

Using Assurance Cases and Boolean Logic Driven Markov Processes to Formalise Cyber Security Concerns for Safety-Critical Interaction with Global Navigation Satellite Systems

Chris W. Johnson

School of Computing Science, University of Glasgow, Glasgow, UK, G12 8RZ.
Tel +44(141)3306053, Fax +44(141)3304913,
Johnson@dcs.gla.ac.uk, <http://www.dcs.gla.ac.uk/~johnson>

Abstract

Satellite-based location and timing systems support a wide range of mass market applications, typically using the GPS infrastructure. Until recently, these applications could not be used within safety-critical interfaces. Limits to the accuracy, availability, integrity and continuity of the space-based signals prevented regulatory agencies from certifying their use. Over the last three months, however, the latest generation of augmented Global Navigation Satellite Systems (GNSS) have been approved for use in safety-related applications. They use a range of techniques to overcome the limitations of previous infrastructures. This means that they can be used as primary navigation tools in a wide range of interactive systems, including aircraft cockpits, railway signalling tools etc. Unfortunately, a range of organisations including the UK Ministry of Defence, have raised concerns about our increasing vulnerability to attacks on these satellite based architectures. These threats are compounded by the difficulty of representing and reasoning about the impact of jamming, spoofing and insider threats for the end-users of safety-critical systems. A sudden loss of navigational support can undermine users' confidence in complex applications and pose a significant threat to distributed situation awareness. We show how formal reasoning techniques can be used to identify the safety and security concerns that jeopardise interaction with future generations of Global Navigation Satellite Systems applications.

Introduction

This paper considers security and safety concerns for the users of Global Navigation Satellite Systems (GNSS). Our concern is partly justified because there is a growing dependency on timing and location information provided by these applications. This creates significant vulnerabilities for many different infrastructures across Europe and North America (RAE, 2011). First generation GNSS architectures, such as GPS and GLONASS, suffer from a number of errors. Some of these problems stem from satellite geometry. If all the satellites are closely grouped together then the benefits of differential signal processing will be reduced. This tends to act as a multiplier for errors induced from other sources. For instance, gravitational forces create subtle changes in the orbit of the satellites within a GNSS constellation. Further problems arise from multipath effects. The signals arriving at a receiver are often reflected from large structures including buildings. This creates inaccuracies of up to 75 meters in urban environments because the reflected signal will take longer to arrive than a direct transmission. Further problems stem from ionospheric effects. Radio waves can be considered to travel at the speed of light in outer space. However, the ionizing effects of solar radiation form layers that refract electromagnetic waves from satellite transmissions. Most end users do not correct for unforeseen changes such as variations introduced by strong solar winds (Köhne and Wößner, 2010).

These errors and the lack of suitable correction mechanisms help to explain why first generation GNSS architectures have not been widely integrated into safety-related systems (Bhatti and Ochieng, 2007, Johnson and Atencia-Yepe, 2011). For example, the International Civil Aviation Organization (ICAO) Required Navigation Performance provides minimum performance criteria for the use of these architectures in a range of applications. These criteria are expressed in terms of **Accuracy**- how correct is the aircraft's position estimate; **Integrity**- the largest aircraft position error can reach without detection; **Availability**- how often can the aircraft use the systems with the desired accuracy and integrity; **Continuity** - the probability that an operation once commenced can be completed and **Time to Alert** - the maximum interval before a performance issue is detected.

Satellite Based Augmentation Systems (SBAS) address these limitations and are intended to support Safety of Life (SoL) applications. These architectures include the North American Wide Area Augmentation System (WAAS) and the Asian Multi-functional Satellite Augmentation System (MSAS) as well as the European Geostationary Navigation Overlay Service (EGNOS). These systems use a number of ground reference stations, to compare known information about the time and their location with the signals received from the satellites to

compute differential corrections and integrity data for each monitored satellite. This information is collated by a smaller number of master stations that then broadcast deviation corrections using a second network of geostationary satellites. The corrections can then be applied to location information derived from the GPS or GLONASS networks. Europe has just certified the EGNOS GNSS for Safety of Life (SoL) applications, including approaches to aircraft runways. Further applications are proposed, for example in railway signalling, where exact information about the speed and position of trains can be used to ensure adequate separation without 'locking' sections of track, as in conventional approaches.

The EGNOS infrastructure uses redundancy to address many of these concerns in SoL applications. For instance, each of the four master stations rotates from being the Master to a Hot-Back-Up and then to be a Cold-Back-Up. The complexity of the underlying design created a host of human factors concerns. In particular, configuration and maintenance posed particular challenges given the potential consequences of errors by the systems engineering teams. The design of this infrastructure was, therefore, guided by the assumption that no single operator error would lead to a loss of integrity. This has significant implications in terms of the security assessments in later sections; it is less clear whether this level of protection might also be offered against deliberate attacks.

