



Protecting personal data in IoT platform scenarios through encryption-based selective disclosure

José L. Hernández-Ramos^{a,*}, Salvador Pérez^a, Christine Hennebert^b, Jorge Bernal Bernabé^a, Benoit Denis^b, Alexandre Macabies^b, Antonio F. Skarmeta^a

^a Department of Information and Communication Engineering, Computer Science Faculty, University of Murcia, Spain

^b CEA-LETI, Grenoble, France

ARTICLE INFO

Keywords:

Attribute-based encryption
Confidentiality
Internet of Things
Smart buildings

ABSTRACT

As the Internet of Things evolves, citizens are starting to change the way they share information and communicate with their surrounding environment, enabling a constant, invisible and sometimes unintended information exchange. This trend raises new challenges regarding user's privacy and personal consent about the disclosure of personal data that must be addressed by flexible and scalable mechanisms. Towards this end, this work introduces the concept of bubble, as a coalition or group of smart objects that can be created according to the relationship between their owners. The proposed approach is based on the use of attribute-based encryption to protect the associated data according to users' preferences, and FI-WARE components for deployment purposes. As a scenario example, the solution is integrated with a radio localization system, in order to protect location data in the context of smart buildings. Finally, this work provides implementation details about the required components, as well as their evaluation on real smart environment scenarios.

1. Introduction

The inclusion of *Information and Communications Technology* (ICT) in our everyday environments has motivated the development of different initiatives to encourage the deployment of new services and applications in the scope of Smart Cities [1]. Much of the success of these scenarios is based on the continuous data sharing from a huge amount of heterogeneous data sources, such as smartphones, transport infrastructures or devices physically deployed in our surrounding environment. These devices are currently enabled to share their sensed data to platforms that extract knowledge from that information, in order to provide more sustainable and efficient cities to citizens.

In this sense, trust perception of users will be key for the success of new services and applications. The volume and sensitivity of the information to be exchanged will require effective and automated mechanisms to make users aware about how their data are shared. Given the huge number of devices producing data, as well as the number of services and citizens that will make use of this information, the application of scalable mechanisms to protect the access to shared information is crucial. Furthermore, such mechanisms should make use of cryptographic approaches that do not require complex key management tasks. Indeed, protecting the access to highly sensitive information is a challenging aspect in scenarios where a large number of heterogeneous devices need to interact, and where different actors are involved with conflicting interests. On the one hand, companies require a huge amount

of information, in order to provide rich experiences and customized services to users. On the other hand, for citizens, maintaining the control on how their data are shared and under what circumstances is a difficult goal to be accomplished.

In order to address these challenges, this work presents the concept of *bubble*, as well as its design and implementation by using specific technologies and mechanisms. A bubble can be defined as a group of smart objects whose associated data are shared under a set of security restrictions according to users' preferences. The proposed approach is based on the use of *Attribute-Based Encryption* (ABE) [2], which provides a high level of flexibility and scalability to address some of the requirements for providing confidentiality to shared data. Specifically, we have made use of the *Ciphertext-Policy Attribute-Based Encryption* (CP-ABE) scheme [3], in which each piece of information is encrypted under a certain logic combination (or policy) of attributes, while a private key is associated with a certain set of these attributes. Thus, an entity is able to decrypt a certain ciphertext if its key satisfies the policy that was used to encrypt it. CP-ABE provides two main properties to be leveraged in these scenarios. On the one hand, unlike symmetric key cryptography, it does not require entities to share cryptographic keys, which could be unfeasible given the huge number of data producers and consumers. On the other hand, it is seamlessly applicable to maintain confidentiality in one-to-many configurations (even in the presence of a central platform), since a piece of information can be encrypted and

* Corresponding author.

E-mail address: jose-luis.hernandez-ramos@ec.europa.eu (J.L. Hernández-Ramos).

shared with groups of entities using a single message. The resulting approach has been deployed on the FI-WARE platform¹ through the integration with the *Orion Context Broker Generic Enabler* (GE),² which is intended to serve as a central platform for sharing information among different entities, and an extension of the IdM Keyrock proposed in [4]. The integration of the CP-ABE scheme into this platform allows to leverage the features already offered by existing GEs, such as the use of the publish/subscribe pattern, and the *OMA Next Generation Services Interface* (OMA-NGSI) [5] as the data format to represent information in such sharing ecosystem. Furthermore, the performance of the resulting approach has been evaluated on a real scenario by considering different software components and practical evaluation aspects. Based on these aspects, the main contributions of this work are:

- Definition of *bubble* as a concept to represent groups of entities, whose data are protected under different combination of attributes by considering potential relationships between their owners.
- Use of CP-ABE to enforce data protection, and integration with existing IdM schemes and tools to foster the realization of an attribute-based data sharing ecosystem.
- Implementation and integration with current and enriched FI-WARE components to promote compatibility and interoperability with existing IoT deployments.
- Implementation of open-source prototypes³ as part of the EU FP7 project SOCIOTAL,⁴ as well as an exhaustive performance evaluation by using different settings and components.

The remainder of this paper is organized as follows: Section 2 presents a description of some works related to our proposal. Section 3 describes the main building blocks of our approach. An overview of the required architecture and main interactions among the identified components is given in Section 4. Such overview is detailed in Section 5 where some usage examples are also provided. The integration of the resulting approach with a radio localization system is described in Section 6 to show the applicability of the proposal on a real scenario. Section 7 describes the implementation details of the required components, and a exhaustive performance evaluation is given in Section 8. Finally, Section 9 concludes the paper with an outlook of our future work in this area.

2. Related work

In recent years, the security challenges and requirements that are derived from common scenarios and use cases in smart cities have attracted significant interest from academia [6]. The huge volume and continuous data exchange, mainly fostered by the strong development of IoT, as well as the sensitivity of the information being exchanged, make users' privacy more difficult to be preserved. Indeed, data protection represents one of the main cornerstones to cope with security and privacy concerns in IoT scenarios [7,8]. This is also significantly reflected by the application of the *General Data Protection Regulation* (GDPR)⁵ as the main legal framework governing data protection aspects in Europe.

In common IoT-enabled scenarios, physically deployed devices (e.g. a sensor) senses and processes data about its surrounding environment, which are shared with groups of users or applications through central data platform. This component is additionally intended to apply data mining techniques [9] to extract knowledge from the information that is transmitted from data sources. However, since it will have access to all the information, this platform becomes an obstacle to ensure that only authorized entities are able to access the information being

shared, which requires to move the enforcement of security preferences closer to data sources. In this sense, the application of flexible and scalable encryption schemes is crucial to ensure only authorized users and applications will have access to shared data.

In parallel with the strong development of data-driven smart cities, *Attribute-Based Encryption* (ABE) [2] has received notable attention, due to its flexibility to provide confidentiality on data outsourcing scenarios, and their advantages compared to the use of symmetric key cryptography [10]. ABE represents the generalization of *Identity-Based Encryption* (IBE), in which users' identities are represented as a set of attributes (e.g. nationality, gender or address). Based on ABE, two alternative schemes were proposed. On the one hand, in the *Key-Policy Attribute-Based Encryption* (KP-ABE) [11], a piece of data is encrypted with a set or list of attributes, while private keys are associated with logical combinations of attributes. On the other hand, in a *Ciphertext-Policy Attribute-Based Encryption* (CP-ABE) scheme [3], data are encrypted under a policy of attributes, while private keys of participants are associated with sets of attributes. This way, the entity responsible for outsourcing a certain piece of information can exercise control over how their data are shared by specifying the combination of attributes that must be satisfied by intended receivers. Currently, there are two main lines of research related to the use of CP-ABE. Firstly, the direct application of CP-ABE presents some practical challenges related to the size of the ciphertexts and keys and the use of expensive cryptographic primitives that are required by the scheme. These limitations have led to the emergence of different theoretical works [12,13], in order to provide more lightweight, expressive and practical alternatives. Secondly, the advantages of CP-ABE are being exploited in different use cases related to the Cloud Computing paradigm [14–16]. These approaches partially fulfill some of the security needs to be addressed in IoT scenarios where information needs to be shared through central data platforms. Furthermore, they do not consider the inclusion of IoT devices acting as data sources, which are usually deployed on smart cities or smart buildings environments.

Regarding the application of the CP-ABE scheme in scenarios where IoT devices are involved, there are some recent proposals that try to accommodate the needs of the scheme to exploit its benefits. The approach proposed in [17] presents a scheme in which IoT devices delegate expensive CP-ABE cryptographic operations in a more powerful device, called assisting node, which is assumed to be trusted. The results of these operations are given to the sensor, which is responsible for computing the ciphertext. The work presented in [18] offers a lightweight KP-ABE scheme to be used in constrained devices. Its main contribution is that the proposed scheme is not based on bilinear pairings, but on the use of *Elliptic Curve Cryptography* (ECC) [19]. Authors in [20] present AGREE, an approach to exploit energy harvesting opportunities to pre-compute and cache suitably chosen CP-ABE-encrypted keys. While their work is focused on WSNs, it could be also considered in the context of IoT devices with tight resource constraints. Moreover, [21] and [22] proposed the integration of the CP-ABE scheme with the *Message Queuing Telemetry Transport* (MQTT) protocol [23], in order to exploit the advantages provided by the publish/subscribe communication pattern. Furthermore, the work proposed by [24] addresses the size of CP-ABE keys, as an aspect limiting its applicability on IoT devices. Towards this end, it proposes a CP-ABE scheme with constant-size secret keys independent of the number of attributes. Additionally, [25] and [26] provide comprehensive performance evaluations of the main CP-ABE operations, by considering different hardware and security levels.

These proposals provide different insights about the application of CP-ABE on IoT-enabled scenarios. However, while these works partially address some of the aspects derived from this integration, the proposed work is intended to provide a more comprehensive and practical approach through the integration with different technologies to realize and implement the concept of *bubble*. Firstly, unlike previous works, the proposed approach is based on our IdM system [4] by considering the inclusion of the System for Cross-domain Identity Management (SCIM)

¹ <http://www.fiware.org/>.

² <http://catalogue.fiware.org/enablers/publishsubscribe-context-broker-orion-context-broker>.

³ <https://github.com/sociotal?tab=repositories>.

⁴ <http://sociotal.eu/>.

⁵ https://ec.europa.eu/commission/priorities/justice-and-fundamental-rights/data-protection/2018-reform-eu-data-protection-rules_en.

standard [27] that provides a common user schema, extension model and well-defined REST API. In this work, the SCIM standard is used to define users' attributes, so that CP-ABE keys, and consequently the definition of bubbles are based on these values. Secondly, the concept of bubble has been defined by considering the OMA NGSI information model [28] to represent it. OMA NGSI specifications provide a common data model and powerful interfaces for finding information, and it is widely used on current IoT deployments. Thirdly, the resulting approach has been deployed on a real testbed and integrated into the FI-WARE platform, in order to leverage the benefits of a CP-ABE based approach on an already established sharing ecosystem. It should be noted that, in contrast to [29] where security aspects are enforced by the central platform, our proposal moves the enforcement of security aspects to the end devices, in order to ensure end-to-end data protection. Such approach could be also exploited by considering the main notions from the so-called *Social Internet of Things* (SIoT) [30], so different relationships between devices' owners can drive the way in which data from these devices are protected. These aspects are further explained in the following sections.

3. Preliminaries

Before the description of the proposed approach, this section gives an overview of the main concepts and mechanisms that have been used in this work.

3.1. SocIoTal bubbles

In order to address scalability requirements in IoT, many current scenarios are based on the massive sharing of data among groups of devices, users and services. These scenarios typically make use of infrastructures that allow communicating entities to remain decoupled. However, the application of security mechanisms is a challenging aspect, specially when the entities involved may be unknown to each other. To address these needs, the SocIoTal project explores the impact of the relationships between people and their smart objects, in order to establish groups or coalitions where information is shared following certain security conditions. In particular, we consider the relationships defined in [30], as a starting point to define security restrictions for a controlled data sharing.

To deal with the dynamism and ephemeral nature of relationships among people and their associated smart objects, SocIoTal envisions the use of two different kinds of groups. On the one hand, a *Community* [29] is defined as a group of users and smart objects that is statically created and managed. Consequently, the user responsible of the bubble has to manage explicitly the joining and leaving actions of other users. On the other hand, a *Bubble* is considered as a group of smart objects whose associated data are shared under a set of security restrictions based on attributes. This way, users are *implicit members* of the bubble as a consequence of satisfying a particular combination of attributes. These bubbles can be created according to preferences that are specified by the owners of such devices, stating different constraints about the relationships that can be established by their smart objects with other entities. As an example, Fig. 1 shows a high-level representation of a bubble. In this case, Alice and Bob maintain a *Ownership relationship* with their corresponding devices, so each user and his/her devices make up an *Ownership bubble*. Furthermore, Alice and Bob work in the same organization, so they also keep a *Co-work relationship*. As a consequence of this relationship, Alice's car is accessible to Bob, while Bob's video camera is also available to Alice. Then, these devices and their owners state a *Co-work relationship enabled bubble*.

As already mentioned, in order to realize the bubbles' foundations, we propose the use of the CP-ABE scheme. Thus, bubble's security restrictions are defined by using specific combinations or policies of attributes. Therefore, users do not need to be explicitly associated to a bubble; they are implicitly linked because of their attributes. In this sense, only users satisfying such policies will be able to access data from the smart objects of a particular bubble.

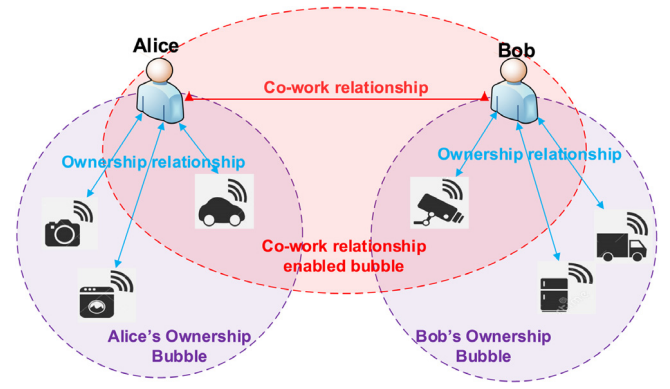


Fig. 1. Bubble high-level representation.

