Pre- and Post-Correlation GNSS Interference Detection within Software Defined Radio

Kevin Sheridan, Yeqiu Ying, Timothy Whitworth, *Nottingham Scientific Limited*

BIOGRAPHIES

Kevin Sheridan is Technical Manager at Nottingham Scientific Limited (NSL) where he works on the development of robust GNSS positioning solutions for critical applications. He has a Ph.D. from University College London.

Yeqiu Ying is a Principal Navigation Engineer at NSL. He leads the development of NSL's software receiver technologies which provide the flexibility to support a range of applications, including radio frequency interference (RFI) detection. He has a Ph.D, from the University of Leeds.

Timothy Whitworth is a Navigation Engineer at NSL where he has successfully developed and tested a series of RF interference detection and characterization algorithms. He has a Ph.D, from the University of Leeds.

ABSTRACT

In recent years, GNSS has been included in a rapidly increasing number of applications in various sectors, including those regarded to be critical as they concern safety and financial transactions. A major threat to the widespread adoption and use of GNSS is its vulnerability to signal interference and jamming, which can severely degrade the GNSS service and impact performance. Effects range from a loss of accuracy to complete denial of GNSS services.

Mechanisms to counter this threat include stricter regulation on the sale, possession and operation of jamming devices; better enforcement of existing legislation; raising public awareness that RF jammers marketed as privacy protection devices can degrade GNSS over much wider areas than advertised and are in most cases illegal, and better robustness within GNSS-based equipment. Interference detection and characterisation can be a tool to enforce regulations and to better understand the threat in order to produce more effective counter-measures.

Nottingham Scientific Limited (NSL) is leading a European Commission-funded project named DETECTOR which will design, develop and validate a low-cost device for detecting GNSS interference and jamming within road transport applications. The device will be capable of operating in a stand-alone mode or as part of a networked solution depending on the needs. Its purpose is to detect and characterise radio frequency interference which can disrupt GNSS-based services. The intention is that the device will be used by police forces, highways authorities, toll operators, ports authorities and governmental organisations to help combat low-cost and do-it-yourself GNSS jamming technologies.

The design of the detection device is based on software defined radio (SDR) technology. A real-time software GNSS receiver enables the continuous monitoring of various metrics of the receiver processing and therefore robust detection of the appearance of interference.

1. IMPACTS OF RADIO FREQUENCY INTERFERENCE

GNSS signals are very susceptible to noise, due to their extremely low power which leads to a very widely documented vulnerability (see for example [1] to [5]). 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.

At many points in 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 because 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 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 below the noise floor when they arrive at the receiver, in the nominal scenario it will have the characteristics of the additive white Gaussian noise (AWGN). However, when 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.

2. DETECTION METHODS

The detection techniques used in this activity exploit 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 in 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 the software receiver of the DETECTOR sensor. The post-correlation algorithms, however, can be use standard measurements such as satellite orbit information and signal to noise ratios (SNR) either from our dedicated software receiver or from the COTS receiver.

Post-Correlation

The implemented post-correlation algorithms rely on statistical tests of the SNR measurements, applying similar techniques to those proposed in [6]. In a well surveyed environment, the SNR measurements under nominal conditions from a static receiver can be characterized as a function of satellite elevation. Based on this information, a reference SNR value for a specific site can be obtained via statistical curve fitting techniques based on the collected measurements over a period of time. Techniques taking into account transmission strength, atmospheric effects and the obstruction environment around the antenna 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 below a certain level can be smoothed out with the remaining part of the clean measurements.

Thresholds for the SNR at each elevation angles can be calculated based on the desired probability of false alarm, P_{fa}. During the online detection phase each epoch of SNR measurements are compared to these thresholds from the reference dataset. Each tracked satellite with a measured SNR and elevation is tested, and a failure is declared if the SNR value is below the threshold. Multiple failures of more than a configured number of satellites within the same epoch lead to the decision that the overall test has failed. 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 "failures".

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 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 is also used to indicate the possible presence of RFI.

DETECTOR uses a combination of all these post-correlation techniques and complements these with further pre-correlation techniques in order to reach a global decision of detection.