The complex and critical nature of satellite based augmentation systems also justifies the use of formal analysis techniques. Numerous mathematical studies were made of the underlying algorithms to establish the correctness of the relationships between the signals received by the ground reference stations and the differential corrections as well as the reliability information broadcast to the eventual end users. However, the reliability analysis was driven by semi-formal techniques based on Failure Modes, Effects and Critical Analysis. The results from these studies were supported by operational observations of test applications to provide evidence that helped to demonstrate conformance with the requirements listed above. The following pages argue that complex interactions between safety and security concerns justify the use of more formal techniques to support the future development of next generation GNSS.

A further motivation for the use of mathematical reasoning is that safety and security requirements extend beyond the underlying SBAS to include the interactive applications that rely on these infrastructures. This extends formal analysis in novel ways; most formal approaches tend to focus on an individual level of abstraction. In contrast, many SBAS applications also employ Receiver Autonomous Integrity Monitoring (RAIM) to detect potential faults in the underlying GNSS measurements. Additional signals, for instance from other satellites arrays, are used to confirm the fixes derived from the underlying SBAS. RAIM services offer significant benefits to end users. For instance, they can be used during critical phases of flight, such as an approach, when the pilot needs to be informed of such inaccuracies as soon as possible (Oliveira and Tiberius, 2008). As we shall see, formal analysis can be used to verify that RAIM warnings support SBAS data during critical operations. However, mathematical verification techniques must also be extended to consider a host of human factors issues. For instance, it is unclear whether end users must understand the underlying reasons why their systems indicate the need to perform a 'go around' when RAIM alerts identify potential inaccuracies in GNSS data.

Security Threats to GNSS Infrastructures

Most of the design concerns that motivated the development of SBAS focussed on safety rather than security requirements. The existing infrastructures remain vulnerable to a range of attacks. An early warning was provided by an approach into New Jersey during December 1997. The crew of a Continental trans-Atlantic flight lost all GPS signals; jeopardising confidence in an array of on-board systems. It was initially believed that this had been caused by an intentional jamming attack. It later turned out to have been the unintended result of a US military test. A 200-kilometer "interference zone" was created by a GPS antenna with a 5-watt signal, stepping through frequencies.

The UK Ministry of Defence (MOD) illustrated the potential threat for maritime navigation (Grant, Williams, Ward and Basker, 2009). A medium powered jamming device generated noise over a pre-defined area of the UK coastline. This study clearly illustrated the impact that the threats to GNSS integrity can have upon the end users of these infrastructures. Particular problems were identified for crews using integrated bridge systems. This technology brings together navigation tools with autopilot control so that a jammed GPS signal could lead to a significant deviation without warning. Even if an alert was issued there are significant barriers to identifying the correct position given the consequent loss of situation awareness. The crews in this trial were all aware that the GPS signals would be jammed. However, multiple simultaneous alarms rapidly increased the crews' workload as they cross-checked navigational information. The impact for on-board systems was

compounded by the impact of jamming for shore-based services. Numerous errors began to undermine the Vessel Traffic Services that provide an overview of coastal areas. Some of the data returned by vessels was based on incorrect GPS fixes that contradicted radar sources.

Many of the vulnerabilities associated with convention GNSS architectures stem from the relatively weak signals that are used. A common analogy is to compare GPS output to using the power of a car headlight across one third of the Earth's surface at more than 20,000km. Most western military organisations can interrupt GNSS signals; simulation software enables planners to identify the optimal allocation and distribution of jamming systems. The military development of satellite navigation jamming devices has been mirrored by the increasing availability of hand held systems that cost little more than \$100 and have a range of several kilometres. These portable technologies can be used in a range of criminal activities – for instance, to disrupt the signals to GPS tracking devices that would otherwise report the location of a stolen vehicle or shipment. It is illegal to offer these devices for sale within the European Union. This is because they cannot comply with the existing Electro-Magnetic Compatibility (EMC) directives; the prohibition was not primarily intended to protect GNSS services. Within the UK, national legislation prevents the operation of a jammer but it is not illegal to own such a device (RAE, 2011).

Further threats illustrate the relationship between underlying systems vulnerabilities and the usability of safety-critical SBAS applications. First generation GNSS infrastructures provide little support for users trying to authenticate signals. This makes it possible to 'spoof' location information through the broadcast of fake GNSS-like signals or the rebroadcast of valid GNSS signals. Signal simulation software can be used to recreate the anticipated GPS signals for a given route using a particular set of waypoints and timing intervals. Coupled with a spoofing transmitter, these simulators can fool the user into thinking that they are following a specified route. The problems of designing the simulator and then integrating it with effective, mobile jamming technologies have created significant barriers to their application for criminal ends. However, these are likely to be eroded in coming years and the potential threats cannot be discounted. The criminal motivation is proportionate to the diversification of GNSS applications – including route monitoring for toll and insurance pricing.

