On the Threat of Systematic Jamming of GNSS

James T. Curran[1], Michele Bavaro[2], Pau Closas[3], Monica Navarro[3]
[1]*Independent Researcher, Cork, Ireland*
[2]*European Commission, Joint Research Centre, Ispra, Italy*
[3]*CTTC: Centre Tecnológic Telecomunicacions Catalunya, Barcelona, Spain*
jamestcurran@ieee.org, michele.bavaro@jrc.ec.europa.eu, pclosas@cttc.es, monica.navarro@cttc.cat

BIOGRAPHY

James T. Curran received a B.E. in Electrical & Electronic Engineering in and a Ph.D. in Telecommunications from the Department of Electrical Engineering, University College Cork, Ireland. He has worked as a research engineer with the PLAN Group in the University of Calgary, and at the JRC in, Italy. He is currently a radionavigation engineer with Serco for ESA in Noordwijk, the Netherlands. His research interests include signal processing, information theory, cryptography and software defined radio for GNSS.

Michele Bavaro received his master degree in Computer Science in 2003 from the University of Pisa. Shortly afterwards he started his work on Software Defined Radio technologies applied to navigation. First in Italy, then in The Netherlands and in the UK he worked on several projects being directly involved with the design, manufacture, integration, and test of radio navigation equipment and supporting customers in the development of their applications. Today he is appointed as Technical Project Officer at the European Commission Joint Research Centre.

Dr. Pau Closas is Head of the Statistical Inference for Communications and Positioning Department and Senior Researcher at the Centre Tecnològic de Telecomunicacions de Catalunya. He received the MSc and PhD degrees in Electrical Engineering from Universitat Politècnica de Catalunya in 2003 and 2009, respectively. He also holds a MSc degree in Advanced Mathematics and Mathematical Engineering from UPC since 2014. During 2008 he was Research Visitor at the Stony Brook University, New York, USA. His primary areas of interest include statistical and array signal processing, estimation and detection theory, Bayesian filtering, robustness analysis, and game theory, with applications to positioning systems and wireless communications.

Dr. Monica Navarro is a Senior Researcher at the Centre Tecnològic de Telecomunicacions de Catalunya within the Communication Systems Division. She received the MSc degree in Telecommunications Engineering from Universitat Politècnica de Catalunya in 1997 and the PhD degree in Telecommunications from the Institute for Telecommunications Research (ITR), University of South Australia, in 2002. From 1997 to 1998 she was a Research Assistant at the Department of Signal Theory and Communications at the UPC, where she worked on the development of fractal shape multiband antennas for wireless cellular communications systems. She has also been part-time lecturer at the Universitat Pompeu Fabra, Barcelona. Her primary areas of interest are on digital communications and signal processing, particularly on iterative information processing, adaptive transmissions and coding techniques, signal processing for synchronization, estimation and detection theory with applications to radio communications systems, including wireless mobile communications, deep-space communications, wireless sensor networks, and positioning applications.

Abstract

This paper presents a study of the threat of malicious interference to GNSS and examines the special case where the jamming device is incrementally more sophisticated than a typical always-on interference source. The concept of a systematic jamming attack is considered, where the interference signal is intentionally synchronized with the GNSS signals, with the intention of causing maximum disruption with the minimum power expenditure. Various attack methodologies are examined for the case of a civilian L1 receiver. It is shown that, depending on the attack strategy, the target signal and the target receiver, data-recovery, navigation and timing can be denied to a user with some tens of decibels less average power than a traditional jamming attack. It is further shown that some attacks may be capable to effectively deny some receiver functionality in a subtle manner such that presence of the malicious interference goes undetected. Key signal and receiver features that expose a vulnerability are identified and some means of improving receiver robustness are provided.

INTRODUCTION

The vulnerability of GNSS to various forms of malicious interference have been widely discussed in recent years, and have considered a wide range of both real and potential attacks. Some of these have included extensive studies of commercially available jamming devices [7, 9] while others have considered the more comprehensive case of spoofing, where the interference takes the form of genuine GNSS signals [4, 12, 13].

Studies of simple jamming attacks have demonstrated that it is relatively easy, given sufficient broadcast power, to deny the use of GNSS to many commercial receivers [6, 10, 15]. However, it has also been shown that given the easily identifiable or periodic nature of simple jamming signals, a receiver can often mitigate the threat, for example, via the use of adaptive filtering or pulse blanking [2, 8]. Furthermore, it has been demonstrated that the jamming signal itself can be readily exploited to identify [11] and locate [3] the jamming source. Recent work on GNSS spoofing have shown that current receivers are vulnerable to a well calibrated spoofing attack [4], and it is clear that many receivers can be manipulated without arousing any suspicion. However, such attacks are highly complicated and require knowledge of the GNSS signals, and the attack scenario, including precise timing and positioning [5].

It is proposed here that some middle ground must exist. A typical jammer is blind to the GNSS signals it overwhelms, and simply relies on power and spectral occupation to deny the GNSS signals to a nearby user. In contrast, a spoofing device must faithfully replicate the characteristics of genuine GNSS signals. An attack will be highly sensitive to alignment of time, phase and power, with the GNSS signals it intends to replace and the effectiveness of the attack is dependent on the fidelity of the signals it broadcasts. However, it is likely that some device might be created which is only slightly more complex than a simple in-car jammer, but which can exploit some information of the GNSS signals it intends to overwhelm in such a way as to increase the efficiency or effectiveness of a jamming-based denial of service attack.

