GPS incident on broadcast networks

Document id:	NIDxxxx
Revision:	A6
Category	Report
Currently responsible:	Magnus Danielson

Abstract

This is a report on how the GPS incident 2016-01-26 affected broadcast networks.

Contact: Magnus Danielson <magda@netinsight.net>

Table of Contents

1 General

1.1 Scope

This report is written on request of Executive Secretariat for the U.S. Civil GPS Service Interface Committee (CGSIC) in order to help the investigations of the GPS event, to illustrate the problems that affected in particular broadcast networks.

It is written with the generous input of Täpp-Anders Skivall at RF Coverage, Didrik Eherenborg at Ehrenborg Networks/Meinberg, Andre Hartmann at Meinberg, Martin Burnicki at Meinberg, Stephen R. Hamilton at CGSIC and Björn Gabrielsson at FOI.

1.2 Revision history

Date	Revision	Responsible	Changes
2016-03-10	PA1	Magnus Danielson	First revision
2016-03-15	PA2	Magnus Danielson	Revision after input from Täpp-Anders Sikvall, Didrik Ehrenborg and Stephen Hamilton, additional text.
2016-03-21	PA3	Magnus Danielson	Revision after input from Stephen Hamilton.
2016-04-05	A4	Magnus Danielson	CGSIC Delivery version.
2016-04-14	A5	Magnus Danielson	Minor adjustment to properly refer to CGSIC
2016-04-20	A6	Magnus Danielson	Minor editorial fix.

Table 1. Revision history

1.3 Abbreviations and acronyms

AGNSS	Assisted GNSS
AGPS	Assisted GPS
BITS	Building Integrated Timing Supply
CGSIC	U.S. Civil GPS Service Interface Committee
DAB	Digital Audio Broadcasting
DVB-T	Digital Video Broadcasting – Terrestrial
EEC	Ethernet Equipment Clock
EGNOS	European Geostationary Navigation Overlay Service
GNSS	Global Navigation Satellite System
GPS	Global Positioning System
GPSDO	GPS Disciplined Oscillator
GSM	Global System for Mobile communications
LTE	Long Term Evolution
MFN	Multi-Frequency Network
MIP	Mega-frame Initialization Packet
NOC	Network Operation Center
NTP	Network Time Protocol
OCXO	Oven-Controlled Crystal Oscillator
OFDM	Orthogonal Frequency-Division Multiplex
PDH	Plesiochronous Digital Hierarchy
PLL	Phase Locked Loop

PMU	Phasor-Measurment Unit
PPS	Pulse Per Second
PRC	Primary Reference Clock
PTP	Precision Time Protocol
SASE	Stand Alone Synchronization Equipment
SBAS	Satellite-Based Augmentation Systems
SDH	Synchronous Digital Hierarchy
SEC	SDH Equipment Clock
SFN	Single-Frequency Network
SNMP	Simple Network Management Protocol
SSM	Synchronization Status Message
SSU	Synchronization Supply Unit
Sync-E	Synchronous Ethernet
TCXO	Temperature Compensated Crystal Oscillator
TDM	Time Division Multiplex
TDMA	Time Division Multiple Access
TETRA	Terrestrial Trunked Radio
UHF	Ultra High Frequency
UMTS	Universal Mobile Telecommunication System
UTC	Coordinated Universal Time
VHF	Very High Frequency
WAAS	Wide Area Augmentation System

1.4 References

[EANTC] EANTC, “Interoperability Showcase 2016 White Paper, 2016, http://www.eantc.de/fileadmin/eantc/downloads/events/2011-2015/MPLSSDNNFV_2016/EANTC-MPLSSDNNFV2016-WhitePaper_Final.pdf

[G781] ITU-T Rec. G.781, “Synchronization layer functions”, 2008-09io4

[G811] ITU-T Rec. G.811, “Timing characteristics of primary reference clocks”, 1997-09

[G.8262] ITU-T Rec. G.8262, “Timing characteristics of a synchronous Ethernet equipment slave clock”, 2010-07

[G.8264] ITU-T Rec. G.8264, “Distribution of timing information through packet networks”, 2008-10

[G.8265.1] ITU-T Rec. G.8265.1, “Precision time protocol telecom profile for frequency synchronization”, 2010-10

[IEEE1588] IEEE 1588-2008, “IEEE Standard for Precision Clock Synchronization Protocol for Networked Measurement and Control Systems”, 2008

[ISGPS200H]IS-GPS-200H., “Navstar GPS Space Segment/Navigation User Interfaces”, 2013-09-24

[KAPLAN] Elliot D. Kaplan and Christopher J. Hegarty editors, “Understanding GPS – Principles and applications (second edition)”, 2006 Artech House, ISBN 1-58053-894-0

[Meinberg] Meinberg, “The impact of the GPS anomaly on January 26th on MEINBERG GPS receivers”, 2016-01-28, https://www.meinbergglobal.com/english/news/global-gps-time-anomaly-on-tue-jan-26th.htm

[RAI] A. Bertella, C. Confalonieri, B. Sacco, A. Scotti, M. Tabone, “GPS timing receivers for DVB-T SFN application: 10 MHz phase recovery”, RAI and RAI Way, 2013

[TR101190] ETSI TR 101 190, “Digital Video Broadcasting (DVB); Implementation guidelines for DVB terrestrial services; Transmission aspects”, 2008-10

[VOLPE] John A. Volpe National Transportation Systems Center, “Vulnerability assessment of the transportation infrastructure relying on the Global Positioning System”, 2001-08-29.

