# SureSpace: orchestrating IoT devices to certify location

João Maria Marques Tiago Instituto Superior Técnico, Universidade de Lisboa joao.marques.tiago@tecnico.ulisboa.pt

# Abstract

Location-aware mobile applications are increasingly popular and useful. However, as more services rely on location, there are concerns that users may misreport their location to gain undue advantages. One way to prevent such location spoofing is to rely on location certification systems. For example, SureThing uses Wi-Fi or Bluetooth beacons and adhoc witnesses to allow a user to make proof of location at a specific time and place. This approach can be extended to smart spaces, such as smart buildings, managed by platforms like DS2OS. In this work, we propose, implement, and evaluate SureSpace, a new system that combines location certification with smart space management, to verify the location of users inside a smart building. The new system relies on existing infrastructure to act as witnesses, while introducing a varied combination of devices and proof techniques to achieve security by diversity in the proofs, and thwart spoofing attacks. The system was evaluated in laboratory conditions, and was shown to be effective using light and audio signals.

# 1 Introduction

Mobile applications have gained popularity in the last decade, given the increased ubiquity and pervasiveness of mobile devices on people's everyday lives. Meanwhile, the Internet of Things (IoT) [12] is emerging as a network of interconnected smart devices that collect and consume information for management, and decision-making purposes in smart environments, depending upon characteristics like network availability or coverage area [6]. Mobile devices can interact with IoT devices for added value services, but security is a critical concern. Context attributes like identity, time and location, need to be trusted to allow for good security decisions when granting access to resources [3].

The use of *location* in mobile applications is increasingly

popular and useful. However, in most cases, location is collected by the device itself, and the user may misreport location to gain undue advantages.

To prevent location spoofing, location certification mechanisms can be deployed to prove that a user is at a specific location (either geographical or logical), at a specific time. SureThing [9] is a recent proof of location framework built with Java on the Android platform that uses location measurements collected with GPS, Wi-Fi and Bluetooth Low Energy (BLE) to certify location. SureThing makes use of *witness* devices to verify and attest to the presence of users in crowded physical spaces with diverse devices. In its current version, it does not take advantage of existing infrastructure available at a location.

With this work, we introduce SureSpace, an extended version of the original architecture of SureThing, designed to engage in smart environments in a secure way. To generate a location proof, *beacons*, which are on-site smart devices, broadcast unique signals meant to be captured by the prover using *witnesses* in the prover device. If the captured signal matches the original signal to a certain threshold, the location proof is deemed valid. Different approaches and techniques are supported for signal processing and matching. To discover, configure and control beacons inside the smart space, we use DS2OS [10], a smart space management framework. It bundles a beacon discovery mechanism to account for the high dynamism of smart environments, allowing beacons to be added or removed in runtime.

#### 2 Background

In this Section, we present the two systems that were composed to build SureSpace: SureThing and DS2OS.

## 2.1 SureThing

SureThing [9] is a location certification system for Android devices that lets users prove their location. Its design and architecture are influenced by other location certification systems, such as APPLAUS [13], and Crepuscolo [7], to the extent that they share some core components. To locate the user, SureThing uses different *location estimation techniques*, such as Geo Proof (geographic location obtained from GPS or ANLP<sup>1</sup>), Wi-Fi Proof (via Wi-Fi fingerprinting), and Beacon Proof (using BLE beacons).

<sup>&</sup>lt;sup>1</sup>The Android Network Location Provider, that uses both cell tower and Wi-Fi to determine the device location.

The goal of location certification is to prevent location spoofing, by requiring proof of a claimed location. Using their Android smartphone, the *prover* says that they are at a specific location, and the system challenges the claim, by asking for *evidence*, and/or one or more *witnesses* at the location. The proof mechanism starts by collecting some sort of *signal*, that is unique to the location at a specific time. This signal is abstract, and it can be both a "natural" signal, like the radiation or noise level, or an artificial signal, such as a radio. The prover collects the signal, with errors, and keeps the evidence for later verification. When the prover wants to make proof of location and time, they are challenged by the verifier. The verifier can ask for the signal or some of its features, and compare them with a template or witnesses observations.

Since the original SureThing paper, there have been other applications implemented with the approach, some with adhoc witnesses, others with fixed witnesses, and also some with beacons that transmit pseudo-random signal sequences. There is also an ongoing work that is providing witness and user *privacy protection* through the use of different privacy mechanisms [8]. Until this work, SureThing lacked a way to leverage *multiple* signals and the smart space orchestration capabilities, required for a robust and diverse proof of location process.

# 2.2 DS2OS

DS2OS [10] is a smart space orchestration framework focused on the development of services for smart spaces. Smart devices deliver valuable data (e.g. light or temperature conditions), and perform useful work (e.g. turn a light or heating on). These capabilities can be used to implement scenarios where smart devices work together towards a common goal. The coordinated management of smart devices, known as smart space orchestration, is possible if smart devices can be interconnected to share data over a network.

In DS2OS, services are logical processes that deliver functionality in a smart space, usually grouped in two categories: *adaptation services*, that provide an interface to a smart device, and *orchestration services*, that implement the logic behind pervasive scenarios. In practice, each smart device requires an adaptation service, and smart devices are coordinated by an orchestration service. Regardless of their category, services have unique identifiers in DS2OS, conveniently named after the nature of the service (e.g. the adaptation service for a temperature sensor *could* go by temperatureadaptationservice).

