

A Trust Framework for Evaluating GNSS Signal Integrity^{*}

Xihui Chen^{1**}, Gabriele Lenzini¹, Miguel Martins², Sjouke Mauw¹, and Jun Pang¹

¹ University of Luxembourg, Luxembourg

²itrust consulting s.à r.l., Luxembourg

Abstract. We develop a novel trust framework based on subjective logic to evaluate the integrity of received GNSS civil signals. We formally define *signal integrity* for the first time in the framework and use it to precisely characterise different spoofing detection methods. Our framework captures the uncertainty during the inference of signal integrity which has been largely ignored or not explicitly specified in the literature. Our framework also gives rise to several natural ways to combine the outputs of various spoofing detection methods on signal integrity. We have validated our framework through experiments using both real and simulated signals and the results show that our framework is effective.

1 Introduction

Global Navigation Satellite Systems (GNSS) have become an essential element in people’s daily lives since the American Global Positioning System (GPS) started to offer free civil signals. Nowadays, almost all smart-phones and other mobile devices on the market are equipped with GNSS receivers. People’s access to their real-time locations has popularised numerous location-based applications. These applications are not restricted to offer services for leisure, such as geo-social networks and points of interest search, but also deployed in safety-critical products, like driverless vehicles and aviation navigation. However, as civil signals are neither signed nor encrypted, there is no way to authenticate their originators. In addition, they are broadcast in the open air with a relatively weak strength. Therefore, civil signals can be easily interfered with or even taken over by false signals, which are called *jamming* and *spoofing*, respectively [2, 3].

In the last decade, a number of scientific experiments and examples have successfully demonstrated that civil signals are vulnerable to spoofing. For instance, in 2012 Humphreys et al. [4] managed to take control of an American unmanned plane by sending faked GPS signals. The experimental results lead to the conclusion that once critical applications are targeted, people’s safety and even homeland security can be practically threatened by spoofing attacks. In such attacks, even if GNSS receivers are tamper-resistant, people still cannot guarantee the correctness of the calculated locations.

It is noted by the Volpe report [5] that there were no practical mitigation methods for spoofing attacks and we believe that it is still the case now, especially for GNSS

^{*} This is a short version of the published paper at CSF’13 [1]. The research was partially supported by the European Space Agency (ESA). We also thank Carlo Harpes for his comments.

^{**} Supported by the National Research Fund, Luxembourg under the project SECLOC 794361.

civil signals. Navigation message authentication is considered as an effective method to prevent spoofing [6]. However, due to the long deployment cycle and high costs this is not a feasible approach in the near future [7]. Instead, researchers have proposed many methods with the aim to *detect* but not to *prevent* spoofing. The general idea is to make use of some observable features that should be present when signals are not spoofed. A spoofing attack is detected if one or more of such features are not observed.

Some low-cost methods are proposed to detect unsophisticated spoofing. For instance, Papadimitratos et al. [8] summarise three spoofing detection tests: location inertial test, clock offset test and Doppler shift test. There are also some methods that make use of more advanced attributes of GNSS signals. For example, Nielsen et al. [9] monitor the correlation between the strengths of two signals from different satellites because the strengths always change independently. Psiaki et al. [10] utilise the correlation between the encrypted military signals received by different receivers as the military signals transmitted by the same satellite should be physically the same even if they cannot be decrypted by civil receivers. The above detection methods are designed under the same principle. Namely, given a signal, a method takes the measurement of a certain attribute of the signal as input, calculates the predicted values and claims the absence of spoofing when the measurement is sufficiently close to the prediction.

Research questions. Although researchers have shown the effectiveness of their (own) detection methods through various ways, we find that the existing spoofing detection methods still suffer from the following problems:

1. The notion of signal integrity has not been formally defined, which leads to ambiguous interpretations. Tippenhauer et al. [7] define spoofing from the viewpoint of localisation results. However in some sophisticated spoofing, the attackers may gradually fool receivers to calculate the planned position and then allow receivers to calculate the right location and time when the attack starts [7].
2. Spoofing detection methods have not been systematically characterised. This leads to incorrect inference of signal integrity from the consistency of measurements with the predicted values.
3. The output of a detection method is always *qualitative*, i.e., whether a signal's integrity is preserved or not, while we believe that it should be *quantitative* by its nature. On one hand, the noise from the environment always influences the receipt of GNSS signals and causes changes on certain attributes. On the other hand, a powerful attacker can generate signals with certain attributes consistent with the prediction. Thus, the consistency of such attributes should not always lead to the conclusion of the signal being integrous.
4. The outputs from different spoofing detection methods might conflict with each other and there exist no algorithms to combine the outputs of different methods into a coherent conclusion. Combining the results of multiple detection methods is necessary as more evidences usually lead to more reliable conclusions.