2 Background

On 2016-01-26 the GPS constellation experienced a significant disturbance of UTC time. This report aims to illustrate its effect on the broadcast network infrastructure. In order to illustrate the chain of events, several sub-systems will be described along with a rough analysis of the effect on each of them.

2.1 Telecommunication

Traditionally telecommunications and broadcast networks only had need for stable frequency. Historically, this requirement has been satisfied by high quality crystal oscillators, rubidium standards and for larger telecommunication, networks cesium standards. Most uses of radio clocks (DFC77, MSF etc) have been to set the time of these systems.

For telecommunication, a hierarchy of clocks have been used, where cesium clocks have been used as Primary Reference Clock (PRC), specified to be within +/- 1E-11 relative frequency error of the SI second [G.811], allowing for the maximum slip-rate of one sample every 70 days. Lower quality clocks are locked to the higher quality clock, often in the form of a hierarchy. Within an operators network, the nodes build a distribution tree, which can be fairly deep, and the clock qualities and PLL properties are standardized in order for the system to have well behaved properties.

For traditional GSM networks, providing network synchronization on the E1 PDH connection used for holding voice calls has sufficed both to provide network synchronization and also reference to the air-interface. Even if GSM uses a TDMA type of access, the phase of the hand terminals relative to the base station is trimmed over the air on a regular basis, and the hand terminals also inherit the frequency from the current base-station. On hand-over, the relative phase-difference between base-stations would be an issue, but several options to mitigate it are provided, so that smooth handover can be assured, clearing the TDM slot from the hand terminal migrating from a base-station and thus allowing for a new call (and new revenue). This way GSM avoids phase alignment of the air interface. UMTS have similar requirements. Improving the phase of the GSM base-stations has shown to improve the performance of the network.

The release of GPS disciplined clocks has provided PRC quality reference at a much lower price than the traditional cesium clocks, and also allowed for phase alignment. Integration with base-stations allows not only phase reference, but also provides AGPS/AGNSS integration to aid hand terminals. TETRA depends on phase alignment of base stations in order to handle hand-over without dropping the call, lacking the GSM set of relative mode operations. It has also allowed some of the modes of LTE to use phase alignment. The trend is clear for higher requirements on frequency stability and phase stability for these applications.

2.2 Broadcast networks

For land-based broadcast networks, a number of transmitter towers are fed signals over a high quality telecommunication network. The traditional analog transmissions of TV and radio required only frequency stability for carrier generation. Analog TV also required relatively good stability for the color carrier and related timing, which is provided by the “house clock” of the production facility. Each TV and radio transmitter gets allocated a VHF or UHF frequency and then, within some distance, this frequency is not used by any other transmitter in order for the transmitters not to cause interference of each others signals, as experienced by a user.

With the development of DAB and DVB-T, the signal encoding now needed to handle multi-path suppression, something achieved using the OFDM technique. This handled some of the multi-path problems unsolved by the analog signals.

The scenario of transmitting the same programme over a larger area then involves multiple transmitters, and in order not to interfere with each other, they use multiple frequencies for transmissions throughout the network, this is also known as Multi-Frequency Network (MFN).

The multi-path suppression capability of the receivers makes them able to handle strong reflections, but also two transmitters transmitting the same signal at approximately the same time at the same frequency. This would allow regions to form where transmitters are frequency and phase aligned and transmitting the same programme on the same frequency at the same time. This scenario is known as Single-Frequency Network (SFN). The benefit of SFN is that overlapping allows receivers to use the combined energy of overlapping transmitters, rather than having to choose either one of them. It also allows a much more generous use of support transmitters to fill in gaps (aka “gap-fillers”) since these will now not require frequency allocations. The higher spectrum efficiency of SFN has allowed the spectrum authorities to re-allocate UHF spectrum from TV-transmissions over to broadband services such as LTE. The increased deployment of SFN in broadcast network thus also increases the need for coordinated and stable phase.

2.3 Power-grids

Power-grids use GPS based timing for the measurement of the voltage and current phasors of the power-grid, as means to monitor the power-grids behavior, covering inter-area oscillations, forces oscillations, resonance modes, islanding, phasing of networks for the closing or opening of breakers, steering of network frequency and phase, load balancing and reactive power management. This is done using Phasor-Measurement Units (PMU), and has become the leading tool to measure the quality of the power-grid. PMUs use the GPS time as a common reference phase for all measurement, and the time is also used to time-stamp each measurement such that data from different locations can be collected, logged and analyzed both in real-time as well as logged for post-mortum analysis. PMU data based analysis of power-grids has entered the control-room, and is starting to see applications as “closing the loop” becomes feasible and automatic steering for stability being deployed. NASPI is the leading PMU forum, in close cooperation with DOE and DOE PNNL lab.

PMU uses +/- 1 us level timing and the GPS event is certainly measurable in those systems, forming for this case a disturbance throughout the time of the event.

While not being reported from our customer base, it is included for completeness.

2.4 Financial systems

Financial systems have increased their timing requirements, such that time-stamping of packets to the trading systems is now performed with 50 ns resolution. High-speed trading pushes delay limits and hence timing of when buy or sell orders comes into the trading system becomes of importance in order to establish a strict order by which buys and sells was done. The SOX law put requirements on financial system in the US, and similar requirements have been made in EC.

2.5 Transportation tracking

In modern shipping, containers holding valuable goods, uses GPS for tracking and mobile systems for reporting. Similarly have fleet management systems become increasingly used, where GPS reports the location of trucks and boats. GPS (and mobile network) jammers have now become used to obstruct such trackers and fleet systems, for the purpose of stealing as well as for the purpose of hiding the trucks whereabouts, as the truck driver may want to do an unsanctioned stop.