Smart devices produce and consume pieces of unstructured information, known as *context*. To become manageable, context is shaped into *context models*, that represent entities in a structured way. Context models have *context nodes* (i.e. attributes) to describe properties of the entity they are linked to. To determine its type, each context node has a type attribute, which is, in fact, a context model itself, usually simpler. To understand this better, consider a context model of a temperature sensor, of type /sensor/temperature. The context model should include, at least, two properties: isOn, a /boolean for the operational status of the sensor, and value, a /number used to represent the read temperature. The temperatureadaptationservice would, then, include a context node of type /sensor/temperature to represent the sensor. In other words, all properties of the sensor's context model would be implicitly included in that context node.

Each service has its own context model, and inter-service communication is achieved through the manipulation of context models. Simply put, a context model is a blackboard: some services write on it, and other services read from it to achieve their objective. To prevent unauthorized operations, context nodes have read and write permissions. Following the example, all services should be able to change the value of isOn, so that the sensor can be turned on. However, value must be read-only to all services but the temperatureadaptationservice, authorized to update the temperature value. In this light, context models are interfaces for services, a sort of service contract that (1) specifies which attributes can be read and/or written, via get and set operations, respectively, and (2) by whom, according to an access control policy. For instance, if an orchestration service wishes to get the current temperature, it would have to call set ('isOn', true) on the context model of the temperatureadaptationservice (to enable the sensor), and then get ('value') to get the actual temperature.

Context models are stored in a distributed system over a peer-to-peer network, known as the Virtual State Layer (VSL). Peers of the network are called Knowledge Agents (KAs), and each agent persists a subset of the context models. To be granted access to the distributed knowledge, services register to a KA of their choice, that becomes responsible for the respective context models. From that moment on, these context models become available to other services, even if connected to a different agent.

One important feature of DS2SOS is that services can subscribe to specific context nodes to receive a notification when the node value changes. This feature is useful, for instance, when a service wants to take different actions depending on the new value. Another key feature of DS2OS is dynamic service discoverability [4, 11], that allows adaptation services to be discovered by different criteria: type of smart device (e.g. temperatureadaptationservice), or type of attribute (e.g. /boolean).

## **2.3** Supported location representations

Geocodes represent geographic locations on Earth, where each location is assigned a unique identifier, used to distinguish between entities. Usually, geocodes are short and human-readable. In SureSpace, we use geocodes to locate both the prover and the orchestrator within the building. Our problem domain allows locations to be represented with coarser granularity, because the focus is on determining if the prover is within the boundaries of the room, and not necessarily at a specific location within that area.

Currently, SureSpace supports two geocode systems: Open Location Code (OLC), and What3words (W3W). By default, SureSpace uses the OLC geocode system to encode locations into *plus codes* that represent squares on the Earth's surface. The longer the code, the smaller the square is, and, in its full length of 11 characters, a plus code represents a  $3 \times 3 m$  square. For example, 8CCGPVP5+GMW describes the reception hall of a building in Lisbon. Shortening the code to



Figure 1: SureSpace architecture.

8CCGPVP5+GM, a larger square containing the reception hall is drawn. This feature is useful to represent larger rooms with a single plus code.

## **3** Architecture

SureSpace is a location certification mechanism designed for smart environments. Its architecture integrates elements from SureThing, that enriches SureSpace with an assortment of techniques and technologies for location certification, and DS2OS, that offers the possibility to configure and orchestrate smart devices in a smart space, as well as the possibility to dynamically discover them in run-time.

#### **3.1** Components

Figure 1 outlines all components of SureSpace, focusing on the communication flows between them. Components are presented following the expected interaction sequence.

The prover is a SureSpace user that engages with the system in order to prove their location. The prover device is the device used by the prover during all interactions with the system.

The Certificate Authority (CA) is the long-term identity provider of all active entities of the system, similar to CAs in use on the Internet for website certification. To be deemed legitimate and to be able to engage with SureSpace, entities must register themselves to the CA. Each entity generates a public/private key pair. The private key is known only to the entity, and is kept safe and secure on their side. The public key is used to generate a certificate signing request in order to apply for a public key certificate. If the request is approved by the CA, a public key certificate is issued and assigned to the requester entity.

An orchestrator is the entry point to SureSpace system because it is the component the prover first reaches out to. An orchestrator is the "mastermind" of any proof of location, responsible for coordinating the process at the highest level, and preventing malicious communication flows coming from unauthorized parties. It implements diverse logical subprocesses that include, for example, the dynamic discovery of new orchestrated rooms, the dynamic discovery of new beacons and their orchestration, and the delivery of accurate information about a specific location proof. Orchestrators run an *orchestration service* to communicate with the adaptation services of the orchestrated beacons.

A Knowledge Agent (KA) is a node in the distributed knowledge network of DS2OS (refer to 2.2). More specifically, KAs are context repositories that persist relevant information used by the orchestrator. To deliver room-level orchestration, each orchestrated room has its own KA, that behaves like a *proxy* to the room (meaning that no rooms share the same KA). Agents hold information about their geographical location by running a *localization service*, so that the orchestrator can discover them by location when looking for new orchestrated rooms. Since beacons register themselves to the closest KA in their vicinity, new beacons are easily discoverable and accounted for. During a proof of location, an orchestrator will need information about engaging beacons (and respective trusted witnesses), like the value of specific attributes. That information becomes available to other KAs, since they are all nodes in the same distributed knowledge network.

