An approach to evaluate trust and reputation of things in a Multi-IoTs scenario

Computing (2020) 102:2257–2298
https://doi.org/10.1007/s00607-020-00818-5

An approach to evaluate trust and reputation of things in a Multi-IoTs scenario

Domenico Ursino*, and Luca Virgili

DII, Polytechnic University of Marche
d.ursino@univpm.it, l.virgili@pm.univpm.it
* Contact Author

Abstract

In the past research, trust and reputation have been investigated for communities of people, for organizations and for multi-agent systems. However, in the last few years, things are becoming increasingly relevant in the Internet scenario and, at the same time, increasingly complex. As a matter of fact, the term "Internet of Things" (hereafter, IoT) is becoming more and more common in both the scientific and the technological contexts. But, if a thing can have a profile and a behavior like a human, it is not out of place to extend the concept of trust and reputation to things and to define ad-hoc approaches for their computation. In this paper, we investigate trust and reputation of a thing in a Multiple IoTs scenario and we propose a context-aware approach to evaluate them. This task is not immediate because it should consider all the peculiarities of a thing compared to a human and all the specificities of a Multiple IoTs scenario compared to a community of people.

Keywords: Trust, Reputation, Internet of Things, Thing's Profile, MIoT, SIoT, Social Internetworking

1 Introduction

We experience the concepts of trust and reputation everyday; for example, when we buy something on an online service provider, in most cases, we do not have enough information about the service and/or the provider. This forces us to accept a "risk of prior performance", like paying for services and goods before receiving them. This information asymmetry can be mitigated thanks to the concepts of trust and reputation. The term *trust* is reported in literature with different nuances; therefore, it could be difficult to understand what it really is. However, as stated in [28], there are mainly two ways of defining trust. The first one is called *reliability trust*. This type of trust can be defined as the subjective probability by which an individual A expects that another individual B performs a given action from which its welfare depends [25]. The second type of trust is called *decision trust*. In this case, there is one party that is willing to depend on something or somebody in a given situation with a feeling of relative security [29]. Both these definitions involve the notions of *dependence* and *reliability* on the trusted entity and a certain *risk* related to a misbehavior of service provider. Starting from the concept of trustworthiness, it is possible to define the one of *reputation*. According to the Concise Oxford Dictionary [1], reputation is "the beliefs or opinions that are generally held about

someone or something". Therefore, the concept of trust is based on personal and subjective events described through factors and evidences. Instead, reputation can be considered as a collective measure of trustworthiness, based on advices or ratings from members of a community.

In the past computer science research, the concepts of trust and reputation have been investigated for communities of people, for organizations, for wireless sensor networks, for vehicular ad-hoc networks, and for multi-agent systems, and a lot of relevant results have been obtained [28, 45, 36, 8, 26, 21, 19, 6].

In the last few years, things are becoming increasingly important in the Internet scenario [9, 10, 11, 3, 4, 5, 33, 7, 14, 34, 27] and, presumably, in the future, the number of objects connected to the Internet will be much higher than the corresponding number of people. As a matter of fact, the term "Internet of Things" is becoming more and more common and, based on it, increasingly complex architectures [24, 23], requiring things to show a smart and social behavior [38, 22], are continuously proposed in literature. Social Internet of Things (hereafter, SIoT [9]), Multiple IoT Environment (hereafter, MIE [10]) and Multi Internet of Things (hereafter, MIoT [11]) are only three of the latest architectures with these characteristics.

This paper aims at providing a contribution in this setting. It starts from the MIoT architecture that, in our opinion, is the most general one proposed for modeling communities of things. In the MIoT model, things are organized in networks called IoTs. A thing can belong to one or more IoTs. Things belonging to more IoTs behave as "bridges" and allow communication and interaction between different IoTs of the MIoT. Things interact with each other through suitable transactions. The analysis of the information content exchanged by a thing with the other ones of the MIoT allows the construction of the thing's profile. The profile of a thing can be further enriched by considering the profiles of the things directly connected to it, according to the homophily principle characterizing social networks [30].

These are the same considerations that underlie the profiles of humans. As a consequence, most of the ideas and results about trust and reputation of a human in a community or of an agent in a multi-agent system can be extended and, possibly, redefined for a thing in a MIoT. Clearly, this extension is not immediate because it must consider all the peculiarities of a thing w.r.t. a human or an agent, and the specifities of a MIoT w.r.t. a community of people or a multi-agent system.