Some of these concerns are being addressed through technological innovation. For instance, spoofing will become far more difficult once Galileo begins to provide encrypted signals for use in safety-related applications. Other threats continue to affect future GNSS architectures. As mentioned, the design of the EGNOS and Galileo ground based systems focused on meeting accuracy, integrity, availability, continuity and time to alert requirements. Software and hardware teams were focused on a series of 'feared events' and failure modes. Algorithmic barriers and standard operating practices, including maintenance procedures, were then created to address these concerns. Deliberate attacks were not part of this analysis. In consequence, a number of 'low level' vulnerabilities may persist. Fortunately, the defences that were created in response to safety concerns also provide protection against potential security threats. The same CRC and error checking techniques that help to identify potential failure modes can also identify a range of attacks. There remains some concern as to whether these defences would offer sufficient protection against 'insider' threats. For instance, most SBAS infrastructures rely on configuration files that enable operators to respond to the failure of particular components. This increases confidence in meeting the safety requirements, cited above. However it also creates opportunities for malicious reconfiguration. Similarly, GNSS infrastructures are typically designed to operate autonomously for short periods of time. Elements of the infrastructure can also be commanded from more than one ground station. This creates a concern that external agents or insiders could spoof legitimate commands or gain temporary control of the infrastructures. This might sound relatively far-fetched. However, the investment in relatively simple attack modes such as those used by STUXNET provides a warning of future vulnerabilities as more and more national infrastructures rely on satellite based navigation and timing information. These potential threats also reinforced the point that security concerns extend beyond the scope of an initial safety analysis into the entire operational life cycle of GNSS architectures through development to deployment and maintenance.

A recent report from the UK Royal Academy of Engineering (2011) argued that 6% of GDP in Western Countries depends on GNSS technology. It went on to criticise the lack of backup technologies. At a national level, agencies should monitor and report on disruption to GNSS signals. At an international level, greater attention should be paid to the vulnerabilities that over-reliance on this technology is creating in the financial markets. Managerial and operational staff should prepare for GNSS outages from ten minutes to a month. The RAE team also argued that GNSS vulnerabilities should be explicitly included in the risk assessments that support critical infrastructures. The limited scope of the RAE report did not, however, identify formal or semi-

formal techniques that might support such analyses. In contrast, the following paragraphs identify a range of tools that might increase the resilience of safety-related SBAS applications.

High-Level Structures: Semi-Formal Dependability Cases

The safety of Satellite Based Augmentation Systems (SBAS) depends on many different forms of evidence, including but not limited to risk assessments, architectural descriptions, development standards, test data, independent audits etc. In previous work, we have described how semi-formal argumentation techniques help to structure these different sources of evidence. In particular, Johnson and Atencia Yopez (2011) use this approach to reason about the impact of security threats for the users of safety-critical GNSS applications. In this approach, the edges of a graph denote that evidence supports a particular goal. These diagrams support discussion between stakeholders who often have very different viewpoints on the validity of particular safety or security arguments. For instance, the Goal Structuring Notation (GSN) uses a *goal or claim* to represent an assertion that can be assessed as either true or false (Kelly and Weaver, 2004). A developer might assert that RAIMs techniques are ‘acceptably safe during low probability continuity failures’. A *solution* can be used to present the evidence that supports a goal or strategy. This is important because it provides a link between the high level argument structure embedded within GSN and the more detailed documentation provided by specific development techniques such as Fault Trees, FMECA, Formal methods etc.