A beacon is a smart device embedded into the trusted infrastructure ready to be used in a proof of location (e.g. a smart bulb). By default, SureSpace is not aware of existing beacons by themselves, since they may not be directly connected to the system. Thus, each beacon requires a proxy to become visible to, and controllable by the system. That proxy is an *adaptation service*, that represents and controls exactly one beacon. During the proof of location, each beacon generates and broadcasts exactly one *signal* meant to be captured by, at least, one corresponding witness in the prover device. A signal is something produced by a beacon that can be received by a corresponding witness. For instance, visible light, emitted by a smart bulb (the beacon), and acknowledged by a light sensor (the witness), could be used as a signal. The definition of signal, however, is left open to avoid narrowing down the possibilities to a small set of conventional signals, paving the way for out-of-the-box ideas. Theoretically, any equipment can be used as a witness, provided it is able to acknowledge signals from a specific beacon. Witnesses in the prover device are *untrusted*, because they are not part of the trusted infrastructure.

A signal is generated based on a set of *quirky properties*, that feed a deterministic signal generator. If these properties are disclosed, the original signal can be easily replicated. Naturally, signals have different characteristics/features, and are susceptible to deterioration induced by multiple factors during their transmission. Moreover, witnesses have limited capabilities, and might not be able to acknowledge all characteristics of the signal, but only a subset of them. Thus, what the witness receives is not the original signal, but a *degraded* representation of some of its characteristics. In some cases, part of the characteristics of the signal can be successfully derived from the analysis and/or processing of the degraded representation. To the derived information we call *proof ambient*, because it represents the witness' knowledge of the original signal.

To determine the legitimacy of the location proof, we

measure the accuracy of the proof ambient by quantifying similarity between the original signal, that was transmitted by the beacon, and the degraded representation, that was received by the witness. To do that, we require a *trusted representation* of the signal from a *trusted witness*, embedded into the reliable infrastructure, capable of acknowledging the same set of characteristics, to ensure the two representations are compatible. In some cases, usually when the infrastructure is not open for modification, trusted witnesses can be *virtual*. Without the need for real devices, virtual witnesses mimic the behavior of receiving the signal through emulation to deliver the same functionality. Alongside with a beacon, corresponding trusted witnesses, be them virtual or not, are represented and controlled by the same adaptation service.

The verifier is the last entity to engage in SureSpace. It measures the reliability of a location proof by assessing the legitimacy of the proof ambient derived by the prover. The verifier implements adequate criteria to compare different representations of the same signal. However, there is no predetermined assessment criteria, because (1) the definition of signal itself remains abstract enough to encompass a variety of beacon, and (2) application-specific criteria might have to be taken into account. Regardless of the implementation details, the verifier must output a boolean value that represents the assessment result. If true, the location proof is deemed reliable, and, thus, accepted.

#### 4 Location certification process

The process encompasses three stages: *pre-authorization*, *proof*, and *verification*.

## 4.1 **Pre-authorization stage**

A proof of location requires the orchestration of a subset of beacons scattered across a smart location, like an orchestrated room. Beacon orchestration requires some context information to be readily available at proof-time (like which beacons are engaging with the proof of location).

The objective of this pre-stage, represented in Figure 2, is to produce an *authorization*, requested by the prover and issued by an orchestrator, that works as a token used to trigger the proof stage, while containing relevant metadata necessarey for the proof of location.

First, the prover device determines its compatibility with the existing beacon-witness mapping (*Device compatibility check* step). A beacon is deemed supported by the prover device if and only if it has, at least, one compatible untrusted witness (usually a sensor, like a light sensor). Consequently, the set of supported beacons depends on the hardware properties of the prover device.

The prover device estimates its location resorting to external mechanisms (*Device location estimation* step). For instance, GPS can be used if the signal is strong enough. Upon locating the prover device, GPS coordinates are then converted to a plus code (refer to 2.3). In the absence of a suitable localization system (like in GPS-constrained environments), scanning an on-site QR code with the geocode of the room might be sufficient to locate the prover within the building. In the end of this step, the prover requests a proof authorization to the orchestrator (*Request authorization* step).



Figure 2: Process diagram for the pre-authorization stage.



Figure 3: Process diagram for the proof stage.

Following the validation of the request, the orchestrator determines which beacons are available at the location reported by the prover (step *Adaptation service discovery*). Different orchestrated rooms offer different beacons, discoverable via adaptation services delivered within that room. To determine which beacons are available at the reported location, the orchestrator (1) lists all orchestrated rooms, (2) determines the closest one to the prover, and (3) discovers which beacons are available in that room.

Finally, the orchestrator selects beacons eligible for the proof of location (*Eligible beacons selection* step). A beacon is eligible if and only if (1) it is supported by the prover device, and (2) available at the prover location. A beacon selection policy might filter eligible beacons, depending on policy criteria (e.g. the security level of the room). In the end of this step, the orchestrator generates an authorization, stores it for future use, and sends it to the prover.

#### 4.2 **Proof stage**

This stage, represented in Figure 3, starts when the prover submits the proof authorization to the orchestrator (step *Submit authorization*).

Beacons generate signals based on a set of quirky properties, used to feed a deterministic signal generator. To minimize the likelihood of broadcasting the same signal twice, a seed value is generated in an unpredictable way (step *Generate random seed*), and it is used to populate quirky properties with pseudorandom values derived from it.

The orchestrators lock the adaptation services of the selected beacons, and applies new pseudorandom values to the quirky properties (step *Configure beacons*). When beacons are ready, they start broadcasting their signal. At the same time, trusted witnesses in the infrastructure start receiving the signal, as well as untrusted witnesses in the prover device. Untrusted witnesses send the information directly to the prover device, that safely stores the information, and trusted witnesses share the information with the orchestrator via the VSL (by updating the corresponding properties of their context models).