Investigating trust and reputation of things in a social context is extremely beneficial. Indeed, it has a lot of applications. Think, for instance, of the detection and isolation of a malicious object, the support of thing cooperation, the detection and the manipulation of thing reliability parameters, the evaluation of quality of services, just to cite a few of them. The presence of a thing's profile allows us to define context-aware notions of trust and reputation, along with suitable approaches for their computation. These notions are well suited to capture and address the complexity of the scenario we are investigating.

This paper is organized as follows: In Section 2, we present two examples motivating the usefulness of our approach. In Section 3, we provide an overview of related literature. In Section 4, we introduce some preliminary notions, in particular the MIoT paradigm and the concept of thing's profile. In Section 5, we describe the proposed approach. In Section 6, we illustrate the experimental campaign that we conducted to test it. In Section 7, we discuss the implications and the possible exploitation of the results obtained through our experiments. Finally, in Section 8, we draw our conclusions and have a look at possible future developments.

2 Motivating Examples

In this section, we present two use cases that allow us to better explain the usefulness of our approach. The former, illustrated in Subsection 2.1, regards the trust and reputation of smart objects in a smart city. It describes how a person can take advantage on the data exchanged by her smart objects and the ones of the city to improve the effectiveness and the efficiency of her activities and, ultimately, the quality of her life. The latter, described in Subsection 2.2, concerns the trust and reputation of smart objects in a shopping center. It illustrates how customers can use the data exchanged by their smart objects and the ones of the shopping centers to improve the effectiveness and the efficiency of the shopping tasks and, ultimately, the overall quality of their purchase experience.

2.1 Trust and reputation of smart objects in a smart city

As a first example case, consider some public areas (such as parks, squares, shopping centers, etc.) in a smart city, and assume that a group of people actively visits them. Each area is equipped with several smart objects for monitoring weather, air quality, traffic conditions, level of noise, etc., along with several actuators, such as smart lamps or information hubs provided as online services. Each person may have several smart devices, such as smartwatches, smartphones, other wearable devices, and so on. People and places can interact with each other through their smart objects [23].

Such a scenario can be modeled through a MIoT $\mathcal{M}$ consisting of a set $\{\mathcal{I}_1, \mathcal{I}_2, \cdots, \mathcal{I}_m\}$ of IoTs, each representing a public area. The set of the objects of $\mathcal{M}$ comprises the smart objects in the public areas and the set of personal devices of people visiting them. If an object o_j of the MIoT is active in the k^{th} public area, it has an instance ι_{j_k} in the IoT $\mathcal{I}_k$. Clearly, when a person with a smart object o_j moves around different public areas, corresponding to different IoTs, o_j will have different instances, one for each IoT.

Each visitor of an area is generally interested in a certain kind of activity; for instance, she could be a fitness runner. The final goal of the MIoT is supporting people to get the best experience from their activities. In this setting, trust and reputation can play a key role in reaching this objective. In the following, we report some possible usage scenarios.

Assume that a person wants to go out for a run. First, she needs to choose the best area for the run, based on weather conditions, traffic and other parameters that she considers relevant. To carry out her choices, she can check data provided by the sensors of each public area of her interest, the information hubs or other trusted runners. The choice of the information sources to consult is usually related to the trust and the reputation of the smart objects present therein. Once a person has performed her choices, she can decide to send this information to the MIoT in order to serve, in her turn, as information provider for the community.

A similar activity flow may happen in several other circumstances in which there is a decision to make, e.g., when a user must choose the best shopping center where she can buy a given object, the best cinema where she can see a movie, etc.

In all these cases, data regarding the choices of a user can be coupled with those registered during the activities she performed as a consequence of these choices (e.g., data coming from personal smartwears) in order to confirm the correctness of the choice or, on the contrary, to alert the other users of the evaluation errors. For instance, imagine a scenario in which a person verifies that the weather was actually much colder than the sensor in the public area seemed to indicate. In this case, the trust of the person in the sensors of that area decreases. This could also lead to a decrease of the

overall reputation of these sensors, thus influencing the decision of the other users. In particular, the reputation decrease of the smart objects of the public area determines how many users are impacted by the negative experience of the user and how much strong this impact is.