Pre-Correlation

Unlike the post-correlation tests, the pre-correlation techniques are very computationally intensive. To run in real-time they 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 clean 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 seek to identify any cases where the evaluated signal has a higher power than the reference signal as this may be 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 a series of laboratory and field tests are being used to determine optimal values considering the usual performance vs. complexity trade-offs.

3. DETECTOR ARCHITECTURE

The DETECTOR system is composed of two major elements: networked DETECTOR field sensors or probes, and a DETECTOR server at the back-office for data storage, processing and analysis. This is illustrated Fig. 1.

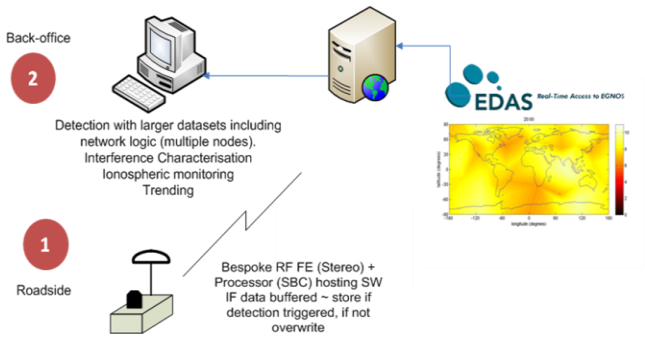

Figure 1: DETECTOR System Overview

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, which can be configured to cover all GNSS frequency bands, performs the GNSS receiver front end processing. A COTS hardware receiver is also included to provide redundancy of measurements. Ethernet and wireless communication modems are included for the data communications. A sample of digital data is stored in a circular buffer. If an interference event is detected the samples are stored, otherwise the buffer is overwritten. All elements of the design are selected to be low-cost, allowing for a scalable solution with large numbers of sensors in future operation.

Figure 2: DETECTOR Sensor

Back-office Server

The back-office collects transmitted digital samples and detection logs from the networked field sensors and performs more comprehensive interference detection and characterization analyses. It also makes use of additional information such as road databases and dynamic motion models to determine if multiple detection events are likely to be due to a single jammer. In addition, it monitors the state of the ionosphere to identify disturbances which could impact many receivers, and prevents this from incorrectly being attributed to intentional interference.

A database of jammer "signatures" will be built up to help develop effective countermeasures. Results can be interrogated to determine trends in the numbers, types and usage patterns of jammers over time.

4. RESULTS

DETECTOR algorithms have been tested in a laboratory and in field trials using dedicated sensors, and also using data available from existing GNSS reference networks. The software can process RINEX and/or NMEA files, and perform post-correlation interference detection. In the UK, like many other countries, it is possible to obtain data from continuously operating GNSS reference stations which have been established to support land surveying and geodetic applications. Several of these are being continually monitored to check for potential interference 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.

Detection Test Results

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

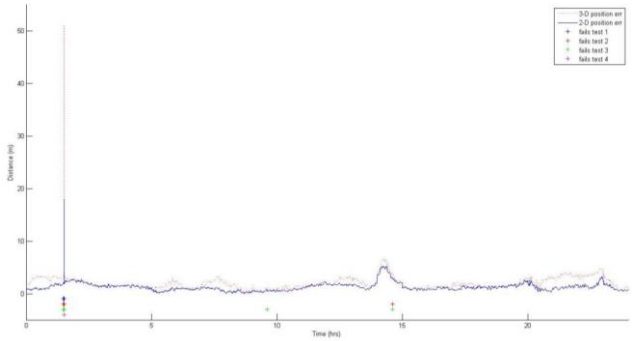

Figure 3: Position error at an urban reference site, with interference flags shown.

Fig. 4 and Fig. 5 show the estimated SNR around these 3 minutes, for the existing reference site data and our own equipment respectively. In both cases it is easy to see how the SNR significantly degrades.

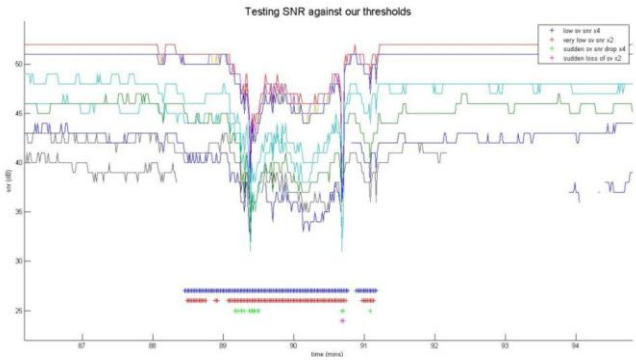

Figure 4: SNR values at reference site during event.