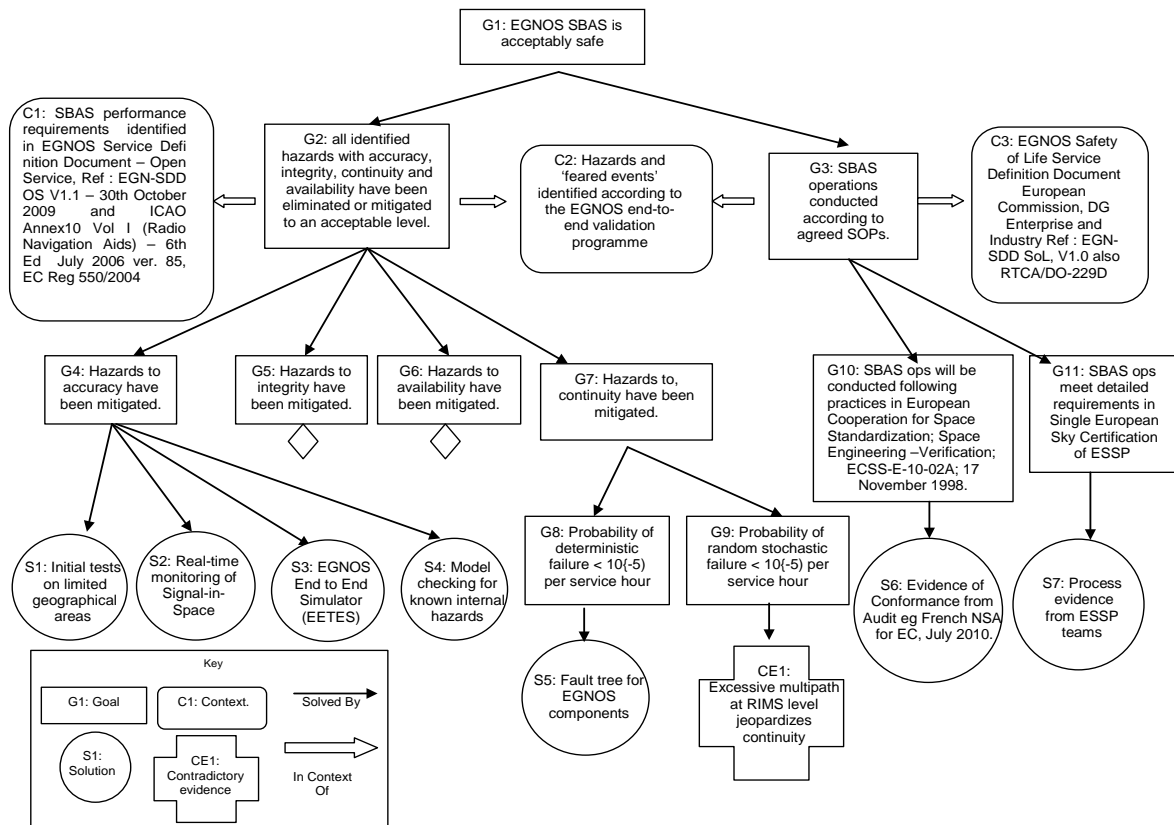


Figure 3: Initial GSN for Satellite Based Augmentation Systems (SBAS)

Figure 1 shows how GSN can be used to argue that an SBAS is acceptable safe. This top level goal can be broken down into sub-claims. In this case, G2 focuses on eliminating or mitigating the hazards that might undermine the ICAO performance requirements in terms of accuracy, integrity, continuity and availability. G3 focuses on the need to operate the SBAS according to the identified Standard Operating Procedures (SOPs). These goals are placed within the context of specification and requirements documents including EC Reg 550/2004, EGN SDD SoL etc. Evidence that these sub goals have been addressed can be derived from a range of tests – initially on limited geographical areas and subsequently by more sustained monitoring of ground stations. The EGNOS End to End Simulator (EETES) also provides evidence of robustness against accuracy concerns. Solution S4 also shows how evidence from formal analysis techniques, including model checking, can be integrated into the safety arguments within a GSN.

The semi-formal, safety case structures used to support EGNOS certification for SoL applications are significantly more complex than that illustrated in Figure 1. A modular approach was, therefore, developed to structure arguments that addressed low level infrastructure concerns through to usability issues (Johnson and Atencia Yopez, 2010). Part A of the safety case explains why the system has been ‘designed, developed and deployed’ in a manner compliant to ICAO Standards and Recommended Practices (SARPS). This was coordinated by the European Commission with support from the European Space Agency as the lead body in the initial design of the EGNOS architecture. In contrast, Part B argues that the SBAS will be operated and maintained to meet the ICAO SARPs by the commercial European Satellite Services Provider (ESSP). Additional safety cases are then required for each of the end-user applications that are built on top of the SBAS SoL architecture.

Goodenough, Lipson and Weinstock (2008) have used GSN to analyse security concerns. In this case, the top level goal demonstrated that the system was acceptably secure. Their approach focussed on relatively low-level threats to software systems; structuring evidence that an implementation will not be vulnerable to buffer overflow attacks. The arguments still relied upon a range of human factors claims. For instance, programmers must be trained to avoid vulnerabilities in their code. Competent programmers must also be used to review the software. Goodenough, Lipson and Weinstock’s work is particular relevant for this workshop; their assurance cases also integrated evidence from mathematically based tools for static code analysis. It is possible to identify a number of different ways in which semi-formal argumentation techniques might integrate security AND safety concerns for interactive systems:

- *Integration within a single dependability argument.* Under this approach, the top level goal would be to demonstrate the dependability of a complex system. A first sub-goal would present the arguments that any implement was acceptably safe. A second sub-goal would then structure the evidence showing that the system was acceptably secure. This approach raises a number of concerns – in particular, it is difficult to show that some security evidence has implications for systems safety and vice versa.
- *Integration of safety concerns into security assurance cases.* This approach would begin by constructing the security arguments pioneered by Goodenough, Lipson and Weinstock (2008). Additional nodes might then be introduced into the diagram to distinguish evidence or arguments about security concerns that might undermine the safety of any implementation. This approach suffers from a range of practical problems – for example, it is possible to identify potential safety concerns with every threat or vulnerability. However, there are additional safety hazards that would not be represented in the combined diagram because they are not strictly related to the original security assurance case.
- *Integration of security threats into safety cases.* We have chosen to adopt a third approach. This begins by developing a conventional safety case, such as that illustrated in figure 1. The second stage is to use conventional forms of threat and vulnerability analysis to identify security concerns that were not identified during the previous step. Additional evidence must then be introduced into the hybrid structure to document any additional mitigation that must be introduced to address these security concerns beyond those that were already considered as part of an initial safety assessment.

Figure 2 shows how safety and security arguments can be integrated into a single graphical structure (Johnson and Atencia Yopez, 2011). The augmented assurance case identifies two safety concerns. The first uses the UK MOD studies to provide evidence that localised disturbances to GPS or GLONASS signals would not be visible to an EGNOS ground station. The threats posed from such interference can be mitigated through the application of the RAIMs techniques mentioned in previous sections of this paper. However, the representation of security and safety arguments within an integrated GSN helps to document the importance of these approaches for the dependability of future applications. Figure 2 also includes a second set of security concerns based around the potential ‘insider threat’ to GNSS infrastructures. These are rarely modelled within simulation environments; however, coordinated attacks by individuals who are familiar with the ground architecture of an SBAS system would undermine many of the defences that are intended to mitigate the impact of individual human errors.




Figure 2: Integrating Security Threats to GNSS Architectures within GSN Safety Arguments

Previous paragraphs have shown how semi-formal assurance cases can be used to represent and reason about the interaction between security and safety concerns in complex, interactive systems. These graphical structures collate the many diverse forms of evidence, including formal analysis, that support dependability arguments. They are particularly appropriate for SBAS and GNSS applications where the quality of end-user interaction is determined by the interaction of multiple infrastructures each supported by the cooperation of many different commercial and regulatory organisations.

A number of caveats can be raised about the use of the hybrid technique, illustrated in Figure 2. For instance, the EGNOS safety case is divided into different parts addressing the design, implementation and operation of the underlying infrastructures. Further components have been developed to demonstrate that particular applications are acceptably safe. It remains to be seen whether it is possible to trace the impact of particular threats and vulnerabilities across these different boundaries to identify the impact that particular cyber attacks might have upon the end users of GNSS data. Further limitations stem from the qualitative and semi-formal nature of most safety cases. This limits the inferences that can be drawn about the potential impact of safety hazards and security threats. In order to look for more quantitative support, we must look more closely at the techniques that can be used to obtain the evidence, which is documented in the nodes of an assurance case.

Lower-Level Analysis: Boolean Driven Markov Processes (BDMP)

Safety argumentation techniques, such as the GSN used in Figures 1 and 2, provide a graphical means of integrating evidence from formal analysis with the products from other safety-related software engineering processes that are required by most development standards and by regulatory agencies. This is important because formal methods are not sufficient in themselves to support the design, implementation and operation of most complex GNSS applications; they must be augmented by risk analysis techniques, by functional testing and usability studies, by environmental and observational data etc. This section goes on to look at one formal approach that can be used to consider both safety and security requirements, while reinforcing these initial links between complementary development techniques.

Boolean Driven Markov Processes (BDMPs) have strong similarities to more conventional fault trees that have been supported the development of many different safety-related interfaces. Bouissou and Bon (2003) provide an overview of the approach and explain one important difference with the BDMP semantics. Boolean Driven

Markov Processes extend conventional Fault Trees with triggers that are represented by dotted arrows similar to the priority AND gates that have been integrated into some Fault Trees. The target of the arrow is activated if the source state occurs. These relationships between the leaf nodes in the BDMP can be formalised within Markov Processes, the use of Boolean logic embedded in the tree notation also helps to reduce the potential state space explosion associated with conventional Markov Processes in applications as complex as the GNSS described in this paper.