3 Sub-systems

3.1 GPSDO

The GPS Disciplined Oscillator (GPSDO), which for telecommunication and broadcast folks often is referred to as the “GPS clock”, typically consists of a GPS receiver and then a separate oscillator and control logic. There exists GPSDOs where the oscillator is tightly integrated into the GPS receiver itself, but most common is the use of a separate GPS OEM module.

The most common GPSDOs use GPS receivers only using the GPS L1 C/A code signal, typically using code phase solutions. Multi-frequency GPSDOs are very rare, effectively ruling out L2C, L5 signals. Multi-system GPSDOs exists, however, typically allowing GPS+GLONASS operation.

The GPS receiver typically operates in a timing mode, where the location of the receiver antenna has been surveyed and hence known. With this knowledge, the GPS receiver converts the observed pseudo-ranges from all observed satellites into time errors for a time-only solution rather than the 3D+T solution commonly used for GPS receivers performing 3D positioning and time solutions. Further, a T-RAIM mechanism is used to remove out-liners of the pseudo-range observations, similar to the traditional RAIM for 3D+T solutions. A GPS receiver capable of and operating in time-only mode can operate on a single GPS satellite, but the use of multiple satellites reduces the confidence bounds on the timing-solution provided.

The solution of time is corrected for the individual GPS satellites error to GPS time. Further, a correction term for the difference between GPS master clock time and UTC, such that receivers can alternatively produce UTC approximation (often referred to as UTC time, but it is not a formal UTC time, just an uncalibrated local replica approximating UTC). For many systems, when they use a GPS clock, they in fact use the UTC time and phase, not that of the GPS clock directly.

The GPS receiver produces a PPS signal the rising edge of which will reflect the clock cycle nearest to the GPS or UTC second. Many GPS receivers also present a PPS offset for the time error (aka “sawtooth correction) of the edge and second marker its timing solution wanted to represent, as the typical clock cycle of GPS receivers has a period being tenths of ns.

The oscillator being controlled is typically a TCXO, OCXO or rubidium clock, typically being of the frequency of 10 MHz. The 10 MHz clock is divided down to a 1 Hz signal, and the phase is compared with the PPS signal using a Time-Interval Counter (TIC). The Time Error (TE) being measured is then used to steer the frequency using either a PLL or Kalman filter. The time-error can be augmented with the PPS offset error to remove that systematic noise. Additional steering logic resets the divisor to align with the PPS pulse. Supervision of the GPS receiver state is also done, such that when there is a lack of GPS signals, the control system enters the holdover mode in which steering from the GPS module is stopped until it has regained signal. The quality of the clock and the quality of the steering then allows the GPSDO to maintain it's frequency and phase for some time to be within suitable limits. The more expensive oscillator, the better the hold-over properties.

Some GPSDOs provide alarm indications. The most basic one is only to say there is GPS reception, indication of which might be visible on a diode (aka LOCK). When no lock is achieved or when lock has been lost, some receivers allows for outputs to be turned off (aka squelch). Alarm relay output may be available and status over serial port may be available as well. Professional clocks also provide Ethernet/IP interfaces, SNMP management, syslog interface, web-page with logs etc. as well as providing NTP time.

“Telco”, “Central Office” or SSU style clocks often include redundant power, receivers, oscillators, output drivers etc. in order to handle partial failures without affecting service.

3.2 Telecommunication network

Within telecommunication networks, such as SDH/SONET, PDH etc., synchronization is needed primarily in order to avoid data loss as the TDM oriented data is being multiplexed. What is strictly needed is syntonization (achieving the same frequency/rate) but is commonly called synchronization in the telecommunication environment, even if this is formally incorrect as synchronization attempts to achieve the same phase, but most telecommunication standards do not do that.

The traditional telecommunication synchronization is concerned with the accumulation of jitter and wander, and how this can be filtered. This forms the basis for clock stability, clock noise and PLL bandwidths for the filtering of this added noise. The separation of added phase noise into jitter and wander reflect mainly it's source and the ability to filter it, so the standard SDH Equipment Clock (SEC) is able to filter jitter (above 10 Hz) but not wander (below 10 Hz). The next level, being supplied by a Synchronization Supply Unit (SSU), typically the central synchronization equipment on a station, is able to filter quite a bit of the wander, having a much more stable clock (often a “telecom rubidium”) and much narrower PLL bandwidth. The telecommunication network then gets its long term stability from the PRC clocks as presented before, but GPS sources into the SSU have become widespread.

The telecommunication equipment then have a SEC clock and uses that for all it's communication. It then can select clocks that are being sent to the station SSU over a SASE/BITS interface in order to get the station clock back in return. The synchronization routing is done with network management tools dedicated to the purpose, while main signaling in the network is done according to the Synchronization Status Message (SSM) system, as described in [G.781].

With the modernization of telecommunication infrastructure, the Synchronous Ethernet (Sync-E) [G.8264] provides the same basic service as traditional telecommunication as a frequency transport, transporting SSM-messages over Ethernet and using a variant of the SEC being the Ethernet Equipment Clock (EEC) [G.8262].

Another recent modernization is the IEEE 1588v2 [IEEE1588] aka Precision Time Protocol (PTP). It's telecommunication profile [G.8265.1] make is also adapt the SSM based infrastructure, EEC style clock etc.

