A Trust Framework for Evaluating GNSS Signal Integrity

Xihui Chen*, Gabriele Lenzini*, Miguel Martins‡, Sjouke Mauw*†, Jun Pang†

*Interdisciplinary Centre for Security, Reliability and Trust, University of Luxembourg
†Faculty of Science, Technology and Communication, University of Luxembourg
‡itrust consulting, Luxembourg

***Abstract*—Through real-life experiments, it has been proved, not only in theory but also in practice, that civil signals of Global Navigation Satellite Systems (GNSS) can be spoofed. Consequently, a number of spoofing detection techniques have been proposed to verify the integrity of GNSS signals.**

In this paper, 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 validate our framework through experiments using both real and simulated signals and the results show that our framework is effective.

I. 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 signal strength. Therefore, civil signals can be easily interfered with or even taken over by false signals, which are called *jamming* and *spoofing*, respectively [1], [2].

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. [3] 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 [4] that there were no practical mitigation methods for spoofing attacks and we believe that it is still the case now, especially for GNSS civil signals. Navigation message authentication is considered as an effective method to prevent spoofing [5]. However, due to the long deployment cycle and high costs this is not a feasible approach in the near future [6]. 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. For instance, under normal circumstances, the strength of GPS signals is rarely above -153.5 dBW. If a received GPS signal has a higher strength, then a detection method claims that the integrity of the signal is not preserved.

Spoofing detection techniques. Some low-cost methods are proposed to detect unsophisticated spoofing [7], [8], [9], [10], [11]. For instance, Papadimitratos et al. [10] summarise three spoofing detection tests: location inertial test, clock offset test and Doppler shift test. Inertial sensors, such as speedometers and altimeters, can be used to predict future locations based on past ones, which are usually close to the real locations. The clock offset test measures the time offset of a receiver's local clock to the system time. As clocks usually drift with a fixed ratio, future clock offsets can be computed and the real offsets should be around them. Doppler shifts are also predictable if the relative velocities of a receiver to the satellites are available.

There are also some methods that make use of more advanced attributes of GNSS signals. For example, Nielsen et al. [12] monitor the correlation between the strengths of two signals from different satellites because the strengths always change independently. Psiaki et al. [13] 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. To the best of our knowledge, the existing detection methods in the literature all belong to this category.

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. [6] define spoofing from the viewpoint of localisation results, i.e., whether a receiver calculates the real location and time. However, this is not completely correct from the perspective of GNSS signals. 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 [6].
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. For example, in the inertial test [10] locations cannot be correctly predicted once the past ones are calculated based on spoofed signals. In such cases, the consistency of current calculated locations does not indicate the integrity of signals.
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. The inconsistency of these attributes does not always come with spoofed signals. 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. As we are not certain about the impacts of noise and the ability of the attackers on tuning signals' attributes, uncertainty in spoofing detection is inherently inevitable and should be quantified.
4) The outputs from different spoofing detection methods might conflict with each other and so far there exist no algorithms to combine the outputs of different methods into a coherent conclusion. Combining the results of multiple detection methods is necessary due to the fact that 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. The main reasons for us to use subjective logic are that it quantifies uncertainty in logic reasoning and provides a series of operators which correspond to logic operators and take uncertainty into account. Remark that our purpose of this paper is not to propose new methods to detect spoofing attacks. Instead, we aim to provide a generic understanding of spoofing detection and develop methods to derive correct conclusions on spoofing detection.

In our framework, we first formalise GNSS systems and receivers, based on which signal integrity is formally defined (Sect. III-C). 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 (Sect. III-E).

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 (Sect. IV). Next, we present a method based on conditional reasoning with subjective logic opinions to evaluate signal integrity for an individual detection method (Sect. V). 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 (Sect. VI). They are designed to capture different assumptions about the attackers' ability to manipulate attributes. In order to validate the effectiveness of our framework, we collect a large dataset of real GPS signals. In spite of the lack of real spoofing scenarios, we simulate the data of spoofed signals in a realistic way. The experimental results show that the framework is rather effective (Sect. VII).

II. Preliminaries

A. GNSS Signals

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 [14]. Navigation data carries information about the obits 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.[1] A

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

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.

A GPS receiver antenna captures signals from the satellites in range. From those signals the receiver calculates a three-dimensional coordinate as follows. A receiver runs replicas of the C/A codes synchronised with those of all the deployed satellites, based on which it separates the signals originated from different satellites and measures their time offsets with the replicas. These offsets are in fact the transition time of the signals. By multiplying with the speed of light, we can obtain the distances to the satellites, which can also be calculated as the Euclidean distance based on the locations of the satellites and the receiver. As navigation data includes the satellites' locations, we have only three variables to solve. Thus with three satellites, we can compute the three-dimensional location in theory. In practice due to the unknown offset between the clocks of the receiver and satellites, a fourth satellite is required.

B. GNSS Signal Spoofing

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 actually.

The above two ways of spoofing have been validated in the literature. Using the first approach, Humphreys et al. [2] implement a simulator which uses a GPS receiver to decode GPS signals and then broadcasts them with arbitrary delays. Tippenhauer et al. [6] theoretically prove that an attacker can spoof multiple receivers at the same time by carefully deploying broadcasting antennas in certain positions. These positions simulate the geometry of satellites. With respect to the second approach, Nighswander et al. [15] implement a simulator which re-broadcasts signals with arbitrary navigation messages. This method can attack multiple receivers more efficiently in larger areas compared with the simulator of Tippenhauer et al. [6] as satellites' geometry is ignored.

C. Subjective Logic

We give a brief introduction to *subjective logic* opinions and the operators on them used in the following discussion. For details we refer readers to its tutorial [16].