3.2. Ciphertext-policy attribute-based encryption

While symmetric key cryptography has been considered as cryptographic solution on emerging IoT scenarios [31], the cumbersome key management and distribution of these approaches hamper a suitable level of scalability to be provided. In contrast, through the use of traditional certificate-based public key cryptography, information can be encrypted to be accessible to a single specific receiver. These limitations require the use of more sophisticated cryptographic approaches in order to provide the bubbles' security functionality previously described.

In this sense, *Identity-Based Encryption* (IBE) [32] is an alternative scheme, in which the concept of identity is not determined by a public key, but a string. Therefore, an entity could make a specific information only accessible to other users whose identity is described by a specific string. Based on this, *Attribute-Based Encryption* (ABE) [2] represents the generalization of IBE, in which the identity is not represented by a single string, but by a set of attributes related to their identity. In ABE, a piece of data can be made accessible to a set of entities whose real, probably unknown identity, is based on a certain set of attributes.

By using ABE foundations, two alternative approaches were proposed. In KP-ABE, data are encrypted under a set of attributes, while secret keys are associated with combinations or policies of attributes. In contrast, in CP-ABE, a piece of data is encrypted under a policy of attributes, while keys are associated with sets of attributes. Therefore, CP-ABE could be seen as a more intuitive way to apply the concepts of ABE, in which an entity can exert greater control over how the information is disseminated to other entities. The CP-ABE scheme [3] consists of four main algorithms:

- **Setup** ($\lambda \rightarrow \{PP, MSK\}$). It takes an implicit security parameter λ as input to generate the public parameters PP that are common to all users of the system (for example, the universe of attributes U), as well as a master secret key MSK , which is used to generate secret keys.
- **KeyGeneration** ($\{MSK, A\} \rightarrow SK_A$). It takes as input the master key MSK and the set A . The result is a private key SK_A .
- **Encrypt** ($\{PP, M, PT\} \rightarrow CT$). It takes the message M , public parameters PP and a decryption policy PT representing subsets of attributes that are allowed to decrypt M . The result of this algorithm is a ciphertext CT containing PT .
- **Decrypt** ($\{PP, CT, SK_A\} \rightarrow M$). The decryption algorithm takes as input the public parameters PP , the ciphertext CT with a PT associated, and a private key SK_A . If the set A satisfies the policy PT , the corresponding entity will be able to decrypt CT with SK_A .

These algorithms are employed to realize the security features of the bubble concept. Section 5 provides a detailed description about the use of CP-ABE within the proposed approach.

3.3. FI-WARE

FI-WARE⁶ is a European middleware platform that promotes the development and global deployment of applications for the Future Internet. FI-WARE delivers a reference architecture, as well as the specification and implementation of different open interfaces called *Generic Enablers* (GEs). Among the FI-WARE GEs, the *Identity Management GE* relies on standard protocols, such as the *Security Assertion Markup Language* (SAML) [33] and OAuth [34], to provide authentication and authorization features, which allows to manage users' access to networks, services and applications. The IdM GE is also responsible for the user profile management, as well as Single Sign-On (SSO) and identity federation across different service domains. Keyrock IdM⁷ is an open source implementation of the FI-WARE's IdM GE, which, in turn, relies on Keystone,⁸ providing further functionality to support the SCIM standard. SCIM is intended to reduce the cost and complexity of user management operations through a common user schema, extension model and REST API with a rich, but simple set of operations. Keyrock has been leveraged in the scope of SocIoTal project to enable the management of users and smart objects' identities, so that security credentials can be associated with the IoT identity attributes defined within Keyrock. Indeed, as will be described, the CP-ABE keys employed within the proposed approach are generated according to these attributes, which in turn, are employed to define the security restrictions for a bubble. For further information about the privacy-preserving IdM system for IoT used in our proposal, the reader is referred to [4].

In addition, the FiWare *Publish/Subscribe Context Broker GE* allows publication of context information by entities, referred as *Context Producers*, so that published context information becomes available to other entities, referred as *Context Consumers*, which are interested in processing the published context information. The Orion Context Broker⁹ is an implementation of this GE, providing the OMA NGSI-9 and NGSI-10 interfaces [5] to access to the broker. The Context Manager component of SocIoTal architecture [35] is a fork of the Orion Context Broker, which has been extended with new functionality such as context management and access control. Furthermore, as described in the following section, the definition of bubbles follows the OMA NGSI Context Management Information Model [36], in order to provide an interoperable approach with existing solutions and deployments.

Based on these three main building blocks, following sections provide a detailed description of the proposed approach.

4. Bubbles scenario architecture and overall functionality

The bubble is the core concept to enable a dynamic and secure data sharing among different users. As already mentioned, a bubble represents a group of devices whose associated data are shared under specific attributes. Before starting the sharing information process within a bubble, it must be modeled and registered so other devices can be added and shared with other users. For this purpose, bubbles are modeled by following the OMA NGSI Context Information Model [36], so they are represented as a *Context Entity* that is intended to group several devices, which are modeled in the same way.

Listing 1: Bubble definition

```
1 {
2   'entities': [{
3     'id': 'bubbleID',
4     'type': 'http://sociotal.namespace.bubble',
5     'isPattern': 'false'
6   }],
```

```
7   'attributes': [{
8     'name': 'BubbleEntity',
9     'value': ['bubbleEntityID'],
10    'type': 'http://sociotal.namespace.bubble.bubbleEntity'
11  }],
12  {
13    'name': 'attributeName',
14    'value': 'attributeValue',
15    'type': 'attributeType'
16  },
17  'metadata': [{
18    'name': 'BubbleRequired',
19    'value': 'true',
20    'type': 'http://sociotal.namespace.bubble.required'
21  }]
22 }
23
24
25
26 }
```

Based on this, Listing 1 shows the main scheme to define a bubble:

- **entities:** the entities (bubbles) to be registered. It includes the bubble's id, the type (<http://sociotal.namespace.bubble>) and the *isPattern* field (lines 3–5), which is always stated to *false* when registering a new bubble.
- **attributes:** it makes reference to a list of attributes related to the bubble(s) to be registered by specifying their name, type and value. In the Listing 1, the attribute *BubbleEntity*, type <http://sociotal.namespace.bubble.bubbleEntity>, is providing the list of entities' id that are part of this bubble (lines 8–10). In addition to specify the list of entities that are composing the bubble, it is necessary to define which attributes are required to get access to the information provided by the bubble's entities (lines 13–21). For this purpose, such attributes include a set of metadata (lines 16–20) that indicate the <http://sociotal.namespace.bubble.required> type. For these attributes, by following the SCIM schema, the list of supported types is:

- urn:ietf:params:scim:schemas:core:2.0:organization
- urn:ietf:params:scim:schemas:core:2.0:department
- urn:ietf:params:scim:schemas:core:2.0:streetAddress
- urn:ietf:params:scim:schemas:core:2.0:locality
- urn:ietf:params:scim:schemas:core:2.0:postalCode
- urn:ietf:params:scim:schemas:core:2.0:country.

Furthermore, it should be noted that other attributes can be employed to restrict how information is to be shared within a bubble.

Based on the previous definition of *bubble* and the approaches described in Section 3, Fig. 2 shows an overview of the proposed scenario. In particular, as already mentioned, the CP-ABE cryptographic scheme has been employed to manage the security aspects associated to the creation and management of bubbles. Furthermore, in the scope of the SocIoTal project, some of the FIWARE GEs have been integrated and extended to realize the intended functionality.

For the sake of clarity, we adopt a producer/consumer approach in which a user or device can act as a data *producer* or *consumer* at any time [37]. Furthermore, the *SocIoTal Platform* is intended to embrace different infrastructure components to cope with key management and data sharing procedures. In particular, the main identified components are:

- **Web User Environment (WUE):** it represents the central point where users define their bubbles through a user-friendly interface.
- **IdM Service:** it acts as a repository of users that are previously registered. In particular, it integrates the FI-WARE Keyrock service, which has been extended to support additional attributes based on the SCIM standard following our previous work [4].
- **CP-ABE Key Manager:** it is responsible for generating the cryptographic material associated to the CP-ABE scheme. Specifically, it generates CP-ABE keys associated to the attributes from the IdM Service for a particular user.
- **Context Manager:** it is the central point where devices and bubbles are registered to enable a secure sharing of devices' data among different users.

⁶ <https://catalogue.fiware.org>.

⁷ <https://catalogue.fiware.org/enablers/identity-management-keyrock>.

⁸ <https://docs.openstack.org/developer/keystone/>.

⁹ <https://catalogue.fiware.org/enablers/publishsubscribe-context-broker-orion-context-broker>.

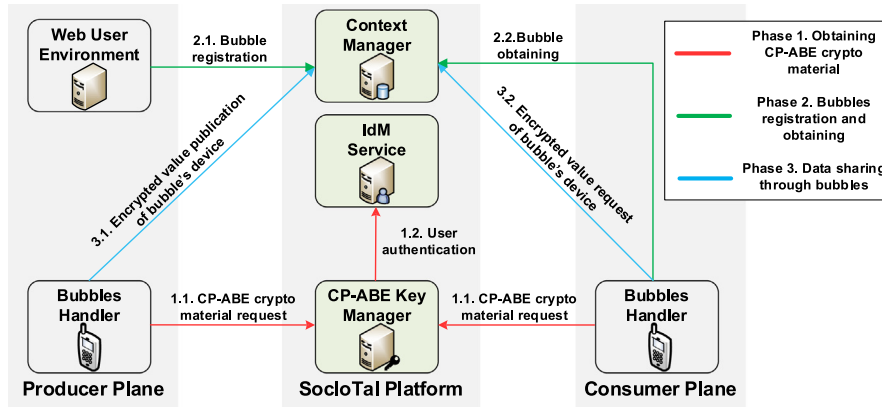


Fig. 2. SocioTal bubbles registration, discovery and operation overview.

- **Bubbles Handler:** it represents a user application to get registered bubbles and to share devices' data by using CP-ABE encryption and decryption algorithms.

In addition to the main components to realize the bubbles approach, previous figure shows an overview of the main interactions among such components for that. Firstly, during the *Phase 1. Obtaining CP-ABE crypto material*, a user employs her *Bubbles Handler* to request a CP-ABE key and the public parameters that are required to encryption and decryption algorithms (step 1.1). Then, the *CP-ABE Key Manager* authenticates the user through the *IdM Service* to get the attributes associated to her (step 1.2). These attributes are used to generate the corresponding CP-ABE key for that user that is delivered together the public parameters. It should be pointed out that this process is only performed once (or until the cryptographic material is revoked), and it is independent to the definition of different bubbles, so key management aspects are simplified. Indeed, changing a bubble's policy does not require the revocation or generation of new cryptographic material.

During the *Phase 2. Bubbles registration and obtaining*, a user employs the *WUE* to register bubbles (step 2.1) by specifying a set of devices to be added and the attributes that are required to share within that bubble. This bubble is registered in the *Context Manager* by following the data model previously described. Then, other users, through the *Bubbles Handler*, try to discover bubbles in the *Context Manager* (step 2.2) by using NGSi9-10 interfaces (e.g. via *queryContext* method).

After previous processes have been completed, in the *Phase 3. Data sharing through bubbles*, a user updates the value of an entity of such bubble, by making use of CP-ABE indicating the policy that is employed to encrypt such value. This policy makes reference to the attributes that are specified in the bubble definition. For example, a CP-ABE policy could be represented as “organization=umu AND department=diic”. The entity is updated with the encrypted value in the *Context Manager* by using NGSi-9/10 interfaces via the *updateContext* method (step 3.1). Finally, a user can get the new encrypted value from the *Context Manager* through the *queryContext* operation; in this case, only those users who satisfy the CP-ABE policy associated to that bubble will be able to decrypt the updated value (step 3.2).

5. Bubbles scenario interactions

After providing a general description of the proposed scenario, this section is intended to delve into the 3 main phases previously defined. Fig. 3 provides a more detailed overview of the required interactions for each phase that are explained below.

5.1. Obtaining CP-ABE crypto material

As described in the previous section, initially the users try to obtain the required cryptographic material to protect their data within a bubble. For this purpose, while it is not shown in the figure, two previous processes are assumed. Firstly, the *CP-ABE Key Manager* has already previously run the *Setup* algorithm (see Section 3.2), so the public parameters *PP* and the master key *MSK* have already been generated. It should be noted that this operation is only required initially to setup the CP-ABE ecosystem. Secondly, users have already been registered in the *IdM Service*; therefore, this service has a set of attributes for each user. For registration purposes, a user can employ the *WUE* as described in [29].

Following the previous notation and Fig. 3, we consider two users acting as a data *Producer P* and a data *Consumer C*. It should be noted the steps associated to the first phase are repeated by each entity. According to it, a user makes usage of the *Bubble Handler* that sends a request to the endpoint where the *CP-ABE Key Manager* accepts requests for generating cryptographic material. This request includes an authentication credential (P's/C's credential) (step 1.1) to ensure that only legitimate users will be able to obtain the appropriate keys. Then, the *CP-ABE Key Manager* queries the *IdM Service* in order to authenticate the requesting user (step 1.2). As already mentioned, the *IdM Service* integrates an extended version of the FIWARE Keyrock IdM that includes support for additional attributes following the SCIM standard. Furthermore, this service supports different authentication mechanisms (step 1.3) including login/password and certificate-based procedures, so that the credential could be instantiated by considering different approaches. It should be noted that this authentication process could be also performed in a privacy-preserving way based on the use of Idemix [38] by following our previous proposal in [4]. In case of a successful authentication process, the *IdM Service* returns the set of user attributes (π/γ) to the *CP-ABE Key Manager* (step 1.4). Then, this entity runs the *KeyGeneration* algorithm (see Section 3.2) to generate a new CP-ABE key that is associated to those attributes. Consequently, the CP-ABE keys for P and C are generated as (step 1.5):

$$KeyGeneration(\{MSK, \pi\} \rightarrow SK_{\pi})$$

$$KeyGeneration(\{MSK, \gamma\} \rightarrow SK_{\gamma})$$

Then, each CP-ABE key is sent together with the public parameters *PP* in the response to the requesting user (step 1.6). For this, both elements are base64-encoded and included in a JSON object following the example in Listing 3. In addition to the CP-ABE key and public parameters, the set of attributes associated to the key is also included in the response, so that users can easily visualize their own attributes.

