GNSS Interference Detection with Software Defined Radio

Yeqiu Ying, Timothy Whitworth, and Kevin Sheridan

Abstract—**Satellite navigation is particularly susceptible to radio-frequency interference. In order to safeguard against and mitigate both unintentional and malicious interference, it is important to have accurate methods to detect it. This paper describes DETECTOR, a GSA-funded project whose objective is to detect and characterize interference for road transport applications. The overall architecture as well as the hardware and software design is described, and the results of one of the real-world tests in an urban environment are shown. During this test several interference sources (jammers) were detected, showing the validity of the design and the necessity of its goals.**

Index Terms—**GNSS, Interference Detection, Jamming**

I. Introduction

IN recent years, Global Navigation Satellite Systems (GNSS) have seen a rapid increase of applications in various sectors. A major threat to the widespread adoption of GNSS concerns the vulnerability of GNSS to signal interference and jamming [1]. Unwanted signals in the GNSS bands can severely degrade the service and impact on the performance [2]. Effects range from a loss of accuracy to complete denial of service. This can lead to catastrophic consequences in safety critical, mission critical, and business critical operations. It is therefore of paramount importance that reliable and robust techniques for interferer/jammer detection can be developed and deployed to protect GNSS infrastructures and services from unintentional and deliberate interference.

Nottingham Scientific Limited leads a GSA-funded project to carry out the design, development, validation, and commercial feasibility assessment for the production of a low-cost GNSS interference and jamming detection solution for deployment within road transport applications. The primary purpose of this product, called “DETECTOR,” is to detect and characterize GNSS jamming equipment being used in road vehicles.

The DETECTOR device is designed based on software defined radio (SNR) technology. A real-time software GNSS receiver enables the continuous monitoring of various metrics of receiver processing, and therefore robust detection of the appearance of interference. Multiple devices deployed over an area can also be networked and connected to a central processing server working cooperatively, which can allow more sophisticated and accurate detection algorithms.

The rest of the paper is organized as follows. Section II describes how radio-frequency interference (RFI) will affect the functioning of GNSS equipments. Section III describes the methods used to detect interference. The architectural design of DETECTOR is discussed in Section IV, and the results of testing are in Section V. Finally the conclusions and future works are described in Section VI.

II. Effects of RFI on a GNSS receiver

The impact of RFI to the GNSS has been studied recently by several research institutes [1] [2], and the main focus is on the impact at the service level. In [3], the impact of the RFI on a low cost GPS receiver has been studied. GNSS signals are very susceptible to noise, due to their extremely low power. Any increase in the noise level at the receiver antenna will adversely affect the performance of GNSS receivers. If the interference level is so high that the receiver electronic components are saturated, the signals might well be unrecoverable. When extra noise is present at the front end, the receiver will encounter the following situations:

1. Low noise will affect measurement accuracy.
2. Medium noise will cause problems with tracking, and make it harder to (re-)acquire satellite signals. Satellites at low elevation may be lost.
3. High noise will completely destroy the receiver’s ability to acquire/track the desired signals.

Along the GNSS receiver processing chain, measurements are available, either internally to the receiver or exported to the application level, which can be used to detect the presence of RFI. One good indicator within the receiver is the gain value of the controllable gain amplifier before the analogue signals are fed into the analogue to digital converter (ADC). This is due to the fact that the input signal to the ADC is required to be matched to the dynamic range of the ADC to guarantee the quantization accuracy. Therefore, within the GNSS receiver implementation, an automatic gain control (AGC) circuit is

Manuscriptreceived June 15, 2012. The DETECTOR project has received funding from the European Community's Framework Programme (FP7/2007-2013) under Grant Agreement No. 277619-2 with the European GNSS Agency (GSA).

Y. Ying, T. C. Whitworth, and K. Sheridan are with Nottingham Scientific Limited, Nottingham, NG2 1RT, UK (phone: +44 (0)115 9682960; e-mail: yeqiu.ying@nsl.eu.com)

normally implemented to automatically adjust the gain value based on the output of the ADC. When the ADC input signal is higher than the nominal level due to the presence of excessive RFI, the AGC will try to lower the gain value of the adjustable gain amplifier, and vice versa. Similarly, the characteristics of the digital signals at the output of the ADC will be changed in the presence of different RFI. Since GNSS signals are buried under the noise floor when they arrive at the receiver, and in the nominal scenario it shall have the characteristics of the additive white Gaussian noise (AWGN). However, when the excessive RFI is present, these characteristics may be changed. Therefore, the digital signals at the output of ADC can be used to detect the presence of RFIs.