Network latency has direct impact on the level of synchronization between orchestrator and prover, since only upon clearance from the orchestrator will the prover initiate the process. If desynchronized, "dead times" may occur in the beginning (beacons are broadcasting, but the prover is waiting for clearance), and in the end of the process (beacons have ceased their activity, but untrusted witnesses remain active). This phenomena *should be* mitigated in the verification stage, without relying on clock synchronization.

## 4.3 Verification stage

Theoretically, a proof of location is accepted if the proof ambient is *complete* enough (with regard to all the signals broadcast by engaging beacons), and *accurate* enough (if the extracted information is in line with the original signal's information).

Even in optimum conditions, signal degradation will decrease the accuracy of the proof ambient. Internal factors (e.g. witnesses sensitivity) and external factors (e.g ambient noise, topology of the orchestrated room) might lead to a misrepresented, yet legitimate, proof ambient. To account for such errors, a *margin of error* is considered when comparing the trusted representation of the original signal with its degraded representation.

To verify the location proof, the prover submits the proof ambient of the different signals to the verifier. A final judgment is yield, either *accepting* or *rejecting* the location proof. Lest compromising the versatility of the verification step, implementation details of the stage are left open.

#### **5** Implementation

### 5.1 Development platform

We developed a prototype of SureSpace as a Java project, using the JDK (Java Development Kit) version 11, and used Maven 3.6.3 to manage the building process and help with dependencies version management. We implemented a custom version of Ardulink 2, a Java solution for coordinating Arduino boards [2] (required for the experimental setup), and implemented the Certificate Authority, the orchestrator, and the verifier from scratch (some code was partially inspired in existing work by João Ferreira [9]). We also implemented all DS2OS services (adaptation, localization, and orchestration services) by extending available templates. To bootstrap the VSL, KAs need to be launched one by one to become peers (refer to 2.2). In practical terms, a KA is bundled as an executable JAR file. However, the code of a KA was no longer compatible with our setup, so we modified it, and recompiled it. The prover was implemented as an Android mobile application, for a richer user-experience, compiled in Java 8 against API 30.

#### 5.2 Cross-entity communication

Depending on their domain, entities communicate using different underlying communication protocols. DS2OS en-



Figure 4: Message structure shared in SureSpace.

tities (KAs and services) communicate with the VSL via REST connectors (using HTTPS as tunneling protocol for added security). Communications between SureSpace entities (CA, orchestrator, and verifier) are supported by gRPC (in Java), following a remote invocation paradigm. gRPC uses Protocol Buffers (protobuf) to provide a platformindependent representation of remote interfaces, and it was chosen because of its efficiency, and loose coupling between clients and servers [1]. Moreover, communications between SureSpace entities share a common payload format, illustrated in Figure 4, that (1) allows for a standardized process of message validation, and (2) fosters the implementation of the desired security properties (integrity, authentication, and non-repudiation). All messages hold information about their source (field sender), and destination (field receiver), to prevent message forwarding. Replay attacks are mitigated by using a securely generated random number (field nonce). For integrity, authentication, and nonrepudiation purposes, a digital signature is generated over the message (field signature) using SHA-512 with the source entity's private key. To reduce the number of interactions with CAs, the source entity's certificate (field certificate) is attached to the body of the message prior to signing.

## 5.3 Entity identification

Each SureSpace entity is assigned a public key certificate that is part of a certificate chain that terminates with the SureSpace Root CA. Public/private key pairs are generated using 2048-bit RSA, and certificates are signed using SHA-512. Moreover, each entity is identified by a hierarchical identifier, unique within the SureSpace domain.

## 5.4 Beaconing technique

To determine if a signal is appropriate for a proof of location, we classify it based on two metrics: *difficulty of replication*, and *difficulty of acknowledgment*. Signals should be difficult to replicate without knowing the quirky properties used for their generation. If signals have noticeable patterns or are reused, a malicious party can easily replicate them. At the same time, signals must be versatile enough to account for common limitations shared by compatible witnesses, so that signals can be easily received by most witnesses.

We propose a technique based on *time fragmentation* to reach an equilibrium between these two metrics. During the

proof stage, each signal is broken into a fixed number of consecutive, same-length fragments, and each fragment is generated based on a set of quirky properties.

For simplicity, we adopt the following notation:

- $b \in B$  denotes beacon b, in the set of supported beacons, B
- *w* ∈ *W<sub>b</sub>*, *b* ∈ *B* denotes witness *w*, in the set of witnesses compatible with beacon *b*, *W<sub>b</sub>*
- *f*<sub>b,i</sub> ∈ *F*<sub>b</sub>, *b* ∈ *B*, *i* ≥ 1 denotes the *i*-th fragment of the set of fragments that compose a signal broadcast by beacon *b*, *F*<sub>b</sub>
- $q \in Q_{f_{b,i}}, b \in B, i \ge 1$  denotes quirky property q, in the set of quirky properties used to generate fragment  $f_{b,i}$ ,  $Q_{f_{b,i}}$
- $S_b = f_{b,1} || f_{b,2} || \dots || f_{b,n}, b \in B, i \ge 1$  denotes a signal with  $n \ge 1$  fragments, broadcast by beacon b

Algorithm 1 proposes a pseudocode of the technique. Since each fragment is very likely to have a different set of quirky properties, we assume that no two fragments share the same set of quirky properties, resulting in the creation of unique signals.