Our contributions. We propose a novel trust framework based on subjective logic to evaluate the integrity of GNSS signals and address the above identified research questions. In our framework, we first formalise GNSS systems and receivers, based on which signal integrity is formally defined. Then we present a generic formal description of spoofing detection methods and classify them based on the relationships

between consistency of attributes and signal integrity. To address the uncertainty in reality, we first take into account the impact of environmental noise and propose a way to obtain an opinion on the consistency of an attribute with its prediction. Next, we present a method based on conditional reasoning with subjective logic opinions to evaluate signal integrity for an individual detection method. In the reasoning, we deal with the uncertainty of the attackers' capability of tuning signals' attributes.

In the end, we propose three algorithms to combine the outputs from different spoofing detection methods. They are designed to capture different assumptions about the attackers' ability to manipulate attributes.

2 Preliminaries

2.1 GNSS Signals and Signal Spoofing

A GNSS system is a constellation of satellites which broadcast navigation signals to the earth. In this paper, we take GPS as a representative due to its popularity. Other systems, such as GLONASS and Galileo, are similar.

GPS satellites are equipped with atomic clocks which are synchronised with the universal time. GPS signals are transmitted in two frequencies f_{L1} and f_{L2} on which navigation data and spreading codes are modulated [11]. Navigation data carries information about the orbits of satellites and spreading codes are used to identify satellites. Each satellite has two unique spreading codes: the coarse acquisition (C/A) and the encrypted precision code (P(Y)). The C/A code is publicly known and encoded in civil signals while the P(Y) code is encrypted and can only be accessed by certified military devices. As we focus on civil applications of GNSS systems, throughout the paper we only consider scenarios where civil signals are targeted by the attackers. Thus, we simply refer to civil signals in the paper as signals.³ A satellite generates its signals by modulating its C/A code and navigation data with the carrier wave of frequency f_{L1} and sends them into the air with a transmitter.

Signal spoofing can be implemented in the following two ways. (a) Because C/A codes are public and no authentication mechanisms protect them, an attacker can construct a signal modulated with a C/A code having arbitrary time offset to the synchronised one. This forgery will lead a receiver to calculate an incorrect distance to the satellite. (b) Since the format of navigation data is also publicly known, an attacker can generate navigation data with arbitrary information but conforming with the format. In this way, the receiver will learn an incorrect location of the satellite. By either or both of these two ways, receivers can be fooled to calculate any locations, no matter where they are. The above two ways of spoofing have been validated in the literature.

2.2 Subjective Logic

Subjective logic opinions. An *opinion* expresses the belief about one or multiple propositions from a space called the *frame of discernment*. An opinion over a frame X is a composite function consisting of three components – a belief function, an uncertainty

³ The P(Y) codes are still part of our signals and can be used to detect specific spoofing attacks.

mass and a base rate function. The belief function assigns belief mass to each proposition in X , which can be interpreted as the positive belief on the truth of the element. It is sub-additive, meaning that the sum of all propositions' belief mass is not larger than 1. Uncertainty mass is the amount of belief that is not assigned as belief mass. The base rate function expresses the *a priori* probability of each proposition in X being true.

Definition 1 (Subjective logic opinion). Let X be a frame $\{x_1, \dots, x_n\}$. An opinion on X can be represented by $w_X = (\mathbf{b}_X, u_X, \mathbf{a}_X)$ where $\mathbf{b}_X : X \rightarrow [0, 1]$ is the belief function, $u_X \in [0, 1]$ is the uncertainty mass and $\mathbf{a} : X \rightarrow [0, 1]$ is the base rate function. Furthermore,

$$\sum_{x \in X} \mathbf{b}_X(x) \leq 1; \quad u_X = 1 - \sum_{x \in X} \mathbf{b}_X(x); \quad \sum_{x \in X} \mathbf{a}_X(x) = 1.$$

The expectation probability of $x \in X$ being true is:

$$E_X(x) = \mathbf{b}_X(x) + \mathbf{a}_X(x) \cdot u_X.$$

When the frame is binomial, e.g., $X = \{x, \bar{x}\}$, the opinion about the truth of x can be denoted as $w_x = (b, d, u, a)$ where $b = \mathbf{b}_X(x)$, $d = \mathbf{b}_X(\bar{x})$, $u = u_X$ and $a = \mathbf{a}_X(x)$ indicating the belief, disbelief, uncertainty and the *a priori* rate about x being true. The expectation probability of x being true is $E(w_x) = b + a \cdot u$.

Conditional belief reasoning. Conditional reasoning offers a way to calculate the truth of a proposition y based on the evidence about another proposition x which has a conditional relation with y . According to the causal relation, we have *deductive* reasoning and *abductive* reasoning. If x (resp., y) is the antecedent, then the reasoning is deductive (resp., abductive). Compared to the probabilistic method, subjective logic takes opinions as input in the reasoning and thus captures the underlying uncertainty.

Deduction and abduction on binomial frames, i.e., $X = \{x, \bar{x}\}$ and $Y = \{y, \bar{y}\}$ have the following notations:

- $w_{y|x}$: conditional opinion on y given x being TRUE;
- $w_{y|\bar{x}}$: conditional opinion on y given x being FALSE;
- w_x : opinion on the proposition x ;
- $w_{y||x}$: opinion on y deduced/abduced from the observation on x .