III. Detection Methods

The proposed detection techniques take the benefits of the flexibility of software GNSS receiver concept, so that the above-mentioned measurements are accessible, some of which are not usually available from commercial off the shelf (COTS) receivers. The detection algorithms are composed of “pre-correlation” and “post-correlation” techniques. The term pre-/post- correlation is defined based on where the algorithms take the measurements along the receiver processing chain, separated by the essential GNSS receiver processing function: correlation. More specifically, the pre-correlation algorithms make use of the digital signals at intermediate frequency (IF) that are available in our software receiver of the DETECTOR sensor. The post-correlation algorithms, however, can be using standard measurements such as satellite orbit information and signal to noise ratio (SNR) measurements either from our dedicated software receiver or from the COTS receiver.

A. Post-Correlation

The implemented post-correlation algorithms rely on the statistical tests of the SNR measurements. It is worth pointing out that similar techniques were proposed in [4]. In a well surveyed environment, the SNR measurements under nominal conditions from a static receiver can be characterised. Based on this information, a reference SNR value can be obtained via statistical curve fitting techniques based on the collected measurements over a certain period of time. This reference is dependent on the orbit information, as well as satellites transmission signal strength information, atmosphere information, and the environment information. Techniques taking into account transmission strength, atmosphere impact and environment geo-location impact such as [4] are implemented to improve the accuracy of the reference curve of SNR against satellites elevation angles. It is desirable that during the period the measurements for computing the reference are collected, there is no RFI present. However, RFI under certain level can still be smoothed out with the remaining part of the clean measurements.

Thresholds for each of the elevation angles can be calculated based on the desired probability of false alarm, P_{fa}, and the reference, and during the online detection phase, each epoch of SNR measurements are compared to the thresholds. Each tracked individual satellites, which provides the SNR and elevation measurement, will be tested, and a failure of this satellite will be declared if the SNR value is below the threshold. Multiple failures of more than a certain number of satellites within the same epoch will lead to the decision of the failure of the tests. Specifically, two tests of this type are performed, with different P_{fa} (one indicating low SNR, one indicating *very* low SNR) and the number of allowable satellite fails.

In addition, differential tests are performed. One differential test checks the SNR value drop over a short period, and if the drop is more than a pre-set threshold, the corresponding satellite will be declared failing. It is likely that the receiver may lose tracking of some of the satellites in the presence of RFI. Therefore, a test checking the loss of tracking of the satellites over a certain window period will indicate the possible presence of RFI.

In DETECTOR, we perform all these above mentioned tests, and the results will be further fused with the pre-correlation techniques and co-operative techniques in order to reach a global decision of detection.

B. Pre-Correlation

Unlike the post-correlation tests, the pre-correlation techniques are very computationally intensive. It is envisaged that in order to run in real-time, they will have to take snapshots of data, rather than the entire captured signal.

Again, the software requires a clean reference to compare against. In this case it could be based on just a few seconds of data, taken within the previous hour. This can be used to get accurate estimates of the histogram and the power spectrum density (PSD).

The tests only consider the case where the evaluated signal has higher power than the reference signal as being caused by interference, i.e., they are one-sided tests against the null hypothesis of no interference present. The parameters such as Fast Fourier Transform (FFT) size, evaluation window size, etc., are all configurable, and later field testing will look to optimize these.

C. Cooperative

In the case of DETECTOR, cooperative detection refers to having several DETECTOR devices/nodes each collecting and testing their own data, and pooling the results for further analysis.

One area of interest in this project is motorway (highway) monitoring. Under these conditions, it should be fairly easy to conceive of some cooperation methods to improve the detection process. An example would be to take a snapshot of the license plates of the cars by a motorway gantry when an interference flag is raised. If one vehicle is present in the snapshots of several gantries on the same stretch of road, it would be likely that a jammer is being used by that vehicle, and authorities could use this result to pursue the case further.

At this stage of DETECTOR cooperation has not been developed, but it will be developed and implemented in the future work.

IV. DETECTOR Architecture

The DETECTOR system is composed of three major elements: networked DETECTOR field sensors, a DETECTOR server at the back-office for results-logging and data fusion, and monitoring systems. This is illustrated Fig. 1.