# 5.4.1 Light signal time fragmentation

Consider a light source (e.g. a LED) used as a beacon in a proof of location,  $b_{light}$ , that can switch between states on and off with period  $P \in [P_{MIN}, P_{MAX}]$ , measured in a convenient unit ( $P_{MIN}$  and  $P_{MAX}$  are, respectively, the lowest and the highest supported periods). In the proof stage, the beacon broadcasts a signal,  $S_{light}$ , split into *n d*-seconds fragments. During each fragment  $f_{light,i}$ , i = 1, 2, ..., n, the beacon switches between states with a pseudorandom period  $P_i$  (the quirky property), derived from the seed, forming a *power-on and power-off* sequence with rates that vary between fragments.

Theoretically,  $P_i$  can take any value in range  $[P_{MIN}, P_{MAX}]$ , and nothing prevents two pseudorandom periods from being equal or close enough to generate similar or indistinguishable fragments. To avoid this, we adopt an approach that (1) prevents reusing the same period and (2) reduces the probability of picking periods too close in the range.

As witness, we consider a light sensor,  $w_{light}$ , capable of measuring light intensity in a convenient unit, at a sampling rate not less than  $P_i^{-1}$ ,  $\forall i$ . Plotting the measurements over time offers a representation of  $S_{light}$  based on one of

```
1 <model type="/complex/beacon">
2 <isOn type="/basic/boolean"/>
3 <switchingPeriod type="/basic/number"/>
4 </model>
```

Figure 5: Simplified context model of a light beacon.

its properties (light intensity). The analysis of that representation may confirm significant variations in light intensity, which are an interpretation of the power-on and power-off sequence.

## 5.4.2 Audio signal time fragmentation

Consider an audio source (e.g a speaker) used as a beacon in a proof of location,  $b_{audio}$ . The beacon can be programmed to play any song out of a predefined set of *m* songs, and  $songId \in \{0, 1, ..., m-1\}$  is the index of the song to be played. This song is any regular song that plays on the radio, and we consider it over any synthesized melody because a song is more easily tolerated by the human ear for extended periods of time. During its activity, the beacon broadcasts a signal,  $S_{audio}$ , with a single *d*-seconds fragment, and a pseudorandom  $songId \in Q_{f_{b_{audio}},1}$  (the only quirky property) is derived from the seed.

As witness, we consider a sound sensor,  $w_{audio}$ , capable of measuring sound amplitude in a convenient unit. Plotting the measurements over time offers a representation of  $S_{audio}$  based on one of its properties (audio amplitude). The analysis of that representation may confirm variations in amplitude and frequency that match the song being played.

## 5.5 Supported beacons

Adding support for a new beacon requires writing its context model with all the configurable properties (the trusted witness requires its own context model too), and implementing the adaptation service, so that the beacon can be discoverable.

Currently, SureSpace supports two beacons (a light beacon, and an audio beacon), and next we go over the steps required to add support for them.

#### 5.5.1 Light beacon and witnesses

Based on the example in 5.4.1, we considered a light beacon capable of switching between states on and off with a configurable period. Figure 5 is a simplified context model of the beacon, where ison is a boolean used to control the beacon (if set to true, the beacon is working), and switching-Period is the time it takes for the beacon to switch between states (in seconds).

As trusted witness, we considered a light sensor capable of measuring light intensity at a configurable sampling rate not less than switchingPeriod<sup>-1</sup>. Figure 6 is a simplified context model of the trusted witness, where intensity is the measured light intensity (in a convenient unit), intensity-SamplingRate is the sampling rate at which the sensor is reading (in Hertz), and isOn is a boolean used to control the witness (if set to true, the witness is working). The orchestration service subscribes the intensity attribute on the trusted witness context model. Every time the attribute value changes because of a new measurement, the orchestration

```
1 <model type="/complex/witness">
2 <intensity type="/basic/number"/>
3 <intensitySamplingRate type="/basic/number"/>
4 <isOn type="/basic/boolean"/>
5 </model>
```

#### Figure 6: Simplified context model of a light witness.

```
1 <model type="/complex/beacon">
2 <isOn type="/basic/boolean"/>
3 <songId type="/basic/number"/>
4 </model>
```

Figure 7: Simplified context model of an audio beacon.

service is notified. At that moment, the value is timestamped (in milliseconds), and stored in the orchestrator to compose the trusted representation of the light signal.

As untrusted witness, we consider any device capable of measuring light intensity to ensure both witnesses acknowledge the same property of the signal (intensity).

5.5.2 Audio beacon and witnesses

Based on the example in 5.4.2, we considered an audio beacon capable of playing a specific song out of a set of songs stored in a raw format (like WAV). Figure 7 is a simplified context model of the beacon, where isOn is a boolean used to control the beacon (if set to true, the beacon is working), and songId is the index of the song to be played.

As trusted witness, we considered a sound sensor capable of measuring the sound amplitude at a specified sampling rate. Figure 8 is a simplified context model of the trusted witness, where amplitude is the measured sound amplitude (in a convenient unit), amplitudeSamplingRate is the sampling rate at which the sensor is reading (in Hertz), and isOn is a boolean used to control the witness (if set to true, the witness is working).

The trusted witness was implemented as fully virtual to demonstrate the feasibility of the approach. Since we have access to the songs in a raw format, it is possible to emulate the behavior of a physical witness. Every amplitudeSamplingRate<sup>-1</sup> s, the virtual witness reads the sound amplitude from the file, and updates the amplitude attribute on its context model, which has been subscribed by the orchestration service. Every time the attribute value changes, the orchestration service is notified. At that moment, the value is timestamped (in milliseconds), and stored in the orchestrator to compose the trusted representation of the audio signal.