Assume we have a causal conditional between x and y , i.e., “if x then y ” (denoted by $x \rightarrow y$) and $w_{y|x}$ and $w_{y|\bar{x}}$ are learned. If we have an observation on x which gives the opinion w_x , then the deduced opinion on y should be calculated by considering both of the situations when x is TRUE and FALSE. In subjective logic, ‘ \odot ’ is used as the operator calculating the opinion on y given w_x and the two conditional opinions $w_{y|x}$ and $w_{y|\bar{x}}$, i.e., $w_{y||x} = w_x \odot (w_{y|x}, w_{y|\bar{x}})$. If we have evidence on y i.e., the opinion w_y , then the opinion on x can be calculated by abductive reasoning. The idea is to calculate $w_{x|y}$ and $w_{x|\bar{y}}$ based on $w_{y|x}$ and $w_{y|\bar{x}}$ using the Bayesian theorem, where the *a priori* probability of x , i.e., a_x , is required. In this way, deductive reasoning can thus be used. In subjective logic, ‘ \oslash ’ is the abductive operator calculating w_x based on $w_{y|x}$, $w_{y|\bar{x}}$ and a_x , i.e., $w_{x||y} = w_y \oslash (w_{y|x}, w_{y|\bar{x}}, a_x)$. We refer the readers to [12, 13] for the details of the implementation of the operators. Conditional reasoning is applicable on multinomial opinions as well.

3 A Trust Framework

3.1 GNSS Systems

A GNSS system consists of a number of satellites which move in certain orbits. We denote by \mathcal{S} the set of running satellites of the GNSS system. Let \mathcal{L} be the set of all geographic coordinates and \mathcal{T} be the set of time points. We use $\xi(S, t) \in \mathcal{L}$ to denote the real location of satellite $S \in \mathcal{S}$ at a given time $t \in \mathcal{T}$.

Satellites broadcast radio signals to the earth. GNSS signals are generated by a fixed procedure such that they have a common pattern. We take GPS signals as an example. A GPS signal includes at least two components: (1) the C/A codes of a deployed satellite (2) a navigation message with ephemeris information. Let Θ be the set of all possible GNSS signals that conform with the pattern. We use the function $sig : \mathcal{S} \times \mathcal{T} \rightarrow \Theta$ to return the signal transmitted by a satellite at a given time.

Natural factors, such as ionospheric scintillation and tropospheric effects, can attenuate signals. Attenuation can cause effects on many attributes of a signal, e.g., carrier phase advance and power decrease. Its impact is determined by the routes that signals take to arrive on the ground. As these routes are subsequently determined by where they reach and when they are generated, we use $\eta(S, \ell, t)$ to denote the attenuation on the signal of $S \in \mathcal{S}$ which is generated at time t and arrives at ℓ . We denote by $\eta(S, \ell, t) \diamond sig(S, t)$ the signal when $sig(S, t)$ reaches the earth. The signal is still an element of Θ as long as the spreading codes and the navigation data are available.

3.2 GNSS Receivers

A GNSS receiver is a device to capture GNSS signals and calculate a location with a localisation algorithm. In fact, a receiver captures the combination of the signals of all satellites in range. Let \mathcal{G} be the set of combined signals and let \uplus be the combination operation on any two signals with the same radio frequency. Then for any $s \in \mathcal{G}$, there exists a set of GNSS signals $\Theta' \subseteq \Theta$ such that $s = \uplus_{sig' \in \Theta'} sig'$. The set \mathcal{G} is closed under the signal combination operation. We use $s(\ell, t) \in \mathcal{G}$ to denote the combined signal received by the receiver located at $\ell \in \mathcal{L}$ at time $t \in \mathcal{T}$.

Given a received signal, the receiver separates the GNSS signals modulated based on their unique features, e.g., C/A codes. This process can be modelled by function $sigCom : \mathcal{G} \rightarrow 2^\Theta$ mapping a received signal to the set of combined GNSS signals.

As the receiver has access to the C/A codes of all satellites, given a GNSS signal in Θ it can identify the satellite whose C/A code is modulated. We call the satellite the *originator* of the signal. We use function $ori : \Theta \rightarrow \mathcal{S}$ to return the originator of any signals. The originator is not always the entity that actually generates the signal as the attackers can generate signals with the same code.

A GNSS receiver implements a localisation algorithm that takes a received signal as input and outputs a coordinate and a time point if possible. We denote the algorithm by $loc : \mathcal{G} \rightarrow \mathcal{L} \times \mathcal{T}$. In practice, the output of a localisation algorithm is of the form of a triple consisting of a coordinate, an accuracy in meters and time. The coordinate and the accuracy define a round area centred at the coordinate with a radius of the accuracy. Since our focus is signal integrity, we assume that localisation algorithms always calculate accurate locations with accuracy zero.