A. Field Sensor

The field sensor device is illustrated in Fig. 2. The sensor is composed of an embedded computer which hosts the software receivers, all the detection algorithms, as well as managing the internal and external communications. A software receiver front end (called Stereo) performs the GNSS receiver front end processing, and a COTS receiver is also included to provide redundancy measurements. Both the Ethernet and the wireless communication modems are included for the data communications.

B. Back-office Server

The back-office collects real-time measurements as well as transmitted digital samples from networked field sensors for more sophisticated fusion and detection and characterization computation, relying on addition information such as road model and dynamic motion model. In addition, an atmospheric monitoring server computes the atmosphere impact correction to reduce the false alarm rate.

V. Results

A. Background

It was possible to use the DETECTOR algorithms with data from external sources. The software can process RINEX and/or NMEA files, and do post-correlation interference detection. In the UK, like many other countries, it is possible to obtain data from continuously operating reference stations which have been established to support land surveying and geodetic applications. Several of these sites have been monitored over a period of weeks/months for unusual events.

Preliminary tests identified a significant number of interference-like events in the data. Based on this, a location on an urban road close to an existing reference station was selected for collecting further data with DETECTOR equipment. Capturing digital samples in addition to more standard SNR and AGC data made it possible to use both post-correlation and pre-correlation detection and characterization techniques.

B. Detection Test Results

After processing data from the reference site, 5 or 6 possible events were identified over a 2-day period. Fig. 3 shows part of this, where the position error is also given in order to highlight how problematic (and potentially dangerous) jamming can be. At the start of the day there is a very clear disturbance, lasting ~3 minutes and causing a 100m positioning error. The remaining figures in this paper all concern this event.

Fig. 4 and Fig. 5 show the estimated SNR around these 3 minutes, for the reference site data and our own equipment respectively. In both cases it is easy to see how the SNR significantly degrades. The reference receiver with its high-grade rooftop antenna has better performance than our receiver in which the antenna did not have a clear sky-view. This limits the effectiveness of post-correlation techniques—it will yield poor values of false-alarm versus missed-detection. However, the pre-correlation methods are very good, and the sensitivity is much higher.

Fig. 6 to Fig. 9 show the increasing power of interference over time. This suggests the jammer was mobile, and was moving towards our site. It is not surprising that most of the satellites were lost at this time (Fig. 5). In Fig. 9, the digital power of the evaluation signal was 5.6 times higher than that of the reference.

A few seconds after the snapshot of Fig. 9, the interference died down. 50 seconds after that though, it built up again. The second peak interference snapshot is shown in Fig. 10. This indicates that this is in fact two instances of jamming within a 3 minute window. By comparing the timing of our data with the reference site, it is also possible to determine the direction of travel of these jammers.

As mentioned, these were not the only events of interference. In the space of 41 hours 22 events separate events were detected.

C. Characterization

Although this paper does not go into it in detail, the DETECTOR project is also concerned with characterizing the interference signals. Fig. 11 and Fig. 12 show the spectrogram for the two interference events of Fig. 9 and Fig. 10 respectively. These plots show the interference to be of the “chirp” type—a continuous wave signal quickly swept through a wide frequency range. The two different signatures confirm that there are two jammers present, just a minute apart.

The DETECTOR software is currently able to characterize a signal automatically, based on; statistical periodicity, time periodicity, duty cycle, a swept signal test, a frequency hopping test, power, and bandwidth. In the case of the data here, the software correctly concludes that both the jammers are sawtooth (up) chirp, the first being continuous, the second pulsed. Note that in all likelihood the second signal will in fact be continuous, but the signal will travel outside the frequency of the pass-band of the receiver, i.e. it is the receiver hardware that is turning the signal from continuous to pulsed.

The other events characterized in our 41 hour data capture include several other chirp signals, some powerful single-tone signals, and a few narrow-band signals.

VI. Conclusion

Initial testing has demonstrated the potential of DETECTOR to effectively detect and characterize RF interference sources. Several real-world jammers were detected and characterized. Pre-correlation techniques have been able to detect events which would likely go undetected using only post-correlation methods. Further testing in controlled environments will be undertaken to evaluate these methods more comprehensively and to identify enhancements to the current baseline.

Solutions developed within the project are expected to be able to detect and characterize jammers being used in road applications. This helps understand the nature of the threat to GNSS services and to develop effective counter-measures.

Acknowledgements

The project partners are Nottingham Scientific Ltd., Università di Bologna, Sanef, Black Holes B. V., and AGIT.

Any views expressed here are entirely those of the authors and do not necessarily represent the project partners or the EC.