Specifically, this work introduces the concept of systematic jamming: where a simple jammer might be equipped with some information of the GNSS signals, and can use this to perform a more sophisticated jamming. For example, it is suggested that the jammer may be equipped with a simple low-cost commercial GNSS receiver [14], such that it would then have access to accurate position and time, and also to satellite ephemerides. With some very basic integration of this information, it might be possible to trigger short and sparse bursts of interference at specific times, such as to deny GNSS to a nearby receiver, and to do so with a very low average power. In this manner, a receiver might be unable to: reliably detect that a jamming attack was ongoing; to effectively mitigate the jamming attack; or to identify or localize the jamming source. In the work that follows, we consider what form such a jammer might take, what the implications for the nearby target receiver might be, and how a target receiver might be equipped to thwart such an attack.

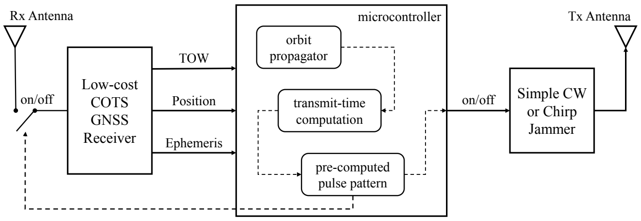
Rx Antenna
on/off
Low-cost COTS GNSS Receiver
TOW
Position
Ephemeris
microcontroller
orbit propagator
transmit-time computation
pre-computed pulse pattern
on/off
Simple CW or Chirp Jammer
Tx Antenna

Figure 1: A block diagram of a hypothetical systematic jammer.

PROBLEM DEFINITION

This work considers the threat that might be posed if a malicious adversary were to add a small amount of added complexity to the typical GNSS jammer, with the intention of providing bursts of interference at specific epochs. A modification to the typical GNSS jammer is envisaged, which includes an on/off keying driven by a microcontroller, as depicted in Fig. 1. The algorithm controlling the keying employs position and timing information sourced from a simple, low-cost GNSS consumer-grade receiver (naturally, it may also implement pulse blanking to avoid self-interference). Based on low-rate position and timing messages, and satellite ephemeris information, a basic orbit propagation can provide accurate estimates of the transmit-time observed on GNSS signals seen in the vicinity of the jammer.

It is proposed that this information might be exploited by an adversary to trigger short pulses of interference which are tightly aligned with certain portions of the navigation message of each satellite. Previous work has demonstrated that, a low duty-cycle pulsed interference, that is appropriately synchronized with the navigation message, can cause disruption to the receiver data recovery process, equivalent to that of an always-on interference [1]. This process requires that the pulse pattern be designed to specifically target weaknesses in the navigation message coding scheme, and it has been shown that for the case of convolutional encoding and block interleaving that a malicious adversary might inflict a denial-of-service (DOS) upon a naïve receiver, using an average interference power 10 to 20 dB lower than the continuous interference case.

Naturally, this offers some distinct advantages to the adversary: a given broadcast power might impose a DOS over a wider geographical area; by broadcasting short sporadic bursts of interference, it may be more difficult for an authority to detect or locate the jamming source; it may also be possible that the interference pattern can be made sufficiently sparse that the target receiver, although experiencing a DOS, might not reliably assert that it is experiencing interference.

This work will examine this threat from two perspectives, that of the malicious adversary, and that of the target receiver. The current Galileo and GPS signals, E1B and L1C/A, respectively, will be studied, seeking to identify how best the adversary might target these signals, and will then analyse to what extent a DOS might be conducted. This study will seek to quantify improvements in jamming effectiveness relative to the continuous interference case. The study will then examine the problem from the perspective of the target receiver, aiming to identify characteristics of a systematic jamming attack that might facilitate either the detection of an ongoing attack, or the mitigation of an ongoing attack.

SYSTEMATIC INTERFERENCE AND DENIAL OF SERVICE ATTACKS

This section provides a study of one specific vulnerabilities of a variety of GNSS signals to a systematic interference. Here, the methodology chosen for the generation of harmful pulse-patterns is based on denying navigation capability of the receiver, rather than denying the signal itself. To produce a position, velocity and time (PVT) solution, a receiver generally needs to extract the ephemeris and the time-of-week (TOW) from each satellite used in the computation of the navigation solution. In many occasions a recently extracted ephemeris may be available and, so, need not be extracted. In other cases, when a coarse position and satellite ephemerides are available, a receiver may need to extract the TOW from one satellite, and can infer the TOW on others. Thus, it is noted that, at the very minimum, a receiver must extract a TOW record at least once, from at least one satellite navigation message, in order to provide a PVT. This section examines the design of interference pulse patterns which might disrupt this process.

Sensitive navigation data

A TOW message is broadcast by all GNSS signals at regular intervals, and generally occupies a very small portion of the overall navigation message. In the case of GPS L1 C/A and GLONASS L1 OF the TOW is broadcast in an unencoded form once per subframe, whereas for Galileo E1B, it is encoded, and broadcast once per pair of pages. As such, the denial of TOW for the GPS and GLONASS signals requires either the denial of the subframe synchronization, or denial of the raw data itself. In the case of Galileo, the TOW might be denied by either denying page synchronization, or by inducing errors in the symbol decoding process. The basic details of the navigation messages, as shown in Fig. 2, are as follows.