3.3 Signal Integrity

When a received signal is free of spoofing, we usually say that the integrity of the signal is preserved, meaning that the signal has not been modified maliciously by the attacker. In other words, an integrous signal is generated by a satellite and without artificial interference, e.g., replaying, before reaching the receiver. Given a received signal, the key point of verifying its integrity is to calculate its reference signal which is supposed not to be spoofed. First, the time between the generation of the reference signal and its arrival at the receiver should be equal to the amount of time required to travel the distance between its originator and the receiver by the speed of light. Second, it should suffer the correct amount of attenuation, e.g., $\eta(S, \ell, t)$, during the transition. We use $|\ell, \ell'|$ to denote the Euclidean distance between two positions ℓ and ℓ' . Based on the above discussion, signal integrity can be formally defined as:

Definition 2 (Signal integrity). *Given a received signal $s(\ell, t)$, we say that $s(\ell, t)$ is integrous if and only if for each $sig' \in sigCom(s(\ell, t))$, there exists $t' \in \mathcal{T}$ such that*

$$(sig' = \eta(ori(sig'), \ell, t') \diamond sig(ori(sig'), t')) \wedge (c \cdot (t - t') = |\xi(ori(sig'), t'), \ell|)$$

where c is the speed of light.

In the following discussion, we use $\mathcal{I}_{s(\ell, t)}$ to denote the proposition that “ $s(\ell, t)$ is integrous” while $\neg \mathcal{I}_{s(\ell, t)}$ represents the negation that “ $s(\ell, t)$ is not integrous”. In practice we cannot use Def. 2 to verify signal integrity by computing the integrous signals and comparing them with the received ones. On one hand, the location of a receiver is under calculation and not available until the integrous signals having been received. Without the location, it is impossible to derive the transmission time of the received GNSS signals and thus the generation time cannot be obtained. On the other hand, the attenuation cannot be measured due to the nature of unpredictability of the environment. Therefore, we cannot learn the set of GNSS signals that should be received.

3.4 Attacker Model

In general, the aim of an attacker is to fool a receiver to calculate a fake location. According to the literature, the attackers have two ways to achieve this purpose – software attacks on receivers [14] and GNSS signal spoofing [7]. In this paper, we focus on the risks coming from signals, as people can protect their receivers against malware but have no control of signals. We assume that the localisation algorithm of a receiver is always well protected and free of misbehaviour. Formally, given a received signal $s(\ell, t)$ if it is integrous then we have $loc(s(\ell, t)) = (\ell, t)$.

The attackers that we consider have similar capabilities in terms of signal transmission to the attackers assumed by Tippenhauer et al. [7]. They have full control of wireless channels by blocking, intercepting, delaying and replaying GNSS signals. Furthermore, we assume that the attackers can manage to make all their signals received by the targeted receivers at any preferred time. For signal generation, we assume that the attackers can generate any GNSS signal in Θ that can be interpreted by receivers.

3.5 Spoofing Detection Methods

A spoofing detection method aims to evaluate the integrity of a given signal. It takes the measurement of a certain attribute of the signal as input and calculates a set of predicted values of the measurement. At last it decides whether the signal is integrous, by comparing the measurement to its predicted values. In the following discussion, we formally characterise spoofing detection methods and classify them.

Given a received signal $s(\ell, t)$ we denote by $Attr(s(\ell, t))$ the set of attributes of $s(\ell, t)$ that can be measured and explored by a spoofing detection method. In this paper, we assume that a spoofing detection method explores only one attribute as it is designed in the literature. The value of an attribute can be measured by a receiver or calculated by other agents. We denote by $m_\alpha(s(\ell, t))$ the value of attribute $\alpha \in Attr(s(\ell, t))$ of $s(\ell, t)$. The domains of the measurements are different between attributes. To be generic, we use $dom(\alpha)$ to denote the domain of α . Note that for the sake of simplicity, we assume that a measurement has just a single value in its corresponding domain, while in practice the measurement of an attribute might be of different forms, e.g., a subset of the domain. Our approach given below can be extended to capture this.

We observe that a spoofing detection method actually realises three sequential steps: generating reference measurement, validating current measurements and assessing signal integrity. We address them one by one in the following.

Step 1: Generate reference measurements. Given an attribute, a spoofing detection first calculates a set of values that should contain its measurement when the received signal is integrous (called *reference set*). Different detection methods have various ways to calculate their reference sets.

We recognise two basic ways. One is to make use of a sufficiently large collection of integrous signals and calculate the set of all values that occur frequently. The other approach is to use the observation that the measurements of some attributes change over time in a fixed pattern. Based on a number of past signals the value of the current signal can thus be computed. Based on the distinction between these two approaches, we can divide spoofing detection methods into two categories – *stateless* and *stateful*. Let $\mathcal{R}_\alpha(s(\ell, t)) \subseteq dom(\alpha)$ be the calculated reference set of attribute α of signal $s(\ell, t)$. Stateless and stateful detection can be formally defined as follows:

Definition 3 (Stateless spoofing detection). Given a received signal $s(\ell, t)$, we say that a spoofing detection method on attribute $\alpha \in Attr(s(\ell, t))$ is *stateless* if $m_\alpha(s(\ell, t)) \in \mathcal{R}_\alpha(s(\ell, t))$ if $s(\ell, t)$ is integrous, where $\mathcal{R}_\alpha(s(\ell, t))$ is calculated by a function $f_\alpha : \mathcal{G} \rightarrow 2^{dom(\alpha)}$, i.e., $\mathcal{R}_\alpha(s(\ell, t)) = f_\alpha(s(\ell, t))$.