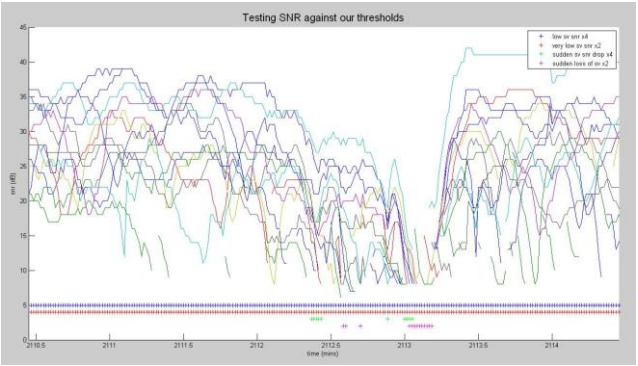

Figure 5: SNR values at from DETECTOR equipment.

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 – using SNR as an indicator of interference becomes unreliable when there can be large fluctuations due to other causes, primarily multipath, which are obvious in Figure 5. However, the pre-correlation methods are still very effective, and the sensitivity is much higher.

Figure 6 shows the changing power of interference over this three minute period. Initially no interference source is present and the power spectral density and histogram matches expected values for nominal conditions. Over time the impact of the interference signal on the receiver increases, with higher power leasing to more samples in the outer bins. This reaches a point at which the signal being evaluated has a digital power more than 5 times that of the nominal reference signal (labeled 1 in figure 6). The trend indicates that the jammer was mobile, and was moving towards our site. It is not surprising that most of the satellites could no longer be tracked at this time. As the jammer moves away the interference level naturally drops. 50 seconds later though, it builds up again (to a peak labeled 2 in figure 6). This indicates that this is in fact two instances of jamming within a 3 minute window.

As mentioned, these were not the only events of interference. In the space of 41 hours 22 separate events were detected, far more than were detected by simply assessing SNR values from the nearby reference station.

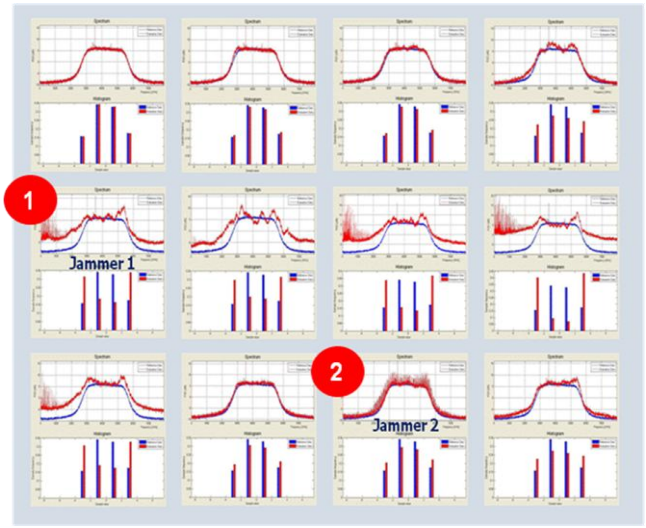

Figure 6: PSD & Histogram over 3 minute period.

Characterization Results

The DETECTOR project is also concerned with characterizing the interference signals. Figure 7 shows the spectrogram for the two interference events. 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.

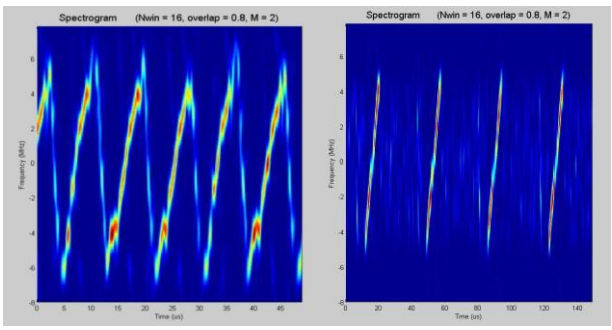

Figure 7: Spectrograms from the two jammers

The DETECTOR software is 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 interference events characterized in our 41 hour data capture include several other chirp signals, some powerful single-tone signals, and a few narrow-band signals. The spectrograms from a sample of these are shown in figure 8. The majority of these look to be attributable to jammers but there are also a number of signatures which are far more likely to be from unintentional man-made sources.