3.3 Broadcast transmitters

The analog transmitters only require the frequency of the carrier to be steered. With the digital transmitters, the symbol rate may also be steered, but it is only with the SFN mode of operation that frequency, symbol rate, phase and transmitted symbols needs to be coordinated among the transmitters.

In SFN mode, the receiver needs to experience sufficiently coherent transmitters. The transmitters needs to output the same symbols at almost the same time at the same carrier frequency. The SFN operation became a possibility when achieving microsecond level timing became feasible with the GPSDO. While timing requirement sometimes is quoted to be +/- 1 us, in practice +/- 5 us is manageable, as the guard interval between the symbols in the OFDM is used to let multi-path “ring-out” and considering that the guard band can be 224 us [TR101190] allowing 5 us in the overall budget for transmitter time errors is feasible.

A particular issue with SFN is to that transmitted symbols needs to be transmitted about the same time. Given the same transport streams, two modulators will output the same symbol stream, so there is a need to ensure that the transport stream is being sent to the modulator at the same time (or rather, somewhat before to compensate each modulators delay). For DVB-T, this is achieved by putting in a marker, a Mega-frame Initialization Packet (MIP) in the transport stream which holds a reference point and a timing reference to the PPS at the SFN adapter. The overall delay to be required is also encoded. Then, as the stream is received at a transmitter site, it enters an SFN sync adapter which then takes the PPS and 10 MHz and measures the experienced network delay, buffers the signal and then outputs the signal such that the reference MFS experiences the right delay. The delay difference is being monitored such that buffer overflow or underflow can be avoided. Similarly the PPS and 10 MHz is supervised for stability as the transmitters internal timing reference is locked to it, providing symbol rate and carrier frequency.

DAB transmitters achieve SFN similarly to how it is achieved for DVB-T.

When operating in SFN mode, the transmitter needs both the timing signal and the transport stream to be correct, or else the transmitter starts to act as a jammer for nearby regions. Thus, when either is missing, or significantly out of spec, the transmitter needs to turn off it's output. This is being logged and reported back to the Network Operating Center (NOC) over SNMP, along with monitoring of timing and data-stream.

In their analysis for the national DVB-T network and the requirements on GPSDOs, RAI research laboratory built a dedicated test setup and measured the hold-over properties of a number of commercially available GPSDOs [RAI]. They have also provided some recommendations for how they need to operate to handle the misconducts that they could find as they intentionally jammed them.

3.4 GPS-free networks

In order to support SFN, an alternative approach to installing GPSDOs at each transmitter site is to let the network transporting the transport stream also provide the timing of PPS and 10 MHz to the transmitters. This is done using Time-Transfer (TT) which, similar to GPS, compensates for the transmission time, using two-way time transfer, thus compensating for delays through the distribution path. Typically the time-transfer system extends the telecommunication synchronization.

A practical benefit for doing this is that installing a GPSDO can come at the same cost as the GPSDO itself. Further, considering the importance of the GPSDOs, it is additional equipment to be supervised, which makes the NOC have to supervise at least three systems (transmission network, transmitters and GPSDOs) instead of only two if a network based timing solution can be chosen. Further, for a large network, the GPSDOs and their antennas will fail regularly somewhere in the network.

Some operators also have the requirements to be GPS-independent, if feasible. This is to be able to handle GPS jamming. Even with this requirement, these networks not strict GPS independent, because the timing source(s) typically GPSDOs with good properties. Having phase coordinated to UTC helps in evaluation, use of GPSDOs as backup source or for additional sites. Also, much of the protection from jamming is achieved as most of the sites will not be exposed to the typical problems. The remaining sites then become more sensitive, but for commercial operators this is usually acceptable. More investment can be made on the few sites that act as timing reference sites in order to make them more tolerable.

For some locations, only relying on GPS may be prohibited by law or local regulations, but in those situations, GPS+GLONASS can usually be accepted. For such solutions, GPS time-scale can be ruled out as the output as the common denominator in time-scale choices is UTC which both systems support.

The GPS-free network thus serves some of the protection objectives, but even if it depends on GPS ultimately, it serves many purposes in how the network is operated.

4 Failure analysis

4.1 GPS/GPSDO problems

On 2016-01-26 SVN23 was decommissioned, triggering a ground control system software bug, resulting in errors in the GPS-time to UTC-time correction polynomial, specifically the A0 offset that became close to 13,7 us offset [Meinberg], in the up-linked broadcast message for a number of satellites. This will affect the UTC timing solution, but not the position or GPS timing solution. This offset is also far outside of the +/- 1 us limits of [IS-GPS-200H].

As described, the GPSDO receives the signal like any GPS L1 C/A receiver, builds the pseudo-ranges from the code-phase observations, and then corrects them according to [IS-GPS-200H]. They get filtered using T-RAIM and a time-solution is presented. This part of the process was not affected by the problem.

As the GPS time is to be converted into UTC time, for PPS and presentation time, the UTC correction fields need to be applied. As we typically receive 8-10 GPS signals, we can arbitrarily choose any of them to use common parameters, which is also what a typical GPS firmware does. Once a subset of satellites broadcasts incorrect UTC corrections, the individual GPS receiver will then make it's arbitrary selection independently and be either on the mark or 13,7 us off the mark. It may then change its preference as the received signal changes and change its preferred solution to another one. This process effectively makes each receiver make individual jumps to and from the correct and offset signals and thus shift its UTC output time with +13,7 us or -13,7 us. This will continue until all GPS satellites broadcast correct signals again. This behavior was also observed at multiple locations.