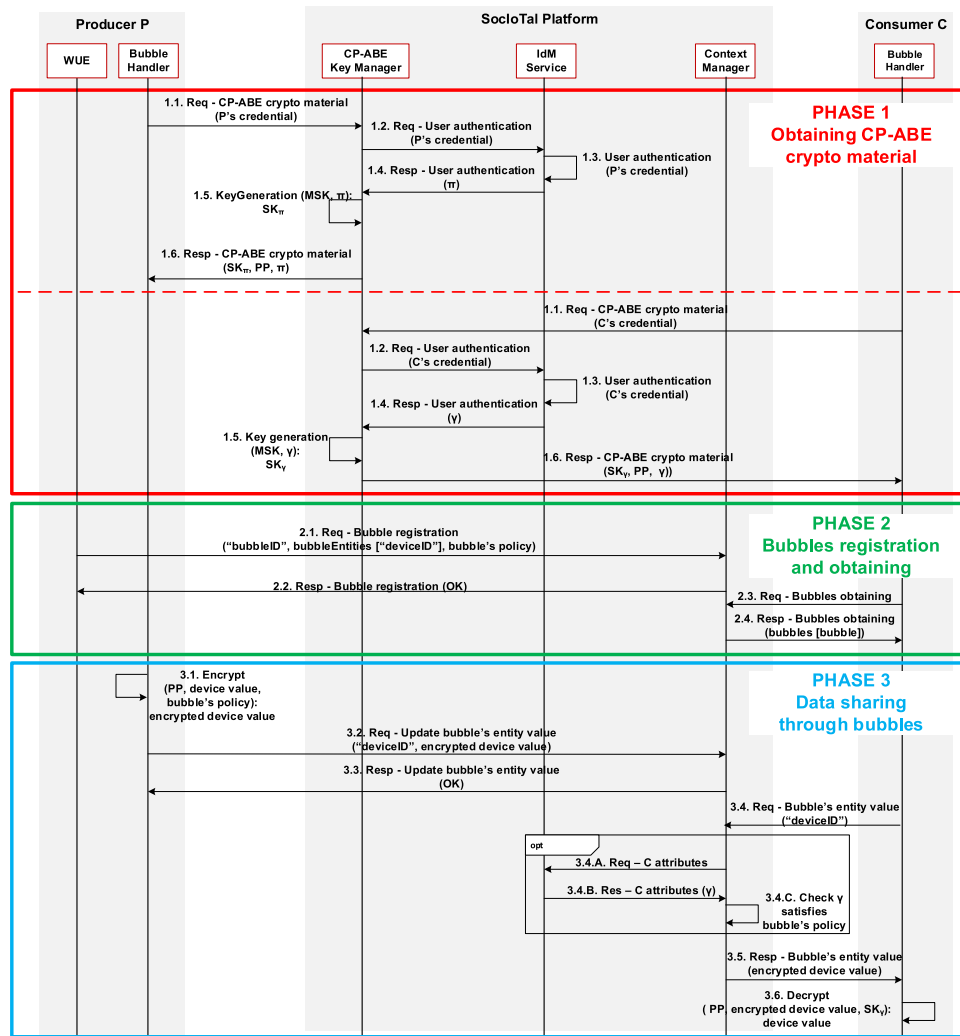


Fig. 3. SocIoTal bubbles registration, discovery and operation overview.

Listing 2: CP-ABE Key Manager response example

```

1 {
2   "cpabekey": "xNlksSd_13ndwe9...",
3   "publicparam": "zcsDko32urnlAdanx...",
4   "attributes": [
5     {
6       "name": "organization",
7       "value": "umu",
8       "type": "urn:ietf:params:scim:schemas:core:2.0:organization"
9     }
10  ],
11  ...
12  ...
13  ]
14 }

```

As mentioned above, this initial phase is only performed when a new key or cryptographic parameters are required (for example, because they have expired). It should be noted that unlike symmetric cryptography based approaches, only public parameters need to be shared by all users. In this sense, the definition or elimination of a bubble does not require the generation or revocation of any cryptographic material, so the aspects of key management are significantly reduced.

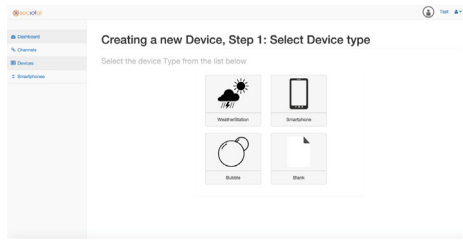
It should be noted that this initial phase is also required in case the cryptographic material is revoked. According to our scenario, CP-ABE keys are associated to the attributes stored in the IdM Service. Therefore, when a user's attribute (for example, her organization) changes, the associated CP-ABE key is not valid anymore. Although the revocation aspects do not represent the focus of our work, different approaches can be considered to address this issue. On the one hand, the use of techniques based on proxy re-encryption [39] has been traditionally used for this purpose, but it requires the definition of new components

in the architecture and an expensive process to update the key of each participant. On the other hand, the use of expiration time attributes can help to mitigate the revocation issue, as suggested by [3]. Indeed, in our previous work, the use of expiration time is associated with symmetric keys that are distributed by using CP-ABE [40]. However, in addition to practical considerations about which amount of time is suitable, it does not address the problem associated to an explicit change of a user's attribute. In our case, considering the proposed architecture, a potential approach is the use of the IdM Service to check the current attributes of a user. In this way, each time a user queries the Context Manager, this entity checks the attributes of the requesting user, as will be further described in Section 5.3.

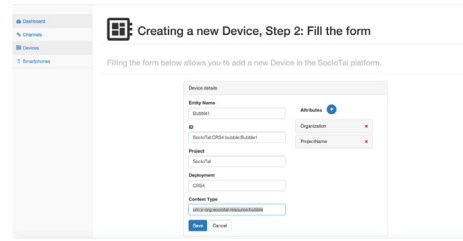
5.2. Bubbles registration and obtaining

When a user has already completed the previous phase, she can use the WUE for the creation and definition of a bubble, associating a set of devices and the required attributes to access devices' data. To do this, the WUE provides a friendly user interface so that a user is able to create a bubble in the same way she registers other devices (e.g. her weather station or smartphone).

Fig. 4 shows an example of WUE screenshots in which a user is enabled to create and define a bubble. As shown in Fig. 4(b), the user can define different details for the bubble (e.g., name, ID or type), as well as the attributes that will be required to get access to the devices data included in the bubble.



(a) SocIoTal bubbles creation



(b) SocIoTal bubbles definition

Fig. 4. WUE examples for creating and defining a bubble.

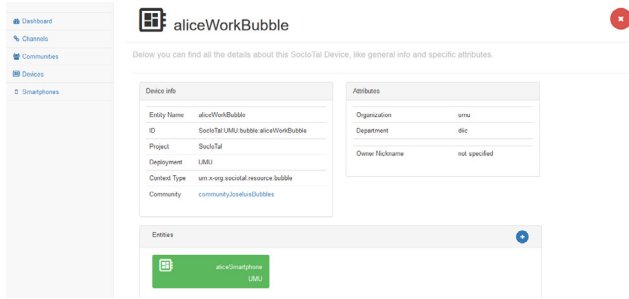


Fig. 5. Bubble definition example view at the WUE.

After creating a bubble, the user can associate her devices (or some of them) to it. Fig. 5 shows an example of a bubble created by the user “Alice” (acting as the *Producer P*) with the name “aliceWorkBubble”. In this case, it is assumed that she has previously registered her smartphone with the name “aliceSmartphone” by using the WUE (see Fig. 4(a)). In addition to the bubble details, such as ID or type, Alice has associated her smartphone with the new bubble. She has also specified the *Organization=umu* and *Department=diic* attributes as the attributes required to access the information generated by devices associated to such bubble. Consequently, all the information provided by the *AliceSmartphone* entity will be accessible only for users who satisfy those attributes.

Following this example, Listing 3 shows the bubble definition that is included as the payload of a NGSI-9 *registerContext* (or a NGSI-10 *updateContext*). This operation is represented by the step 2.1 in Fig. 3, where *bubbleID* is “aliceWorkBubble”, the list of *bubbleEntities* is only composed by “AliceSmartphone”, and the *bubble’s policy* is represented by the conjunction of attributes that are marked with type “<http://sociotal.namespace.bubble.required>” (i.e., *organization=umu AND department=DIIC*)

Listing 3: Bubble registration example

```

1  "contextRegistrations": [{
2    "entities": [{
3      "id": "AliceWorkBubble",
4      "type": "http://sociotal.namespace.bubble",
5      "isPattern": "false"
6    }],
7    "attributes": [{
8      "name": "BubbleEntity",
9      "value": ["AliceSmartphone"],
10     "type": "http://sociotal.namespace.bubble.bubbleEntity"
11   }],
12   {
13     "name": "organization",
14     "value": "umu",
15     "type": "urn:ietf:params:scim:schemas:core:2.0:organization",
16     "metadata": [{
17       "name": "BubbleRequired",
18       "value": "true",
19       "type": "http://sociotal.namespace.bubble.required"
20     }],
21   },
22   {
23     "name": "department",
24     "value": "diic",
25     "type": "urn:ietf:params:scim:schemas:core:2.0:department",
26     "metadata": [{
27       "name": "BubbleRequired",

```

```

28   "value": "true",
29   "type": "http://sociotal.namespace.bubble.required"
30 }
31 ]
32 }

```

In this case, based on Section 3.1, we assume that Alice has a *Ownership relationship* with “aliceSmartphone”, since she is her owner. The definition of the attributes “*Organization=umu*” and “*Department=diic*” can be seen as the specification of a *Co-work relationship*, so that users working within the *diic* department at *umu* will be able to access the data generated by Alice’s smartphone.

Listing 4: Bubbles query message example

```

1  {
2    "entities": [
3      {
4        "type": "urn:x-org:sociotal:resource:bubble",
5        "isPattern": "true",
6        "id": "x"
7      }
8    ],
9    "attributes": []
10 }

```

When the bubble is registered, the Context Manager sends a message indicating the process was successfully performed (step 2.2). Then, a user acting as the *Consumer C*, accesses the Context Manager in order to obtain information about the registered bubbles. Listing 4 shows an example of the payload contained within a OMA NGSI-10 *queryContext* method to get all entities of type “urn:x-org:sociotal:resource:bubble” (step 2.3).

Listing 5: Bubbles response message example

```

1  { "contextResponses": [
2    {
3      "statusCode": {
4        "reasonPhrase": "OK",
5        "code": "200"
6      },
7      "contextElement": {
8        "id": "SocioTal:UMU:bubble:aliceWorkBubble",
9        "attributes": [
10         {
11           "name": "Organization",
12           "value": "umu",
13           "type": "http://sensorml.com/ont/swe/property/OrganizationName",
14           "metadata": [
15             {
16               "name": "BubbleEntity",
17               "value": "AliceSmartphone",
18               "type": "http://sociotal.namespace.bubble.bubbleEntity"
19             }
20           ]
21         },
22         {
23           "name": "Department",
24           "value": "diic",
25           "type": "urn:ietf:params:scim:schemas:core:2.0:department",
26           "metadata": [
27             {
28               "name": "BubbleEntity",
29               "value": "AliceSmartphone",
30               "type": "http://sociotal.namespace.bubble.bubbleEntity"
31             }
32           ]
33         }
34       ]
35     }
36   ]
37 }

```

```

41 }
42 ],
43 "type": "urn:x-org:sociotal:resource:bubble",
44 "isPattern": "false"
45 }
46 }
47 ]

```

Following the example, Listing 5 shows the response to the previous request in which the bubble information “aliceWorkBubble” is obtained (step 2.4). In this way, a user is able to visualize the bubbles registered in the Context Manager, as well as the attributes required to be able to access the data of the associated devices.

5.3. Data sharing through bubbles

Once a user has registered a bubble through the WUE, she can use her Bubble Handler to share protected data by using the Context Manager. In this way, following the previous example, Alice (acting as the *Producer P*) decides to protect the location information that is detected by her smartphone. To do this, the Bubble Handler obtains information about the bubbles registered by Alice in the Context Manager. Since the device “aliceSmartphone” is associated with the bubble “aliceWorkBubble”, the Bubble Handler detects that the device’s data must be protected by using the bubble’s policy (“Organization = umu” AND “Department = diic”). It should be noted that this functionality is not reflected in the sequence diagram, since the Bubble Handler will be usually stored this information, so the interaction with the Context Manager is not required. The Bubble Handler runs the *Encrypt* algorithm (see Section 3.2) to encrypt the location data by using the mentioned policy as (step 3.1):

$Encrypt(\{PP, devicevalue, ("Organization = umu"$
 $AND "Department = diic")\} \rightarrow encrypteddevicevalue)$

where device value is represented by the location data, following the example. Then, the resulting ciphertext *encrypted device value* is then included within an NGSI-10 *updateContext* message, as shown in the Listing 6 that is sent to the Context Manager (step 3.2). This message is in turn answered by the Context Manager indicating the value was successfully updated (step 3.3).