Annotation	Triggered Markov Representation	Semantics
------------	---------------------------------	-----------



Figure 3 introduces the annotations that provide a bridge between the fault tree components and the underlying Markov Processes (Bouissou and Bon, 2003). Figure 4 goes on to apply the approach to model N-2 safety-related failures involving GNSS implementations. Each of the diagrams illustrates how multiple independent failures can undermine the confidence and situation awareness of end-users by interrupting the provision of navigation and timing information. In Figure 4 a) there is a probability λ that the RAIMs secondary system will fail to provide a correct location when it is needed following the loss of a primary SBAS fix. In contrast, Figure 4 b) considers the situation in which the RAIMs fault may occur before the loss of SBAS; this might happen during a ‘fail silent’ problem or as a result of incorrect configuration, maintenance etc. Figure 4 c) illustrates how these two different scenarios can be combined within the same BDMP model.

Piètre-Cambacédès and Bouissou, (2010) have recently extended Boolean Driven Markov Processes to represent and reason about security threats for complex systems. This approach borrows many of the concepts that have been embedded with the attack trees that are themselves based on concepts from Fault Tree notations. As before, the intention is to associate Markov Processes with the leaf nodes of the tree structures. These leaves are used to represent the actions of potential attackers; they capture important state changes that describe whether an attack has been executed or not, whether it can be detected or not once it has begun and so forth. The Markov processes describe the probability of that node being true; this creates the opportunity to develop an executable semantics from the notation in which a state transition in the underlying process will trigger the leaf node to become true. Figure 5 summarises the security extensions to the BDMP modelling language.

Annotation	Semantics
	

Figure 6 illustrates an important strength of the BDMP approach. The probability of an attack on both the RAIM and SBAS implementation might be very low. However, the UK government has recently acknowledged that greater consideration needs to be paid to what are known as N-2 vulnerabilities. These characterise low probability, high consequence failures of multiple infrastructures where there may be hidden interdependencies – for instance, the insider threat from a common infrastructure provider (software or hardware) cannot be discounted for both RAIMs and SBAS technologies.

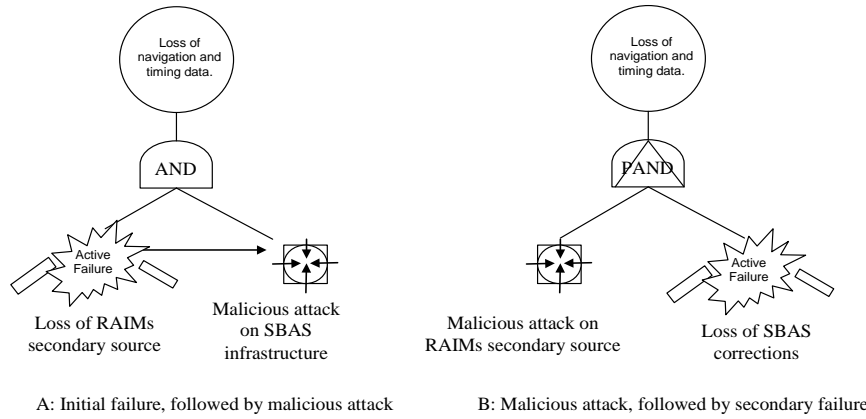


Figure 7: BDMP Hybrid Modelling of Security Attacks and Accidental Vulnerabilities

Figure 7 extends the BDMP approach to consider a hybrid vulnerability in which external attacks exploit wider failures in a national critical infrastructure. In diagram A, the attacker waits until there is a failure in the RAIM infrastructure before mounting an attack on the SBAS infrastructure. This does not seem a plausible scenario – given that the attacker would need to know that the failure had occurred and that the users of an SBAS application would continue operation without the support of the RAIM implementation before they launched their attack. This formalisation is important because it provides the basis for discounting certain lines of attack providing this decision can be validated by a broad range of stakeholders. In contrast, diagram B models a scenario in which the attacker disables the RAIM implementation and allows the system to continue operating until there is a subsequent fault in the SBAS application. Again the plausibility of the attack mode might be questioned; however, this scenario describes the insertion of a latent threat that has significant potential to compromise the safety and security of many complex systems.




Figure 8: BDMP Integrated Modelling of Security Attacks and Accidental Vulnerabilities

Figure 8 builds on the BDMP extensions to construct a model that integrates accidental faults with malicious attacks on both the RAIMs and SBAS infrastructures. As can be seen, the RAIMs for an end-user application might be disabled either through a fault or through an attack. This would lead to a loss of navigation and timing information if it were followed by the failure of an SBAS infrastructure either through an attack or fault. The same consequences would arise if the loss of the SBAS were to be followed by a failure on demand for the RAIM implementation.