GPS: The L1 C/A preamble is an 8 bit sequence (160 ms) transmitted every 6 seconds. The GPS parity is composed of 6 bit (120 ms) transmitted every 600 ms (navigation data word). Checking the consistency of two subsequent preambles, as well as the 10 parity checks in between, is a commonly accepted mean of synchronizing to the 6 seconds boundary.

Galileo: The E1B signal transmits a plain 10 symbol synchronization sequence (40 ms) every second. It is interesting to see that GPS and Galileo synchronization sequences hardly overlap in time. The Galileo message CRC is FEC encoded and then spread by an interleaver. The E1B receiver de-interleaves the data and runs a Viterbi decoder to retrieve the 120 bit/sec of I/NAV. The identification of a word results in resolving a 2 seconds time ambiguity. In addition Galileo E1C is a pilot channel which allows continuous alignment to the 100 ms boundary. Galileo time and GPS time are essentially aligned for what concerns this analysis.

GLONASS: The L1 OF signal transmits a trimmed pseudo-random sequence of 300 ms duration as time mark at the end of each 2-seconds page. GLONASS time implements leap second and is aligned to UTC for what concerns this analysis.

Design of Interference Pulse Patterns

The object of this section is to identify an interference signal that will deny the extraction of the TOW from the above signals using the least amount of energy possible such that the target receiver either remain unaware of the jamming attack; might be unable to effectively mitigate the jamming signal. To simplify the problem somewhat, the jamming signal is restricted to be an on-off-keying of a chirp interference signal, having *on* pulses of fixed duration equal to some integer milliseconds.

Two particular examples are explored here: GPS L1 C/A which is subjected to pulsed interference across the broadcast TOW, and the case of Galileo E1B, which is subject to pulsed interference across a series of symbols spaced according to the symbol interleaver, and are depicted in Fig. 3. The GPS pulse pattern has been aligned with the 17-bit TOW and consists of 6 20 ms pulses, each separated by a 20 ms. The Galileo pulse pattern consists of 15 4 ms pulses, spaced according to the Galileo 8 × 30 block interleaver, such that all 12 pulses appear consecutively once the received symbol stream has been deinterleaved.

This particular choice of pulse patterns is somewhat arbitrary, and has been selected based on some simple experiments. A more thorough design might carefully weigh the choice of number of pulses, pulse duration, and instantaneous interference power, to find a pattern which provides the highest probability of inducing bit errors, with the minimum probability of being detected. This, of course, will depend greatly on the particular monitoring techniques of the receiver - most notably the C/N_0 estimator, and tracking look design.

To align these pulse patterns with the received GNSS signals, they are broadcast with a fixed delay relative to

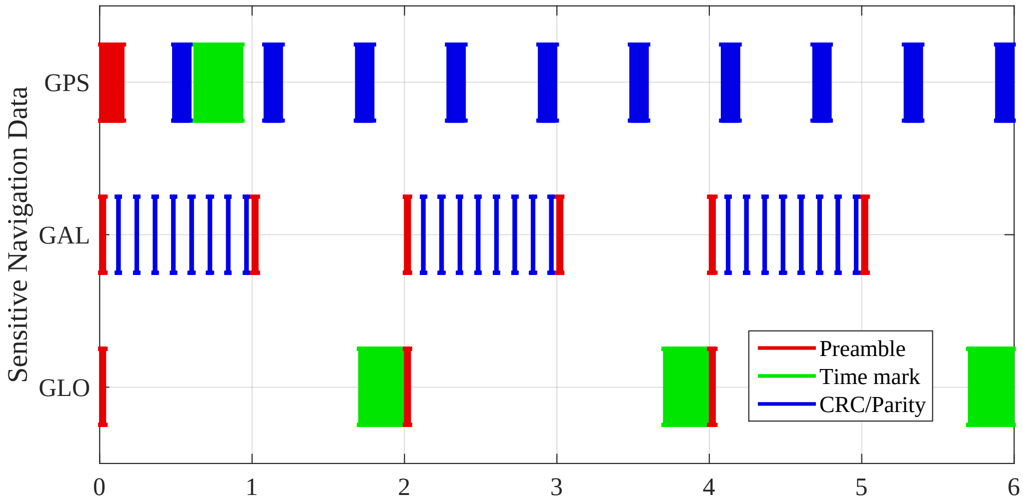

Figure 2: Position in time of various portions of sensitive data contained in each of the GPS, Galileo and GLONASS navigation messages.