Listing 6: Context entity update message example

```

1 { "contextElements": [
2   {
3     "type": "urn:x-org:sociotal:resource:smartphone",
4     "isPattern": "false",
5     "id": "SocioTal:UMU:smartphone:aliceSmartphone",
6     "attributes": [
7       {
8         "name": "Location",
9         "value": "TzVPV3NCSIVONIVwMHERSUZxbC9PUT09...",
10        "type": "http://sensorml.com/ont/swe/property/Location",
11        "metadata": [
12          {
13            ...
14          }
15        ]
16      }
17    ]
18  },
19 ],
20 "updateAction": "UPDATE"
21 }

```

At this point, it should be noted that even though the information is shared through the Context Manager, this entity does not have access to the information provided by the different entities, since data are encrypted on an end-to-end basis. In fact, Fig. 6 shows the encrypted value corresponding the smartphone’s location when it is queried through the WUE.

After that, a user (acting as a *Consumer C*) tries to access Alice’s position. From the message obtained in the previous phase (see Listing 5), she knows that the “aliceSmartphone” device is associated with the “aliceWorkBubble” bubble. Therefore, she can use his Bubble Handler to obtain the data associated to it by using the NGSI-10 *queryContext* method with the content shown in Listing 7 (step 3.4).

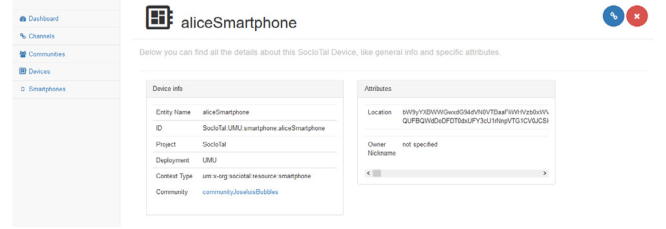


Fig. 6. Alice’s smartphone view at the WUE when location data is encrypted.

Listing 7: Device’s encrypted data query message example

```

1 {
2   "entities": [
3     {
4       "type": "urn:x-org:sociotal:resource:smartphone",
5       "isPattern": "false",
6       "id": "SocioTal:UMU:smartphone:aliceSmartphone"
7     }
8   ],
9   "attributes": []
10 }

```

Considering our architecture, as already mentioned in Section 5.1, we have added optional steps to deal with issues associated to attributes delegation and revocation. In particular, when the Context Manager receives the previous message, it queries the IdM service to get the current attributes of the requesting user (step 3.4.A). Then, the set of attributes (γ) is sent back to the Context Manager (step 3.4.B). Since the policy of a bubble is still accessible to the Context Manager, it can check that the set of attributes satisfy the bubble’s policy (step 3.4.C). Accordingly, if this condition is not satisfied, the Context Manager does not provide the requested data. Instead, it returns an error message, so that the requesting user is redirected to the CP-ABE Key Manager, in order to update her CP-ABE key, so that the *Phase 1* is carried out again. Otherwise, if the verification process is successful, the Context Manager returns the message in Listing 8 in which the location information of the device is encrypted (step 3.5). In fact, the location value corresponds to the output of the CP-ABE encryption that is encoded in base64. Once the above message is obtained, the *Consumer C* uses her Bubble Handler to execute the CP-ABE Decrypt algorithm, (see Section 3.2) by using her private key and parameters obtained in the first phase, as well as the encrypted location value, as follows (step 3.6).

$Decrypt(\{PP, encrypteddevicevalue, SK_{\gamma}\} \rightarrow devicevalue)$

This way, in case the user’s CP-ABE key satisfies the bubble’s policy (“Organization = umu” AND “Department = diic”) she will be able to get the value of *locationData* (i.e., *device value*)

Listing 8: Encrypted entity response example

```

1 { "contextResponses": [
2   {
3     "statusCode": {
4       "reasonPhrase": "OK",
5       "code": "200"
6     },
7     "contextElement": {
8       "id": "SocioTal:UMU:smartphone:aliceSmartphone",
9       "attributes": [
10        {
11          "name": "Location",
12          "value": "TzVPV3NCSIVONIVwMHERSUZxbC9PUT09...",
13          "type": "http://sensorml.com/ont/swe/property/Location",
14          "metadata": [
15            {
16              ...
17            }
18          ]
19        },
20        {
21          "name": "owner",
22          "value": "Alice",
23          "type": "identifier:UUID"
24        }
25      ]
26    },
27    "type": "urn:x-org:sociotal:resource:smartphone",
28    "isPattern": "false"
29  }
30 ],
31 }

```

It should be noted that, even considering the additional interactions to cope with delegation/revocation aspects, two main advantages of the use of CP-ABE are still leveraged. Firstly, while the Context Manager is able to access the information related to a bubble (including the associated policy), it cannot access the data that is shared between users through such bubbles. Secondly, the potential problems associated with an uncontrolled delegation of attributes are also mitigated due to the verification of attributes for each request to the Context Manager. In addition, as already mentioned, changes in the policy of a bubble do not require the revocation or update of the cryptographic material. While it is not the focus of the paper, part of our future work is related to the integration of more efficient revocation schemes [41] that could be potentially analyzed in the context of IoT environments.

6. Use case: sharing location data through bubbles in indoor environments

Following with the example of sharing location data, in this section we describe a scenario in which a user's location data are shared through the proposed approach in a smart building. Indeed, according to the *European Alliance of Companies for Energy Efficiency in Buildings* (EuroACE [42]), we spend over 90% of our time in buildings. Consequently, we consider the uncontrolled disclosure of people's location data in such environments can facilitate tracking tasks that can potentially harm their privacy through the inference of their daily habits. Furthermore, it should be noted that typical GPS-based approaches are not practical in these scenarios due to the lack of signal.

Towards this end, the proposed localization system is intended to be used in indoor environments (e.g. a smart building) and it is based on the use of two different low-level devices. The first ones act as *anchors*, which are static and calibrated. The second ones are mobile low-level devices worn by people. They act as *sinks of Received Signal Strength Indicators* (RSSI) measured at each receiving of a beacon message sent by the anchors. Furthermore, they are connected to the Context Manager via a local gateway through a wireless communication channel. These sinks are able to acquire raw time-stamped location-dependent information (either absolute information or relative information with respect to their neighbors) that can be interpreted further by any localization engine (i.e., in terms of ranging, positioning or position tracking) embedded in the gateway. Thanks to dedicated algorithms, the coordinates of the moving sink nodes are estimated and subsequently transmitted through OMA NSGI-9/10 to the Context Manager where they are stored as contextual information in the broker database. In this scenario, the WSN nodes are IEEE 802.15.4-enabled devices using OpenMote hardware. The embedded application is Contiki-based. The user position is derived from its distance to anchor nodes. Such distance is computed by using RSSI. The user carries a special device called the mobile node, along with a smartphone for graphical feedback. The goal is to display the real-time user's location on a map on the smartphone.

The location process follows these steps:

1. **Collection of the RSSI and computation of the relative distances.** The mobile device broadcasts a beacon packet to its in-range neighbors (anchor nodes represented by blue discs with a signal symbol). The neighbors acknowledge the beacon with a reply packet. The mobile node stores the various RSSIs as they arrive from neighbors.
2. **Deduction of the 2D coordinates.** When all RSSIs are received, they are encapsulated in a message sent to the gateway (GW, green disc) using the existing WSN topology; e.g. in Fig. 7, the route to the gateway starts with node A. The gateway proceeds to compute the estimated location of the device based on the collected RSSIs, using a specific algorithm.

Furthermore, the computed location is pushed to the Context Manager by using OMA NSGI-9/10 messages. Then, the third-party "Location application", hosted in the cloud, receives the updated location of the user/device, and finally, the graphical front-end running on the user's smartphone updates the position (protocol is application-specific).

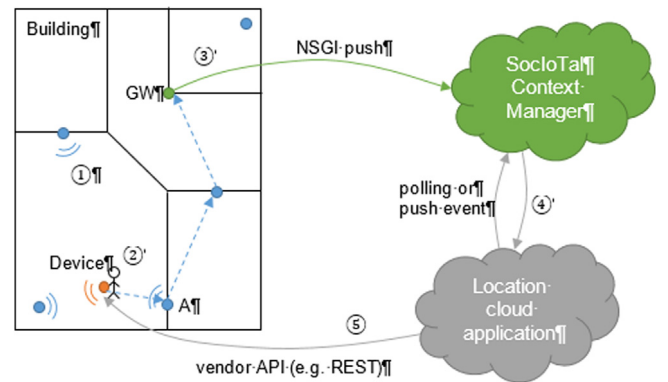


Fig. 7. The radio localization scenario.

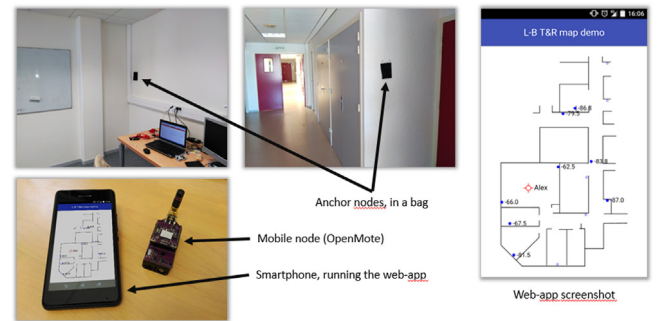


Fig. 8. Deployment of the radio localization demonstrator in an indoor environment.

6.1. Collection of the RSSI and computation of the relative distances

The collection of Received Signal Strength Indicator (RSSI) measurements is deployed in the WSN (see Fig. 8) using the openmote technology. Based on a meshed network topology, the RSSI collected by the mobile sink openmote node are sent to the gateway in order to evaluate the corresponding distance from the sink to the anchors. The radio signal is secured at the MAC layer thanks to the security features provided by the Contiki 3.x embedded OS in the openmotes and the linux 4.4.4 kernel deployed in the local gateway according to the techniques depicted in [43].

Wireless connectivity. Based on the scenario described, N_a fixed anchors are disseminated at known locations in the environment of interest and so-called "blind" mobile nodes to be positioned. In location-enabled IEEE 802.15.4 networks, the ranging functionality usually relies on received power measurements through the RSSI. In terms of RSSI availability, possible radio links depend at least on the power loss suffered over the transmission range and on the receiver power threshold. It is an intrinsic capability of the hardware device. Accordingly, for a given position occupied by the mobile node, the exploitable 1-hop neighborhood $H = \{i/P_{R_{x,i}} > S\}$ is thus defined as a set of reachable anchors, if $P_{R_{x,i}}$ is the received analog power over the link with respect to the i th anchor and S is the equivalent RSSI detection threshold. The entries $\{I_i = 1 \text{ if } i \in H\}$ of the related connectivity (or adjacency) vector directly account for physically feasible RSSI-enabled links.

Received Signal Strength Measurements. On the receiver side, the analog Received Signal Strength (RSS) is generally expressed as follows:

$$P_{R_x}(d) = P_{R_{x,ref}} - 10\alpha \log_{10} [d/d_{ref}] + X_s \quad (1)$$

where $P_{R_{x,ref}}$ is the average power received at the reference distance d_{ref} (generally equal to 1 m), α is the path loss decay exponent and X_s is a random term with standard deviation σ_s . $P_{R_{x,ref}}$ accounts for the

transmit power, transmit and receive antenna gains, and power path loss at the reference distance.

In static scenarios, X_s traditionally accounts for centered Gaussian shadowing. Assuming mobile nodes and/or mobile scatterers, additional small-scale fading (e.g., Rayleigh) is observed around the shadowed received power. However, in most of practical scenarios, these rapid power fluctuations can be significantly mitigated out, and hence neglected, after averaging over several consecutive power measurements so as to draw mostly their deterministic range-dependent component, i.e., considering for ranging $\overline{P_{R_x}}(t) = \frac{1}{K} \sum_{k=1}^K P_{R_x}(k, t)$ instead of $P_{R_x}(k, t)$ at time stamp t . Obviously, the size K of the averaging sequence must be judiciously chosen with respect to the channel coherence time (e.g. [44]), given a mobility scenario (typically, pedestrian in our case).

RSSI-based range measurements and related errors. Based on the previous RSSI measurements collected between pairs of devices, one can determine the corresponding relative distances, given a priori path loss model parameters. Using the median estimator for instance (e.g. [44–46]), which requires only little prior knowledge about channel model parameters but provides an intermediary level of performance, the range measurement is defined as:

$$\tilde{d} = \exp(M) \quad (2)$$

where $M = \frac{(P_{R_x,ref} - \overline{P_{R_x}}) \log(10)}{10\alpha} + \log(d_{ref})$.

In [45], the estimation variance affecting \tilde{d} was shown to be:

$$\sigma_d^2 = d^2 \exp(s^2) [\exp(s^2) - 1] \quad (3)$$

where $s = \frac{\sigma_s \log(10)}{10\alpha}$.

Optionally, the knowledge of ranging variance can be advantageously incorporated within any subsequent positioning step fed by weighted range measurements. Practically, in real-life systems, this variance can be approximated as σ_d^2 by incorporating the estimated distance \tilde{d} instead of the actual distance d in the corresponding formulas (e.g. [45]).

Preferably, due to the highly parametric nature of the retained ranging approach, the path loss model parameters must be a priori calibrated in the operating environment and/or, whenever possible, per anchor [46]. This is for instance obtained empirically through standard data fitting, after collecting received power measurements at known distances. In our case, $P_{R_x,ref} = -62$ and $\alpha = 3.5$ have been determined as optimal parameters.

6.2. Deduction of the 2D coordinates

Positioning algorithm. The considered variant for the implemented positioning algorithm relies on a range-based non-linear *Weighted Least Squares* (WLS) optimization procedure. Following a standard centralized synchronous approach [47], the 2D Cartesian coordinates of a given mobile node are determined as the solutions of the following problem:

$$(\hat{x}_m, \hat{y}_m) = \underset{(\hat{x}, \hat{y})}{\operatorname{argmin}} \sum_{i=1}^{N_a} \frac{I_i}{\sigma_{d,i}^2} (\tilde{d}_i - \hat{d}_i)^2 \quad (4)$$

where \hat{x}_m and \hat{y}_m are the estimated 2D coordinates of the mobile node, $\hat{d}_i = \sqrt{(x_i - \hat{x})^2 + (y_i - \hat{y})^2}$ is the estimated distance between the mobile node and the i th anchor, and I_i represents the force, measurement of the power signal, of the RSSI.