Definition 4 (Stateful spoofing detection). Given a received signal $s(\ell, t)$, we say that a spoofing detection method on attribute $\alpha \in Attr(s(\ell, t))$ is *stateful* if for a given a set of past signals $N = \{s(\ell_1, t_1), \dots, s(\ell_n, t_n)\} (\forall_{s(\ell_i, t_i) \in N} t_i < t)$, $m_\alpha(s(\ell, t)) \in \mathcal{R}_\alpha(s(\ell, t))$ if $s(\ell, t)$ is integrous and $s(\ell_i, t_i)$ is integrous for any $s(\ell_i, t_i) \in N$, where $\mathcal{R}_\alpha(s(\ell, t))$ is calculated by a n -ary function $f_\alpha : \mathcal{G}^n \rightarrow 2^{dom(\alpha)}$, i.e., $\mathcal{R}_\alpha(s(\ell, t)) = f_\alpha(s(\ell_1, t_1), \dots, s(\ell_n, t_n))$.

In a stateless spoofing detection method a reference set is computed based on the received signal whose integrity is under evaluation. The reference set in a stateful detection method relies on some past signals. The integrity of the past signals determines

the correctness of the reference set to be computed in a stateful detection method. In the definitions, we rely on the casual relation that a measurement falls in its reference set is caused by the fact that the signal is integrous. However, the related works in the literature usually take the opposite but incorrect direction, i.e., the integrity of a signal is concluded from the measurements of its attributes.

Step 2: Validate measurements. After calculating the reference set, the spoofing detection method checks whether the input measurement is in the reference set. If it is the case, we say that the measurement is *valid*. We use $\mathcal{V}_{s(\ell,t)}^\alpha$ to represent the proposition that “ $m_\alpha(s(\ell,t))$ is valid”.⁴

In practice, a reference set predicts a measurement considering an average intensity of natural environment interference on signal during transmission. This can lead to incorrect validity of measurement in the cases where the interference (abnormally) deviates from the average. This means that the measurement should be valid once the interference is normal. If we can learn how much the deviation of the current interference is from the average, then there will be a way to obtain the corresponding value to the average case. However, the impact of the interference cannot be measured. Therefore, it is undesirable to have a definite conclusion that a measurement is invalid once it is out of the reference set. Instead, since subjective logic opinions can allow us to capture the uncertainty caused by the environmental interference, we express the conclusion of a detection method on the validity of $m_\alpha(s(\ell,t))$ by an opinion. It is denoted by $w_{\mathcal{V}_{s(\ell,t)}^\alpha}$ and called the *validity opinion* of $s(\ell,t)$ on attribute α .

Step 3: Assess signal integrity. At last, a spoofing detection method assesses the integrity of received signals based on the validity of the measurements.

The output of a spoofing detection method is usually qualitative in the literature, which is not correct in reality. This is mainly because: 1) unpredicted environmental interference on signals leads to uncertainty of measurement validity; 2) there does not exist a definite causal relationship from measurement validity to signal integrity. For instance, some attackers can generate signals with valid measurements if they have access to powerful simulators. In such situations measurements are valid but signals are spoofed. False negative/positive ratios are thus defined to estimate the frequency of such situations and assess the performance of the detection in the literature.

In our approach, we use a subjective logic opinion to capture the uncertainty about the integrity of a signal. Given $s(\ell,t)$, we denote the opinion on its integrity by $w_{\mathcal{I}_{s(\ell,t)}}^\alpha$ and call it an *integrity opinion*.

Summary. Based on the above discussion, upon the receipt of the measurement of an attribute α , we can summarise the three steps that a spoofing detection method sequentially performs as follows:

1. Calculate the reference set $\mathcal{R}_\alpha(s(\ell,t))$;
2. Evaluate the validity of $m_\alpha(s(\ell,t))$ according to $\mathcal{R}_\alpha(s(\ell,t))$, i.e., $w_{\mathcal{V}_{s(\ell,t)}^\alpha}$;
3. Infer the opinion on the integrity of $s(\ell,t)$ based on $w_{\mathcal{V}_{s(\ell,t)}^\alpha}$, i.e., $w_{\mathcal{I}_{s(\ell,t)}}^\alpha$.

⁴ The notion of valid measurement is (implicitly) used by almost all existing spoofing detection methods. We formally define it in this paper.

In the literature, the calculation of reference sets in the first step has been extensively discussed. We proceed with how to obtain the validity of measurements in the second step (Sect. 4) and how to derive the integrity of signals in the third step (Sect. 5).

4 Deriving Validity Opinions

We give a method to calculate the validity opinion of an attribute given a received signal by taking into account the environmental interference by developing a function mapping $m_\alpha(s(\ell, t))$ and $\mathcal{R}_\alpha(s(\ell, t))$ to the opinion $w_{\mathcal{V}_{s(\ell, t)}^\alpha}$ for any signal $s(\ell, t)$.