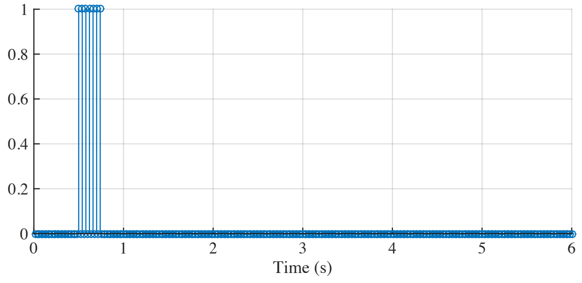

(a) GPS L1 C/A

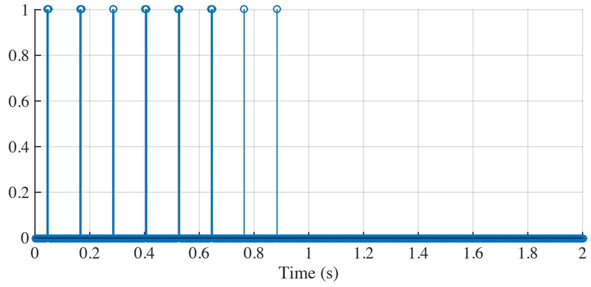

(b) Galileo E1B

Figure 3: Example pulse patterns for the systematic jamming of the GPS L1 C/A and Galileo E1B navigation messages.

the edge of a GPS 6 second boundary. All GNSS satellites broadcast their messages in synchronous, and all have are observed at a range between 18,000 and 24,000 km, depending azimuth and elevation, this fixed delay was set to 67 ms, or approximately 20,000 km. Note that the maximum variation between nearest and furthest satellite results in a misalignment of less than 20 ms, and so the pulse pattern applied to the GPS L1 C/A message will still overlap completely with the 17 bit TOW message. Similarly, owing to the nature of the block interleaver used for Galileo E1B, when the pulse pattern is shifted relative to the encoded symbols, provided they still overlap with a single page, the will deinterleave to a continuous stream.

ANATOMY OF A SYSTEMATIC JAMMER

Until very recently, the only widely available transceiver option existing for radio amateurs and navigation/telecommunication engineers was the Ettus product line: the USRPs. More recently the technological advances in the integration of RF components into single multi-modal chips (mostly driven by the 3G/4G and DTV market) have enabled the design of relatively simple, highly versatile low cost SDR peripherals. A comprehensive review of such hardware is not appropriate here. However worth mentioning are Michael Ossman's HackRF One and the Nuand's bladeRF as both were available in the laboratory and were used to demonstrate the concepts illustrated above. The most relevant specifications for these two devices are in Tab. 1

Designing and building a sophisticated jammer would be a controversial research activity and would pass the message that synchronised jamming needs specialist expertise and skilled personnel, when instead it is relatively simple. Using a general purpose transceiver reading samples

	HackRF One	bladeRF
Freq span	30 MHz - 6 GHz	300 MHz - 4.2 GHz
Bandwidth	20 MHz	28 MHz
Bits	8 I&Q ADC/DAC	12 I&Q ADC/DAC
Interface	USB 2.0 HS	USB 3.0
Radio	RFFC5072+MAX2837	Lime Semi LMS6602D
Baseband	CPLD + MCU	FPGA + CPU
Clock	In, out, 10ppm	In, out, 0.5ppm VCTCXO
Trigger	No (but added)	Yes (recently)

Table 1: Specification of the transceivers used in the tests

from disk, a 6 second long IF file can be created containing bursts of wideband noise at specific positions in time. Such IF file can be played back in a loop for as long as the reference clock of the transceiver maintains good synchronisation with the navigation data bits to be jammed. A real-time orbit propagator is not necessary if synchronisation errors of about 10 ms can be tolerated, as an average travel time of e.g. 67 ms can be assumed for all satellites. A low quality frequency source will maintain ms-level synchronisation in several minutes of operation, but a common TCXO will not build significant skew before a few hours. The experimented approach however present the additional practical requirement to avoid streaming packet losses. Each packet can contain a few ms of data, depending on the chosen bandwidth and quantisation thus a few packet losses will significantly disrupt the synchronisation with the live SIS. Both the hackRF and the bladeRF were designed mainly for telecommunications where the uninterrupted streaming is not as important as it is in radionavigation.

Synchronisation of the Jammer with GNSS-Time

In order to allow proper synchronisation of a common PC peripheral it is necessary to stay away from high abstraction layers and act as close as possible to the hardware. In fact, the USB or Ethernet busses have too unpredictable latency in non-real-time Operating Systems. The ms-level time keeping of the PC clock is already a non-negligible task without bespoke software and low-latency interconnection to a dedicated time-server (e.g. timing equipment exposing a NTP server). Then, the absence of high-priority interrupts in the OS can invalidate the approach altogether. On the other hand the clock accuracy obtainable using directly the CMOS signal, even from a mass-market receiver, is about 50 ns.

A trigger for the bladeRF was added to the stock firmware released on June 2016. At the time of writing the hackRF did not support triggering but, as it is an open hardware and software design, such feature was promptly implemented. One of the GPIOs of the CPLD was used to hold the FIFO data from being dispatched to the DAC until a rising edge of the trigger signal would be registered, as shown in Fig. 4.

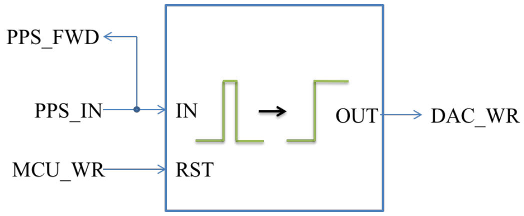

Figure 4: Intuitive diagram of the gate feature implemented in the CPLD (pulse-to-level translator)

A trigger-forward line was also inserted in order to simplify chaining multiple transmitters. The MCU firmware and user-level application were accordingly modified to cope with the trigger feature and to avoid buffer underruns on the disk when streaming. As the CPLD is clocked at twice the DAC speed the resolution of trigger is limited by the clock period of the transmitted signal, so assuming 10 MHz bandwidth to about ±25 ns which is comparable with the GPS module accuracy and well below what is needed for synchronized navigation data jamming.