As for optimization solvers, many solutions can be considered, such as Sequential Quadratic Programming (SQP), Linearized Least Squares (LLS) through Taylor Series Expansion (TSE), Gradient Conjugate (GC), Newton-Raphson, or the popular iterative Gradient-Descent (GD) (e.g., [48,47]), etc. In the following, we consider the latter implementation for simplicity.

Accordingly, given an initial guess $(\hat{x}_m(t, 0), \hat{y}_m(t, 0))$ at time stamp t and iteration $j = 0$, one can simply compute the WLS solution for any iteration $j > 0$, as follows:

$$\begin{aligned} \hat{x}_m(t, j) &= \hat{x}_m(t, j-1) + \delta_x(t, j), \\ \hat{y}_m(t, j) &= \hat{y}_m(t, j-1) + \delta_y(t, j). \end{aligned} \quad (5)$$

where

$$\begin{aligned} \delta_x(t, j) &= \sum_{i=1}^{N_a} \beta_i (\tilde{d}_i(t) - \hat{d}_i(t, j-1)) \frac{\hat{x}_m(t, j-1) - x_i(t, j-1)}{\hat{d}_i(t, j-1)}, \\ \delta_y(t, j) &= \sum_{i=1}^{N_a} \beta_i (\tilde{d}_i(t) - \hat{d}_i(t, j-1)) \frac{\hat{y}_m(t, j-1) - y_i(t, j-1)}{\hat{d}_i(t, j-1)}. \end{aligned} \quad (6)$$

with β_i a scaling descent coefficient, for instance proportional to $1/\tilde{\sigma}_{d,i}$ or even more simply equal to 0.1 in our case and $\hat{d}_i(t, j-1)$ the latest estimated distance equal to:

$$\sqrt{(x_i - \hat{x}_m(t, j-1))^2 + (y_i - \hat{y}_m(t, j-1))^2}$$

Different stopping rules can be applied, either based on a maximum tolerated number of iterations (e.g., under constrained computational complexity) or if the norm of the descent progress does not exceed an a priori threshold (e.g., given an a priori spatial resolution to express the solution).

As for initialization, one can use the weighted centroid of all available anchors or alternatively, an arbitrary subset of anchors leading to the strongest RSSI readings.

6.3. Integrating the radio localization system into the bubbles approach

In order to describe the integrated scenario by using the bubbles concept, Fig. 9 shows the sequence diagram with the required steps to share location data by using the proposed localization system. Following with our example, Alice wishes to share her location data obtained from the proposed system to those users in the building that satisfy certain identity attributes. For this purpose, she makes use of a Mobile Node that sends the received RSSIs to the Gateway deployed in the building, so that Alice's position can be calculated. Then, this information is shared through the Context Manager by making use of the bubble concept.

At this point, it should be pointed out that we have considered two alternatives of the main scenario in which the Bubbles Handler can delegate the CP-ABE encryption operation to the platform. To do this, we have considered the addition of a new component, named *CP-ABE Assistant*, which is responsible for encrypting the information before sending to the Context Manager. As already demonstrated in [26,25], CP-ABE requires expensive cryptographic operations related to the use of pairings. Therefore, in certain situations, the delegation of these operations can alleviate the burden of the end nodes, especially in cases where a huge amount of data needs to be protected. Additionally, we have considered that Alice has associated her Mobile Node (named "Sociotat:UMU:mobilenode:aliceMobilenode") to the bubble "AliceWorkBubble" in the Context Manager, by following the process described in Section 5.2.

With these premises, the description of the sequence diagram is as follows. As already described in Section 5.1, during the *Phase 1*, users try to obtain the required CP-ABE cryptographic material from the CP-ABE Key Manager. This stage could be carried out at any time before the following interactions. In this case, we have considered Alice and Bob act as a data *Producer* and *Consumer*, respectively. Furthermore, it is also considered that Alice has defined the bubble "AliceWorkBubble" associating her device "Sociotat:UMU:mobilenode:AliceMobilenode" to it, with a specific policy by following the steps of *Phase 2*. Then, the integration with the localization system is reflected in *Phase 3*. In this case, the Mobile Node worn by Alice sends the set of RSSIs, which are obtained from the reception of beacon messages sent by the anchors. The set of RSSIs is sent to the Gateway (step 3.1) through the WSN topology. After computing the estimated location (step 3.2), the Gateway sends this information to Alice's Bubble Handler (step 3.3). As already mentioned, at this point two alternatives are considered. In the first case (Option 1), Alice's Bubbles Handler runs the CP-ABE encryption algorithm to encrypt Alice's location by using the AliceWorkBubble's policy, and the CP-ABE parameters that are obtained during the Phase 1 (step 3.4). In the second case (Option 2), the CP-ABE

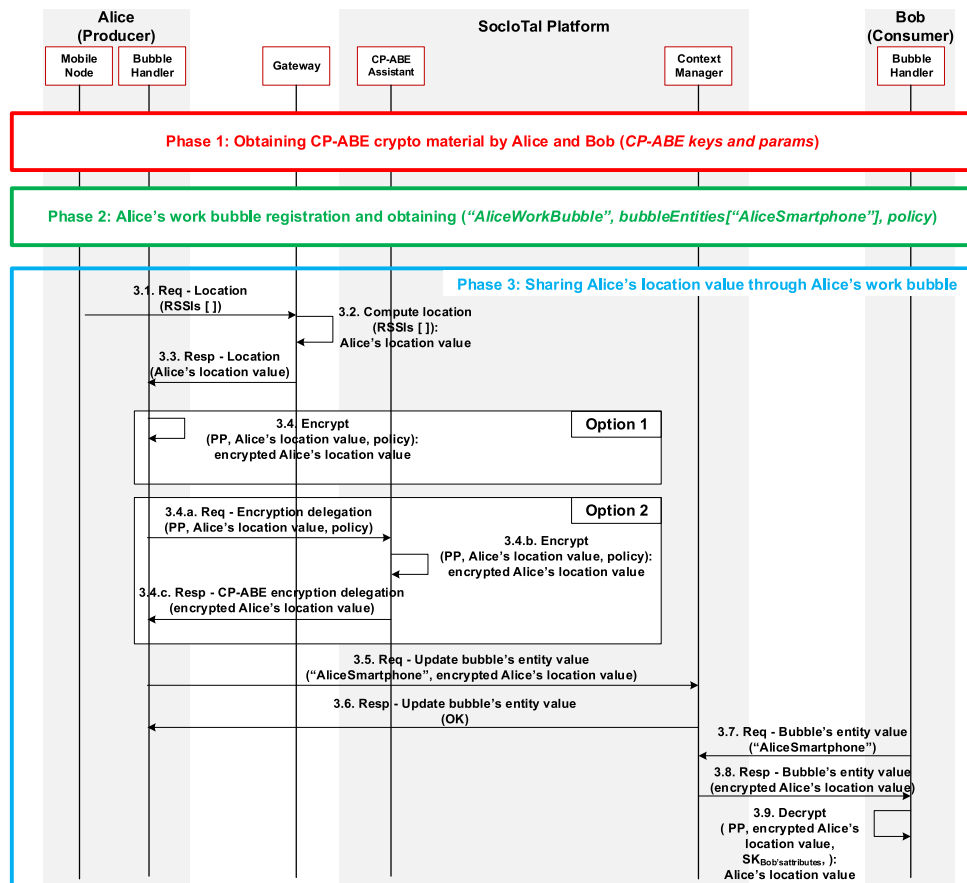


Fig. 9. Main interactions for sharing location data through bubbles by using the proposed localization system.

encryption is delegated to the CP-ABE Assistant service that is provided by the platform. Towards this end, the Bubble Handler component sends the encryption parameters to such service (3.4.a), so that the encryption operation is performed in the same way (3.4.b). Once the encryption has been completed, the result is sent to the requesting Bubble Handler (step 3.4.c).

Then, Alice's Bubbles Handler sends a NGSI *updateContext* message to the Context Manager, in order to update Alice's location of the entity "SocIoTal:UMU:mobilenode:AliceMobilenode" with the encrypted value previously obtained (step 3.5). This message is answered with a confirmation message from the Context Manager (step 3.6). Since then, each query to this entity will return the encrypted value of Alice's location. To do this, Bob can use the NGSI *queryContext* message to get the information associated to "AliceWorkBubble". It should be noted that these steps are optional, since Bob's Bubbles Handler could have previously obtained the information associated to this bubble. Then, this component queries the values associated to "SocIoTal:UMU:mobilenode:AliceMobilenode" through a NGSI *queryContext* method and it obtains the encrypted value of Alice's location (steps 3.7–3.8). In order to decrypt such value, the Bubble Handler runs the CP-ABE decryption algorithm by using Bob's CP-ABE key and the public parameters (step 3.9) that are previously obtained. Additionally, once Alice's location is decrypted, Bob can use the Web-App client (shown in Fig. 8) to show the Alice's location in the map of the building.

7. Use case implementation

Taking into account the previously identified components for the use case, this section is intended to describe the implementation details of each of them. It should be noted that for the implementation of the proposed scenario, we have used the existing implementation of the

WUE and Context Manager components. A more detailed description of each component can be found at GitHub.¹⁰

7.1. SocIoTal platform

As described in Section 4, the platform aims to embrace different components to enable the creation and management of bubbles in the proposed scenario. By taking into account the addition of the CP-ABE Assistant component for our use case, the implementation details for each component are provided.

CP-ABE key manager. It is a HTTPS server implemented in Java servlets accepting requests for obtaining the corresponding CPABE keys and cryptographic parameters. As already mentioned, these keys are associated with the attributes stored in the IdM Service. Towards this end, the implementation integrates a simple Keyrock client library to recover such attributes from that service. For CP-ABE functionality, two different libraries have been employed in order to compare their performance. On the one hand, we have made use of the library provided by [49], which is a Java implementation built on top of the Java Pairing Based Cryptography library (jpBC) [50]. On the other hand, we have deployed the required functionality for CP-ABE aspects through the use of the library provided in [51], which is a C implementation (based on the PBC library [52]) that was developed by the authors of the CP-ABE scheme.

IdM service. This service has been implemented as a modified deployment of the FIWARE Keyrock IdM system. In particular, the KeyRock server has been extended to consider attributes of the SCIM standard that was not included in the original version. It should be noted that a more complete version of this service including privacy-preserving

¹⁰ <http://github.com/sociotal/SOCIOTAL/wiki>.

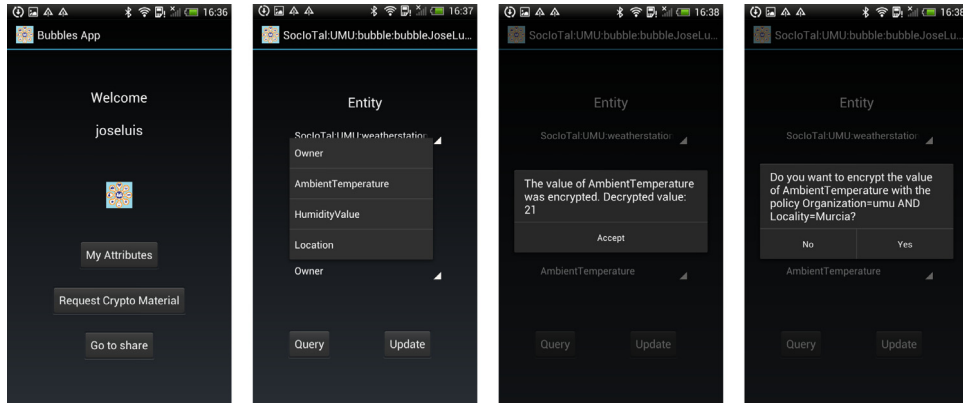


Fig. 10. Bubbles Handler screenshot examples.

features has been extended in the scope of the SocIoTal project, as described in [4].

CP-ABE assistant. It is a HTTPS server implemented in Java servlets accepting requests for encrypting data by using CP-ABE. Like in the case of the CP-ABE Key Manager, the already mentioned CP-ABE libraries have been employed in order to compare their performance in the encryption process.

7.2. Bubbles handler

This represents the main component to manage the secure data sharing aspects through the bubbles, as well as to obtain the required CP-ABE cryptographic material. It integrates a Keyrock Client library that provides a basic API for identity management by implementing a client to interact with the IdM Service for authentication purposes. It also integrates an OMA-NGSI client library in order to ease the management of the messages that are sent to the Context Manager, based on the use of the *gson* library.¹¹ Like in the case of the previous components, we have tested two different libraries to perform the required CP-ABE encryption/decryption operations [49,51].

It should be noted that, while these libraries can be executed in any Java-compatible device, Fig. 10 shows some examples of screenshots related to the deployment on Android devices. Additional information about this component can be found at Github¹²

7.3. Gateway

Implementation of the radio-location metrics on the gateway. The implementation of the location-dependent attribute algorithm related to Section 6.1 in Java is detailed on Listing 9. It could be implemented in any device that supports Java. For SocIoTal trials, it runs in the local Gateway that first estimates the distance of each anchor from the sink node from the collected RSSI values.

Listing 9: Code for the computation of the distances

```
1 private double distance(double rssi){
2     //d = (d_ref x 10^((rssi_ref - rssi) / (10 x a))^2
3     return Math.pow(REF_DISTANCE * Math.pow(10, (REF_RSSI - rssi) / (10 *
4         ALPHA)), 2);
5 }
```

Practical implementation of the positioning algorithm. Moreover, as practical implementation of the algorithm described in Section 6.2, the fixed anchors are calibrated and their relative coordinates in the given environment is estimated as depicted by the code detailed on Listing 10. Given this calibration, Listing 11 depicts the Java implementation of the Gradient-Descent technique used to estimate the relative 2D coordinates of a sink node in the real environment.

