizing emissions is as follows:

- an NTP client – in our case running on a Raspberry Pi 4 single board computer – actively monitors the time shared over the network to identify the second prior to the start of the next session,
- after a delay of 20 ms, the enable pin of the SDR emitter is activated, allowing to trigger the PRN sequence when the next PPS rising edge is detected,
- assuming the 1-PPS edges are aligned on both ends, the next-PPS starts the transmission and both emission are synchronous.

This strategy however leads to a chicken and egg problem: TWSTFT is designed for clock comparison and steering, but with this approach requires that both PPS edges are aligned already. Lacking the alignment condition, the initial TWSTFT comparison will lead in ± 5 ns alignment, which will end up cancelling if the 1-PPS control loop aligns the edges. In this sense, TWSTFT can be considered as a time to delay converter with a factor of 5 ns/s scaling between the broadcast signal synchronization and the receiver clock synchronization.

IV. RECEIVER DRIFT

The opensource, openhardware implementation of the TWSTFT emitter aims at being hardware agnostic and vendor independent since the BPSK modulated PRN sequence at the IF frequency is fully implemented in the FPGA, and the receiver part only requires dual coherent channel analog to digital converters at baseband if the IF to baseband transposition is performed externally, or with sufficient bandwidth

to sample the IF signal and software transpose to baseband. These features are readily found on a multitude of opensource, openhardware platforms including the Fairwaves (now Lime Micro) XTRX, Pluto+, Ettus Research B210 or X310. Our current experiments rely on the latter for reception and on a dedicated FPGA board – Digilent CMOD-A7 – for ease of development and synthesis speed. Using the same hardware on all transceivers might however cancel some common mode source of drift during the TWSTFT analysis: such an effect is highlighted in this section.

In the current X310 usage, the SDR receiver is idle until 30 seconds the next session is expected to start. Because the GNU Radio scheduler and SDR hardware require some time to initialize, recording starts about 20 s before emission is started, keeping some margin for unexpected initialization delays but making sure the first loopback bit locally emitted is recorded. Each session lasts 5 minutes, then the recording is switched off and the receiver returns to an idle state. The power consumption of the FPGA and the ADC is not the same whether the receiver is idle or active, and a warmup is observed as a drift (Fig. 6). Since the same X310 in the same casing is used on both ends of the link, the same drift is observed on both loopback signals which cancel when computing the time delay differences. Identifying such an issue is fundamental towards hardware agnostic solutions since different hardware platforms exhibit different warmup trends with different impacts on fine time delays.




Fig. 6. Loopback signal delay recorded at both ends of the link, exhibiting both the same delay drift during warmup since the same hardware is used on both ends of the link. The 20 ps offset was introduced when displaying the curves for clarity.

In order to avoid warmup drift, the solution of continuously collecting samples and broadcasting over a UDP network has been investigated: in this datagram communication approach, the samples are lost if the data collection is occurring out of a communication time slot, and a UDP client records the samples when the communication time slot starts. Rather than using the datagram UDP, we have used the ZeroMQ Publish/Subscribe with the same result. As a result of continuously

running the FPGA and ADC at full speed, the temperature of the components has stabilized and the source of the delay drift removed.

Fine tuning the broadcast power to maximize the ADC range without saturating is needed to achieve utmost signal to noise ratio. Interestingly, the X310 SDR receiver fitted with BasicRX daughter boards – balun transformers feeding the radiofrequency signals to the ADC – does not saturate the input as would be indicated by GNU Radio when reaching normalized sample values of $[-1 : +1]$, but the output to input signal amplitude slowly collapses as the input amplitude is increased, similar to the saturation point of amplifiers characterized with the IP1 and IP3 offset of the output gain to the nominal gain (Fig. 7). From this figure, it can be seen that an input power above 6 dBm leads to an output amplitude that diverts from the linear relationship between input and output power expected from a linear behavior of the ADC.




Fig. 7. Output amplitude (arbitrary units from the sample IQ data) v.s input amplitude on the X310 BasicRX daughter board. Notice that the output amplitude does not reach 1 even though saturation is seen above 6 dBm input power.

This loss of linear behavior of the ADC leads to an increase of the noise on the time of flight measurement. Fig. 8 exhibits the standard deviation on the time delay measurement during 5 minute sessions as a function of input power, indicated as attenuation so that the samples on the left of the chart saturate the ADC and on the right of the chart lead to a linear behavior. Inset (bottom-left) are the time-domain measurements leading to the standard deviation measurements displayed on the main chart. The time of flight measurement dramatically drops when the linear behavior of the ADC is reached above 15 dB attenuation of the input signal, assuming the temperature drift has been compensated for (upper measurement at 18 dB attenuation when the ADC was started cold). Sub-10 ps standard deviation time of flight is demonstrated on this wired link representative of the loopback connection between the PRN generator and the X310 reference channel.




Fig. 8. Standard deviation as a function of input power, and inset the time-domain curve leading to the standard deviation estimates.

V. CONCLUSION

The TWSTFT measurements carried out with SDR Tx/Rx platforms developed in OP and LTFB, over a period of more than 3 weeks, with two 5-min measurement sessions during each odd UTC hour, lead to initial results already slightly better than those obtained with commercial modems, as shown in Fig. 1. This validates the developments obtained in laboratory, as well as their robustness. The next objective is to work on performance and its improvement, including implementing Delay Locked Loop and Phase Locked Loop tracking.

All development software and gateway is available from https://github.com/oscimp/amaranth_twstft for analysis and reproduction of the setup as well as experimenting with different modulation schemes (QPSK) and pseudo-random sequences.

ACKNOWLEDGEMENTS

These developments were supported by the FIRST-TF and Oscillator Instability Measurement Platform (OscIMP) grants, as well as by Laboratoire National d'Essais (LNE). The SAW delay line was fabricated in the MIMENTO cleanroom, a facility part of the Renatech network. We are indebted to Claudio Calosso (INRIM, Italy) for fruitful discussions and cooking amazing pizzas.

REFERENCES

- [1] J.-M. Friedt, C. Droit, G. Martin, and S. Ballandras, "A wireless interrogation system exploiting narrowband acoustic resonator for remote physical quantity measurement," *Review of scientific instruments*, vol. 81, no. 1, 2010.
- [2] J.-M. Friedt and G. Goavec-Merou, "Time of flight measurement with sub-sampling period resolution using software defined radio," in *13th Annual GNU Radio Conference*, Tempe, AZ, USA, September 2023.
- [3] J.-M. & al. Friedt, "Development of an opensource, openhardware, software-defined radio platform for two-way satellite time and frequency transfer," in *2023 Joint Conference of the European Frequency and Time Forum and IEEE International Frequency Control Symposium (EFTF/IFCS)*. IEEE, 2023, pp. 1–4.
- [4] J. Achkar, E. Meyer, B. Chupin, F. Meyer, O. Chiu, M. Lours, and J.-M. Friedt, "Two-way satellite time and frequency transfer using an opensource, openhardware software-defined radio platform," in *4th URSI AT-RASC*. Gran Canaria: URSI, May 2024, pp. 1–4.