Our main idea is to find an appropriate function degrading the belief on the validity of a measurement in terms of its distance to the reference set. The intuition behind this is that environmental interference with larger variation from the average is less common. The larger the variance is, the farther away that a measurement is from the reference set and thus the less probable that the measurement is valid. There are two necessary elements in the above observation, namely, the distance of a measurement to the reference set and the degradation function.

Distance of measurements to reference sets. Suppose that the distance between any two elements in $\text{dom}(\alpha)$, e.g., x and x' , is given as $\|x - x'\|$. The calculation and domains of the distances may vary between attributes. In this paper, we assume that the distances are normalised into real numbers. The distance of a measurement from a reference set is assigned zero if it is in the set. Otherwise, it is set as the minimum distance of the measurement to the values in the reference set.

Degradation function. The degradation function should be smooth and be compatible with the probability distribution of the environmental interference suffered by the given signal. Note that the choice of the distribution influences the accuracy of the validity opinion and should be carefully assessed with extensive analysis, e.g., using sufficiently large number of samples. We observe that the measured values of most attributes mentioned in the literature fit normal distributions best, e.g., signal strengths and clock offsets. Although some attributes may fit different distributions, in our framework we mainly take the normal distribution as an example to define the degradation function. The main idea can be adapted to other distributions.

5 Inferring Signal Integrity

We show how to derive the integrity opinion of a signal based on the measurement validity of one of its attributes, by studying the causal relationships between measurement validity and signal integrity. Since stateless and stateful methods have different causal relationships, they require different methods to derive integrity opinions. We describe the method for stateful spoofing detection.

5.1 Stateful Spoofing Detection

In a stateful spoofing detection method, e.g., on attribute α , a reference set is calculated based on a set of past signals. For the sake of simplicity, we assume that a stateful

detection method only makes use of one past signal. However, our method given below can be generalised to other cases.

For a signal $s(\ell, t)$, let $s(\ell', t')$ ($t' < t$) be the past signal based on which $\mathcal{R}_\alpha(s(\ell, t))$ is calculated. According to Def. 4, we can see that a reference set is computed in a specific way such that once past signals and the signal to be verified are both integrous, the corresponding measurement is valid. This gives rise to the following conditional relation for signal $s(\ell, t)$:

$$\mathcal{I}_{s(\ell', t')} \wedge \mathcal{I}_{s(\ell, t)} \rightarrow \mathcal{V}_{s(\ell, t)}^\alpha.$$

We cannot derive the integrity opinion $w_{\mathcal{I}_{s(\ell, t)}}^\alpha$ using the method given for stateless spoofing detection methods due to the involvement of the integrity of the past signals. In probability theory, if we can learn the joint probabilities $p(\mathcal{I}_{s(\ell', t')}, \mathcal{I}_{s(\ell, t)})$ and $p(\neg\mathcal{I}_{s(\ell', t')}, \mathcal{I}_{s(\ell, t)})$, then the probability $p(\mathcal{I}_{s(\ell, t)})$ can be calculated by summing them up. This calculation is called *marginalisation*. In subjective logic if we learn the beliefs on $\mathcal{I}_{s(\ell', t')} \wedge \mathcal{I}_{s(\ell, t)}$ and $\neg\mathcal{I}_{s(\ell', t')} \wedge \mathcal{I}_{s(\ell, t)}$, then the opinion on $\mathcal{I}_{s(\ell, t)}$ can be computed in a similar way. Let I be the following multinomial frame made of $\mathcal{I}_{s(\ell', t')}$ and $\mathcal{I}_{s(\ell, t)}$:

$$I = \{ \mathcal{I}_{s(\ell', t')} \wedge \mathcal{I}_{s(\ell, t)}, \neg\mathcal{I}_{s(\ell', t')} \wedge \mathcal{I}_{s(\ell, t)}, \mathcal{I}_{s(\ell', t')} \wedge \neg\mathcal{I}_{s(\ell, t)}, \neg\mathcal{I}_{s(\ell', t')} \wedge \neg\mathcal{I}_{s(\ell, t)} \}.$$

Let w_I be the multinomial opinion on I . Using the above causal relationship, we can calculate w_I based on the measurement validity through the abduction reasoning. As w_I contains the beliefs on $\mathcal{I}_{s(\ell', t')} \wedge \mathcal{I}_{s(\ell, t)}$ and $\neg\mathcal{I}_{s(\ell', t')} \wedge \mathcal{I}_{s(\ell, t)}$, we can compute the integrity opinion on $\mathcal{I}_{s(\ell, t)}$. Specifically, the calculation can be described in the following two steps:

1. Compute w_I based on $w_{\mathcal{V}_{s(\ell, t)}^\alpha}$. The computation is an abductive reasoning from $\mathcal{V}_{s(\ell, t)}^\alpha$. Let $w_{\mathcal{V}_{s(\ell, t)}^\alpha|I}$ be the set of *a priori* conditional opinions on $\mathcal{V}_{s(\ell, t)}^\alpha$ when each proposition in I is true, i.e., $\{w_{\mathcal{V}_{s(\ell, t)}^\alpha|x} \mid x \in I\}$. This calculation is as follows:

$$w_I = w_{\mathcal{V}_{s(\ell, t)}^\alpha} \overline{\odot} (w_{\mathcal{V}_{s(\ell, t)}^\alpha|I}, \mathbf{a}_I).$$