Listing 10: Code for the anchor calibration

```
1 static{
2     //Fill in the anchor map
3     anchorMap.put("00124b00060d85b3", new Coord(0.00, 0.00));
4     anchorMap.put("00124b00060d84c4", new Coord(26.00, 0.00));
5     anchorMap.put("00124b00060d8475", new Coord(46.00, 0.00));
6     anchorMap.put("00124b00060d8473", new Coord(54.00, 65.00));
7     anchorMap.put("00124b00060d8480", new Coord(25.00, 65.00));
8     anchorMap.put("00124b00060d84ba", new Coord(0.00, 65.00));
9     anchorMap.put("00124b00060d84f5", new Coord(0.00, 38.00));
10    anchorMap.put("00124b00060d84aa", new Coord(23.00, 34.00));
11 }
```

Listing 11: Code for the computation of the coordinates

```
1 public Coord computeLeastMeanSquare() {
2     //Mapping of anchor address <=> mean RSSI with this anchor
3     Map<String, Double> meanRssi = meanRssi.size();
4     if (size < 3)
5         //We cannot compute localization with less
6         //than three anchors
7         return null;
8
9     final List<String> macs =
10         new ArrayList<>(meanRssi.keySet());
11     //Choose the first anchor as an arbitrary reference
12     //for computations
13     final String firstMac = macs.get(0);
14     final Coord firstCoord = anchorMap.get(firstMac);
15     final double firstDistance =
16         distance(meanRssi.get(firstMac));
17
18     //Fill the matrices
19     Matrix a = new BasicIDMatrix(size, 2);
20     Matrix h = new BasicIDMatrix(size, 1);
21     for(int i = 1; i < meanRssi.size(); i++){
22         final String mac = macs.get(i);
23         final Coord coord = anchorMap.get(mac);
24
25         // A = [ x_i - x_0  y_i - y_0 ]
26         //      [ x_k - x_0  y_k - y_0 ]
27         a.set(i, 0, coord.x - firstCoord.x);
28         a.set(i, 1, coord.y - firstCoord.y);
29
30         // h = [ x_i^2 - x_0^2 + y_i^2 - y_0^2 + d_0^2 - d_i^2 ]
31         //      [ x_k^2 - x_0^2 + y_k^2 - y_0^2 + d_0^2 - d_k^2 ]
32         //      [ ----- ]
33         //      hprime
34
35         final double hprime = Math.pow(coord.x, 2) -
36             Math.pow(firstCoord.x, 2) +
37             Math.pow(coord.y, 2) -
38             Math.pow(firstCoord.y, 2);
39         h.set(i, 0, hprime, firstDistance -
40             distance(meanRssi.get(mac)));
41     }
42 }
```

¹¹ <https://github.com/google/gson>.

¹² <https://github.com/sociotal/SOCIOTAL/wiki/SocioTal-Group-Manager>.

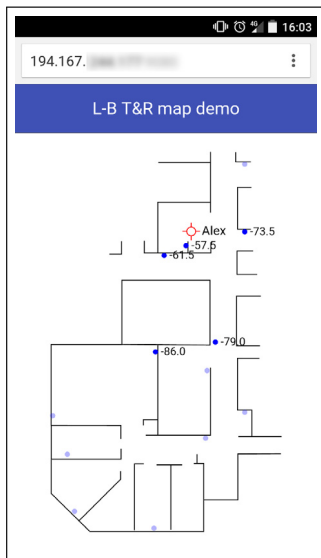


Fig. 11. Web App screenshot.. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Once estimated on the local gateway thanks to this implementation of the gradient-based method, the time-stamped 2D $(\tilde{x}_m(t), \tilde{y}_m(t))$ coordinates produced at time t can be sent to the Bubbles Handler component for encryption purposes before sending to the Context Manager.

7.4. Web app environment

As part of the deployment of the radio localization system (see Fig. 8), a localization Web App has been developed and it is shown in Fig. 11. The Web App provides a map of the building where the active anchors – those whom the RSSI is taken into account by the sink node – are colored in blue. The inactive anchors are colored in pastel blue and the sink node (in this case, worn by Alex) makes reference to a person moving on the building. His displacement is followed on the map thanks to the web app. This way, in case of a successful CP-ABE decryption process, a user can employ this app to visualize the position of other users in the building.

8. Evaluation and performance results

After describing the implementation details of the different components, this section provides a description of the specific hardware, as well as the results obtained from the different processes previously explained.

8.1. Hardware features

For the radio localization process, the OpenMote-CC2538 hardware technology is used for the low-level devices. The OpenMote-CC2538 is a System-on-Chip embedding a Cortex-M3 microcontroller and a CC2520-like radio transceiver from Texas Instrument compliant for 2.4 GHz IEEE 802.15.4–2006 standard. The microcontroller runs up to 32 MHz and includes 32 Kbytes of RAM and 512 Kbytes of Flash, as well as the usual peripherals. It embeds a battery monitor, temperature and lighting sensors, as well as a 3-axes accelerometer. This sensor mote is designed for IoT applications. For this experiment, the latest version of Contiki 3.x is embedded as Operating System. In the following the metrics used to estimate the coordinate of a sink node from the collected RSSI measure from the fixed anchors are detailed both with the analytic formula and the corresponding implementation in the local gateway. The experiments have been performed with the OpenMote hardware

that are narrow-band IEEE 802.15.4 devices operating at 2.4 GHz in the ISM band. Furthermore, the gateway functionality was implemented on an Intel(R) Xeon(R) CPU E3-1220 v3 @ 3.10 GHz.

Moreover, the Bubbles Handler has been implemented on an Android device Samsung Galaxy J3 (6) with a 1.5 GHz quad-core processor, 1.5 GB of RAM and version 5.1 of Android (Lollipop). In addition, the platform components (besides the Context Manager and the WUE) were deployed on an Intel Core i5-5257U 2.70 GHz processor and 2 GB of RAM with Ubuntu 14.

8.2. Considerations about bubbles scenario evaluation

As previously described, we have made use of two different implementations for the CP-ABE functionality. Since the Bubbles Handler is deployed on an Android device, it should be noted that, for the C library, we have followed a similar approach to [26] by using *Android Native Development Kit* (NDK) and CMake¹³ tools, in order to accommodate C code into the Android application. By using both implementations, we have made use of the following aspects:

- **Hardware.** As described in Section 6, the CP-ABE encryption operation could be executed on the Bubbles Handler or the CP-ABE Assistant within the platform. Consequently, evaluation results consider different hardware devices for the CP-ABE encryption process.
- **Performance parameters.** In addition to the delay required for the main CP-ABE algorithms, energy consumption is also considered for evaluation purposes in the case of the Bubbles Handler, which is deployed on Android devices. For this purpose, we have made use of the tool Powertutor [53], which is also used in previous works [26].
- **Number of bubbles' attributes.** As stated at Section 4, we use up to 6 different SCIM attributes to be employed for restricting the access to data within a specific bubble. While additional attributes could be used, it would make attribute management more difficult to be performed.
- **Security level.** Both implementations use type A pairings based on the curve $y^2 = x^3 + x$ over the field F_p for some prime $p = 3 \bmod 4$. The security strength of the scheme can be tuned by modifying two parameters: the size of the field p , and the prime order q [26,25]. In our case, evaluation is performed by considering 80-bit (i.e. $|p| = 512$, $|q| = 160$) and 112-bit (i.e. $|p| = 1024$, $|q| = 224$) security levels.
- **Scalability.** While previous aspects are already considered by [26] and [54], we assess the impact of the concurrent execution of CP-ABE cryptographic algorithms. Beyond the use of location data, we consider this aspect crucial due to the potential huge amount of data to be protected on an IoT-enabled scenario.

Additionally, it should be noted that for CP-ABE encryption and decryption evaluation, only CP-ABE policies with AND gates and 1-level policy tree are considered. This is due to the implicit bubble definition in which all the attributes in the bubble are used to restrict the access to data within a specific bubble. While a more expressive (and consequently, more complex) definition could be used, the proposed bubble specification aims to provide a simple but still consistent view so that end users are enabled to control the disclosure of their devices' data. In addition, each experimentation result was obtained as the mean value of 10 executions, and confidence intervals were calculated by considering a 95% confidence level.

¹³ <https://cmake.org/>.

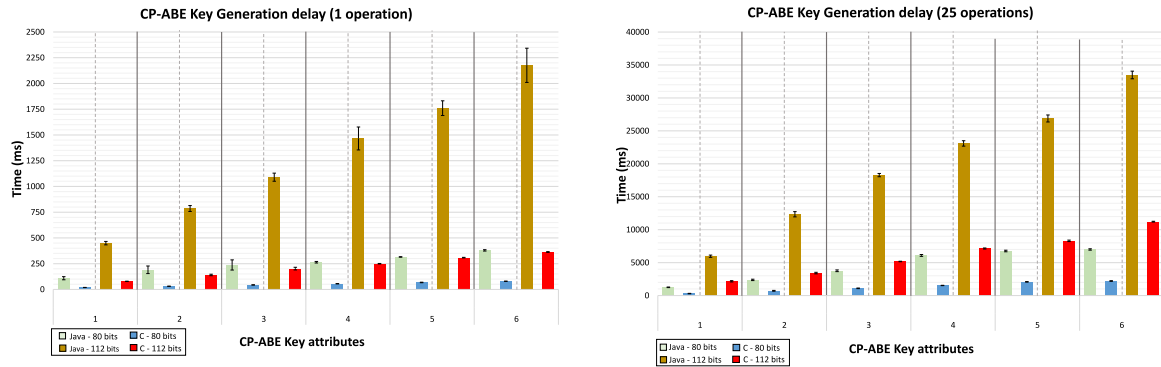


Fig. 12. CP-ABE Key Manager performance.

8.3. Obtaining CP-ABE cryptographic material

Taking into account these premises, Fig. 12 shows the delay required for both implementations in the case of CP-ABE key generation that is performed on the CP-ABE Key Manager. Indeed, this process makes reference to the step 1.5 in Fig. 3). The x -axis represent the number of attributes that is associated to the generated key. Firstly, for a 1-operation workload, the implementation based on [49] requires 108 ms to generate a 1-attribute key (379 ms for a 6-attribute key) in case of 80-bits security level. This delay grows until 449 ms and 2176 ms when considering a 112-bits security level. This time is considerably reduced by using the implementation based on [51] in which the delay obtained is 78 ms and 368 ms for this security level.

Secondly, for a 25-operation workload, the difference is more considerable; the Java library requires 1253 ms to generate 1-attribute key (7097 ms for a 6-attribute key) in case of 80-bits security level (5973 ms/33484 ms for 112-bits). For the C library, this delay is reduced to 2167 ms/11221 ms for 112-bits security level, which represents a 63,72% and 66,49% decrease respectively. While it represents a considerable better performance, it should be noted that this process is not expected to be done frequently, since CP-ABE keys are associated to users' attributes, which does not affect possible dynamic changes with CP-ABE policies or the creation of new bubbles in our scenario.

8.4. Detecting and sending location data to gateway

Regarding the interactions related to the localization system, once the data collected by the Mobile Node, aggregated messages including the RSSI measures are sent to the Gateway (step 3.1 in Fig. 9). This operation takes few milliseconds (between 1 ms and 100 ms) in a scope that trigs on the mote transmission and includes the receiving and the time of raising the IEEE 802.15.4 stack in the Linux environment on the gateway.

8.5. Data processing on the gateway

The data processing (step 3.2 in Fig. 9) consists of performing the distances for each collected RSSI using the barycenter method. This processing takes few μ s as depicted in the table below. For these measurements, the statistics are obtained from a collection of 1 million of samples uniformly distributed on the considered interval. Once the distances evaluated, the gradient descent method is used to estimate the 2D coordinates of the sink mote. This technique takes significantly more time than the least mean square method, as depicted in Table 1. Nevertheless, the precision, the consistency and the reproducibility of this technique is higher and justifies its use to provide to the end user of relevant estimation of the sink localization.

8.6. Encrypting location data through bubbles

Taking into account the proposed alternatives for CP-ABE encryption (and delegated encryption), in this section we compare both implementations when they are executed on different hardware.

On the one hand, Fig. 13 shows the delay required for the encryption by fluctuating the workload (from 1 to 5 operations). In particular, these results refer to the step 3.4 in Fig. 9. In case of a 1-operation workload, the implementation based on the Java library requires 1232 ms/5283 ms for 1 and 6-attribute policies when a 80-bits security level is considered. This time is around 3799 ms/20 006 ms in the case of a 112-bits security level, representing a 208,36%/278,69% increase. For this operation and the same security level, the C library requires 351 ms/1355 ms, which means a 90,76%/93,23% decrease regarding the previous library. Furthermore, in case of a 5-operation workload, the implementation based on C represents a 92,37%/95,33% decrease by considering the same configuration.