The links between BDMP diagrams and fault trees enable the use of minimal cut set algorithms to identify the different scenarios that might compromise the operation of complex, safety-critical systems. Integrated diagrams, such as that illustrated by Figure 8, can be also drive the quantitative analysis of attack and failure scenarios using the underlying Markov Processes that support each of the lead nodes. Their analysis is intended to derive the probability of an undesired event within a given period of time or the mean time in which this event will occur. Piètre-Cambacédès and Bouissou (2010) show how these models can be used to suggest potential vulnerabilities in these integrated models. Their work focuses on the quantitative analysis of pollution incidents. However, their results can be extended to our case study. In systems whose mission time is short it is better to conduct a malicious attack on both infrastructures even though the probability of success might be relatively small. The probability of detection is also constrained by the limited time exposure. Over longer periods of operation, or if an attack takes longer to mount, it is better for the attacker to use a hybrid approach. For example, malicious code might be inserted into a RAIMs implementation and triggered during ‘accidental’ problems in the operation or configuration of a GNSS. This attack strategy could be reversed if there was a higher probability of detection associated with the RAIMs compared to the SBAS infrastructure. In either case, the relatively low probability of detection for malicious code inserted by an ‘insider’ remains a significant concern across most critical infrastructures.

In order to support this quantitative, integrated use of BDMP for security and safety assessments we must define values for in-operation (λ) and on-demand failures (γ). It is also important to model the probability that an attack action will be successful (β). The probability of active failure for the SBAS in Figure 8 can be derived from existing safety case evidence, mentioned previously. The probability of the active loss and the on-demand failure of RAIMs implementations can be derived from application specific safety assessments, e.g. supporting ICAO Required Navigation Performance criteria. However, the remaining values are less clear cut. How can one assess the probability of a successful attack on a RAIMs implementation or an SBAS infrastructure? It is for this reason that we have not attempted to calculate the derived values from the Markov Processes; there are further concerns over attempts to accurately assess the existing vulnerabilities of a critical component of European infrastructures. Some national agencies have argued that the managers of key systems must act as though state sponsored cyber attacks have a probability of one. In other words, the design, management and operation of critical systems must be guided by the assumption that they will be the target of external agencies. This simplifies some aspects of the quantitative assessments, security analysts can focus on the probability of detection before any intrusion has the potential to affect the end users of the underlying infrastructure. The hybrid BDMP analysis in Figure 8 provides a framework for this more detailed work.

Conclusions

Previous generations of Global Navigation Satellite Systems could not be used for safety-related applications. Ionospheric interference and multipath signals combined to create problems of accuracy, integrity, availability and continuity. The underlying architectures provided no guarantees about the time taken to alert the end users to potential problems. These issues have all recently been addressed through the development of ‘Safety of Life’ extensions for Satellite Based Augmentation Systems. In consequence, the European Commission has supported the certification of these infrastructures for applications ranging from precision approaches to runways through to advanced railway signalling systems. Similar provisions have been made for the use of the North American North American Wide Area Augmentation System (WAAS) and the Asian Multi-functional Satellite Augmentation System (MSAS).

A range of tools and techniques have been used to support the development and operation of SBAS infrastructures. These include conventional risk assessment methods, such as HAZOPS and FMECA. They also include a range of empirical tests both on the ground and space based components. Formal analysis has been applied to represent and reason about the algorithms used to calculate necessary corrections to the GPS and GLONASS signals. Usability tests have also supported the systems engineering and configuration tools used in the operation of SBAS infrastructures. Graphical argumentation techniques have been used to gather this

evidence into a coherent structure that demonstrates augmentation infrastructures are acceptably safe within a range of different contexts.

At the same time that satellite based augmentation systems have been approved for location and timing information in safety-critical interfaces, a range of organisations including the US Department of Defense and UK Ministry of Defence; have raised concerns about increasing civil vulnerability to attacks on the underlying infrastructures. These vulnerabilities are compounded by a range of human factors concerns; the more that end users begin to rely on navigation and timing infrastructures then the greater the consequences will be for any interruption to those services. Previous studies by the MoD and by the FAA have shown that crews suffer a significant loss of situation awareness when there is any interruption to GNSS infrastructures. The integration of navigation and timing data into a wide range of applications also undermines operators' confidence in a host of interactive systems, not simply those that report on their present location. For example, studies of GNSS jamming have shown that relatively cheap devices undermine a host of maritime bridge information systems and shore based traffic monitoring applications.

Further concerns arise because it is unclear how to represent and reason about the safety concerns that are created by the diverse security threats to GNSS architectures, including spoofing and the insider threat to both space and ground based systems. Such security concerns invalidate many of the assumptions that support the provision of critical services. One approach would be to extend the application of argumentation techniques such as GSN from safety-related applications to represent security argumentation. Several examples have been developed to show how this can be done for a range of software applications. However, this suffers from a number of limitations. In particular, it can be difficult to represent and reason about the impact that security threats might have upon underlying safety arguments. We have, therefore, extended previous approaches to show how security threats might be used to challenge the evidence that supports arguments about GNSS Safety of Life applications. The intention is to provide an integrated, risk-based approach to the identification of attack scenarios that can help assess the resilience of safety cases to security threats across the life cycle extending from design to maintenance and operation.