2. Compute $w_{\mathcal{I}_{s(\ell, t)}}^\alpha$ based on w_I . Suppose $w_I = (\mathbf{b}, u, \mathbf{a})$ and $w_{\mathcal{I}_{s(\ell, t)}}^\alpha = (b, d, u, a)$,

$$\begin{aligned} b &= \mathbf{b}(\mathcal{I}_{s(\ell', t')} \wedge \mathcal{I}_{s(\ell, t)}) + \mathbf{b}(\neg\mathcal{I}_{s(\ell', t')} \wedge \mathcal{I}_{s(\ell, t)}); \\ u &= u; \quad d = 1 - b - u; \\ a &= \mathbf{a}(\mathcal{I}_{s(\ell', t')} \wedge \mathcal{I}_{s(\ell, t)}) + \mathbf{a}(\neg\mathcal{I}_{s(\ell', t')} \wedge \mathcal{I}_{s(\ell, t)}). \end{aligned}$$

The base rate vector \mathbf{a}_{HI} expresses the *a priori* probability distribution on the four propositions in I . Note that $\mathcal{I}_{s(\ell', t')}$ and $\mathcal{I}_{s(\ell, t)}$ are independent as the signals $s(\ell, t)$ and $s(\ell', t')$ do not depend on each other and can be generated by two different sources. As $s(\ell', t')$ is a past signal, we assume that its integrity opinion has already been calculated, i.e., $w_{\mathcal{I}_{s(\ell', t')}}^\alpha$. The expectation probability of $\mathcal{I}_{s(\ell', t')}$, i.e., $E(w_{\mathcal{I}_{s(\ell', t')}}^\alpha)$, is thus the *a priori* probability of $\mathcal{I}_{s(\ell', t')}$ being true. Recall that we set $a(\mathcal{I}_{s(\ell, t)})$ to 0.5 to express the absence of any knowledge about $\mathcal{I}_{s(\ell, t)}$ being true. We can calculate \mathbf{a} as follows:

$$\begin{aligned} \mathbf{a}(\mathcal{I}_{s(\ell', t')} \wedge \mathcal{I}_{s(\ell, t)}) &= E(w_{\mathcal{I}_{s(\ell', t')}}^\alpha) \cdot 0.5; \\ \mathbf{a}(\mathcal{I}_{s(\ell', t')} \wedge \neg\mathcal{I}_{s(\ell, t)}) &= E(w_{\mathcal{I}_{s(\ell', t')}}^\alpha) \cdot 0.5; \\ \mathbf{a}(\neg\mathcal{I}_{s(\ell', t')} \wedge \mathcal{I}_{s(\ell, t)}) &= (1 - E(w_{\mathcal{I}_{s(\ell', t')}}^\alpha)) \cdot 0.5; \\ \mathbf{a}(\neg\mathcal{I}_{s(\ell', t')} \wedge \neg\mathcal{I}_{s(\ell, t)}) &= (1 - E(w_{\mathcal{I}_{s(\ell', t')}}^\alpha)) \cdot 0.5. \end{aligned}$$

Some *a priori* conditional opinions are applied during the inference of signal integrity. They should be assessed properly to guarantee the correctness of integrity opinions. We propose an approach to determine their values in the following section.

5.2 Determining the Conditional Opinions

We can divide the conditional opinions used in Sect. 5.1 into two classes according to whether spoofed signals are involved – *integrous signal based (isb)* and *spoofed signal based (ssb)*. Specifically, the opinions $w_{\mathcal{V}_{s(\ell,t)}^\alpha | \mathcal{I}_{s(\ell,t)}}$ and $w_{\mathcal{V}_{s(\ell,t)}^\alpha | \mathcal{I}_{s(\ell',t')} \wedge \mathcal{I}_{s(\ell,t)}}$ belong to the former class while the later class includes $w_{\mathcal{V}_{s(\ell,t)}^\alpha | \neg \mathcal{I}_{s(\ell,t)}}$, $w_{\mathcal{V}_{s(\ell,t)}^\alpha | \mathcal{I}_{s(\ell',t')} \wedge \neg \mathcal{I}_{s(\ell,t)}}$, $w_{\mathcal{V}_{s(\ell,t)}^\alpha | \neg \mathcal{I}_{s(\ell',t')} \wedge \mathcal{I}_{s(\ell,t)}}$ and $w_{\mathcal{V}_{s(\ell,t)}^\alpha | \neg \mathcal{I}_{s(\ell',t')} \wedge \neg \mathcal{I}_{s(\ell,t)}}$.

Determining isb conditional opinions. In practice, reference sets should be carefully chosen to ensure that the number of spoofed signals that have valid measurements should be small while most integrous signals have valid measurements. Reference sets do not contain all possible values that an integrous signal should have and there are situations where an integrous signal has an invalid measurement. The *isb* opinions express how likely these will not happen. Given the calculation of reference sets, we can estimate *isb* opinions by counting the frequency of valid measurements in a sufficiently large dataset of integrous signals.