Figure 8: Simplified context model of an audio witness.

As untrusted witness, we consider any device capable of measuring sound amplitude to ensure both witnesses acknowledge the same property of the signal (amplitude).

# **5.6** Verifier implementation

To verify a location proof, the verifier quantifies similarity between different representations of the same signal: a degraded one (from untrusted witnesses in the prover device), and a trusted one (from trusted witnesses in the infrastructure). If more than one beacon is used (and, thus, more than one signal is involved), individual similarity estimates are weighted for a final similarity estimate.

Based on the beacons we support, we use the MATLAB Engine API for Java to quantify similarity. Next, we detail the approach used for comparing representations of the same signal, for both light signals, and audio signals.

## 5.6.1 Light signal representations similarity estimation

Representations may be sampled at different rates, impeding their comparison. We bring them to a common rate by upsampling the representation with the lowest frequency, using linear interpolation. This process produces an approximation of the representation that would have been obtained by sampling at a higher rate.

Because clock synchronization is not a requirement, potential delays between representations may exist. To align them without relying on timestamps, we use correlation to determine where representations overlap the most, and then align them.

At last, we normalize both representations, and calculate the linear correlation coefficient between them,  $corr_{light 1}$ , given by Equation 1

$$r = \frac{\sum_{i=1}^{n} (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^{n} (x_i - \bar{x})^2} \sqrt{\sum_{i=1}^{n} (y_i - \bar{y})^2}}$$
(1)

where *r* is the linear correlation coefficient, *n* is the number of samples in the representations (they have the same size),  $x_i$  and  $y_i$  are the sample points, and  $\bar{x}$  and  $\bar{y}$  are the means of the samples. This coefficient measures the linear relationship between the two representations, and is used as an estimate of their similarity.

# 5.6.2 Audio signal representations similarity estimation

Audio and light signals are fundamentally different, and so are their representations. For that reason, we cannot follow the previous approach. Instead, after normalization, we use dynamic time warping to resample and align the representations. This algorithm stretches the two representations onto a common set of instants, such that the sum of the Euclidean distances between corresponding points is smallest. Then, we calculate the linear correlation coefficient between the aligned representations, *corr<sub>audio</sub>* 1, and use it as a first estimate of their similarity.

To improve the estimate, we calculate the power spectrum of the two representations. Simply put, a power spectrum is a frequency-domain interpretation of an audio signal representation because it describes the distribution of power



Figure 9: Experimental setup components, and orchestrated area.

(sound amplitude) into frequency components. In this context, this information is relevant because each song has a different time-frequency structure. Thus, we calculate the linear correlation coefficient between the power spectra of the two representations, *corr<sub>audio 2</sub>*, and use it as a second estimate of their similarity.

The two estimates are weighted for a final similarity estimate, given by Formula 2

$$w_1 \times corr_{audio\ 1} + w_2 \times corr_{audio\ 2}, w_2 = 1 - w_1 \qquad (2)$$

Weights  $w_1$  and  $w_2$  must be tuned based on a training set, since they are necessarily beacon- and witness-dependent. 5.6.3 Combined signals similarity estimation

For a final similarity estimate, individual similarities are estimated according to Formula 3

$$w_{3} \times corr_{light \ 1} + \\ w_{4} \times (w_{1} \times corr_{audio \ 1} + w_{2} \times corr_{audio \ 2}), \qquad (3) \\ w_{4} = 1 - w_{3}$$

Weights  $w_3$  and  $w_4$  need to be tuned.

# 6 Evaluation

In this Section, we present the experimental setup used to evaluate the SureSpace prototype, describe the evaluation criteria, and discuss the evaluation results.

## 6.1 Experimental setup

For the experimental setup, we used inexpensive equipment, with less accuracy, but more representative of commodity equipment that we expect to find in a smart building. We dropped room-level orchestration, and opted for a smaller, yet representative, orchestrated area shown in Figure 9.

Recall notation from 5.4, where *b* denotes a supported beacon, and *w* a corresponding witness. Based on the supported beacons, we used a Grove Chainable RGB Led V2.0 as light beacon,  $b_{light}$ , a Grove Light Sensor V1.2 as trusted light witness,  $w_{light}$ , and a JBL GO 2, connected to a Grove MP3 V2.0 module, as audio beacon,  $b_{audio}$ . The trusted audio witness,  $w_{audio}$ , is fully virtual, eliminating the need for physical equipment. For connectivity reasons,  $b_{light}$ ,  $w_{light}$ , and the Grove MP3 V2.0 module were all connected to a Grove Base Shield V2.0 for an Arduino Uno board. The

prover device was a Huawei Mate 20 Pro Android smartphone, shipped with Android 10, equipped with a built-in ambient light sensor, the untrusted light witness,  $w'_{light}$ , and a microphone, the untrusted audio witness,  $w'_{audio}$ . During the proof of location, the prover device is steady in the center of the orchestrated area, as depicted.

We built a dataset by running 80 location proofs under the same *controlled scenario*. Each location proof delivered *four* signal representations (two representations per signal, a trusted one and a degraded one). To be assessed, a location proof is submitted to the verifier (refer to 5.6).

In some cases, the verifier may classify location proofs incorrectly, either by accepting a location proof that should be rejected (false positive), or, conversely, by rejecting a location proof that should be accepted (false negative). To evaluate SureSpace, we consider (1) the false positive rate (*FPR*), which is the percentage of all negatives that still yield positive, (2) the false negative rate (*FNR*), which is the percentage of positives that yield negative, and (3) the success rate, which is the percentage of correct classifications. *FPR* and *FNR* are inversely proportional to the success rate, and the verifier should focus on minimizing them.