Testing for synchronizing of two transmitters was performed by generating two wideband CDMA sequences and playing them through a combiner into a single channel receiver. The orthogonality of the sequences allows reception of both without interference by means of a matched filter as well as measurement accuracy limited by the bandwidth of the recording device. In this case, a GNSS-like PRN was generated and used to test the triggering of a pair of hackRF devices.

Synchronization with the true time was further tested by generating a simulated satellite signal using actual broadcast navigation data and orbit of one particular satellite signal visible at the time, but modulating it with a different PRN, one that was not visible at the time. The playback was triggered using a GPS-aligned PPS, and this simulated signal was combined with the a live feed from a roof-top antenna, and broadcast into a standard GNSS receiver. The receiver produced a set of pseudoranges, one for each visible satellite, along with a pseudorange for the simulated signal. The principle of the test is that the simulated signal

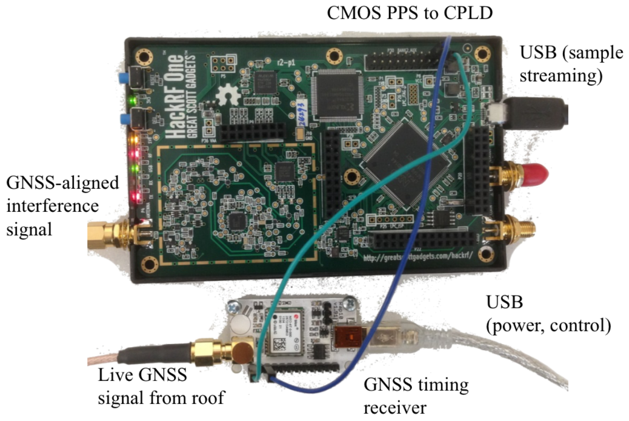

Figure 5: Prototype systematic jammer constructed using a hackRF One and a μBlox timing receiver for PPS generation.

would appear as with a new SVID, but exhibit the same pseudorange as the satellite from which its navigation data and orbit were copied. Indeed, this simulated range was aligned well with the corresponding genuine signal, indicating that the trigger did launch the transmission at a PPS edge. However, it was observed that the range diverged rapidly due to the poor clock quality of the transmitter. This indicated that it would be necessary to periodically re-synchronize the transmission with GPS time.

LIVE TESTING WITH COTS RECEIVER

This section briefly describes results of a simple systematic interference test conducted on a COTS GNSS receiver. The prototype systematic jammer was constructed using a single hackRF One, which derived synchronisation with GPS time via a uBlox timing receiver, which delivered a rising edge on a trigger once every 30 seconds, as depicted in Fig. 5. Note that although this device delivered a very precise timing reference, the systematic jamming attack does not necessarily require such accuracy, indeed the GNSS propagation delay is approximated with an error of up to 10 ms. Therefore, a 1 to 10 ms accurate reference derived from a wired or wireless network, being WiFi or a 3G mobile network, would suffice.

The test consisted of a conductive combination of a life GNSS feed from a roof mounted antenna with a systematic interference signal. The unit under test was a uBlox M8N receiver and was configured to deliver raw observations to a host PC for post processing. The bench-top configuration is shown if Fig. 6, where the interference signal is combined directly with the live GNSS feed, and is subsequently split between the UUT and a spectrum analyser. A series of tests were conducted to identify the minimum instantaneous power that could be applied which would reliably deny the extraction of the TOW from the target signal. In the conductive test, the live GNSS signals arrive to the combiner having been amplified by both the antenna, and an in-line amplifier, and so the interference signal was broadcast at a correspondingly increased power level. In a live broadcast attack, this increase in power would not be necessary, but the interference signal power would have to be adjusted according to the free space path loss between the transmit antenna and the target receiver.

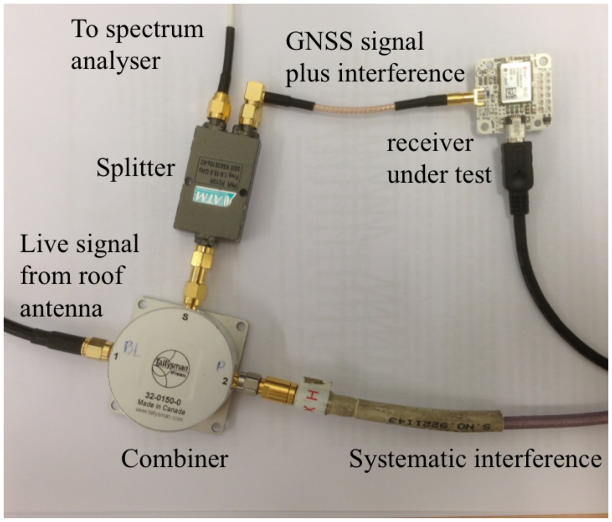

Figure 6: Configuration of the live testing of a uBlox M8N receiver, subject to systematic interference. The experiment consisted of the conductive combining of interference with a live signal feed from a roof mounted antenna.