References

[1] C. Dixon, M. Dumville, etc, “GNSS vulnerabilities: testing the truth,” in *Coordinates Magazine*, March 2012.

[2] M. Thomas. “Global navigation space systems: reliance and vulnerabilities,” *The Royal Academy of Engineering GNSS Vulnerability Report*

[3] S. Storm van Leeuwen, “Electromagnetic interference on low cost GPS receivers”, *National Aerospace Laboratory Report*, 2008

[4] R. |Thompson, A. Dempster, etc., “Detection of RF interference to GPS using day-to-day CN0 differences,” *International Symposium on GNSS*, Oct. 2010

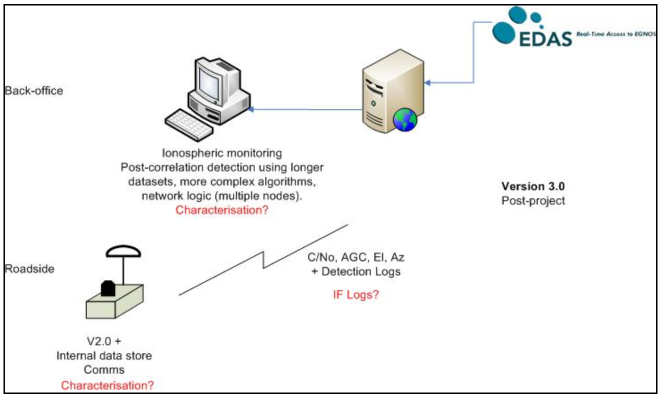

Fig. 1. DETECTOR system architecture.

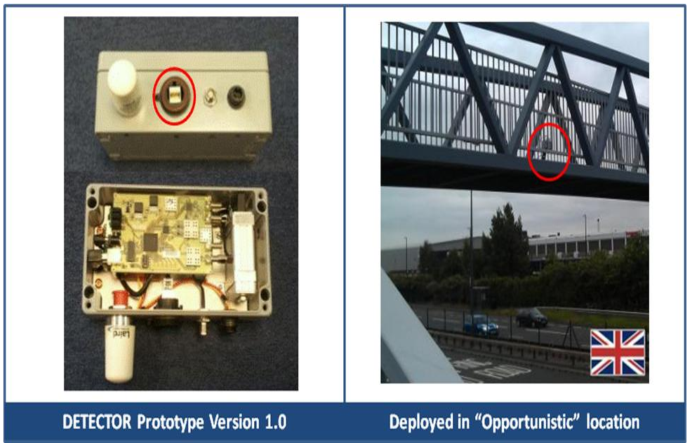
DETECTOR Prototype Version 1.0
Deployed in “Opportunistic” location

Fig. 2. DETECTOR field sensor

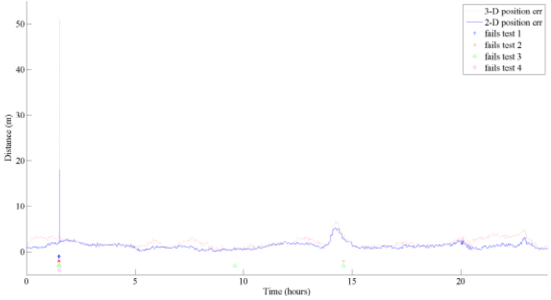

Fig. 3. Position error post-processed from an urban reference site, with interference flags marked below, taken from the log file from DETECTOR. There is a significant event towards the start of the day.

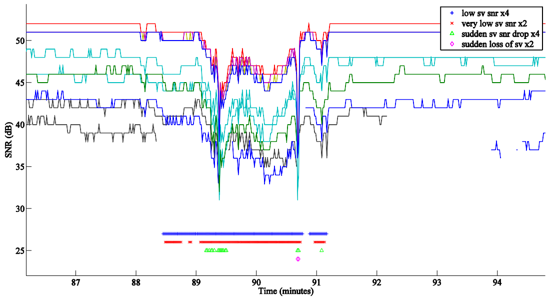

Fig. 4. The SNR of the satellites at the reference site, during the significant event at the start of Fig. 3, again shown with interference flags.

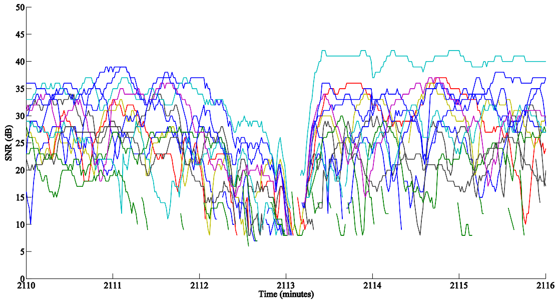

Fig. 5. The SNR of the satellites seen by our antenna around the same time as Fig. 4. The time on the x-axis is different because the equipment was already running for over a day by this point.