The integration of security concerns into safety cases helps to sketch the potential consequences of a malicious attack on an underlying SBAS. A key benefit is that the safety case provides a means of collating the diverse sources of evidence from design, testing and analysis summarised above. However, this evidence must be derived using other tools and techniques. It is for this reason that this paper has also presented a means of analysing the more detailed interactions between the security and safety of GNSS. In particular, we have shown how Boolean Driven Markov Processes help to avoid some of the state explosion limitations of conventional Markov techniques using extensions to the well-known Fault Tree notation. An important benefit of this approach is that it has already been developed to consider attacks on safety-related applications. However, it has not previously been applied to consider failure modes and security threats to critical infrastructures, such as the SBAS described in this paper.

Much remains to be done. In particular, the quantitative application of the Markov processes relies upon estimates of the likelihood of successful attacks on elements of the SBAS infrastructure. This, in turn, must be validated by expert judgement across a range of stakeholders – including the designers and operators of the infrastructure components, the end users of the GNSS applications as well as national intelligence agencies similar to the UK Centre for the Protection of National Infrastructure. The elicitation of these estimates reinforces the close integration between human factors and formal analysis that is a recurring theme in this work and in any study to address the interactions between security and safety.

Acknowledgement

The work described in the paper has been supported by the UK Engineering and Physical Sciences Research Council grant EP/I004289/1. Special thanks are due to Amaya Atencia Yepez, GMV who provided valuable comments on an early draft of this paper – any remaining errors remain those of the author.

References

U.I. Bhatti and W.Y. Ochieng, Failure Modes And Models For Integrated GPS/INS Systems. The Journal of Navigation, 60(2):327–348, 2007.

M. Bouissou and J.-L. Bon, A New Formalism that Combines the Advantages of Fault Trees and Markov Models: Boolean Logic Driven Markov Processes, Reliability Engineering and System Safety journal, (82)2:149-163, 2003.

EUROCONTROL, Safety Case Development Manual, Technical report DAP/SSH/091, Brussels, Belgium, 2006.

J. Goodenough, H. Lipson and C. Weinstock, Arguing Security - Creating Security Assurance Cases, Software Engineering Institute, Carnegie Mellon University, 2008. Available as of 19/3/2011 on <https://buildsecurityin.us-cert.gov/bsi/articles/knowledge/assurance/643-BSI.pdf>

A. Grant, P. Williams, N. Ward and S. Basker, GPS Jamming and the Impact on Maritime Navigation. The Journal of Navigation, 62(2): 173-187, 2009.

C.W. Johnson and A. Atencia Yopez, Safety Cases for Global Navigation Satellite Systems' Safety of Life (SoL) Applications. In H. Lacoste-Francis (ed.), Proceedings of the Fourth International Association for the Advancement of Space Safety, Huntsville Alabama, NASA/ESA, Available from ESA Communications, ESTEC, Noordwijk, The Netherlands, ISBN 978-92-9221-244-5, ESA Technical report SP-680, 2010.

C.W. Johnson and A. Atencia Yopez, Mapping the Impact of Security Threats on Safety-Critical Global Navigation Satellite Systems, To Appear in the Proceedings of the 29th International Systems Safety Society Conference, Las Vegas, USA, August 2011.

L. Piètre-Cambacédès and M. Bouissou, Modelling Safety and Security Interdependencies with BDMP (Boolean logic Driven Markov Processes). IEEE International Conference on Systems Man and Cybernetics (SMC), 10-13 Oct. 2010, 2852 – 2861, 2010.

L. Piètre-Cambacédès and M. Bouissou, Attack and Defence Dynamic Modelling with BDMP (Boolean logic Driven Markov Processes). Proceedings of the 5th International Conference on Mathematical Methods, Models and Architectures for Computer Network Security, LNCS 6258, St Petersburg, Russia, pp86-101, September 2010.

T P Kelly and R A Weaver, The Goal Structuring Notation - A Safety Argument Notation. In Proceedings of the Dependable Systems and Networks 2004 Workshop on Assurance Cases, July 2004.

RAE, Global Navigation Space Systems (GNSS): Reliance and Vulnerabilities, Royal Academy of Engineering, London, UK, 2011. Available as of 19/3/2011 on: http://www.raeng.org.uk/news/publications/list/reports/RAoE_Global_Navigation_Systems_Report.pdf