Denial of GPS L1 C/A PVT

In the first test, the ability of the systematic jammer to deny observations and a PVT from GPS L1 C/A was examined. The experimental setup described above was used, and the pulse pattern depicted in Fig. 3 (a) was used. The prototype jammer was powered up and allowed to initialize and align with GNSS time. Next the receiver under test was issued a cold-start command its behaviour was observed. The test was repeated with progressively increasing interference power until a power level was established at which the receiver was unable to produce a PVT, which was observed to occur at an *instantaneous* interference to noise floor level of approximately 30 dB, as depicted in Fig. 7 A trace of the 11 GPS satellites being tracked by the receiver are shown in Fig. 8, where it can be seen that the received C/N_0 for the L1 C/A signal ranges from 49 to 35 dBHz, but experiences brief reductions in power of approximately 6 dB. During the entire test, the receiver was unable to provide a sufficient set of observations and ephemerides such that a PVT could be computed. Unfortunately it was not possible to gain enough visibility into the internal receiver functionality to determine exactly which informa-

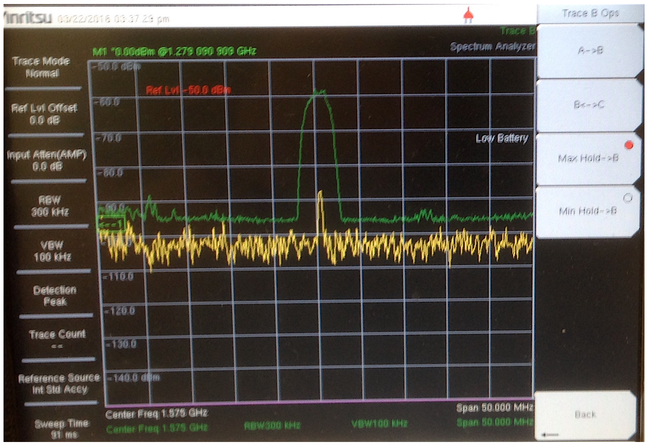

Anritsu 03/22/2016 03:37:29 pm

M1 *0.00dBm @1.279 090 909 GHz

Ref Lvl -50.0 dBm

Trace Mode Normal

Ref Lvl Offset 0.0 dB

Input Atten(AMP) 0.0 dB

RBW 300 kHz

VBW 100 kHz

Detection Peak

Trace Count --

Reference Source Int Std Accy

Sweep Time 91 ms

Center Freq 1.575 GHz

Span 50.000 MHz

Center Freq 1.575 GHz RBW300 kHz VBW100 kHz Span 50.000 MHz

Figure 7: Screenshot of the spectrum analyzer attacked to conductive test setup, depicted in Fig. 6, showing the instantaneous power spectrum (yellow) and max-hold of the spectrum (green) which is indicative of the instantaneous interference-to-noise-floor ratio.

tion was successfuly extracted. It would have been helpful to understand whether ephemeris, almanac, health status and other variables were available, or whether the annihilation of the TOW and subsequent CRC failure rendered all data unavailable. Nonetheless, the results confirm that it is possible to deny a GPS L1 C/A based PVT via the targeted jamming of just a small portion of the navigation message. Beyond the results presented here, a similar systematic interference test was conducted and configured to run continuously over 24 hour period, such that the receiver experienced a complete change in the visible constellation. Again, it was found that the receiver was unable at any point to provide a PVT.

Denial of Galileo E1B PVT

The second test conducted was designed to assess the ability of the systematic jammer to deny observations and a PVT from the Galileo E1B signals. The pulse pattern was further changed to that of Fig. 3 (b) and an experimental setup similar to the GPS case was used, however, due to the low availability of healthy Galileo satellites, the live GNSS feed from the roof antenna was replaced with a simulated signal sourced from a Spectracom GSG-6 Series Multi-Constellation simulator. In this case the pulse pattern significantly more distributed in time, being spread relatively evenly across the I/NAV odd page. This particular pulse pattern was shaped according to the interleaving pattern, rather than being aligned with a particular data word, with the intention that once it is deinterleaved, it will appear as a continuous stream of symbol errors arriving at the decoder. It is expected that this will have a greater impact on the performance of the decoder, than would a set of sparse errors. Interestingly, the ability of this approach to deny the navigation message is relatively insensitive to its alignment with the beginning of the page. Provided the complete set of pulses are received within one page, they will be de-interleaved into a continuous stream. A screenshot from one of the tests is shown in Fig. 9 which includes a trace from 8 Galileo and 9 GPS satellites. As expected, the Galileo E1B message has been denied by the systematic interference. Two interesting observations were made during this test. Firstly, it was noted that the reception of the GPS L1 C/A signal was relatively unaffected. 8 of the 9 satellites report useful observations, and the receiver steadily provided a GPS-based PVT. It is suspected that the single GPS satellite not providing observations is in fact a false acquisition. The second particularly striking observation is that the C/N_0 reported by the receiver under test does not exhibit any significant variation either for GPS or for the Galileo satellites. A consistently high C/N_0, in the range of 48 to 49 dBHz, was reported for all Galileo satellites, despite the fact that the receiver was unable to extract navigation data from any of them.