It is worth pointing out the relevance of the smart object reputation in this context. As a matter of fact, some smart objects of the MIoT could assume the role of reliable information hubs for the whole MIoT if their reputation is particularly strong and durable over time.

Trust and reputation may also have an important role in the detection and the management of possible anomalies characterizing one or more devices in the network. As an example, assume that a weather sensor in a public area is malfunctioning; in this case, all the objects relying on its data will be affected by this anomaly. First of all, this leads to a decrease of its trust and reputation. Furthermore, if one or more other trustworthy weather devices are present in the same area, they could help the whole MIoT to determine the sensor malfunction, to avoid the propagation of its effects and, finally, to repair it.

2.2 Trust and reputation of smart objects in a smart shopping center

Another possible scenario, where trust and reputation play an important role, is a big shopping center consisting of several buildings, each of them dedicated to specific product typologies, such as food, clothing, do-it-yourself, electronic devices, and so on. In this context, smart devices can be modeled by a MIoT $\mathcal{M}$ consisting of a set $\{\mathcal{I}_1, \mathcal{I}_2, \cdots, \mathcal{I}_m\}$ of IoTs, one for each building. The set of the objects of $\mathcal{M}$ consists of the set of the smart sensors present in each building (including video surveillance, temperature sensors, fire sensors, presence sensors, etc.) and the set of personal devices of visitors (including smartphones, tablets, smartwatches, etc.).

Each object o_j that interacts with the ones of the k^{th} building has an instance ι_{j_k} representing it in $\mathcal{I}_k$. Clearly, when the owner of an object o_j, such as a smartphone, moves throughout the buildings of the shopping center, o_j will have different instances associated with the different buildings of the center.

Here, a smart system of the shopping center could push offers to the enabled customer devices based on proximity, past preferences, habits, and so on. Analogously, based on the knowledge provided by the smart objects and the sensors dispersed in the shopping center, a personal device can suggest its owner the most comfortable and promising places to visit during her stay in the shopping center.

In this scenario, each person connected to the MIoT is interested in a certain kind of activity, somehow related to shopping. Indeed, users can play several roles ranging from vendors, suppliers or customers.

While a customer visits the building of a shopping center, her device may constantly locate the nearest ones and query for interesting products or offers. In the meantime, it could query other customers' smart objects (for instance, wearable devices) to measure her vital parameters in order to evaluate her pleasure in checking the products of a shopper. This can represent feedback information that the device supplies to the MIoT. Furthermore, a personal device of a customer can act as a personal shopper providing her with suitable suggestions. It interacts with the other objects of the MIoT, considers the offers of the shops, elaborates this information through machine learning algorithms, makes some proposals to its customer, registers her feedback and transmits them to the other devices in order to improve the quality of its recommendations.

Assume, now, that a customer wants to go out for shopping. First, she needs to locate the

best building to start with. This activity can be carried out by contacting her devices that act as personal shopper or by checking the preferred destinations of "special" customers (for instance, the most influential ones) or, again, by detecting the most comfortable shops. Once the desired knowledge has been obtained, the device can process it to make its suggestions. After the customer has made her choices and has performed her shopping activities, she can share information about her experience. In this way, she and/or her devices can become information providers for other customers. In this scenario, trust and reputation play an important role. For instance, the reputation of each smart object determines how many devices (and, ultimately, how many people) it can influence and how strong its influence is.

As in the previous scenario, an important issue to investigate and address is the presence of possible anomalies. The impact of an anomaly depends on several factors; the reputation of the affected objects is certainly one of the most important of them. As an example, given an anomaly of the device acting as a personal shopper, for instance the loss of historical data on product prices, the corresponding suggestions might not be the most convenient ones for its owner. In this case, the anomaly will certainly have a high impact on the device's owner. Furthermore, it can have an impact, even if smaller, on all the other objects (and, ultimately, on the corresponding customers) that it can reach and influence. The extension and the strength of the impact of an object o_j on an object o_q depend on the value of the trust of o_q in o_j, on the overall reputation of o_j and on the decrease of the reputation of o_j.

3 Related Literature

In computer science research, there is a plenty of papers addressing trust and reputation. Each of them proposes a model to handle these concepts from different points of view. However, as in most cases, the efficiency and the effectiveness of each model depend on the environment where it works.