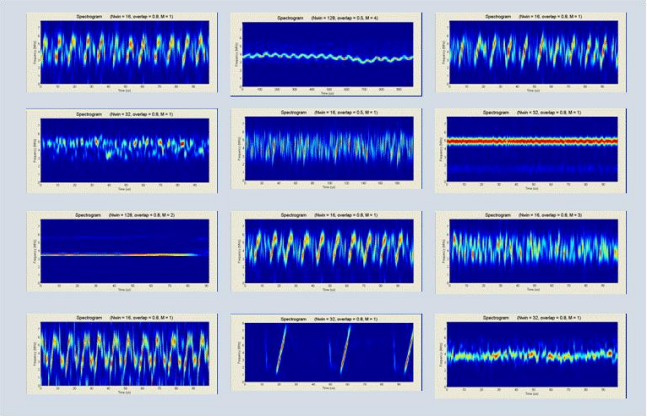

Figure 8: Spectrograms from a sample of interference events

5. FURTHER TESTING

In addition to the investigation reported above, DETECTOR has been tested in a variety of scenarios, and more tests are planned in order to further refine the system design. In June 2012, NSL participated trials in Sennybridge, Wales, with the support of the Ministry of Defence and the Defence Science and Technology Laboratory (DSTL). In these trials, jammers with known and configurable characteristics were operated at a remote site. A variety of test were performed with the jammers static and the detection equipment in a moving vehicle, then the situation was reversed with jammers moving and detection equipment static. In all cases the DETECTOR solution was able to detect and correctly characterize the jammer. Figure 9 shows the plan of the test site, with figure 10 showing a screenshot from the analysis software illustrating the impact of the interference signal.

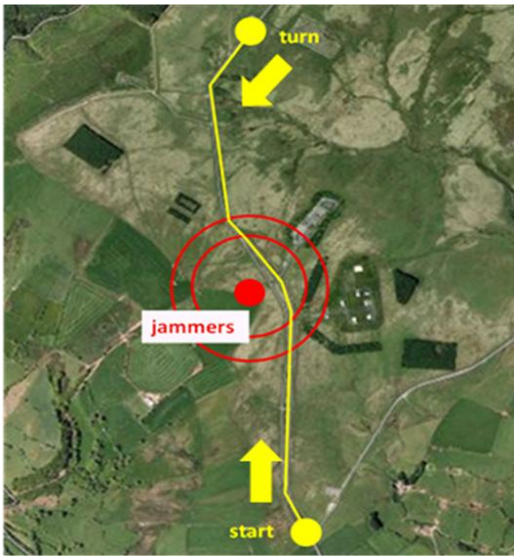

Figure 9: Plan of Jammer test site, Sennybridge, Wales.

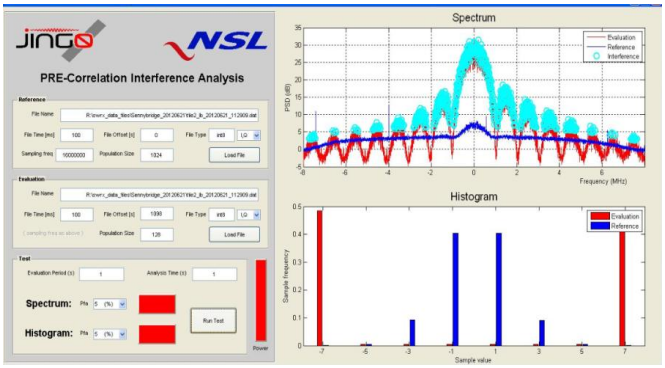

Figure 10: Software screenshot showing the PSD and histogram in the presence of a known jammer.

DETECTOR has also been tested using the specialist laboratory facilities of the Institute for the Protection and Security of the Citizen (IPSC), one of the EC's Joint Research Centres, in Ispra, Italy (Figure 11). In initial tests a sample of commercially available jammers (Figure 12) were operated in an anechoic chamber with the proposed DETECTOR equipment (Stereo RF Front End) and a more comprehensive spectrum analyzer recording and analyzing the signals.