The GPSDO also trains it's oscillator, and the usual shift around is in a handful of ns ranges, achieving around 70-100 ns in modern receivers. For the intended application, this is quite enough for many applications. However, as the PPS now shifts 13,7 us we need to steer our oscillator to that phase, which means a drastic frequency shift is needed to slowly shift it in place. As the GPS now selects another satellite with a correct offset, the process restarts with the opposite sign. Thus, the produced 10 MHz and PPS are moving around during the full event. Comparing two GPSDOs does not make sense either, as they will make independent decisions and do these shifts at different times.

Depending on the advancement of the GPSDO algorithm, the GPSDO may output or may not output alarms. Experience shows that some GPSDOs leave much room for improvement in this regard.

While it may seem obvious to try to filter out “good” from “bad”, this has not previously been a design objective for these receivers. The GPS system has behaved reasonably good for a long time.

The GPSDOs of Net Insight illustrated this behavior clearly, with severe side-consequences. Several of our customers reported the same problems. We issued a message to our customers, but received very little additional response. Meinberg customers have reported issues, and the same issues where most likely experienced by all vendors, as this is not a vendor specific issue. The EANTC test [EANTC, page 18] illustrates this fact clearly.

4.2 Telecommunication networks

Similar to the GPSDO oscillator training, each node in the telecommunication network that directly or indirectly derives its timing from a GPSDO being affected will not be able to track in the phase shifts smoothly. Both phase and frequency will be affected. Due to the loop bandwidth, there will be smoothing, however, as much of the phase-step will be smoothed already by the GPSDO. The main effect will be time-lag, causing the phase shift behave like a wave through the network. The variation of delay can cause data-loss in the cases that buffer levels are too close to the margin, but for many nodes the difference will be so small as they track along that they will not be greatly affected. However, adaptation on the output interface can often have trouble compensating for such variations.

During the EANTC testing [EANTC, page 18], the failure affected the full test-setup, causing the long term testing to fail. It should be made clear that all vendors were affected. As this is the testing of the ITU-T recommendations for all future 4G and 5G networks, it illustrates how national infrastructure has become highly sensitive to these problems.

The network testing inside of Net Insight was greatly affected as our GPSDOs started to misbehave.

Since wireless telecommunication is a significant part of civilian society, it has grown to become a critical system in terms of larger national crisis. Events affecting large parts of the telecommunication capability affects emergency numbers (i.e. 112/911) as well as basic communication. Civilian network capabilities have taken over for many of the previously dedicated networks and land-line systems in a range of alarm, health care, information spreading etc. Regulators have requirements on the availability, and this event may be of interest for any long-term system availability analysis and robustness requirements.

4.3 Broadcast transmitters

As the typical broadcast transmitter running SFN experiences things, it monitors it's transport stream as received from the telecommunication network and compares the transport stream arrival time with that of the given timing.

As timing sways in broadcast transmitters having local GPSDO as reference, the transmitter may detect the timing problem, but most significantly, as the transport stream slowly moves around, the timing reference which will move relatively fast. As the transmitter normally expects the timing reference to be stable and network to change, it will report this as instability of the transport stream. We had several TV and radio broadcast customers that initially reported this as a problem of the transport network. One customer even started to see problems where timing was independent of the transport networks, so they concluded themselves that the problems where more common to their GPS use rather than their transport network use.

Transmitters that get their timing from the network can experience significantly lower levels of difference between their reference and the network signal. Their changes have significant amounts of common mode, as they have the same source. However, due to differences in output properties, there will be somewhat different responses over time and phase-shift, so some of it will be detectable. However, as the transmitter follows the timing, it too will experience much of this as common mode and for such customers we had reports that their GPSDOs had problems.

Outside of our own customer base, BBC issued a statement reporting that their DAB-transmissions were affected by the event.

Public TV and radio broadcasts are part of any country's national resources, being critical infrastructure and these systems need to be robust, especially at times of significant national crisis. Continuity of DVB-T and DAB transmissions, alongside other SFN transmissions such as analog FM, is a significant availability issue and in this case several such networks were affected. This is also why GPS-free networks are required in several countries. Further, there is strict availability requirements on these infrastructures. For these reasons, this is why this event has serious implications on such transmitter networks, as several broadcast networks where affected and capability to transmit without disturbance effectively lost.

5 Previous events

This is not the first GPS issue. Multiple signal anomalies have occurred over the years where the transmitted signal have caused issues one way or another. In addition, a number of jamming or spoofing cases exists, but they are somewhat out of scope here, as the focus is on the transmitted signals and their interpretation.

Some additional type of problems is included, to further illustrate the set of vulnerabilities that the full system experience.

5.1 1024 week roll-over

The traditional GPS L1 C/A signal [ISGPS200H] only has a 10-bit field to identify the GPS week. This provides a 19,6 year calender cycle, which has it's start with GPS week 0 started on 1980-01-06T00:00:00Z. As GPS receivers have been manufactured, the roll-over time has come closer, so receivers have used simple roll-over compensation based on some random GPS-week and then correct the GPS week number accordingly, thus shifting the GPS week roll-over further into the future. In form of pseudo-code, this form of correction becomes:

```
if (gps_week < 500)
          gps_week = gps_week + 1024;
```

Thus, if the GPS-week received is in the range 500-1023, it is interpreted as GPS-week 500-1023, while if the received GPS-week is in the range of 0-499, it is interpreted as GPS-week 1024-1523. The trouble with this correction method is when GPS-week 1524 is transmitted as GPS-week 500, at which time the presented time jumps 19,6 years backwards in time. Such events have been shown to occur regularly for different values, and the errors range from just displaying the incorrect date to failure to maintain lock.