As reported in [28], there are some features that allow the cataloguing of general trust and reputation models proposed in literature. These are: *(i)* trust classes, *(ii)* categories of trust semantics, *(iii)* reputation network architectures, and *(iv)* reputation compute engines. As for this last issue, there are several families of approaches to compute trust and reputation. For instance, some possible families could be based on: *(i)* the sum or the average of ratings, *(ii)* fuzzy operators, *(iii)* "flow" models computing trust and reputation scores through looped or long chains. For example, Google's PageRank [32] belongs to this last family.

Before deeping on trust and reputation in the IoT scenario, it is necessary to spend some words on the computation of these measures in an online service provisioning [46, 35], which is the first Internet context where these concepts have been applied. One of the most famous reputation systems is the eBay's one [35]. In this system, after each purchase, the buyer and the seller have the opportunity to rate each other as positive, negative or neutral. The architecture is centralized and the central authority computes the reputation score of each participant as the sum of positive and negative ratings. Even if this system is primitive and can be quite misleading, it seems to have a strong positive impact in the marketplace. On the other hand, there are experts' sites where pools of individuals are willing to answer questions in their areas of expertise. For instance, AskMe is one of these sites; in this system, a participant has to pay a fee to take part to the corresponding network.

Another important contribution in trust management is reported in [15]. In this paper, the authors introduce a framework, called Socialtrust, aiming at analyzing the reliability of information exchanged

in online social networks. In order to perform trust computation, Socialtrust considers three factors, namely: *(i)* the trust group feedback, *(ii)* the difference between the user's perceived quality and the trust concept, and *(iii)* the tracking of user behavior. The authors also describe the application of Socialtrust to MySpace profiles.

In [39], the authors propose an approach to find users in a social network, who are able to spread a specific information as far as possible. In particular, a company selects a set of people, who are willing to send advertisements to their friends in order to get discounts or free goods. If a user is highly respected by her friends, her advertisements will be probably spread over the social network.

Finally, another notable reputation system is the PageRank [32]. In this case, the collection of hyperlinks to a given page can be exploited to evaluate the reputation score of that page.

After a description of trust management in online service provisioning, we examine the transpositions of all these concepts to the IoT context. A well-defined reputation system is really relevant in IoT, because it is necessary to manage the services provided by objects. As previously described, the concept of trust encompasses factors like the goodness of a service or the reliability and the availability of an object.

The relevance of trust management in an IoT context is investigated in [47]. Here, the authors show how trust management can favor data fusion and mining, privacy and information security.

In literature, it is possible to catalogue trust and reputation approaches operating in IoT scenarios according to the type of architecture that the authors decided to develop. Based on it, three different kinds of model can be recognized, namely: *(i)* centralized, *(ii)* semi-centralized, and *(iii)* distributed ones.

An example of a centralized model is proposed in [37]. This model can evaluate the context in which objects work. In this architecture, there is a node called Trust manager, which handles all the information related to the trustworthiness of agents in the IoT. The evaluation approach consists of five phases, namely: *(i)* information gathering; *(ii)* entity selection; *(iii)* transaction; *(iv)* reward or punish, and *(v)* learning. Roughly speaking, when an object requires a service, it asks to the trust manager a list of trustable nodes offering that service. Then, after the transaction completion, it sends a report to the trust manager that contains a positive score (reward) or a negative one (punish) regarding the service provider.

Another example of a centralized system is described in [18], where the authors propose a model called Trusted Resource Sharing (TRS). TRS has three main components, namely Trust, Usage and Relation. The Trust component, which is the one of interest to this paper, is developed through a centralized architecture, in which there is an entity managing trust and resource policies. Trust is evaluated for both objects and resources available in the network. However, IoT is expected to exponentially grow in the next future, so a central authority could be a bottleneck. Indeed, a failure on the Trust Manager can block the whole network. Furthermore, the Trust Manager could also be attacked to change the trust and the reputation of each participant.

A further centralized system is described in [42]. Here, the authors propose a trust model called "REK", developed in a SIoT context. In order to evaluate trustworthiness, REK leverages two indicators, namely "Experience" and "Reputation". Both these indicators are modeled using mathematical tools and are extracted from previous interactions among entities in the SIoT environment. The "Experience" component is modeled by means of PageRank [32].