Subjective logic opinions. In subjective logic, 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 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. It can be interpreted as the perceived imprecision of the probability estimates. 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 $\{x_1, \ldots, x_n\}$. An opinion on X can be represented by $w_X = (\vec{b}_X, u_X, \vec{a}_X)$ where $\vec{b}_X : X \to [0,1]$ is the belief function, $u_X \in [0,1]$ is the uncertainty mass and $\vec{a} : X \to [0,1]$ is the base rate function. Furthermore,*

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

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

$$\vec{E}_X(x) = \vec{b}_X(x) + \vec{a}_X(x) \cdot u_X.$$

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 $b = \vec{b}_X(x)$, $d = \vec{b}_X(\bar{x})$, $u = u_X$ and $a = \vec{a}_X(x)$ 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 has been discussed in both binary logic and probability calculus. It 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 $x \to y$) 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, '$\circledcirc$' 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., $w_{y\|x} = w_x \circledcirc (w_{y|x}, w_{y|\bar{x}})$. 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, $\overline{\circledcirc}$ is the abductive operator calculating w_x based on $w_{y|x}$, $w_{y|\bar{x}}$ and a_x, i.e., $w_{x\|y} = w_y \overline{\circledcirc}(w_{y|x}, w_{y|\bar{x}}, a_x)$. We refer the readers to [17], [18] for the details of the implementation of the operators.

Conditional reasoning is applicable on multinomial opinions as well. Suppose two multinomial frames X and Y. Assume conditional opinions $w_{Y|X}$ and $w_{Y|\overline{X}}$ are available. Note that $w_{Y|X} = \{w_{Y|x} \mid x \in X\}$ and $w_{Y|\overline{X}} = \{w_{Y|\bar{x}} \mid x \in X\}$ where $w_{Y|x}$ (resp., $w_{Y|\bar{x}}$) represents the conditional opinion on Y given that x is TRUE (resp., FALSE). The opinion on Y based on observations on X (i.e., w_X) can be calculated by deductive reasoning, i.e., $w_{Y\|X} = w_X \circledcirc w_{Y|X}$. Likewise, the opinion on X based on observations on Y can be calculated by abductive reasoning, i.e., $w_{X\|Y} = w_Y \overline{\circledcirc}(w_{Y|X}, \vec{a}_X)$ where $\vec{a}_X$ is the *a priori* distribution on X.

III. A Trust Framework

In this section, we propose a trust framework to evaluate signal integrity.

A. 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. The formats of locations and time points are out of our discussion since different formats can be converted from one to another. For instance, the coordinate N25°07.450' is represented in degrees and minutes while it can also be of the form of only degrees, i.e., 25.124167. 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 $sig : \mathcal{S} \times \mathcal{T} \to \Theta$ 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.

B. 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 in it based on their unique features, e.g., C/A codes. This separation process can be modelled by function $sigCom : \mathcal{G} \to 2^{\Theta}$ 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 $ori : \Theta \to \mathcal{S}$ to return the originator of any signals. Note that by the originator of a signal we only mean that the originator's spreading code is modulated in the signal, implying that, whenever it is received, the receiver would think it is from the satellite. The originator is not always the entity that actually generates the signal as the attackers can also 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 $loc : \mathcal{G} \to \mathcal{L} \times \mathcal{T}$. 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. For the same reason, we also omit the implementation difference between receivers. The notations mentioned are summarised in Tab. I.

C. 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

Table I
THE NOTATIONS AND FUNCTIONS.

$\mathcal{S}$	set of running satellites of the GNSS system;
$\mathcal{T}$	set of time points;
$\xi(S,t)$	position of satellite S at time t;
Θ	set of GNSS signals;
$sig(S,t)$	GNSS signal transmitted by satellite S at time t;
$\eta(S,\ell,t)$	attenuation of the signal leaving S at t to reach ℓ;
$\mathcal{G}$	set of combined GNSS signals that can be captured ;
$sigCom(s)$	set of GNSS signals combined in $s \in \mathcal{G}$;
$ori(sig)$	satellite whose C/A code is modulated in sig;
$loc(s)$	location and time calculated using received signal s.

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.

D. 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 [15] and GNSS signal spoofing [6].

Software attacks on receivers target at the localisation algorithms implemented on receivers. Infected by malware, the receiver can be forced to calculate incorrect coordinates. GNSS signal spoofing is to feed a receiver with simulated signals such that even the correct localisation algorithm cannot compute the right location.

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. [6]. 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.

With regard to signal generation, we assume that the attackers can generate any GNSS signal in Θ that can be interpreted by receivers. However, the attackers cannot generate the military signals due to the encrypted P(Y), but it can intercept and replay them.

E. 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. For instance, the values of attributes, e.g., signal strength and Doppler shift, are calculated by receivers while others, e.g., power correlation of signals from two satellites, are not provided directly by receivers. 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 easily extended to capture these cases.

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 $f_\alpha : \mathcal{G} \to 2^{dom(\alpha)}$, 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 $N = \{s(\ell_1,t_1),\ldots,s(\ell_n,t_n)\}$ ($\forall_{s(\ell_i,t_i)\in N}\ t_i < t$), $m_\alpha(s(\ell,t)) \in \mathcal{R}_\alpha(s(\ell,t))$ 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 $f_\alpha : \mathcal{G}^n \to 2^{dom(\alpha)}$, i.e., $\mathcal{R}_\alpha(s(\ell,t)) = f_\alpha(s(\ell_1,t_1),\ldots,s(\ell_n,t_n))$.*

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 $\mathcal{V}^\alpha_{s(\ell,t)}$ to represent the proposition that "$m_\alpha(s(\ell,t))$ *is valid*".[2]

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 $w_{\mathcal{V}^\alpha_{s(\ell,t)}}$ 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 $w^\alpha_{\mathcal{I}_{s(\ell,t)}}$ 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., $w_{\mathcal{V}^\alpha_{s(\ell,t)}}$;
3) Infer the opinion on the integrity of $s(\ell,t)$ based on $w_{\mathcal{V}^\alpha_{s(\ell,t)}}$, i.e., $w^\alpha_{\mathcal{I}_{s(\ell,t)}}$.

In the literature, the calculation of reference sets in the first step has been extensively discussed. In this paper, we take it as given. We proceed with how to obtain the validity of measurements in the second step (Sect. IV) and how to derive the integrity of signals in the third step (Sect. V).