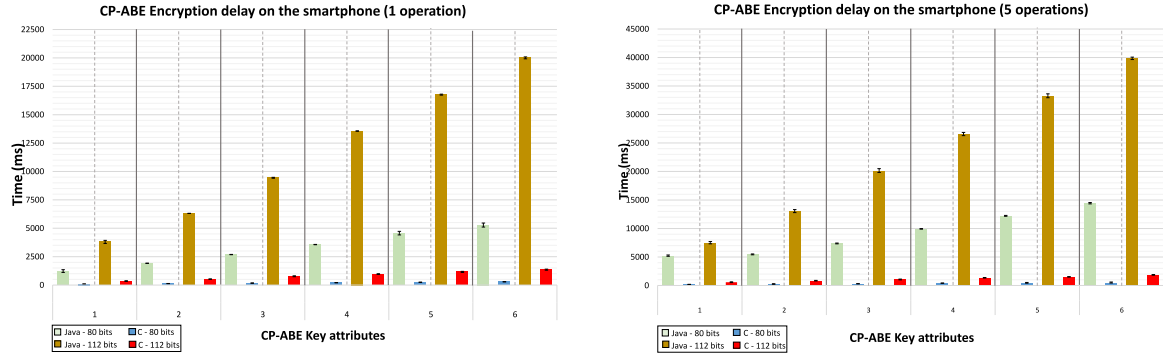
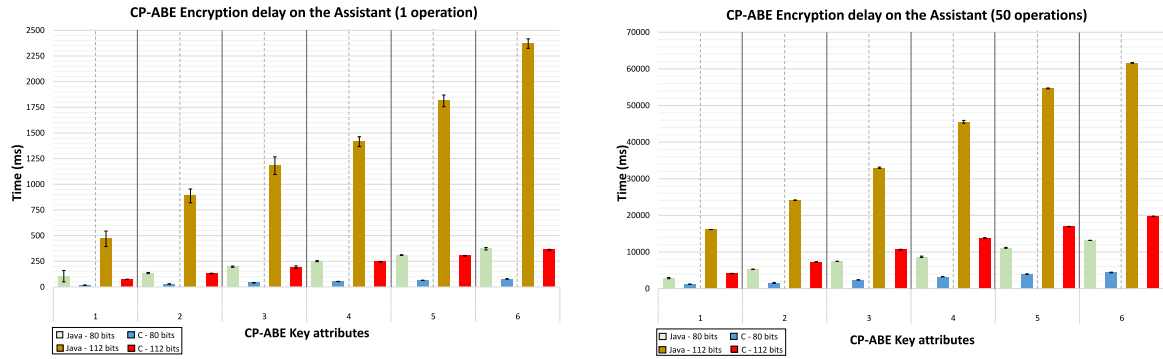
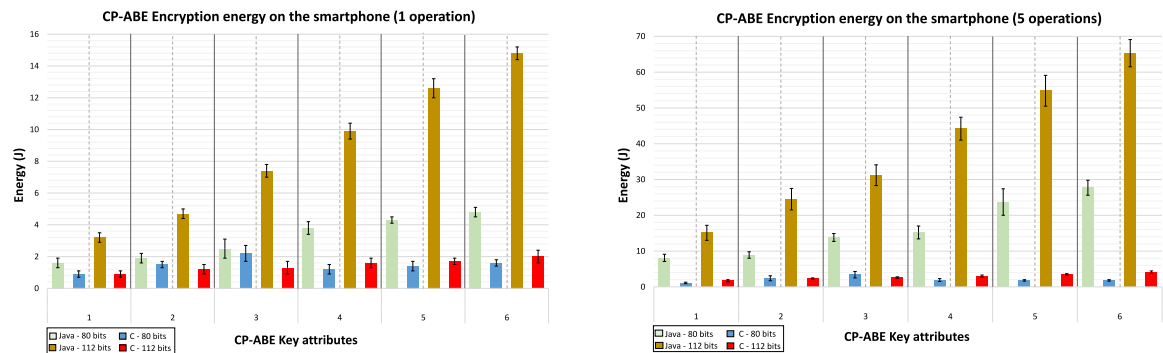
On the other hand, Fig. 14 depicts the performance of the CP-ABE encryption operation when it is executed by the CP-ABE Assistant (3.4.b). It should be noted that, in this case, we have considered a workload up to 50 operations, since this component will be in charge of encrypting data coming from different devices in the building. In this case, for a 1-operation workload, the Java implementation needs 468 ms/2370 ms for 1 and 6-attribute policies when a 112-bits security level is considered. In this sense, the C library requires 74 ms/363 ms, so a 84,19%/84,68% decrease regarding the previous library. It should be noted that, since CP-ABE encryption is performed more efficiently in the platform, this decrease is less significant than in the case of the smartphone. In case of a 50-operation workload, the Java library requires 16 126 ms/61 639 ms for 1 and 6-attribute policies and a 112-bits security level. For the C library, this delay is reduced to 4125 ms/19 741 ms, so a 74,42% and 67,97% decrease is obtained respectively. In general, while evaluation results are obviously better in the case of the CP-ABE Assistant, it should be noted that this entity would be able to access the data to be encrypted. Consequently, depending on the security and performance requirements of a specific scenario, the inclusion of this component could be appropriate or not.

Moreover, the aspects related to energy consumption are a major concern for scenarios where certain devices (e.g. a smartphone) are employed. Because of it, we provide a set of experimental results regarding the energy consumption for the CP-ABE encryption operation in case of the Bubbles Handler is running this operation. These results are shown in Fig. 15 by considering a workload between 1 and 5 operations. According to it, the encryption operation requires between 3,2 J and 14,8 J in the case of the Java implementation for 1 and 6-attributes policies, and a 112-bits security level. In case of the C implementation, this consumption is reduced to 0,9 J and 2 J meaning a 71,88% and 86,49% decrease. In addition, for a 5-operations workload, the Java library consumes 15,1 J/65,3 J, while the C implementation requires 1,7 J/4,1 J, which represents an 88,74%/93,72% decrease.

Table 1

Performance of the location algorithms in the Gateway.

Location computations (Values in μ s)					
Barycenter method	1000 000 samples	min 1.0	max 5328.0	mean 1.18	stddev 7.4
Least mean square method	1000 000 samples	min 7.0	max 54536.0	mean 7.8	stddev 57.2
Gradient descent method	10 000 samples	min 2432.0	max 89195.0	mean 2569.3	stddev 1042.2

**Fig. 13.** CP-ABE Encryption delay on the Bubbles Handler.**Fig. 14.** CP-ABE Encryption delay on the CP-ABE Assistant.**Fig. 15.** CP-ABE Encryption energy consumption on the Bubbles Handler.

8.7. Decrypting location data through bubbles

When encrypted data are obtained through the Context Manager, they need to be decrypted by using the CP-ABE decryption algorithm (step 3.9 in Fig. 9). For this operation, Fig. 16 shows the time required

by the Bubbles Handler by considering a workload of 1 and 5 operations. In this case, unlike the encryption, we have considered the decryption process should not be delegated to any other entity, since the CP-ABE key would be disclosed, so that such entity would have access all the information being shared.

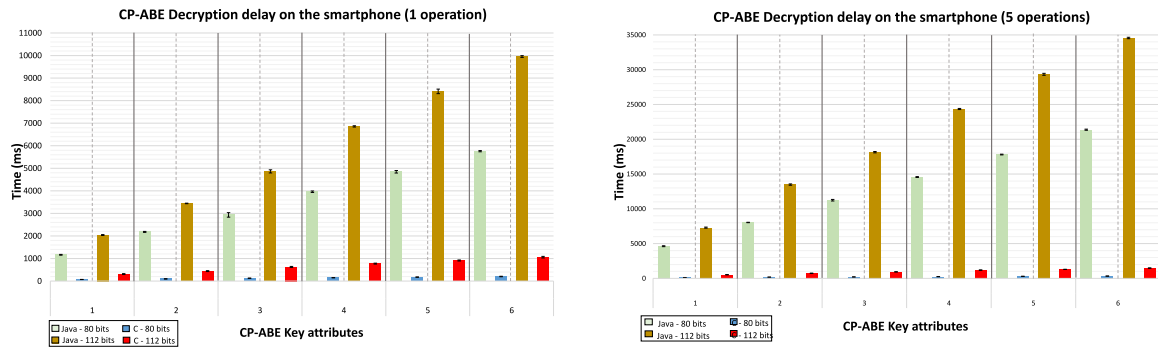


Fig. 16. CP-ABE Decryption delay on the Bubbles Handler.

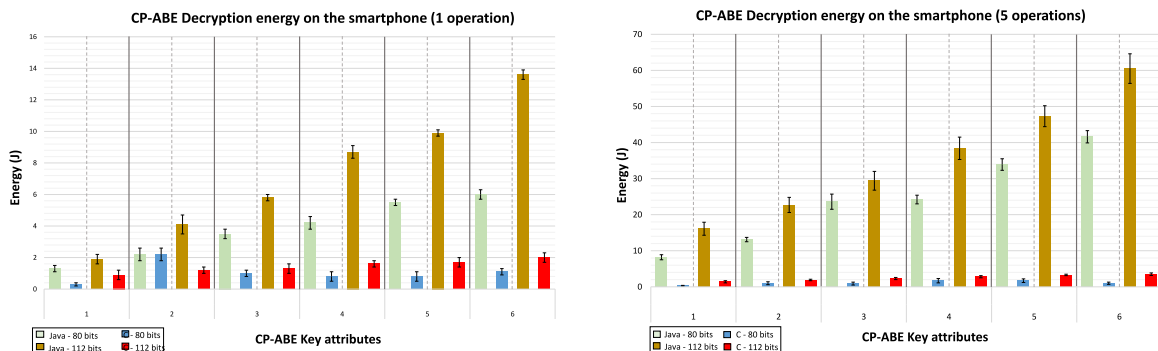


Fig. 17. CP-ABE Decryption energy consumption on the Bubbles Handler.

According to results, in case of a 1-operation workload, the Java-based implementation library requires 1165 ms/5758 ms for 1 and 6-attribute policies when a 80-bits security level is considered. This delay is increased to 2041 ms/9956 ms in the case of a 112-bits security level, representing a 75,19%/72,91% increase. For this operation and the same security level, the C library requires 309 ms/1059 ms, which means a 73,48%/89,36% decrease regarding the previous library. Furthermore, in case of a 5-operation workload, the implementation based on C obtains 526 ms/1476 ms (7297 ms/34559 ms in case of the Java library), which represents a 92,79%/95,73% respectively.

For energy consumption (Fig. 17), like in the case of encryption, values remain almost constant in case of the implementation based on the C library. Indeed, it requires 0,3 J/1 J for 1 and 6-attribute policies when a 80-bits security level is considered. This consumption is around 0,9 J/1,8 J in the case of a 112-bits security level. In case of the Java implementation, the energy consumption grows up to 1,3–6 J for a 80-bits security level and 1,9–13,6 J by considering a 112-bits level. Finally, in case of a 5-operation workload and this security level, the Java library consumes 16,1–60,5 J, while the implementation based on C requires 1,4–3,5 J, which represents a 1050% a 91,3%/94,21% decrease by considering the same configuration.

Such results demonstrate the proposed approach could be used even when a huge amount of personal data needs to be protected through the use of bubbles. Indeed, as the number of concurrent operations grows, the performance improvement by using the library from [51] is still more significant, providing practical results even when considering 6-attributes policies and a 112-security level. With the proposed approach, while the integration of CP-ABE on smartphone devices requires heavy cryptographic operations to be performed on end devices, confidentiality is ensured on an end-to-end basis. However, it should be noted that the Bubbles Handler component could be instantiated by a different entity (beyond the use of a smartphone). In this case, the delegation of

the CP-ABE encryption into the assistant component could be suitable at the expense of reducing the security restrictions in the scenario. In this sense, the application of such delegation approach could be analyzed and compared by considering more constrained IoT devices, such as in [54].

9. Conclusions and future work

Current trends in smart cities promote the development of new services based on sharing massive amounts of information, representing the cornerstone of a data-driven economy. In these scenarios, the specification and enforcement of security preferences on shared data is an extremely challenging aspect, which must be addressed by comprehensive approaches considering aspects such as flexibility, scalability, heterogeneity and usability. This work has presented the concept of bubble as a way to protect devices' data by considering potential relationships between their owners. Toward this end, the use of the CP-ABE scheme provides significant benefits in terms of key management to provide confidentiality even when this information needs to be shared through central data platforms. The resulting approach has been further deployed on different hardware components, as well as the FI-WARE platform in order to demonstrate their applicability on a real environment. In order to complement our proposal, future work is mainly focused on the integration of anonymous signature schemes to protect the integrity of shared data within a bubble in a privacy-preserving way.

Acknowledgments

This work has been partially funded by the project ANASTACIA, which has received funding from the European Union's H2020 Research and Innovation Programme under Grant Agreement 731558, and CHIST-ERA PCIN-2016-010.