To ensure the reproducibility of the experiments, the duration of the proof stage was set to 30 seconds in all proofs of location (a reasonable value in a human time-scale that works in a possible meeting room scenario). Regarding the audio component, beacon  $b_{audio}$  could choose between 20 different predefined songs, all sampled at 44.1 kHz (refer to 5.4.2). Regarding the light component, light signals were broken into two 15-seconds fragments, and beacon  $b_{light}$ could switch between states with a pseudorandom period  $P \in [0.5, 7.5]$  (refer to 5.4.1). The upper bound of the range is 7.5s (=  $\frac{15s}{2}$ ) to ensure that light signal fragments generate, at least, one complete on-off sequence (i.e. the beacon is powered on, and then it is powered off for the same amount of time at least once). The lower bound of the range is 0.5sto ensure compatibility with all light witnesses (based on the information we collected,  $w'_{light}$ , the Android ambient light sensor, is the light witness that reports the lowest sampling rate of  $\approx 2Hz$ ).

# 6.2 Optimal weight tuning

We started by dividing our dataset into two subsets: a training set, and a test set. Following the Pareto principle, the training set accounted for 20% of the dataset (i.e. 16 location proofs), and the test set accounted for the remaining (i.e. 64 location proofs).

We started by tuning weights  $w_1$  and  $w_2$  using the training set (refer to 5.6.2). We crossed audio signal representations from all location proofs in the training set, ending with 256 combinations (16 × 16), from which 16 should be accepted (the number of legitimate location proofs), and 240 should be rejected (the number of fabricated location proofs). We varied  $w_1$  (recall that  $w_2 = 1 - w_1$ ) to find the optimal combination that would minimize FPR + FNR. Figure 10 plots the sum as a function of  $w_1$ . The local minima is at  $w_1 = 0.661$ , which means that  $\langle w_1, w_2 \rangle = \langle 0.661, 0.339 \rangle$  offers the best success rate (FPR = 33.33% and FNR = 37.50%). Replac-



Figure 10: Optimal value of  $w_1$ .



Figure 11: Optimal value of  $w_3$ .

ing  $w_1$  and  $w_2$  in Formula 2, the audio signal representations similarity estimate is given by Formula 4

$$0.661 \times corr_{audio\ 1} + 0.339 \times corr_{audio\ 2} \tag{4}$$

Then, we tuned weights  $w_3$  and  $w_4$  (refer to 5.6.3). Following the same approach, we crossed audio and light signals from all location proofs in the training set, and varied  $w_3$  (recall that  $w_4 = 1 - w_3$ ) to find the optimal combination that would minimize FPR + FNR. Figure 11 plots the sum as a function of  $w_3$ . The local minima is at  $w_3 = 0.436$ , which means that  $\langle w_3, w_4 \rangle = \langle 0.436, 0.564 \rangle$  offers the best success rate (FPR = 7.50% and FNR = 6.25%). Replacing all weights in Formula 3, the final similarity estimate is given by Formula 5

$$0.436 \times corr_{light 1} + 0.564 \times (0.661 \times corr_{audio 1} + 0.339 \times corr_{audio 2})$$
(5)

## 6.3 Approach effectiveness

We validated our approach by testing our signal representation similarity estimate against the test set. We crossed light and audio signal representations from all location proofs in the test set, ending with 4096 combinations ( $64 \times 64$ ), from which 64 should be accepted (the number of legitimate location proofs), and 4032 should be rejected (the number of fabricated location proofs). In the end, the verifier classified location proofs correctly in 94.78% of the cases, with rates *FPR* = 5.06% and *FNR* = 15.63%.

For demonstration purposes, consider the two light signal representations of a location proof accepted by the verifier. Figure 12 plots the normalized light intensity read by both witnesses in different colors (the trusted witness,  $w_{light}$ , and the untrusted witness,  $w'_{light}$ ). It becomes evident that the light signal is split into two fragments, with the second fragment starting at around 15 *s*. At that moment, the rate at which  $b_{light}$  changes between states on and off decreases.

The trusted representation is sharper and has more spikes than the untrusted representation. This difference can be ex-



Figure 12: Light signal representations overlapping.



Figure 13: Audio signal representations overlapping.

plained by the very different sampling rates at which both witnesses work. Based on the information we collected during the experiments,  $w_{light}$  works at  $\approx 14$  Hz, while  $w'_{light}$  works at  $\approx 2$  Hz. For that reason, the trusted witness is aware of the commence of the second fragment, while the untrusted witness misses it, since it occurs in-between samples. Notwithstanding, representations overlap, suggesting both witnesses were exposed to the same light conditions.

Additionally, consider the two audio signals representations of a location proof accepted by the verifier. Figure 13 plots the normalized sound amplitude read by both witnesses in different colors (the virtual trusted witness, waudio, and the untrusted witness,  $w'_{audio}$ ). Although both representations overlap, the trusted representation has stronger and neater variations in sound amplitude, while the untrusted representation looks more sketchy, with periods of constant sound amplitude. Once again, this is a consequence of sampling the same sound at different sampling rates. Based on the information we collected during the experiments, waudio works at  $\approx$  30 Hz, while  $w'_{audio}$  works at  $\approx$  12 Hz. The power spectra of both representations, depicted in Figure 14, shows prominent peaks in magnitude occurring around the same frequencies. This suggests that both witnesses were capturing the same song being played by  $b_{audio}$ .



Figure 14: Audio signal representations spectra overlapping.