IV. Deriving Validity Opinions

In this section, we give a method to calculate the validity opinion of an attribute given a received signal by taking into account the environmental interference. Essentially, we

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

develop a function mapping $m_\alpha(s(\ell,t))$ and $\mathcal{R}_\alpha(s(\ell,t))$ to the opinion $w_{\mathcal{V}^\alpha_{s(\ell,t)}}$ 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 $dom(\alpha)$, 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, i.e., $\| x - x' \| \in \mathbb{R}$. 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. Let $d_\alpha(s(\ell,t))$ be the distance between $m_\alpha(s(\ell,t))$ and $\mathcal{R}_\alpha(s(\ell,t))$. Then it can be defined as follows:

$$d_\alpha(s(\ell,t)) = \begin{cases} 0 & m \in R \\ \min\limits_{v \in R} \| m - v \| & m \notin R \end{cases}$$

where $m = m_\alpha(s(\ell,t))$ and $R = \mathcal{R}_\alpha(s(\ell,t))$.

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 the following we take the normal distribution as an example to define the degradation function. The main idea can be adapted to other distributions. Assume $w_{\mathcal{V}^\alpha_{s(\ell,t)}} = (b, d, u, 0.5)$. The base rate is set to 0.5 so as to express that we have no preference. The other three parameters can be computed as follows:

$$b = e^{-\frac{d_\alpha(s(\ell,t))^2}{2 \cdot var^2}}; \quad d = 1 - b; \quad u = 0$$

where var represents the variance required by the original normal distribution and it determines how fast b drops along with $d_\alpha(s(\ell,t))$. The uncertainty u can be interpreted as the confidence in the existence of the normal distribution. As we have already assumed its existence, we assign 0 to uncertainty u.

We can determine the value of var if a distance and the corresponding belief are given. In our method, we take the maximum distance allowed for a measurement and assign the minimum belief to it. Let d_{max} be the maximum allowed distance to the reference set and b_{min} be the corresponding minimum belief. We can calculate var as follows:

$$var = \frac{d_{max}}{\sqrt{-2 \cdot \ln b_{min}}}.$$

V. Inferring Signal Integrity

In this section, we show how to derive the integrity opinion of a signal based on the measurement validity of one of its attributes. We study the causal relationships between measurement validity and signal integrity, based on which conditional reasoning can be used. Since stateless and stateful methods have different causal relationships, they require different methods to derive integrity opinions.

A. Stateless Spoofing Detection

In a stateless spoofing detection method, e.g., on attribute α, a reference set is calculated in such a way that as long as a signal is integrous, its measurements must be valid (see Def. 3). Therefore, given a signal $s(\ell,t)$, the following conditional relationship holds:

$$\mathcal{I}_{s(\ell,t)} \rightarrow \mathcal{V}^\alpha_{s(\ell,t)}.$$

The validity opinion $w_{\mathcal{V}^\alpha_{s(\ell,t)}}$ has already been calculated based on the methodology given in Sect. IV. Thus the integrity opinion of $s(\ell,t)$ can be considered as the abduced opinion on the validity of the measurement.

In the abduction, we need two *a priori* conditional opinions on the measurement validity when the signal is integrous or spoofed and the *a priori* probability that the signal is integrous before its reception. Let $w_{\mathcal{V}^\alpha_{s(\ell,t)}|\mathcal{I}_{s(\ell,t)}}$ and $w_{\mathcal{V}^\alpha_{s(\ell,t)}|\neg\mathcal{I}_{s(\ell,t)}}$ be the opinions on the validity of the measurement when the signal is integrous or spoofed, respectively. We set the base rate $a(\mathcal{I}_{s(\ell,t)})$ to 0.5 to indicate no *a priori* knowledge about the integrity of the signal. It is a conservative choice as we want to eliminate the interference of artificial preference as much as possible. Using the abduction operator in subjective logic (i.e., $\overline{\circledcirc}$), we can calculate the opinion on the truth of $\mathcal{I}_{s(\ell,t)}$ as follows:

$$w^\alpha_{\mathcal{I}_{s(\ell,t)}} = w_{\mathcal{V}^\alpha_{s(\ell,t)}} \overline{\circledcirc} (w_{\mathcal{V}^\alpha_{s(\ell,t)}|\mathcal{I}_{s(\ell,t)}}, w_{\mathcal{V}^\alpha_{s(\ell,t)}|\neg\mathcal{I}_{s(\ell,t)}}, a(\mathcal{I}_{s(\ell,t)})).$$

B. 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 $s(\ell,t)$:

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

We cannot derive the integrity opinion $w^{\alpha}_{\mathcal{I}_{s(\ell,t)}}$ 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 $\mathcal{I}_{s(\ell,t)}$:

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

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

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

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

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

The base rate vector $\vec{a}_{HI}$ 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., $w_{\mathcal{I}_{s(\ell',t')}}$. The expectation probability of $\mathcal{I}_{s(\ell',t')}$, i.e., $E(w_{\mathcal{I}_{s(\ell',t')}})$, 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 $\vec{a}$ as follows:

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

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

C. Determining the Conditional Opinions

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

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.

We take $w_{\mathcal{V}^{\alpha}_{s(\ell,t)}|\mathcal{I}_{s(\ell,t)}}$ as an example to illustrate the calculation which can be extended straightforwardly to the opinions used in stateful spoofing detection. Let SC be the collection of integrous signals and $P \subseteq SC$ be the set of samples whose measurements of α are valid. Let $w_{\mathcal{V}^{\alpha}_{s(\ell,t)}|\mathcal{I}_{s(\ell,t)}}$ be (b, d, u, a). The base rate a expresses the *a priori* probability about the truth of $\mathcal{V}^{\alpha}_{s(\ell,t)}$ when the received signal is integrous. We set it to 0.5 when we have no knowledge about $\mathcal{V}^{\alpha}_{s(\ell,t)}$. Then the belief, disbelief and uncertainty can be computed by

$$b = \frac{|P|}{|SC|+2}, \quad d = \frac{|SC/P|}{|SC|+2}, \quad u = \frac{2}{|SC|+2}.$$

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 more powerful an attacker is, the more likely that the measurements of their spoofed signals remain valid.