The value of 500 has been shown to occur in some receivers, while other offsets have been shown for other receivers. This is a re-occurring issue. This problem have affected both civilian and military installations. Affected applications include the fields of telecommunication, broadcast, economical transactions (stock trade being stopped) and astronomy.

This is a deficiency of the signal rather than receivers, while the design of receivers control how well they can handle the case. One approach to overcome this ambiguity is to allow the user to enter the date, where the current year is sufficient to identify which 1024-multiple is the correct one, and a complementary solution is to maintain a standard battery-backed RTC clock, which will solve most of the issues as long as there is an operating battery, at which time a fall-back to manually entering the date can be used.

Modern signals uses a 13-bit value for GPS week, providing a 8192-week roll-over scheme. Modern receivers utilizing L2C, L5 or L1C signals can use this scheme for correction. Similarly other UTC sources can be used to assist for this ambiguity resolution.

5.2 GPS health code

During this GPS event, a particular vendor's aviation GPS experienced failure as each new GPS satellite was set healthy for the first time. Due to the expenses of re-validating aviation systems, at least initially, no adjustment was made. This affected a passenger aircraft on route, causing them to loose GPS position. The pilots had to resort to other means of navigation.

The expenses for re-validating the receiver was prohibitive to get proper operation, which is a worse situation. The quality assurance system should not be the cause of delaying such adjustments.

5.3 PRN31

The traditional description of NAVSTAR GPS identify it as a 24-satellite system. The original signal structure did not handle almanac for more than 24 satellites. It was later extended to 32, but then the use up to 30 seemed like the next limit. In practice, all 32 PRN codes assigned for satellites have been used, as the older satellites have graced us with their ability to operate well beyond their scheduled life-time. However, the description of the system has caused assumptions to be encoded into GPS receiver firmware. One such assumption being made has to do with the operational status of PRN31 which affected the old PLGR and MAGR receivers. As PRN31 was activated in orbit, these receivers where affected.

5.4 PRN32

A particular interesting case was the activation of PRN32. One fine Monday morning, all the GPS receivers of a particular telecommunication operator failed, within minutes, throughout its nation-wide network. The operator of the network received complaints from their customers and discovered that their GPS receivers had hanged. It required manual power-cycling of the receivers, they would sometimes lockup or crash, but within hours they crashed again. This was distributed over around one hundred transmitter sites and required manual intervention on each site, thus making it painstaking. Unfortunatly, manual intervention had little improvement to offer. As a result, their mobile communication service completely failed. The GPS receiver vendor offered no real assistance, claiming ionospheric disturbances. This claim was, however, easy to verify as being false by looking at a magnetometer record, only small increase in activity could be shown, and only 24 hours after the event hit. However, looking at the NANU announcements, it could be correlated to the inclusion of a new GPS satellite as PRN32, and as this satellite came over the horizon it could be tracked.

The particular bug is an interesting illustration of software memory bug. For each satellite we track, we maintain state, so we need to have memory allocated. We can choose to do this for each GPS channel of the receiver, or we can choose to do this for each PRN code. Consider now that we do this for the PRN code, and we have 32 satellites, we can declare this as:

```
struct gps_sat_state_s gps_state[32];
```

thus allocating 32 pieces of gps_sat_state_s as indexed 0-31. However, if elsewhere in the code, the PRN1 is indexed as 1, PRN2 is indexed as 2 etc., PRN32 will be indexed as 32, which is outside of the index-range 0-31. Thus, whenever the receiver starts to track PRN32, it starts to write into a part of the memory not allocated for that purpose. This memory can have any use, but if this happens to be the stack of the processor for instance, the receiver will hang as a result. This is the mostly likely scenario here. Whenever PRN32 is visible, it has a chance of being selected and tracked, and when it was, such as during this GPS event, the receiver crashed.

In this particular case, the GPS vendor was not helpful in upgrading the firmware of the receivers, as it was long out of their support cycle, and most likely, the infrastructure to maintain and release software has been dismantled and is long gone. It also illustrates the problem that GPS receivers can have built-in easter-eggs that unforeseen in their designed life-cycle, and that their operational lifetime can be much longer than often considered. They sit there, they keep doing a fine job, but then one day, given a certain set of circumstances, they fail. In this case they downed the full network and the network's service. The network operator was forced to acquire new GPS receivers and deploy them throughout the network, replacing all the old ones. The end result being that the network was nonoperational for weeks, with several customers very unhappy, as it served as a backup-network for several systems. The cost for the operator was loss of traffic, extensive overtime for staff, and punishment bills for undelivered service.

5.5 GPS/GNSS jamming

While jamming and spoofing is natural scenario in the military world, where electronic warfare is part of the strategic and tactical aspects with measures, counter-measures as well as counter-counter-measures, this is not as common in the civilian world, or rather, used to be. GPS was designed with a certain amount of jamming and spoofing resistance [KAPLANio4], some of this resistance is only available to keyed receivers, being able to fully utilize the Y-code and now new M-code. In addition, military receivers have been designed with jamming resistance as a active concern. Jamming scenarios is tested on regular basis in both lab and open door events.

Most civilian receivers only use the traditional L1 C/A signal, which do contain some jamming suppression in the form of the code, but due to the signal conditionings is relatively easy to build jammers for, as most receivers have no real protection for in-band jamming. Very cheap receivers and antennas can even have bad selectivity such that out-of band signals can jam them, but they work relatively well as the surrounding band is relatively quiet. Only a handful of civilian receivers have been designed with jamming in mind, but receivers are now appearing with jamming detection capabilities.