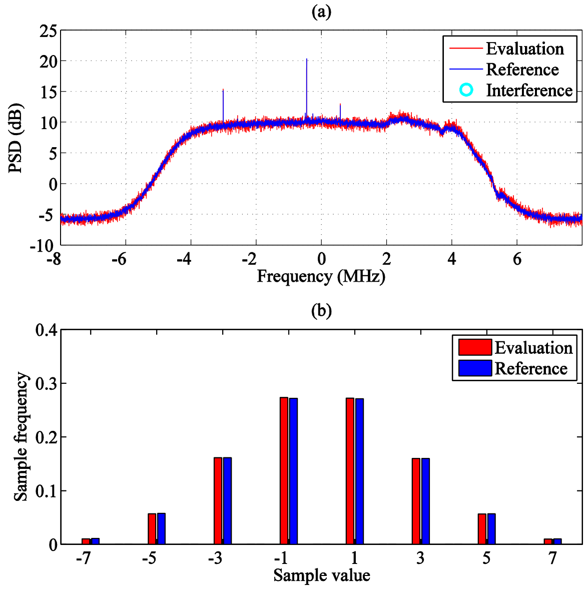

Fig. 6. The pre-correlation figures under “normal” conditions. Part (a) shows the spectrum, and part (b) shows the histogram. The evaluation data in (a) is a little noisier than the reference, since it used fewer windows. The two spikes seen in (a) are from self-interference, and thus do not indicate jamming.

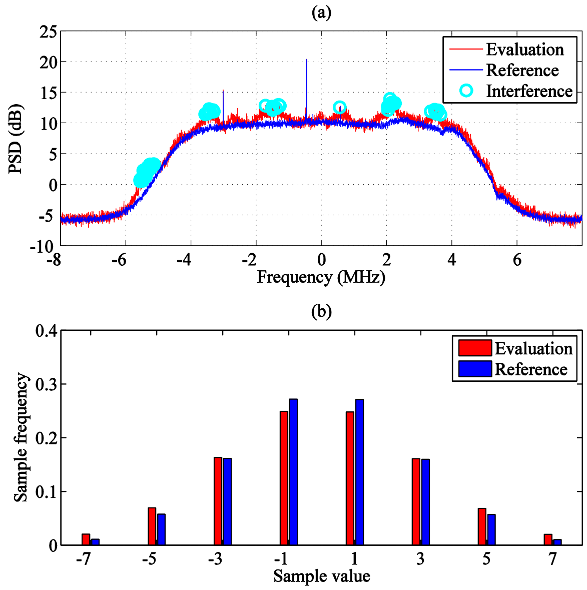

Fig. 7. The pre-correlation results around the start of the event in Fig. 5. The spectrum (a) of the evaluation data shows components above the reference, high enough to trigger detection. The histogram (b) has more samples in the outer bins, indicating more power.

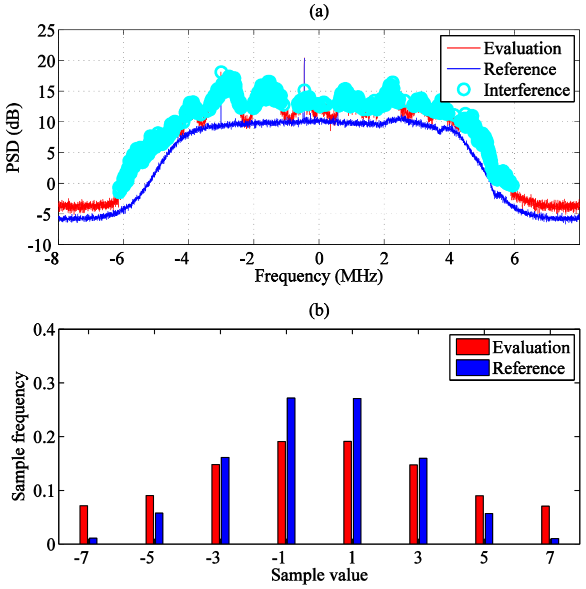

Fig. 8. The pre-correlation results 30 seconds after Fig. 7. The disturbance is worse, and the power is getting higher. This suggests that the jammer is moving closer to the receiver.