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.

There are many restrictions for the attackers to overcome in order to preserve the validity of a measurement, e.g., signal simulators, deployment environment and the availability of equipment. A spoofing attack demanding a simulator of 10,000 euros is harder than the ones which need simulators of 1,000 euros. The difficulty to meet a requirement can be divided into levels. For instance, the prices of equipment can be assigned to levels from *low* to *high*. Meanwhile, the importance of requirements also varies.

Let $Req = \{rq_1, \ldots, rq_k\}$ be the set of requirements and $W = \{w_1, \ldots, w_k\}$ be the set of corresponding importance where $\sum_{1 \leq j \leq k} w_j = 1$. For $rq_i \in Req$, we assign one of the five scores $\{0.2, 0.4, 0.6, 0.8, 1\}$, i.e., $score(rq_i)$. Sometimes, we do not have expertise for every requirement. When we have no idea about the requirement, we set $score(rq_i)$ to 0. The sum of weighted assigned scores can be interpreted as the votes against a successful spoofing attack while the unassigned scores can be seen as the neutral votes. Take $w_{\mathcal{V}^{\alpha}_{s(\ell,t)}|\neg\mathcal{I}_{s(\ell,t)}}$ for example. Let it be (b, d, u, a), then

$$b = \sum_{score(rq_i) \neq 0} score(rq_i) \cdot w_i; \quad d = 1 - b - u;$$
$$u = \sum_{score(rq_i) = 0} w_i.$$

We set a as 0.5 to indicate the absence of any preference.

VI. Combining Integrity Opinions

A received signal has a set of attributes that can be measured and explored by spoofing detection methods. According to Sect. V, given a signal a detection method will calculate its integrity opinion. However, the integrity opinions can be different from each other. This is because:

- The conditional opinions used in spoofing detection methods are different. This leads to different integrity opinions even if the validity opinions are the same.
- Unpredictable environmental interference can cause an integrous signal to have incorrect validity opinions for certain attributes. This subsequently causes incorrect integrity opinions.
- Some attackers are able to tune some attributes of their generated signals so that the corresponding measurements remain valid. This fools the spoofing detection methods to output incorrect integrity opinions.

Thus, a combined integrity opinion is needed to deal with the difference. Furthermore, with more evidences taken into account, the combined opinions will be more reliable. The combination is very useful for location-based applications as they can customise their services based on signal integrity and take proper actions whenever spoofing is detected.

In this section, we propose three algorithms to combine the integrity opinions according to different user security requirements. A combination algorithm can be seen as a function taking a set of individual integrity opinions as input denoted by $\mathcal{W}_{\mathcal{I}_{s(\ell,t)}}$, and outputting a combined integrity opinion denoted by $w_{\mathcal{I}_{s(\ell,t)}}$. Before presenting the algorithms, we start with how to construct the set of integrity opinions, i.e., $\mathcal{W}_{\mathcal{I}_{s(\ell,t)}}$.

Recall that stateful spoofing detection methods make use of the integrity opinions on past signals. Assume that integrity opinions can be combined, we have two types of integrity opinions – combined opinions, e.g., $w_{\mathcal{I}_{s(\ell,t)}}$ and those given by individual stateless methods, e.g., $w^{\alpha}_{\mathcal{I}_{s(\ell,t)}}$. As a consequence, a stateful detection method can output two kinds of integrity opinions – *global* and *local*. A global integrity opinion is calculated using combined opinions on past signals, while a local integrity opinion is based on opinions given by a single stateless method. Given a signal, we thus have two sets of integrity opinions to combine – *global opinion set* and *local opinion set*, denoted by $\mathcal{W}^{glo}_{\mathcal{I}_{s(\ell,t)}}$ and $\mathcal{W}^{loc}_{\mathcal{I}_{s(\ell,t)}}$, respectively. In this section, we use $\mathcal{W}_{\mathcal{I}_{s(\ell,t)}}$ to have a generic description for our algorithms. It can be substituted by either $\mathcal{W}^{glo}_{\mathcal{I}_{s(\ell,t)}}$ or $\mathcal{W}^{loc}_{\mathcal{I}_{s(\ell,t)}}$ in implementation.

A. The Veto Algorithm

In safety-critical applications, failing to detect a spoofing attack can lead to severe consequence. In such situations, false alarms of spoofing are affordable but false claims of integrity are not. To meet this requirement, our idea is to give a spoofing alarm as long as one of the deployed spoofing detection methods gives an opinion indicating spoofing. We choose the integrity opinion with the minimum belief in the integrity of the signal as the combined opinion.

We introduce a relation to compare the belief in the integrity of a given signal expressed by two integrity opinions, i.e., $\preceq \subseteq \Omega \times \Omega$ where Ω is the set of all binomial opinions. An integrity opinion has less belief in the integrity of a signal than another if its expectation probability is smaller or it has a larger uncertainty when their expectation probabilities are equivalent. The relation $\preceq$ is formally defined as follows:

Definition 5 ($\preceq$)**.** *Given two binomial subjective opinions $w = (b, d, u, a)$ and $w' = (b', d', u', a')$, we say that w is not larger than w' (denoted by $w \preceq w'$) if*

$$E(w) < E(w') \vee (E(w) = E(w') \wedge u \geq u') \vee w = w'.$$

Recall that $\mathcal{W}_{\mathcal{I}_{s(\ell,t)}}$ is the set of integrity opinions output by spoofing detection methods. The calculation of the combined integrity opinion $w_{\mathcal{I}_{s(\ell,t)}}$ is straightforward. Let

$Veto : 2^{\Omega} \rightarrow \Omega$ be the Veto function, then we have

$$w_{\mathcal{I}_{s(\ell,t)}} = Veto(\mathcal{W}_{\mathcal{I}_{s(\ell,t)}}) \text{ s.t.}$$
$$(w_{\mathcal{I}_{s(\ell,t)}} \in \mathcal{W}_{\mathcal{I}_{s(\ell,t)}}) \wedge (\forall w \in \mathcal{W}_{\mathcal{I}_{s(\ell,t)}}, w_{\mathcal{I}_{s(\ell,t)}} \preceq w).$$

Note that past signals are mandatory for stateful spoofing detection methods to derive integrity opinions. When there are no sufficient opinions available for a stateful spoofing detection method, we set its integrity opinion in $\mathcal{W}_{\mathcal{I}_{s(\ell,t)}}$ to $(1, 0, 0, 0.5)$ to eliminate its impact on the combined opinion.