It should be noted that the vast majority of civilian receivers only do GPS L1 C/A code-tracking. Relatively few provides smoothed carrier-phase or carrier-phase measurements. Very few do any form of multi-frequency, and this is typically in the GIS application where dual-frequency receivers do GPS L1 C/A in addition with L1 and L2 semi-codeless P(Y) tracking. The GIS applications now see increase advancement in receivers, with L2C, L5 tracking, GLONASS L1 and L2 tracking etc. However, outside of these applications, price sensitivity have made GPS L1 C/A only receivers dominate. In a jamming scenario, jamming in the GPS L1 would lock out essentially all GPS users, but some may still get some signal on L2C but for their purpose they don't get the needed precision due to loss of dual frequency capability. T ripple frequency receivers now exists for the GIS segment, where L5 can be used, however the reference network does not really provide the needed support for L5.

Many simple GPS receivers uses 1-bit samplers, which have documented [KAPLAN] poor performance when being jammed. The 1.5 bit receivers with good AGC control [KAPLAN] have somewhat better performance. GPS receivers is essentially blocked either by LNA overloading or lack of dynamic capabilities in the analog front-end and ADC conversion.

Many civilian GPS receivers does not have choke-ring antennas or other antennas with deep nulls in equatorial and below horizon orientation. This is typically only used for multi-path suppression, under the assumption that reflections comes from below or near the antennas equatorial plane. This also provides some margin for ground-based or distant jammers.

The relatively poor jamming resistance of civilian GPS receivers make self-oscillating antenna amplifiers a threat, as illustrated by the Moss landings incident. The use of GPS for tracking has also caused some people either concerned about their privacy, those with criminal intent or those that just don't want their employers to see all they do eager to consider the jamming ability. This have caused a market for cheap and available GPS jammers.

This have caused several incidents such as that reported by FAA in New Jersey, but also it has become in accelerated use by criminals in order to steal cars, boat engines and similar expensive goods, by jamming GPS they disable the GPS tracker on them, allowing for stealing and untraced transport to other country where they can disable the tracker and sell in modified form.

Another example is the jamming of a port, which prohibited the location of containers, such that the cranes moving containers could not locate the right containers, thus halting the operations of the port. By having people run around the port knocking on all the trucks there asking whoever it was running the jammer to turn it of, they where able to have the jammer turned of and the port could open its operations again.

Unintentional jamming also occurs from other systems, such as TV-transmitters who's third overtone is the L1 signal. Another example was when a paging systems transmitter antenna was sitting 2 m from a GPS antenna, and the field-strength from the 21.3 MHz 50W transmitter was sufficient to saturate the LNA of the antenna, achieving full blocking.

Thus, GPS jamming has become a civilian concern, and whenever intentional or unintentional jamming occurs, those operating a critical infrastructure of some sort needs to have made their system robust in the sense that it needs to detect the jamming condition and have some suitable form of counter-action in order for the system to survive, preferably with little or no impact on the availability and performance.

Multi-frequency/Multi-system GPS/GNSS jammers is now available, due to the cheap technology of jamming a band, and that jamming various mobile network frequencies is needed to eliminate mobile communication, extending this into the multiple GNSS systems is a natural extension. Thus, just becoming independent on GPS L1 may not be enough to become jammer resistant.

The civilian GPS vendors have only recently been able to participate in open door jamming events. The jamming scenarios is not always adapted to the fixed GPSDO scenario, which is relevant to much of the fixed infrastructure.

The vulnerability assessment known as the Volpe report [VOLPE], done for DOT, in 2001 remains a good reading, but few consider jamming as part of the threat to many infrastructures.

Augmentation over the network could overcome some of the jamming problems. This assumes sufficient timing stability over the network, network wide monitoring and augmentation of both network and GPS/GNSS receivers. However, few networks provide sufficient infrastructure to achieve this.

5.6 GPS/GNSS Spoofing

While spoofing, the generation of false signals, is a known technique from the military side, it is essentially not considered in the civilian context. The Volpe report [VOLPE] shows how a GPS signal simulator can be used for the attack. Today the cost of GPS or GNSS constellation simulators have dropped, and they have uses outside of GPS/GNSS receiver testing, as they can be used to improve or provide GPS/GNSS navigations in buildings and tunnels. The use of GPS/GNSS in a wide range of devices provides a good platform to aid navigation as we have increased use for it. These tools can however act as spoofers either due to misconfiguration or test configuration. The motivation to use GPS/GNSS simulators instead of re-transmitting the GPS/GNSS signal is that a much better signal to noise can be maintained in the system, since the signal is being generated fresh rather than being picked up from the air close to the noise of the antennas.

The replay of signals also acts as a method to confuse receivers.

The work by University of Texas, Austin show that a relatively cheap setup allows for the spoofing attack to gently “take over” the attacked receiver and then dominate it and make it gently drift away. Cooperative work with the PNNL DOE lab have shown that this also works with GPSDO and PMU measurements, where the PMU measurement tracks the attackers intended time, thus illustrating not only the GPS receiver, but the system attack. This was presented by PNNL at the NASPI GPS vulnerability workshop.

5.7 Installation failures

A further concern which is far from the broadcast GPS signal is a range of installation failures. Antennas that fill up with water due to missing or broken porous plug where water can escape have killed GPS antennas. Lightning might break them, and use of lightning arresters is very sparse. Poor choice of cable, increased loss in cables have occurred. Failing OCXOs also contribute to loss of service.