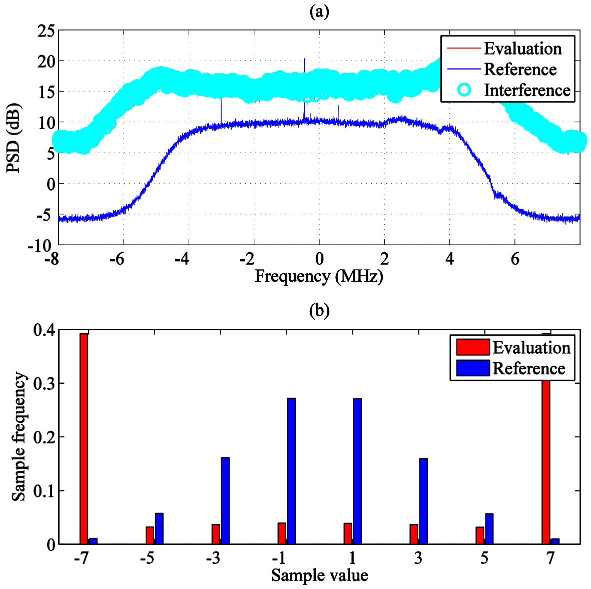

Fig. 9. The pre-correlation results when the interference was at its maximum. The jammer must have been very close to the receiver (the antenna was 25—30 meters from the road).

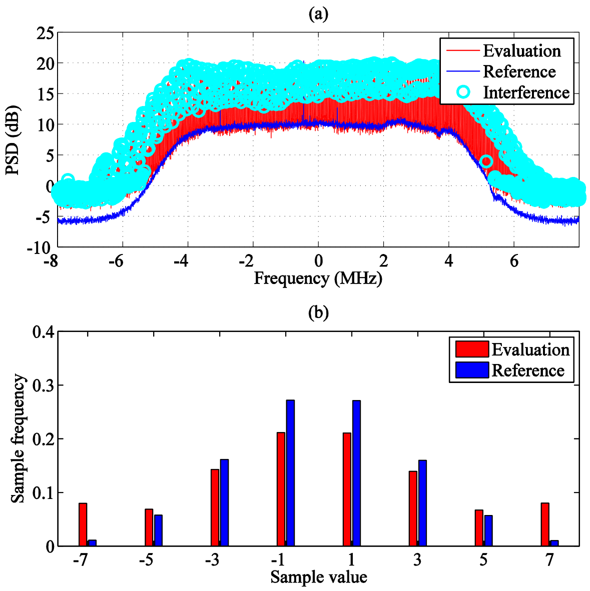

Fig. 10. The pre-correlation results a minute after Fig. 9. This is almost certainly from a *distinct* source.

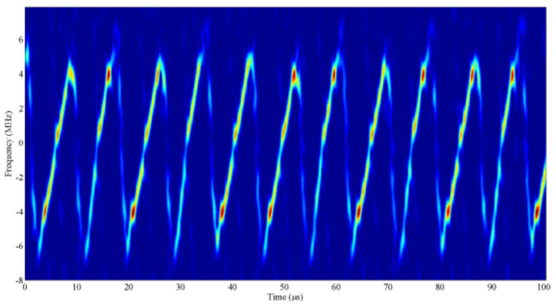

Fig. 11. The spectrogram at the same time as Fig. 9. It shows a sawtooth chirp signal, with a period of ~9µs, comprising ~6µs of sweep, and ~3µs of fly-back. The frequency span is roughly ±7MHz around the L1 centre-frequency (1575.42 MHz).

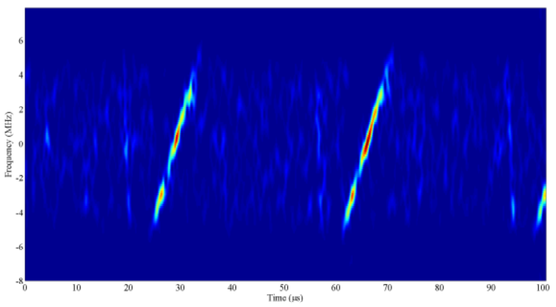

Fig. 12. The spectrogram at the same time as Fig. 10. It shows what looks like a pulsed sawtooth chirp signal, however we may infer that it sweeps continuously, but goes out of the frequency range of the receiver's front end. The sweep rate is lower than the previous jammer, and the period is ~37µs.