B. The Consensus Algorithm

Recall that in a subjective logic opinion, the uncertainty mass can be interpreted as a confidence measurement on the correctness of the probability expectation. Given an integrity opinion, the smaller the uncertainty is, the more likely that its expectation probability of signal integrity is correct. Based on this understanding, the integrity opinions with less uncertainty should play a more important role in the combined opinion.

Intuitively, more evidences should lead to more reliable conclusions. This means that when more integrity opinions are combined, we should have more confidence in the correctness of the combined opinion. In other words, the combined opinion should have less uncertainty mass.

We make use of the opinion fusion operator $\oplus$ [16], which is also called the consensus operator, to combine integrity opinions. The definition of the operator can be found in Appendix A. Consensus is also called *cumulative* fusion and applicable when the evidences giving rise to the opinions are independent. Since the measurements of an attribute do not affect another attribute, we can assume that they are independent. Moreover, the fused opinion simply meets our expectation for the combined opinion, which can be derived straightforwardly from its definition. First, in the fused opinion, a larger proportion of the belief mass comes from the opinion with less uncertainty. Second, more opinions will lead to less uncertainty mass in the fused opinion. Let $Consensus : 2^{\Omega} \rightarrow \Omega$ be the corresponding function of the Consensus algorithm. Then we have

$$w_{\mathcal{I}_{s(\ell,t)}} = Consensus(\mathcal{W}_{\mathcal{I}_{s(\ell,t)}}) = \oplus_{w \in \mathcal{W}_{\mathcal{I}_{s(\ell,t)}}} w.$$

When there are no sufficient past integrity opinions for certain stateless spoofing detection methods, their integrity opinions are set to the vacuous opinion with uncertainty being 1. It is the neutral element of the consensus opinion, so it will have no impacts on the combined opinion.

C. The Combined Algorithm

From their descriptions, it is clear that (1) the Veto algorithm is conservative in the sense that it can lead to more false alarms of spoofing; (2) while the Consensus algorithm can better reduce uncertainty it can lead to more false claims of integrity due to its use of the opinion fusion operator. To achieve a balance of the two situations, we combine the features of the two algorithms and develop a new algorithm. Different from the Veto algorithm, we do not always choose the integrity opinion with the smallest expectation probability to conclude a spoofed signal. Instead, we consider the opinions not only with sufficiently small expectation probabilities and but also with sufficiently small uncertainty. We call such integrity opinions *VETO* opinions.

Definition 6 (((σ, θ)-VETO opinions)**.** *Let $w = (b, d, u, a)$ be an integrity opinion and $\sigma \in [0, 1)$ and $\theta \in [0, 1)$ be the thresholds of the expectation probability and the uncertainty, respectively. It is said to be a VETO opinion if*

$$E(w) \leq \sigma \wedge u \leq \theta.$$

For each individual detection method, σ and θ can be set to different values. Let σ_α and θ_α be the predefined thresholds for the spoofing detection method on attribute α. When combining the opinions from a number of detection methods, if there exist multiple VETO opinions then their consensus is calculated and output as the combined opinion. Otherwise, if there is no VETO opinion, the Consensus algorithm is called to calculate the combined opinion. This new algorithm is called Combined as shown in Alg. 1.

Algorithm 1 The Combined Algorithm

```
1:  Input: W_{I_{s(ℓ,t)}}
2:  Output: w_{I_{s(ℓ,t)}}
3:  Init: w_{I_{s(ℓ,t)}} ← (0, 0, 1, 0.5);
4:  for w^α_{I_{s(ℓ,t)}} ∈ W_{I_{s(ℓ,t)}} do
5:      if isVETO(w^α_{I_{s(ℓ,t)}}, σ_α, θ_α) then
6:          w_{I_{s(ℓ,t)}} ← w_{I_{s(ℓ,t)}} ⊕ w^α_{I_{s(ℓ,t)}};
7:      end if
8:  end for
9:  if w_{I_{s(ℓ,t)}} = (0, 0, 1, 0.5) then
10:     w_{I_{s(ℓ,t)}} ← Consensus(W_{I_{s(ℓ,t)}});
11: end if
```

The combined integrity opinion is initially set to the vacuous opinion. The function isVETO(w, σ, θ) returns true if w is a (σ, θ)-VETO opinion and false otherwise. We start with looking for VETO opinions in $\mathcal{W}_{\mathcal{I}_{s(\ell,t)}}$ and compute the consensus of them if there exist any (line 4-8). If there are no VETO opinions, $w_{\mathcal{I}_{s(\ell,t)}}$ will remain unchanged (line 9) as the uncertainty of a VETO opinion is always smaller than 1 (see Def. 6). Then we compute the consensus of all integrity opinions (line 10).