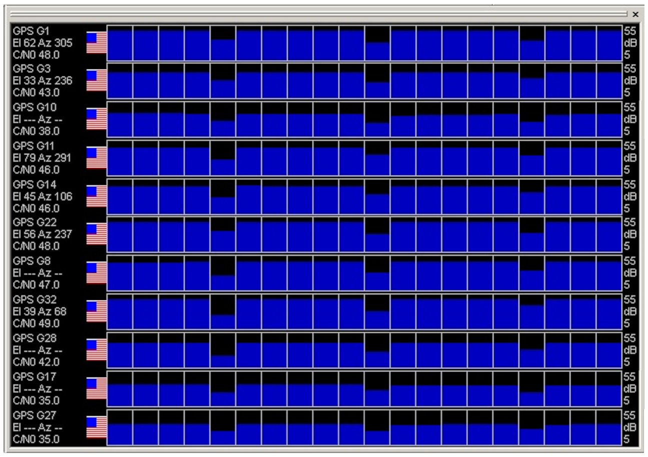

GPS G1
El 62 Az 305
C/N0 48.0

GPS G3
El 33 Az 236
C/N0 43.0

GPS G10
El --- Az --
C/N0 38.0

GPS G11
El 79 Az 291
C/N0 46.0

GPS G14
El 45 Az 106
C/N0 46.0

GPS G22
El 56 Az 237
C/N0 48.0

GPS G8
El --- Az --
C/N0 47.0

GPS G32
El 39 Az 68
C/N0 49.0

GPS G28
El --- Az --
C/N0 42.0

GPS G17
El --- Az --
C/N0 35.0

GPS G27
El --- Az --
C/N0 35.0

Figure 8: Screenshot from the uBlox uCenter™during a systematic interference attack on GPS L1 C/A.

A few interesting conclusions are drawn from these results. We note that is is possible to deny the use of one kind of GNSS signal, in this case, Galileo E1B, while leaving the other, in this case GPS L1 C/A, relatively unaffected, even when the occupy the same RF band. This appears to be due to the relative orthogonality of the navigation message structures, owing to their significantly different symbol periods, 4 ms and 20 ms, and the fact that one employs FEC while the other does not. It is also clear that the observation of C/N_0 may not be a useful means of interference detection, given that the C/N_0 level observed on the GPS and Galileo signals was virtually identical, yet the impact of the interference on the receivers ability to process the signal is drastically different.

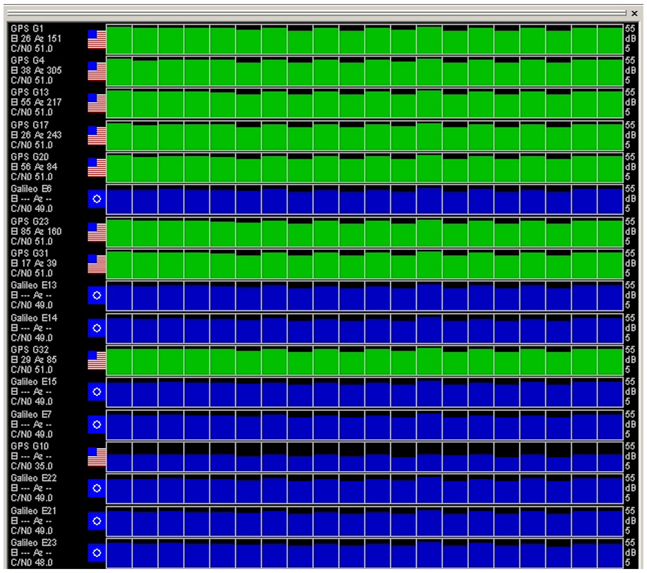

Figure 9: Screenshot from the uBlox uCenter™during a systematic interference attack on Galileo E1B.

Power, energy and synchronisation

The probability of a bit or symbol error occurring is a very nonlinear function of the instantaneous interference power (considering both the complimentary error function, and the nonlinearity of a receiver's front-end), and, as has been shown in the previous experiment, a certain instantaneous power must be applied before any errors are observed in a single test. Increasing the interference power beyond this point will not increase the probability that a bit error has been induced beyond 0.5. To achieve a more reliable denial of the navigation message, more symbols must be targeted, where the probability that the message is corrupted is given by:

$$P_{\text{Err}} = 1 - 0.5^{N_{\text{Pulse}}}. \tag{1}$$

where N_{Pulse} denotes the number of corrupted symbols. This probability tends to unity quite rapidly. Naturally, the total interference energy required increases as a linear function of the number of symbols:

$$E_{\text{Int}} = P_{\text{Int}} N_{\text{Pulse}} T_{\text{Pulse}}, \tag{2}$$

where T_{Pulse} is the pulse periods, being equal to the symbol or bit period. An astute adversary will tune this energy effecting a tradeoff between the likelihood that the navigation message is denied, and the likelihood that the interference power will alert the receiver to the attack, enabling some type of mitigation. Of course, other factors must be considered, including the temporal smoothing that the receiver might apply to signal and noise power estimates, where long averaging of C/N_0 will tend to hide short and sporadic bursts of interference, as shown in the case of Galileo E1B.

Signal	T_{Pulse}	N_{Pulse}	T_{Patt}	K_{Syst}
GPS L1 C/A	20 ms	6	6 s	17 dB
Galileo E1B	4 ms	15	2 s	15 dB

Table 2: Effective reduction in required average interference power when employing systematic jamming.