Determining ssb conditional opinions. The *ssb* opinions are related to spoofing scenarios. They express the opinions on the validity of measurements when some related signals are spoofed. They also describe the power of attackers with regard to tuning attributes when false signals are generated. The method of deriving *isb* opinions is applicable if we have samples of spoofed signals. However, as far as we know there is no publicly available dataset of spoofed signals. Instead, we propose an alternative method estimating *ssb* opinions based on the efforts required for the attackers to generate signals with valid measurements. Intuitively, the more efforts that are required, the less likely that the measurements of spoofed signals are valid.

6 Combining Integrity Opinions

A received signal has a set of attributes that can be measured and explored by spoofing detection methods. Given a signal, a detection method will calculate its integrity opinion. However, the integrity opinions can be different from each other. This is mainly because (1) the conditional opinions used in spoofing detection methods are different; (2) unpredictable environmental interference can cause an integrous signal to have incorrect validity opinions for certain attributes; (3) attackers are able to tune some attributes of their generated signals so that the corresponding measurements remain valid.

We propose three algorithms to combine the integrity opinions according to different security requirements. The *Veto* algorithm gives a spoofing alarm as long as one of the detection methods gives an opinion indicating spoofing, while the *Consensus* algorithm combines integrity opinions to reduce uncertainty by making use of the opinion fusion operator \oplus [15]. The *Combined* algorithm combines their features to achieve a balance between false positives and false negatives.

7 Concluding Remark

We have implemented a prototype consisting of a measurement calculator, a series of spoofing detection methods and an integrity opinion combiner. The measurement calculator is connected to a receiver and used to read the basic measurements that can be calculated by the receiver. It also computes the measurements that cannot be offered by the receiver, e.g., Doppler ratio. The measurements are distributed to the detection methods which calculate the individual integrity opinions. To validate the effectiveness of our framework, we collect a large dataset of real GPS signals and the experimental results show that the framework is rather effective. More details can be found in [1].

References

1. Chen, X., Lenzini, G., Martins, M., Mauw, S., Pang, J.: A trust framework for evaluating gnss signal integrity. In: Proc. 26th IEEE Computer Security Foundations Symposium (CSF), IEEE Computer Society (2013)
2. Warner, J.S., Johnston, R.G.: A simple demonstration that the global positioning system (GPS) is vulnerable to spoofing. *Journal of Security Administration* **25**(19) (2002)
3. Humphreys, T.E., Ledvina, B.M., Psiaki, M.L., O'Hanlon, B.W., Kintner, P.M.: Assessing the spoofing threat: Development of a portable GPS civilian spoofer. In: Proc. 21st Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS), Institute of Navigation (2008) 2314–2325
4. Mixon, M.: Todd Humphreys' research team demonstrates first successful GPS spoofing of UAV. <http://www.ae.utexas.edu/news/archive/2012/> (2012)
5. Carroll, J.V.: Vulnerability assessment of the U.S. transportation infrastructure that relies on the global positioning system. *The Journal of Navigation* **56**(2) (2003) 185–193
6. Kuhn, M.G.: An asymmetric security mechanism for navigation signals. In: Proc. 6th Workshop on Information Hiding (IH). Volume 3200 of LNCS., Springer (2004) 239–252
7. Tippenhauer, N.O., Pöpper, C., Rasmussen, K.B., Capkun, S.: On the requirements for successful GPS spoofing attacks. In: Proc. 18th ACM Conference on Computer and Communications Security (CCS), ACM Press (2011) 75–86
8. Papadimitratos, P., Jovanovic, A.: GNSS-based positioning: Attacks and countermeasures. In: Proc. IEEE Military Communications Conference (MILCOM), IEEE CS (2008)
9. Nielsen, J., Broumandan, A., Lachapelle, G.: Spoofing detection and mitigation. *GPS World* (September 2010)
10. Psiaki, M., O'Hanlon, B., Bhatti, J., Shepard, D., Humphreys, T.: GPS spoofing detection via dual-receiver correlation of military signals. *IEEE Transactions on Aerospace and Electronic Systems* (2013) To appear.
11. Borre, K.: In: A Software-Defined GPS and Galileo Receiver. Applied and Numerical Harmonic Analysis (2007)
12. Jøsang, A.: Conditional reasoning with subjective logic. *Multiple-Valued Logic and Soft Computing* **15**(1) (2009) 5–38
13. Jøsang, A., Pope, S., Daniel, M.: Conditional deduction under uncertainty. In: Proc. 8th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU). Volume 3571 of LNCS., Springer (2005) 824–835
14. Nighswander, T., Ledvina, B.M., Diamond, J., Brumley, R., Brumley, D.: GPS software attacks. In: Proc. 19th ACM Conference on Computer and Communications Security (CCS), ACM Press (2012) 450–461
15. Jøsang, A.: Subjective logic (book draft). available at http://folk.uio.no/josang/papers/subjective_logic.pdf (2012)