VII. Validation

A. Implementation

We implement a prototype for our framework, which consists of three major components – the measurement calculator, a series of spoofing detection methods and an integrity opinion combiner. The measurement calculator is

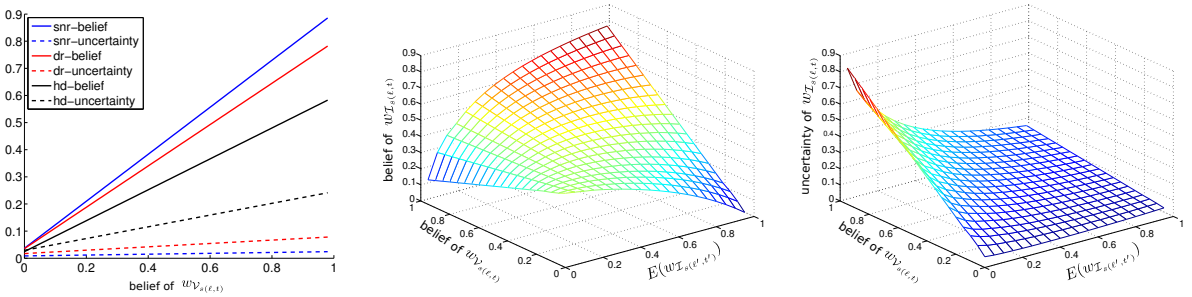

(a) The stateless methods. (b) The stateful method (belief). (c) The stateful method (uncertainty).

Figure 1. The integrity opinions.

Table II
THE PARAMETERS USED IN STATELESS DETECTION.

methods	reference set	d_{max}	ssb opinion
snr	$[0, 62.5]$	2.7	(0.1, 0.8, 0.1, 0.5)
dr	$[1.2829, 1.2837]$	0.004	(0.2, 0.7, 0.1, 0.5)
hd	$[0, 3.5]$	10	(0.4, 0.5, 0.1, 0.5)

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 spoofing detection methods which calculate the individual integrity opinions.

In our prototype, we implement four spoofing detection methods which explore the following attributes, respectively:

- Doppler ratio (dr) between the Doppler shifts of the civil signal and the military signal in a received signal.
- Signal-to-noise ratio (snr) between the power of the signal and the noise in the given RF bandwidth, which is expressed in decibels (dB).
- Height difference (hd) between the height in a calculated coordinate and the real height corresponding to the latitude and longitude in the coordinate, which is expressed in metres.
- Clock offset (cf) – the time difference between the local clock of a receiver and the universal time, which is measured in seconds.

The first three spoofing detection methods are stateless while the last one is stateful as it predicts the clock offset based on one past offset and the drift speed of the local clock. In detail, suppose that the clock offset at t' is cf and the drift speed of the clock is v_{drift}. Then the predicted clock offset at time t is $preCF(t) = off + v_{drift} \cdot (t - t')$.

To learn the reference sets and related parameters, we use a dataset of 160,000 samples of integrous signals which are collected with a professional receiver JAVAD ALPHA2. Each record of the dataset stores the measurements of a signal. We choose a reference set that allows 98% of the samples to have valid measurements. Recall d_{max} is the maximal allowed distance of a measurement from the reference set. It is assigned to a value so that only 5% samples have larger distance. The corresponding minimum belief, i.e., b_{min}, is uniformly set to 0.05. Tab. II lists the parameters used in each stateless detection method.

The reference set of the clock offset at time t is composed of the values between $preCF(t) - 1 \times 10^{-8}$ and $preCF(t) + 1 \times 10^{-8}$. The maximum distance is set as $3 \times 10^{-8} s$ so as to ensure 5% signals with larger distance. With respect to the isb conditional opinions, as about 98% samples have valid measurements, they are set to $(0.98, 0.02, 0, 0.5)$. For the ssb conditional opinions, we assign them a preliminary opinion based on our knowledge. In our implementation, they are set to an identical opinion $(0.1, 0.8, 0.1, 0.5)$.

B. The Experimental Setup

To validate our framework, we prepare three datasets of signal measurements. The first one is called *integrity dataset* storing the measurements of 25,531 integrous signals. These samples are collected using the same GPS receiver as but independently from the dataset used for parameter evaluation. The second is a *spoofed dataset* and synthesised based on the integrity dataset to simulate spoofed signals. This is because no spoofed signals are publicly available. The third dataset is a *mixed dataset* with both integrous signals and spoofed signals.

The spoofed and the mixed datasets contain synthesised records for spoofed signals. The main idea to synthesise such records is to make use of the fact that the attributes of spoofed records have values deviating from those of integrous signals. Furthermore, the amount of the deviation is determined by the attackers in terms of their capabilities to tune spoofed signals. A more powerful attacker will generate signals with less deviation. We take a simple assumption that the attackers' capabilities follow the normal distribution during the construction of the spoofed dataset. To compute a record of a spoofed signal, given an item in the integrity dataset and an attribute, we first decide whether to change its value based on the corresponding *a priori* ssb conditional opinion. If yes, an extra distance is calculated following a normal distribution with d_{max} as the mean and the same variance used in the validity calculation. This extra distance

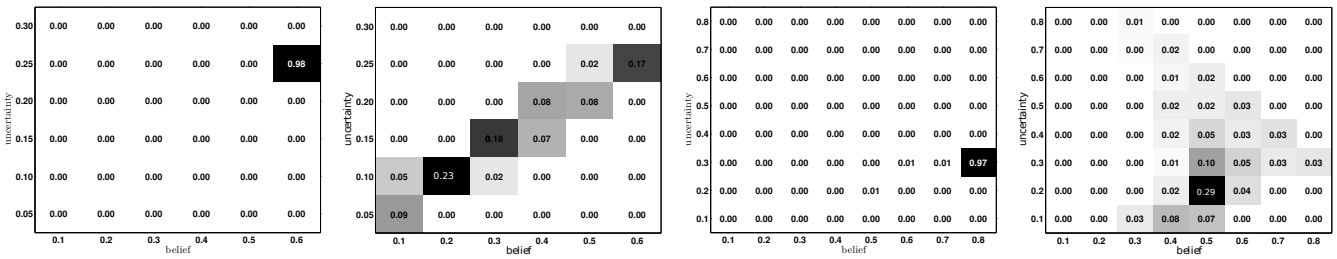

(a) hd (integrity dataset). (b) hd (spoofed dataset). (c) cf+hd (integrity dataset). (d) cf+hd (spoofed dataset).