As far as semi-centralized architectures are concerned, [41] developed a model operating on a SIoT. It introduces three new components into the SIoT, namely Trust Agent, Trust Broker and Trust

Analysis and Management. The proposed model focuses on the social parameters of IoT. Specifically, the authors examine honesty, cooperativeness, and community-interest. These parameters contribute to the building of a knowledge (reported as Knowledge Trust Management - KTM) useful to evaluate the reputation of an agent. Knowledge grows in a huge way over time. In order to manage this big amount of data, the authors propose a fuzzy-based model, capable of representing attributes in vague terms, like "low" or "high", "bad", "acceptable" or "good". KTM is a good way to build up a consistent reputation system capable of adapting to different situations.

As for distributed architectures, there are several works proposing distributed models to compute trust and reputation. One of the first models belonging to this category is described in [16]. Here, the authors study a community of sensors in a Wireless Sensor Network. This model builds two types of reputation, namely direct reputation, computed through personal observations, and indirect one, based on the recommendations of other nodes. Analogously to the model of [41], the one of [16] is based on fuzzy theory. It represents a starting point for future developments in the computation of trust and reputation among IoT devices. However, it considers only a specific IoT environment, with a limited number of measures. The model proposed in [12] is an extension of the one introduced in [16]. Here, the authors describe a dynamic protocol, aiming at addressing the limits presented in previous ones. The concept of trust is modeled through three elements: honesty, cooperativeness and community-interest (these are the same elements described in [41]). As it is a distributed model, each node updates the trust ratings of the other nodes by collaborating with them.

Another distributed system is described in [44]. Here, trust and reputation are computed by analyzing each device from three different viewpoints, namely sensor, core and application ones. Each of these layers has its own trust information. In the sensor layer, trust denotes which node must be contacted for a service. Instead, in the core layer, trust is used to select a set of networks and routes, through which data can be sent. Finally, in the application layer, trust is exploited to evaluate which candidate method of data processing or which storage service are trusted. These three trust scores are, then, composed through fuzzy logic.

A further interesting distributed system is described in [17]. Here, the authors develop a technique to compute trust on an IoT based on the Service Oriented Architecture. This technique aims at selecting feedbacks using rating similarities, communities of interest and social contacts. The whole model is based on a collaborative filtering approach. To guarantee scalability, each node saves trust information only about a subset of nodes of interest and performs a minimum computation to update these values.

Finally, [31] introduces two trust and reputation models for the SIoT environment. The former is called *subjective trustworthiness*. In this case, each node computes the trustworthiness of its neighbors based on its experience and the one of them. Trust computation considers five viewpoints, namely: *(i)* direct opinion; *(ii)* indirect opinion; *(iii)* long-term opinion; *(iv)* relationship factor; *(v)* direct opinion in the credibility. The latter is called *objective trustworthiness*. It is defined in a Peer-to-Peer (P2P) scenario, in which information of each node is visible and managed through special nodes, called Pre-Trusted Objects. In this case, trust computation considers four points of view, namely: *(i)* long-term opinion; *(ii)* short-term opinion; *(iii)* relationship factor in credibility; *(iv)* intelligence in credibility. In both models, the credibility of a node is used to evaluate the opinion of other nodes. Therefore, these models give a high weight to recommendations made by "good" friends and a low weight to feedbacks provided by "bad" friends.

4 Preliminary concepts

4.1 The MIoT paradigm

In this section, we provide an overview of the MIoT paradigm, described in detail in [11], because it is the reference one for our definitions of trust and reputation of a thing in a multi-network scenario.

A MIoT $\mathcal{M}$ consists of a set of m IoTs. Formally speaking: $\mathcal{M} = \{\mathcal{I}_1, \mathcal{I}_2, \cdots, \mathcal{I}_m\}$, where $\mathcal{I}_k$ is an IoT.

Let o_j be an object of $\mathcal{M}$. We assume that, if o_j belongs to $\mathcal{I}_k$, it has an instance ι_{j_k}, representing it in $\mathcal{I}_k$. The instance ι_{j_k} consists of a virtual view (or, better, a virtual interface) representing o_j in $\mathcal{I}_k$. For example, it provides all the other instances of $\mathcal{I}_k$, and the users interacting with it, with all the necessary information about o_j. Information stored in ι_{j_k} is represented according to the format and the conventions adopted in $\mathcal{I}_k$.