In terms of the total interference energy, or the average interference power required to render the PVT unavailable, one can consider the effective gain achieved by applying a systematic interference signal, relative to a continuous interference signal. Essentially expressing the average duty-cycle of the interference, it can be computed as the ratio of *on* time to *off* time:

$$K_{\text{Syst}} = -10log10\left(\frac{N_{\text{Pulse}} T_{\text{Pulse}}}{T_{\text{Patt}}}\right), \tag{3}$$

where T_{Patt} is the repetition period of the interference pattern, being 6 seconds for GPS L1 C/A and 2 seconds for Galileo E1 B. The interference configuration for both the GPS L1 C/A and Galileo E1B are summarised in Tab. 2, where it is suggested that the effective gain of applying systematic jamming, as opposed to continuously broadcast jamming, is in the region of 15 to 17 dB. Moreover, although the results here have been generated using a tightly synchronised transmitter (with an error in the region of some tens of nanoseconds), the principle of operation of the systematic jammer would permit synchronisation errors in the region of 1 to 10 ms. Notably, at this level of timing error, the jammer may no longer need to avail of position information.

CONCLUSION

The literature to date has primarily considered the two extremes of GNSS vulnerability, being either a very simple jamming attack, or a very complicated spoofing attack. However, there appears to be a middle-ground, which is very accessible to a malicious attacker, as it only requires commercial, off-the-shelf components, and some basic integration; yet it can pose a significant threat to a naïve receiver implementation. This increased threat comes at a very small increased attack cost and complexity, and has the potential to disrupt many location-based services, by imposing an undetectable partial (data recovery) or full (position and timing) denial-of-service.

Preliminary results suggest that this attack methodology is feasible and, under certain conditions may be quite effective when targeting a naïve receiver. The extent to which this attack methodology can be used has been studied and the various advantages and disadvantages of certain GNSS signal structures have been identified. Further studies will aim to quantify the increased risk posed by such an attack and to identify receiver strategies which may offer improved resilience.

REFERENCES

[1] J. Curran, M. Navarro, M. Anghileri, P. Closas, and S. Pfletschinger. Coding aspects of secure gnss receivers. *Proceedings of the IEEE*, 2016.

[2] F. Dovis. *GNSS Interference Threats and Countermeasures.* Artech House, Boston, 2015.

[3] D. Fontanella, R. Bauernfeind, and B. Eissfeller. In-car gnss jammer localization with a vehicular ad-hoc network. In *Proceedings of the 25th International Technical Meeting of The Satellite Division of the Institute of Navigation*, pages 2885–2893, September 2012.

[4] T. E. Humphreys, J. Bhatti, D. Shepard, and K. Wesson. The texas spoofing test battery: Toward a standard for evaluating gps signal authentication techniques. In *Proceedings of the 25th International Technical Meeting of The Satellite Division of the Institute of Navigation*, pages 3569–3583, September 2012.

[5] Todd E. Humphreys, B. M. Ledvina, M. L. Psiaki, B. W. O'Hanlon, and P. M. Kintner. Assessing the spoofing threat: Development of a portable gps civilian spoofer. In *Proceedings of the 21st International Technical Meeting of the Satellite Division of The Institute of Navigation*, pages 2314–2325, September 2008.

[6] M. Johnson and R. Erlandson. Gnss receiver interference: Susceptibility and civil aviation impact. In *Proceedings of the 8th International Technical Meeting of the Satellite Division of The Institute of Navigation*, pages 781–791, September 1995.

[7] T. Kraus, R. Bauernfeind, and B. Eissfeller. Survey of in-car jammers - analysis and modeling of the rf signals and if samples (suitable for active signal cancellation). In *Proceedings of the 24th International Technical Meeting of The Satellite Division of the Institute of Navigation*, pages 430–435, September 2011.

[8] J. Samson L. Musumeci and F. Dovis. Performance assessment of pulse blanking mitigation in presence of multiple distance measuring equipment/tactical air navigation interference on global navigation satellite systems signals. *IET Radar, Sonar and Navigation*, 8(6):647–657, July 2014.

[9] R. H. Mitch, R. C. Dougherty, M. L. Psiaki, S. P. Powelland B. W. O'Hanlon, B. W. Bhatti, and T. E. Humphreys. Signal characteristics of civil gps jammers. In *Proceedings of the 24th International Technical Meeting of The Satellite Division of the Institute of Navigation*, pages 1907–1919, September 2011.

[10] B. Motella, S. Savasta, D. Margaria, and F. Dovis. An interference impact assessment model for gnss signals. In *Proceedings of the 21st International Technical Meeting of the Satellite Division of The Institute of Navigation*, pages 900–908, September 2008.

[11] NSL. Spirent detector, Accessed 2016.

[12] O. Pozzobon, C. Sarto, A. Dalla Chiara, S. Pozzobon, G. Gamba, M. Crisci, and R. T. Ioannides. Developing a gnss position and timing authentication testbed gnss vulnerability and mitigation techniques. In *InsideGNSS article*, January 2013.

[13] Spirent. Simsafe, Accessed 2016.

[14] u blox. M8 concurrent gnss timing modules, Accessed 2016.

[15] M. Wildemeersch and J. Fortuny-Guasch. A laboratory testbed for gnss interference impact assessment. In *Proceedings of the 22nd International Technical Meeting of The Satellite Division of the Institute of Navigation*, pages 49–54, September 2009.