Figure 2. Integrity Opinions of individual detection method.

is added to the distance of the original measurement and the resulted distance is used to calculate the validity opinion.

C. Experimental Results

Bounds of integrity opinions. Fig. 1 shows the change of integrity opinions with validity opinions. Fig. 1(a) shows how the belief and uncertainty of an integrity opinion evolve with the belief of a validity opinion in the stateless methods. A general observation is that belief and uncertainty both increase linearly as the beliefs of the validity opinions grow. However, different methods have different output opinions, which are determined by their *a priori* conditional opinions. In our setting, the method snr calculates an integrity opinion with the largest belief and the smallest uncertainty than the other two spoofing detection methods.

Concerning the cf stateful method, since it requires past signals, its calculated integrity opinions should change along with two parameters – the expectation probability of the past signal's integrity and the beliefs of validity opinions. Fig. 1(b) and Fig. 1(c) show the beliefs and uncertainty of the integrity opinions when the two parameters have various values. The maximum belief value occurs when they are both 1.0 while the minimum belief is obtained when the past signal is spoofed and the current signal has a valid measurement. The maximum uncertainty is computed when the past signal is spoofed and the current measurement is valid. We use Tab. III to summarise the bounds of belief and uncertainty of integrity opinions for each method.

Table III
BELIEF & UNCERTAINTY BOUNDS OF INTEGRITY OPINIONS.

methods	min(b)	max(b)	min(u)	max(u)
snr	0.03	0.86	0.01	0.02
dr	0.09	0.78	0.02	0.07
hd	0.02	0.58	0.03	0.24
cf	0.01	0.80	0.03	0.95

Integrity opinions of spoofed and integrous signals. We study what integrity opinions spoofing detection methods calculate when signals are spoofed and integrous. To achieve this, we make use the spoofed and integrity datasets.

We divide integrity opinions into classes according to their beliefs and uncertainty. Each cell in the diagrams in Fig. 2 corresponds to a class of opinions whose beliefs and uncertainty are bounded in certain intervals. The number labelled in each cell is the proportion of calculated integrity opinions which fall in the corresponding class. The grey level of a cell also indicates the proportion. The darker it is, the larger the proportion is. In Fig. 2 we choose hd and cf as examples to show the distribution of integrity opinions when all signals are spoofed or integrous. Note that the cf method uses the integrity opinions of past signals given by the hd detection. We have two major observations. First, the integrity opinions of spoofed signals have much smaller beliefs and uncertainty compared to those of integrous signals. In Fig. 2(a), we can see that 98% of the opinions given by hd on integrous signals have beliefs larger than 0.5 and uncertainty less than 0.3. However, when signals are spoofed, the beliefs of about 50% integrity opinions drops below 0.2 and the uncertainty becomes smaller than 0.15 (see Fig. 2(b)). The opinions computed by the cf detection follow a similar pattern. Second, different methods give different opinions even for the same signals. The opinions on both datasets given by the two methods rarely overlap.

Integrity opinion combination. We use the mixed dataset to validate the performance of the combination algorithms. Intuitively, a combined algorithm is effective if it can calculate large beliefs for integrous signals and small beliefs for spoofed signals. In the mixed dataset, we have 4,748 spoofed signals out of total 25,531 samples (about 18.6%). Fig. 3 shows the results of our three algorithms. They all successfully distinguish spoofed signals from integrous ones (with certain errors). The Veto algorithm assigns smaller beliefs and larger uncertainty to both spoofed and integrous signals, as it is rather conservative when compared with the other two methods. The Consensus algorithm assigns 77% of the signals with beliefs larger than 0.9, and assigns 14% of the signals with beliefs less than 0.2 meaning that about 4.6% of the spoofed signals are not detected. The Consensus algorithm gives uncertainty less than 0.05 to almost all the signals. It is interesting to see that the Combined algorithm gives more balanced results. When signals are integrous, a belief of 0.9 is mostly assigned which is the same as the Consensus algorithm. Meantime, for

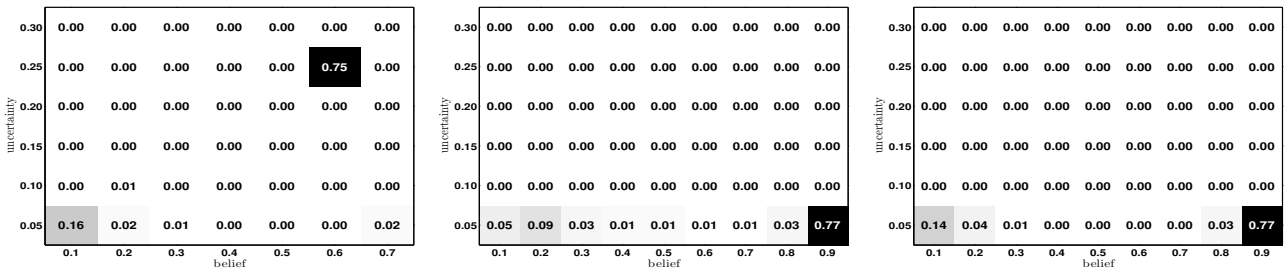

(a) Veto. (b) Consensus. (c) Combined.

Figure 3. Combined opinions of integrous signals.

spoofed signals, it assigns a belief of 0.1 to most of them, which is comparable to the Veto algorithm and much better than the Consensus algorithm. The observations follow the design principles of the algorithms. In practice, the choice of a combination algorithm depends on applications.

VIII. Conclusion and Future Work

In this paper, we proposed a trust framework to evaluate the integrity of GNSS signals. We identified a few problems with existing spoofing detection methods in the literature and addressed them within our framework. First, we clarified the concept of signal integrity and gave a formal definition, which is the first attempt to the best of our knowledge. Second, we precisely characterised spoofing detection methods and extracted the causal relation between measurement validity and signal integrity. We then proposed an approach to derive signal integrity while capturing its uncertainty in a natural way. Last but not least, we presented three ways to combine opinions from various detection methods and validated our work through experiments.