Further tests are now planned in which jammers will be placed in different locations within two vehicles and the detection devices will be operated at a wide range of relative elevations, azimuths, and distances to the vehicles. These tests are designed to how local obstruction conditions, such as opening a car window, influence the effective range of the jammers. Considering both the range over which a jammer will cause significant disruption to other GNSS receivers, and the range over which the signal can be reliably detected and characterized.

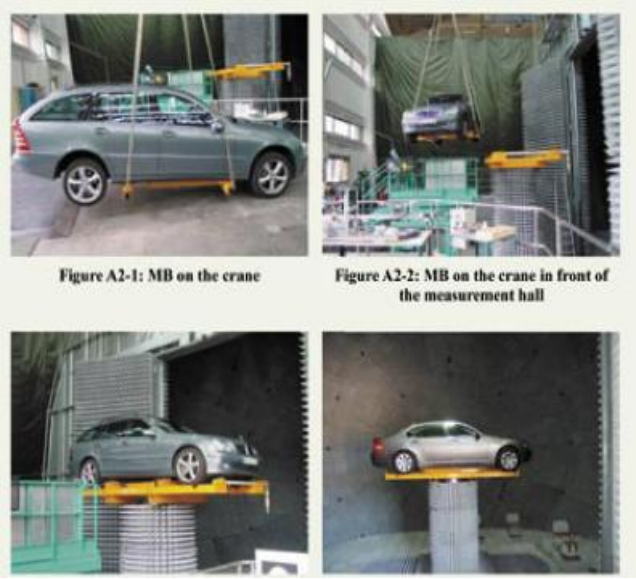

Figure 11: Previous tests at IPSC, JRC (Ispra, Italy).

Figure 12: Jammers used in laboratory testing.

In these tests, DETECTOR was again able to reliably detect and characterize a range of typical jammers. Controlled testing provides an opportunity to better understand the jamming signals produced by various low-cost devices in a clean environment, which is a critical step in developing effective counter-measures. Figure 13 shows the spectrogram of four of the jammers tested. It is interesting to note that many of the jammers tested will disrupt GNSS signals on multiple frequencies, including the proposed Galileo PRS.

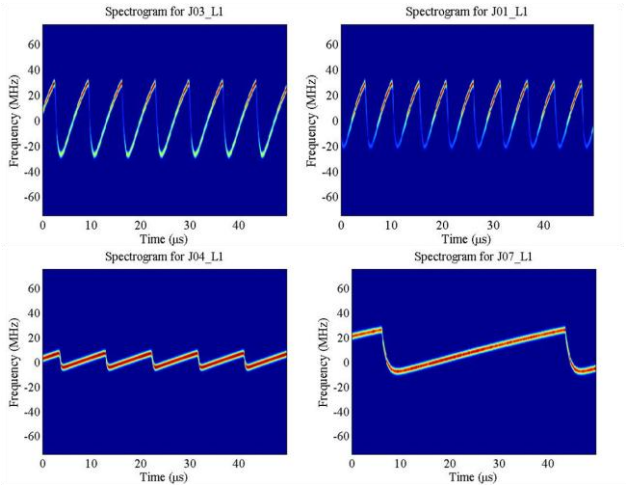

Figure 13: Spectrogram of jammers in laboratory testing.

All the testing opportunities reported so far also make it possible to perform a variety of design trade-offs to support the development of an operational solution. The amount of digital data which is collected, analyzed, communicated and stored is an important design issue. These tests have helped identify the length of sample, bandwidth and quantization level (number of bits) required for reliable detection and characterization. The flexibility of the Stereo RF Front End makes it an ideal platform for performing these types of investigation. For example, by configuring the two available RF chains differently and using them simultaneously, it has been possible to directly compare results using a 2 bit quantization with 6 bits (3I, 3Q) on the other chain. These tests explain why the number of bits differs in figures 6 and 10.