Mounting of antennas require care, among others not so ice builds, that ice builds in mounting holes so it cracks the case open. Further, where snow can build, the antenna should be built such that snow can slide of the antenna cover. GPS antennas can also be popular places for birds to land and sit, and bird spilling can form a cover over the antenna.

Location of the antenna should attempt to be multi-path free and with a good view of the sky. Locating antennas with much of the horizon obscured, deep in urban valleys of buildings is not optimal. One installation saw the loss of satellites, and as they came to the site they discovered that they where building a bridge over the antenna.

Installation near other antennas, in the field of a microwave link or other sources of strong fields is a risk.

Installation where hot steam passes by has been shown to cause trouble. At one location the janitor put the christmast tree on the top of the building, feeling that the GPS antenna mounting was a good place to mound the tree. As the snow falled, the branches went down and no signal received, as the snow melted, the signal came back.

Installation with too long cables can cause too high loss on the cables, making the signal to noise (S/N) ratio to small for tracking, this can lead to unstable results.

6 Recommendations

6.1 Common robustness specifications

Commercial GPS users, such as telecommunication and broadcast network operators, were exposed to a robustness attack from the GPS system itself. The robustness of GPS has not previously been a major issue, and hence there is a weakness in GPS receivers. In order to enable the market to have a relatively transparent view of what robustness measures have been taken in a particular product, it would be good if a number of robustness checks can be described. Profiles or full set can then be required by operators as they buy or upgrade their networks. It also enables the GPS/GPSDO vendors to have equal access to such countermeasures, alarms and monitoring features. The GPS operations can also benefit from this, as it helps to create an additional safety-net.

The GPS signal has considerable redundancy, both within a signal from a particular satellite and between satellites. Additionally, SBAS (such as WAAS and EGNOS) can then improve on this redundancy in a significant manner.

As future receivers are developed, redundancy between different signals (L1 C/A, L1C, L2C, L5) can to some degree, be utilized. But for the a foreseeable future, L1 C/A only receivers can be expected to dominate. Similarly, other GNSS system might aid and, to some degree, this already exists today to some degree.

A number of consistency check that can be applied in receivers can be done. A number of likelihood analyses can be made. For example, orbit parameter changes have certain limits. The GPS or UTC time does not jump or ramp outside of certain limits, etc. Majority decisions and consistency detections is possible to utilize. To some degree, SBAS corrections can be used to validate orbital changes.

Additional robustness can be achieved using stable references, but as receivers can be compared over the network they can act as redundancy for each other. This would not fully handle the given event, but it would have significantly reduced it's impact if properly implemented.

For many systems, the worst situation is silent malfunction. Some GPSDOs can only be monitored by looking at the LOCK LED to be lit, and then you have to physically be there. Telecommunications systems have an elaborate set of built-in tools to discover, monitor and convey various error statuses. The GPSDOs also to be incorporated into this such that the many warning signs can be detected, false alarms avoided, significant problems escalated, turning off signals on significant failures. This helps when building large complex systems, as alternative timing sources may be used.

6.2 Up-load verification

Produced GPS upload signals need to be decoded and verified with robustness checks similar to those used for receivers, both isolated and with use of the predicted properties of orbit, time etc. This provides an additional verification prior to upload the satellites.

6.3 Interface Specification clarification

Management and documentation within the interface specifications can be improved such that future behavior of the system can be anticipated. For instance, the full use of PRN-numbers, active satellites, GPS-week wrap, etc. can be documented such that receiver manufacturers and software developers expect that such parameters may occur in the future.

6.4 Promote improved infrastructure robustness

The improvement of civilian infrastructure robustness includes failure mode awareness, increased robustness of receivers, improved robustness of networks and improved performance monitoring and alarm systems.

Increased robustness of receivers includes both detection of anomalies as well as better use of redundancies. Additional signal sources such as WAAS/EGNOS based SBAS, as well as promotion of cheap civilian multi-frequency receivers that goes beyond the classic L1 C/A signal to include L2C, L5, L1C as well as aiding from semi-codeless L1/L2 P(Y) tracking could help. It is unfortunate that the use of such improved receivers has been hindered by high prices of receivers and antennas.

For many of these purposes, a high degree of transparency of receiver state allows the extended supervision, comparison and augmentation. This should also include improved jammer and spoofing detection methods. Only high-end receivers have capability to deliver such data, and provide standard interfaces for such observables, such as RINEX.

6.5 Recommendations for professional users

While this report highlights some of many issues relating to use of GPS/GNSS receivers in various professional uses, such as telecommunications, broadcast, power-grid and financial systems, of which some has high availability requirements and may even be considered critical infrastructure. This report has the primary goal to provide some insight of the professional use as being affected by the GPS incident, but does not aim to provide a full report of recommendations for such professional users. Such a report should be produced with specific recommendations for this audience.

Professional users should however consider that a number of improvements in their robustness may be required. Operating of GPS/GNSS receivers without monitoring of their performance, means to alarm of failures, fall-back upon failure and a consequence analysis of what will fail and what remedies should be applied. Operation of GPS/GNSS receivers beyond their support life-cycle means that their failure can force them to replace all receivers with affected services being down, possibly for weeks.

A particular danger is that many systems now include GPS/GNSS receivers embedded, or delivered with the system, but their function has little or no documented impact on the functionality of the system. Upon procurement and installation of such system, it has been seen that “there is no need for synchronization” because receivers have been built-in, such systems installed in tunnels have show improper function.