So far, the framework is only validated through simulated spoofed signals. It is interesting to evaluate it in real spoofing scenarios. Moreover, our prototype can be improved and we plan to incorporate more detection methods.

Acknowledgements

Xihui Chen is supported by the National Research Fund, Luxembourg under the project SECLOC 794361. This work was partially supported by the European Space Agency (ESA) under the project "Developing a prototype of Localisation Assurance Service Provider (LASP)", with the contract number 4000102584-10-NL-HE.

We especially thank Carlo Harpes and itrust consulting, Luxembourg for initiating the LASP project and for their valuable comments. We also thank the anonymous reviewers for their comments that help improve our manuscript.

References

[1] J. S. Warner and R. G. Johnston, "A simple demonstration that the global positioning system (GPS) is vulnerable to spoofing," *Journal of Security Administration*, vol. 25, no. 19, 2002.

[2] T. E. Humphreys, B. M. Ledvina, M. L. Psiaki, B. W. O'Hanlon, and P. M. Kintner, "Assessing the spoofing threat: Development of a portable GPS civilian spoofer," in *Proc. 21st Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS)*. Institute of Navigation, 2008, pp. 2314–2325.

[3] M. Mixon, "Todd Humphreys' research team demonstrates first successful GPS spoofing of UAV," http://www.ae.utexas.edu/news/archive/2012/, 2012.

[4] J. V. Carroll, "Vulnerability assessment of the U.S. transportation infrastructure that relies on the global positioning system," *The Journal of Navigation*, vol. 56, no. 2, pp. 185–193, 2003.

[5] M. G. Kuhn, "An asymmetric security mechanism for navigation signals," in *Proc. 6th Workshop on Information Hiding (IH)*, ser. LNCS, vol. 3200. Springer, 2004, pp. 239–252.

[6] N. O. Tippenhauer, C. Pöpper, K. B. Rasmussen, and S. Capkun, "On the requirements for successful GPS spoofing attacks," in *Proc. 18th ACM Conference on Computer and Communications Security (CCS)*. ACM Press, 2011, pp. 75–86.

[7] J. S. Warner and R. G. Johnston, "GPS spoofing countermeasures," *Homeland Security Journal*, 2003.

[8] H. Wen, P. Y.-R. Huang, J. Dye, A. Archinal, and J. Fagan, "Countermeasures for GPS signal spoofing," in *Proc. 18th Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS)*. Institute of Navigation, 2005, pp. 1285–1290.

[9] P. Y. Montgomery, T. E. Humphreys, and B. M. Ledvina, "Receiver-autonomous spoofing detection: Experimental results of a multi-antenna receiver defense against a portable civil GPS spoofer," in *Proc. 22th Technical Meeting of The Institute of Navigation*, 2009, pp. 124–130.

[10] P. Papadimitratos and A. Jovanovic, "GNSS-based positioning: Attacks and countermeasures," in *Proc. IEEE Military Communications Conference (MILCOM)*. IEEE CS, 2008.

[11] A. Jafarnia-Jahromi, A. Broumandan, J. Nielsen, and G. Lachapelle, "GPS vulnerability to spoofing threats and a review of antispoofing techniques," *International Journal of Navigation and Observation*, vol. 2012, 2012.

[12] J. Nielsen, A. Broumandan, and G. Lachapelle, "Spoofing detection and mitigation," *GPS World*, September 2010.

[13] M. Psiaki, B. O'Hanlon, J. Bhatti, D. Shepard, and T. Humphreys, "GPS spoofing detection via dual-receiver correlation of military signals," *IEEE Transactions on Aerospace and Electronic Systems*, 2013, to appear.

[14] K. Borre, *A Software-Defined GPS and Galileo Receiver*. Applied and Numerical Harmonic Analysis, 2007.

[15] T. Nighswander, B. M. Ledvina, J. Diamond, R. Brumley, and D. Brumley, "GPS software attacks," in *Proc. 19th ACM Conference on Computer and Communications Security (CCS)*. ACM Press, 2012, pp. 450–461.

[16] A. Jøsang, "Subjective logic (book draft)," available at http://folk.uio.no/josang/papers/subjective_logic.pdf, 2012.

[17] ——, "Conditional reasoning with subjective logic," *Multiple-Valued Logic and Soft Computing*, vol. 15, no. 1, pp. 5–38, 2009.

[18] A. Jøsang, S. Pope, and M. Daniel, "Conditional deduction under uncertainty," in *Proc. 8th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU)*, ser. LNCS, vol. 3571. Springer, 2005, pp. 824–835.

APPENDIX

THE CONSENSUS OPERATOR ($\oplus$)

Definition 7 (Cumulative fusion operator)**.** *Let w^A and w^B be opinions respectively held by agents A and B over the same frame $X = \{x_j \mid j = 1, \ldots, l\}$. Let $w^{A\diamond B}$ be the opinion such that*
Case I: For $u^A \neq 0 \vee u^B \neq 0$:

$$b^{A\diamond B}(x_j) = \frac{b^A(x_j)u^B + b^B(x_j)u^A}{u^A + u^B - u^A u^B}$$
$$u^{A\diamond B} = \frac{u^A u^B}{u^A + u^B - u^A u^B}$$

Case II: For $u^A = 0 \wedge u^B = 0$:

$$b^{A\diamond B}(x_j) = \gamma^A b^A(x_j) + \gamma^B b^B(x_j), \quad u^{A\diamond B} = 0$$

where

$$\gamma^A = \lim_{u^A \to 0, u^B \to 0} \frac{u^B}{u^A + u^B}, \gamma^B = \lim_{u^A \to 0, u^A \to 0} \frac{u^B}{u^A + u^B}.$$

Then $w^{A\diamond B}$ is called the cumulatively fused bba of w^A and w^B, representing the combination of independent opinions of A and B. By using the symbol '$\oplus$' to designate this belief operator, we define $w^{A\diamond B} \equiv w^A \oplus w^B$.