# 6.4 Resistance to attacks

We consider that SureSpace has successfully resisted an attack when it rejects an illegitimate location proof. In other words, we are interested in the number of false positives, and, thus, consider FPR as the metric to evaluate SureSpace's resistance to attacks.

Based on the attacker model, we consider two attackers that try to prove their false presence by manipulating signals:

- A1 Receives the signal from exactly one random beacon, but not from the others (beacons may have different radius of action);
- A2 Combines legitimate signals representations to derive a synthesized proof ambient from them.

Attacker A1 can either (1) receive the light signal but not the audio signal, or (2) receive the audio signal but not the light signal. Based on Formula 5, we confirm that signals have similar weights (0.436 vs 0.564), so we expect location proofs to be rejected if an attacker only provides the representation of one of the signals. To simulate attacker A1, we used the 4032 fabricated location proofs that should be rejected by the verifier, and, according to the scenario, (1) or (2), we discarded one of the signal representations when submitting the location proof to the verifier. In both scenarios, we determined FPR = 0.00 % — it means that the verifier successfully rejected all location proofs fabricated by the attacker.

In 6.3, we combined legitimate signal representations from different location proofs to fabricate new illegitimate location proofs. In fact, we were already simulating attacker A2, and determined FPR = 5.06 %. In fact, combining signals from different location proofs is not an effective attack because the verifier is capable of telling they were generated for different location proofs (since the likelihood of broadcasting the same signal is low).

## 7 Conclusion and future work

In this paper, we presented SureSpace, a location certification system designed for smart environments. It leverages the capabilities of both SureThing, that defines procedures and techniques for proofs of location, and DS2OS, that provides control over diverse smart devices for orchestration purposes. SureSpace relies on smart devices that broadcast preconfigured signals, and witnesses that capture these signals, to certify location at a specific time and place, without requiring the presence of other users. The system was evaluated in laboratory conditions with inexpensive Arduinocompatible equipment. It was shown to be effective using light and audio signals, with a success rate up to 94.78%. Moreover, we evaluated SureSpace's resistance to attacks by simulating the capabilities of an attacker.

As future work, we consider the possibility of implementing more complex methods for optimal weight tuning, that could possibly improve the effectiveness of SureSpace.

We also consider a new way to verify location proofs based on a *challenge-response* approach. Instead of disclosing the complete captured signal to the verifier, the prover could be challenged to answer questions about specific components of the location proof. For example, and considering the light signal, we could consider questions like "For how long was the beacon in the on state?".

Moreover, SureSpace could use a privacy protection mechanism to protect the prover's identity. [5] proposes an approach based on short-term identifiers, or *pseudonyms*, used in communications. The long-term identifier is only disclosed under defined conditions, allowing mapping the pseudonym to the real identity in some cases (e.g. when the orchestrator wants to validate the prover identity).

Finally, SureSpace could be evaluated in an actual smart environment, using available smart devices. The adoption of SureSpace would also require the development of an improved end-user mobile application. Then, it could be used in real-world applications, like a check-in app for physical meetings.

### 8 References

- [1] About gRPC | gRPC. https://www.grpc.io/about/. Accessed: 2020-12-23.
- [2] Ardulink 2. https://github.com/Ardulink/Ardulink-2. Accessed: 2020-12-21.
- [3] C. A. Ardagna, M. Cremonini, E. Damiani, S. D. C. di Vimercati, and P. Samarati. Supporting Location-Based Conditions in Access Control Policies. In *Proceedings of the 2006 ACM Symposium on Information, Computer and Communications Security*, ASIACCS '06, page 212–222, New York, NY, USA, 2006. Association for Computing Machinery.
- [4] G. Aures and C. Lübben. DDS vs. MQTT vs. VSL for IoT. Network, 1, 2019.
- [5] N. Bißmeyer, J. Petit, and K. M. Bayarou. CoPRA: Conditional pseudonym resolution algorithm in VANETs. In 2013 10th annual conference on wireless on-demand network systems and services (WONS), pages 9–16. IEEE, 2013.
- [6] A. Buckman, M. Mayfield, and S. BM Beck. What is a smart building? Smart and Sustainable Built Environment, 3(2):92–109, 2014.
- [7] E. S. Canlar, M. Conti, B. Crispo, and R. Di Pietro. Crepuscolo: A collusion resistant privacy preserving location verification system. In 2013 International Conference on Risks and Security of Internet and Systems (CRiSIS), pages 1–9. IEEE, 2013.
- [8] J. Costa and M. L. Pardal. A Witness Protection for a Privacy-Preserving Location Proof System, 2020.
- [9] J. Ferreira and M. L. Pardal. Witness-based location proofs for mobile devices. In 17th IEEE International Symposium on Network Computing and Applications (NCA), Nov. 2018.
- [10] M.-O. Pahl. Distributed Smart Space Orchestration. Dissertation, Technische Universität München, München, 2014.
- [11] M.-O. Pahl and S. Liebald. A Modular Distributed IoT Service Discovery. In 2019 IFIP/IEEE Symposium on Integrated Network and Service Management (IM), pages 448–454. IEEE, 2019.
- [12] K. K. Patel, S. M. Patel, et al. Internet of Things IOT: definition, characteristics, architecture, enabling technologies, application & future challenges. *International journal of engineering science and computing*, 6(5), 2016.
- [13] Z. Zhu and G. Cao. APPLAUS: A privacy-preserving location proof updating system for location-based services. In 2011 Proceedings IEEE INFOCOM, pages 1889–1897. IEEE, 2011.