6. VALIDATION AND DEPLOYMENT

Once all aspects of the DETECTOR device are fully consolidated and a final end-to-end prototype is in place, it will be validated at the "AutomotiveGATE" in Germany (figure 14). The Aldenhoven Testing Center (ATC) is a proving ground for positioning systems for road applications. Galileo satellite signals are emulated and transmitted using a series of pseudolites which ensures that three or more Galileo signals are available at any position inside the test bed. This allows the additional benefit of Galileo satellites to be evaluated.

For DETECTOR validation the important aspect of this testbed is that jammers, receivers and detectors can be operated in a realistic road environment. The testbed includes road sections representative of urban, sub-urban and autobahn conditions. Jammers, receivers and interference detectors can be installed on vehicles and on roadside gantries in various combinations. For example, the impact on a GNSS receiver from a brief exposure to a jammer which is being used in a vehicle travelling in the opposite direction can be assessed. Equally, the ability to detect this jammer in slow and fast moving traffic can be

assessed. The effectiveness of detectors installed on overhead gantries rather than at the roadside can be examined, taking into account alternative traffic conditions which could block the path between a jammer and detector.

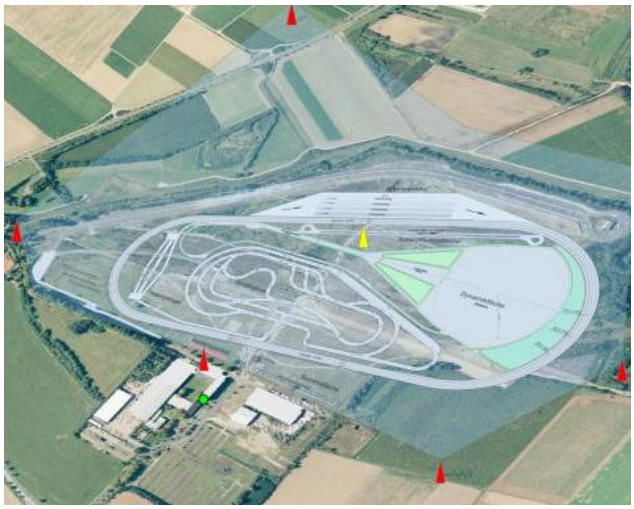

Figure 14: Overview of the AutomotiveGATE test site (red icons indicate locations of pseudolite transmitters).

DETECTOR has been designed with the aim of having multiple sensors deployed on road gantries. Following the validation of the prototype it will be installed at a Sanef operated site in northern France (Figure 15).

Figure 15: Site for initial sensor installation

CONCLUSIONS

DETECTOR is being developed to provide an effective, low-cost means to detect and characterize RF interference sources which degrade GNSS performance in road applications. Initial testing has demonstrated the potential of DETECTOR to reliably detect and characterize jammers in real-world conditions, controlled field tests and in the laboratory.

Pre-correlation techniques, which are made possible through the use of a flexible RF Front End and software receiver, have been able to detect events which would likely go undetected using only post-correlation (SNR) methods. Characterization allows the type of interference signal to be identified which gives a good indication of whether it is unintentional interference or a jammer. Signal characteristics helps understand the nature of the threat to GNSS services and to develop effective counter-measures.

ACKNOWLEDGMENTS

The work presented in this paper has been funded in part under the EC FP7 programme through the European GNSS Agency (GSA). This support is gratefully acknowledged. 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. M. Wildemeersch and J. Fortuny-Guasch (2010). “RadioFrequency Interference Impact Assessment on Global Navigation Satellite Systems.” EC JRC Security Technology Assessment Unit, EUR 24242 EN, January 2010.

2. Bauernfeind R., Kraus T., Dotterbock D., Eissfeller B., Lohnert E. and Wittmann E. (2011) “Car Jammers: Interference Analysis.” GPS World, October 2011.

3. M. Thomas (2011) “Global navigation space systems: reliance and vulnerabilities,” The Royal Academy of Engineering GNSS Vulnerability Report. March 2011.

4. S. Storm van Leeuwen (2008), “Electromagnetic interference on low cost GPS receivers”, National Aerospace Laboratory Report, 2008.

5. H. Kuusniemi (2012) “Effects of GNSS Jammers on Consumer Grade Satellite Navigation Receivers” ENC 2012, April 2012.

6. R. Thompson, A. Dempster, *et al*.(2010), “Detection of RF interference to GPS using day-to-day CN0 differences,” International Symposium on GNSS, October 2010.