References

- [1] A. Caragliu, C.D. Bo, P. Nijkamp, Smart cities in Europe, *J. Urban Technol.* 18 (2) (2011) 65–82, <http://dx.doi.org/10.1080/10630732.2011.601117>.
- [2] A. Sahai, B. Waters, Fuzzy identity-based encryption, in: *Lecture Notes in Computer Science*, Springer, Berlin, Heidelberg, 2005, pp. 457–473, http://dx.doi.org/10.1007/11426639_27.
- [3] J. Bethencourt, A. Sahai, B. Waters, Ciphertext-policy attribute-based encryption, in: 2007 IEEE Symposium on Security and Privacy, IEEE, 2007, <http://dx.doi.org/10.1109/sp.2007.11>.
- [4] J.B. Bernabe, J.L. Hernandez-Ramos, A.F.S. Gomez, Holistic privacy-preserving identity management system for the Internet of Things, *Mob. Inf. Syst.* 2017 (2017) 1–20, <http://dx.doi.org/10.1155/2017/6384186>.
- [5] M. Bauer, E. Kovacs, A. Schulke, N. Ito, C. Criminisi, L.-W. Goix, M. Valla, The context API in the OMA next generation service interface, in: 2010 14th International Conference on Intelligence in Next Generation Networks, IEEE, 2010, <http://dx.doi.org/10.1109/icin.2010.5640931>.
- [6] A. Martinez-Balleste, P.A. Pérez-Martínez, A. Solanas, The pursuit of citizens' privacy: a privacy-aware smart city is possible, *IEEE Commun. Mag.* 51 (6) (2013) 136–141, <http://dx.doi.org/10.1109/MCOM.2013.6525606>.
- [7] M. Frustaci, P. Pace, G. Aloí, G. Fortino, Evaluating critical security issues of the IoT world: Present and future challenges, *IEEE Internet Things J.* (2017) 2483–2495, <http://dx.doi.org/10.1109/jiot.2017.2767291>.
- [8] A. Sedrati, A. Mezrioui, Internet of Things challenges: A focus on security aspects, in: 2017 8th International Conference on Information and Communication Systems, ICICS, IEEE, 2017, <http://dx.doi.org/10.1109/iacs.2017.7921973>.
- [9] C.-W. Tsai, C.-F. Lai, M.-C. Chiang, L.T. Yang, Data mining for Internet of Things: A survey, *IEEE Commun. Surv. Tutor.* 16 (1) (2014) 77–97, <http://dx.doi.org/10.1109/surv.2013.103013.00206>.
- [10] J.-M. Bohlí, R. Kurpatov, M. Schmidt, Selective decryption of outsourced IoT data, in: 2015 IEEE 2nd World Forum on Internet of Things, WF-IoT, IEEE, 2015, <http://dx.doi.org/10.1109/wf-iot.2015.7389146>.
- [11] V. Goyal, O. Pandey, A. Sahai, B. Waters, Attribute-based encryption for fine-grained access control of encrypted data, in: *Proceedings of the 13th ACM Conference on Computer and Communications Security - CCS*, ACM Press, 2006, <http://dx.doi.org/10.1145/1180405.1180418>.
- [12] R. Ostrovsky, A. Sahai, B. Waters, Attribute-based encryption with non-monotonic access structures, in: *Proceedings of the 14th ACM Conference on Computer and Communications Security - CCS*, ACM Press, 2007, <http://dx.doi.org/10.1145/1315245.1315270>.
- [13] N. Attrapadung, J. Herranz, F. Laguillaumie, B. Libert, E. de Panafieu, C. Ràfols, Attribute-based encryption schemes with constant-size ciphertexts, *Theoret. Comput. Sci.* 422 (2012) 15–38, <http://dx.doi.org/10.1016/j.tcs.2011.12.004>.
- [14] G. Wang, Q. Liu, J. Wu, Hierarchical attribute-based encryption for fine-grained access control in cloud storage services, in: *Proceedings of the 17th ACM Conference on Computer and Communications Security - CCS*, ACM Press, 2010, <http://dx.doi.org/10.1145/1866307.1866414>.
- [15] M. Li, S. Yu, Y. Zheng, K. Ren, W. Lou, Scalable and secure sharing of personal health records in cloud computing using attribute-based encryption, *IEEE Trans. Parallel Distrib. Syst.* 24 (1) (2013) 131–143, <http://dx.doi.org/10.1109/tpds.2012.97>.
- [16] Z. Wan, J. Liu, R.H. Deng, HASBE: A hierarchical attribute-based solution for flexible and scalable access control in cloud computing, *IEEE Trans. Inf. Forensics Secur.* 7 (2) (2012) 743–754, <http://dx.doi.org/10.1109/tifs.2011.2172209>.
- [17] L. Touati, Y. Challal, A. Bouabdallah, C-CP-ABE: Cooperative ciphertext policy attribute-based encryption for the Internet of Things, in: 2014 International Conference on Advanced Networking Distributed Systems and Applications, IEEE, 2014, <http://dx.doi.org/10.1109/inds.2014.19>.
- [18] X. Yao, Z. Chen, Y. Tian, A lightweight attribute-based encryption scheme for the Internet of Things, *Future Gener. Comput. Syst.* 49 (2015) 104–112, <http://dx.doi.org/10.1016/j.future.2014.10.010>.
- [19] D. Hankerson, A.J. Menezes, S. Vanstone, *Guide to Elliptic Curve Cryptography*, Springer Science & Business Media, 2006.
- [20] G. Bianchi, A.T. Caposelle, C. Petrioli, D. Spenza, AGREE: exploiting energy harvesting to support data-centric access control in WSNs, *Ad Hoc Netw.* 11 (8) (2013) 2625–2636, <http://dx.doi.org/10.1016/j.adhoc.2013.03.013>.
- [21] M. Singh, M. Rajan, V. Shivraj, P. Balamuralidhar, Secure MQTT for Internet of Things (IoT), in: 2015 Fifth International Conference on Communication Systems and Network Technologies, IEEE, 2015, <http://dx.doi.org/10.1109/cnsn.2015.16>.
- [22] D. Thatmann, S. Zickau, A. Forster, A. Kupper, Applying attribute-based encryption on publish subscribe messaging patterns for the Internet of Things, in: 2015 IEEE International Conference on Data Science and Data Intensive Systems, IEEE, 2015, <http://dx.doi.org/10.1109/dsdis.2015.52>.
- [23] Information technology. Message Queuing Telemetry Transport (MQTT) v3.1.1, <http://dx.doi.org/10.3403/30328634u>.
- [24] F. Guo, Y. Mu, W. Susilo, D.S. Wong, V. Varadharajan, CP-ABE with constant-size keys for lightweight devices, *IEEE Trans. Inf. Forensics Secur.* 9 (5) (2014) 763–771, <http://dx.doi.org/10.1109/tifs.2014.2309858>.
- [25] X. Wang, J. Zhang, E.M. Schooler, M. Ion, Performance evaluation of attribute-based encryption: Toward data privacy in the IoT, in: 2014 IEEE International Conference on Communications, ICC, IEEE, 2014, <http://dx.doi.org/10.1109/icc.2014.6883405>.
- [26] M. Ambrosin, M. Conti, T. Dargahi, On the feasibility of attribute-based encryption on smartphone devices, in: *Proceedings of the 2015 Workshop on IoT Challenges in Mobile and Industrial Systems - IoT-Sys*, ACM Press, 2015, <http://dx.doi.org/10.1145/2753476.2753482>.
- [27] K. Grizzle, E. Wahlstroem, C. Mortimore, System for Cross-domain Identity Management: Core Schema, Tech. Rep., RFC ed., Sep., 2015, <http://dx.doi.org/10.17487/rfc7643>.
- [28] E. Kovacs, M. Bauer, J. Kim, J. Yun, F.L. Gall, M. Zhao, Standards-based worldwide semantic interoperability for IoT, *IEEE Commun. Mag.* 54 (12) (2016) 40–46, <http://dx.doi.org/10.1109/mcom.2016.1600460cm>.
- [29] I. Eliceigui, C. López, L. Sánchez, J. Lanza, L. Muñoz, A. Pintus, A. Manchinu, A. Serra, Design and implementation of a cloud-based platform for unleashing the personal and communal Internet of Things, *Mob. Inf. Syst.* 2017 (2017) 1–14, <http://dx.doi.org/10.1155/2017/2164072>.
- [30] L. Atzori, A. Iera, G. Morabito, M. Nitti, The social Internet of Things (SIoT)—When social networks meet the Internet of Things: Concept, architecture and network characterization, *Comput. Netw.* 56 (16) (2012) 3594–3608, <http://dx.doi.org/10.1016/j.comnet.2012.07.010>.
- [31] O. Garcia-Morchon, S.L. Keoh, S. Kumar, P. Moreno-Sanchez, F. Vidal-Meca, J.H. Ziegeldorf, Securing the IP-based internet of things with HIP and DTLS, in: *Proceedings of the Sixth ACM Conference on Security and Privacy in Wireless and Mobile Networks - WiSec*, ACM Press, 2013, <http://dx.doi.org/10.1145/2462096.2462117>.
- [32] D. Boneh, M. Franklin, Identity-based encryption from the weil pairing, in: *Advances in Cryptology — CRYPTO 2001*, Springer, Berlin, Heidelberg, 2001, pp. 213–229, <http://dx.doi.org/10.1007/3-540-44647-8.13>.
- [33] N. Ragouzis, J. Hughes, R. Philpott, E. Maler, P. Madsen, T. Scavo, Security Assertion Markup Language, SAML, V2.0 Technical Overview, OASIS Committee Draft, vol. 2, 2008.
- [34] D. Hardt, RFC 6749: The OAuth 2.0 Authorization Framework, Internet Engineering Task Force, IETF, 2012, <http://dx.doi.org/10.17487/rfc6749>.
- [35] J.B. Bernabe, I. Eliceigui, E. Gandrille, N. Gligoric, A. Gluhak, C. Hennebert, J.L. Hernandez-Ramos, C. Lopez, A. Manchinu, K. Moessner, M. Nati, C. O'Reilly, N. Palaghias, A. Pintus, L. Sanchez, A. Serra, R. van Kranenburg, SocIoTal — The development and architecture of a social IoT framework, in: 2017 Global Internet of Things Summit, GloTS, IEEE, 2017, <http://dx.doi.org/10.1109/giots.2017.8016286>.
- [36] Open mobile alliance (oma) ngsl context management v1.0, 2010. URL http://www.openmobilealliance.org/release/NGSI/V1_0-20120529-A/OMA-TS-NGSI_Context_Management-V1_0-20120529-A.pdf.
- [37] J.L. Hernandez-Ramos, J.B. Bernabe, A. Skarmeta, ARMY: architecture for a secure and privacy-aware lifecycle of smart objects in the Internet of My Things, *IEEE Commun. Mag.* 54 (9) (2016) 28–35, <http://dx.doi.org/10.1109/mcom.2016.7565269>.
- [38] J. Camenisch, E.V. Herreweghen, Design and implementation of the idemix anonymous credential system, in: *Proceedings of the 9th ACM Conference on Computer and Communications Security - CCS*, ACM Press, 2002, <http://dx.doi.org/10.1145/586111.586114>.
- [39] S. Yu, C. Wang, K. Ren, W. Lou, Attribute based data sharing with attribute revocation, in: *Proceedings of the 5th ACM Symposium on Information, Computer and Communications Security - ASIACCS*, ACM Press, 2010, <http://dx.doi.org/10.1145/1755688.1755720>.
- [40] S. Perez, J.L. Hernandez-Ramos, S.N. Matheu-Garcia, D. Rotondi, A.F. Skarmeta, L. Straniero, D. Pedone, A lightweight and flexible encryption scheme to protect sensitive data in smart building scenarios, *IEEE Access* 6 (2018) 11738–11750, <http://dx.doi.org/10.1109/access.2018.2801383>.
- [41] J. Li, W. Yao, J. Han, Y. Zhang, J. Shen, User collusion avoidance CP-ABE with efficient attribute revocation for cloud storage, *IEEE Syst. J.* 12 (2) (2018) 1767–1777, <http://dx.doi.org/10.1109/jsyst.2017.2667679>.
- [42] E.P. Paper, Smart buildings: Energy efficiency first! 2017, http://euroace.org/wp-content/uploads/2015/10/EA_Smart_Buildings_Feb_2017_Final.pdf. (Online, Accessed 20 December 2017).
- [43] C. Hennebert, J.D. Santos, Security protocols and privacy issues into 6LoWPAN stack: A synthesis, *IEEE Internet Things J.* 1 (5) (2014) 384–398, <http://dx.doi.org/10.1109/jiot.2014.2359538>.
- [44] P. Coronel, S. Furrer, W. Schott, B. Weiss, Indoor location tracking using inertial navigation sensors and radio beacons, in: *The Internet of Things*, Springer, Berlin, Heidelberg, pp. 325–340, <http://dx.doi.org/10.1007/978-3-540-78731-0.21>.
- [45] M. Laaraiedh, S. Avrillon, B. Uguen, Enhancing positioning accuracy through direct position estimators based on hybrid rss data fusion, in: *VTC Spring 2009 - IEEE 69th Vehicular Technology Conference*, IEEE, 2009, <http://dx.doi.org/10.1109/vetecs.2009.5073539>.
- [46] J.J. Nielsen, H. Nouredine, N. Amiot, B. Uguen, M. Laaraiedh, I. Raos, I. Arambasic, V. Savic, S.Z. Bello, J. Dominguez, B. Denis, R. Raulefs, Assessment of cooperative and heterogeneous indoor localization algorithms with real radio devices, in: 2014 IEEE International Conference on Communications Workshops, ICC, IEEE, 2014, <http://dx.doi.org/10.1109/iccw.2014.6881190>.
- [47] B. Denis, L. He, L. Ouvre, A flexible distributed maximum log-likelihood scheme for UWB indoor positioning, in: 2007 4th Workshop on Positioning, Navigation and Communication, IEEE, 2007, <http://dx.doi.org/10.1109/wpnc.2007.353616>.

- [48] G. Destino, D. Macagnano, G. Abreu, B. Denis, L. Ouvry, Localization and tracking for LDR-UWB systems, in: 2007 16th IST Mobile and Wireless Communications Summit, IEEE, 2007, <http://dx.doi.org/10.1109/istmwc.2007.4299285>.
- [49] Y. Wang, CP-ABE java implementation, 2013, URL <https://github.com/junwei-wang/cpabe>.
- [50] A.D. Caro, V. Iovino, jPBC: Java pairing based cryptography, in: 2011 IEEE Symposium on Computers and Communications, ISCC, IEEE, 2011, <http://dx.doi.org/10.1109/iscc.2011.5983948>.
- [51] J. Bethencourt, A. Sahai, B. Waters, cpabe toolkit, 2011, URL <http://acsc.cs.utexas.edu/cpabe/>.
- [52] B. Lynn, PBC library, 2006, URL <http://crypto.stanford.edu/pbc>.
- [53] L. Zhang, B. Tiwana, Z. Qian, Z. Wang, R.P. Dick, Z.M. Mao, L. Yang, Accurate online power estimation and automatic battery behavior based power model generation for smartphones, in: Proceedings of the Eighth IEEE/ACM/IFIP International Conference on Hardware/Software Codesign and System Synthesis - CODES/ISSS, ACM Press, 2010, <http://dx.doi.org/10.1145/1878961.1878982>.
- [54] M. Ambrosin, A. Anzanpour, M. Conti, T. Dargahi, S.R. Moosavi, A.M. Rahmani, P. Liljeberg, On the feasibility of attribute-based encryption on Internet of Things devices, IEEE Micro 36 (6) (2016) 25–35, <http://dx.doi.org/10.1109/mm.2016.101>.