$\mathcal{M}$ can be represented by means of a graph-based notation. In particular, each IoT $\mathcal{I}_k \in \mathcal{M}$ can be modeled as a graph $G_k = \langle N_k, A_k \rangle$. In this case: N_k is the set of the nodes of G_k; there is a node n_{j_k} for each instance $\iota_{j_k} \in \mathcal{I}_k$, and vice versa. Since there is a biunivocal correspondence between a node and an instance, in the following, we use these two terms interchangeably. A_k is the set of the arcs of G_k; there is an arc $a_{jq_k} = (n_{j_k}, n_{q_k})$ if there exists a link from n_{j_k} to n_{q_k}.

Finally: $\mathcal{M} = \langle N, A \rangle$. Here:

- $N = \bigcup_{k=1}^{m} N_k$;
- $A = A_I \cup A_C$, where $A_I = \bigcup_{k=1}^{m} A_k$ and $A_C = \{(n_{j_k}, n_{j_q}) | n_{j_k} \in N_k, n_{j_q} \in N_q, k \neq q\}$.

A_I is the set of the inner arcs (hereafter, *i-arcs*) of $\mathcal{M}$; they link instances (of different objects) belonging to the same IoT. A_C is the set of the cross arcs (hereafter, *c-arcs*) of $\mathcal{M}$; they link instances of the same object belonging to different IoTs.

Now, we can introduce the concept of neighborhood of an instance ι_{j_k} in $\mathcal{I}_k$. Specifically, the neighborhood $nbh(\iota_{j_k})$ of ι_{j_k} is defined as:

$$nbh(\iota_{j_k}) = nbh^{out}(\iota_{j_k}) \cup nbh^{in}(\iota_{j_k})$$

where:

$$nbh^{out}(\iota_{j_k}) = \{\iota_{q_k} | (n_{j_k}, n_{q_k}) \in A_I, |tranSet_{jq_k}| > 0\}$$

and

$$nbh^{in}(\iota_{j_k}) = \{\iota_{q_k} | (n_{q_k}, n_{j_k}) \in A_I, |tranSet_{qj_k}| > 0\}$$

In other words, $nbh(\iota_{j_k})$ comprises those instances directly connected to ι_{j_k} through an incoming or an outgoing arc, which shared at least one transaction with ι_{j_k}.

Furthermore, in a MIoT, *a set MD_j of metadata* can be associated with an object o_j. Our metadata model refers to the one of the IPSO (Internet Protocol for Smart Objects) Alliance [2]. Specifically MD_j consists of three subsets, namely: *(i)* MD_j^D, i.e., the set of *descriptive metadata*; *(ii)* MD_j^T, i.e., the set of *technical metadata*; *(iii)* MD_j^B, i.e., the set of *behavioral metadata*. All details about these metadata can be found in [11].

Given a pair of instances ι_{j_k} of o_j and ι_{q_k} of o_q in $\mathcal{I}_k$, our model registers the set $tranSet_{jq_k}$ of the transactions from ι_{j_k} to ι_{q_k}.

The set $tranSet_{jq_k}$ is defined as:

$$tranSet_{jq_k} = \{Tr_{jq_{k_1}}, Tr_{jq_{k_2}}, \cdots, Tr_{jq_{k_v}}\}$$

A transaction $Tr_{jq_{k_t}} \in tranSet_{jq_k}$ is represented as:

$$Tr_{jq_{k_t}} = \langle reason_{jq_{k_t}}, source_{jq_{k_t}}, dest_{jq_{k_t}}, start_{jq_{k_t}}, finish_{jq_{k_t}}, success_{jq_{k_t}}, content_{jq_{k_t}} \rangle$$

Here: *(i)* $reason_{jq_{k_t}}$ denotes the reason why $Tr_{jq_{k_t}}$ occurred, chosen among a set of predefined values; *(ii)* $source_{jq_{k_t}}$ indicates the starting node of the path followed by $Tr_{jq_{k_t}}$; *(iii)* $dest_{jq_{k_t}}$ represents the final node of the path followed by $Tr_{jq_{k_t}}$; *(iv)* $start_{jq_{k_t}}$ denotes the starting timestamp of $Tr_{jq_{k_t}}$; *(v)* $finish_{jq_{k_t}}$ indicates the ending timestamp of $Tr_{jq_{k_t}}$; *(vi)* $success_{jq_{k_t}}$ denotes whether $Tr_{jq_{k_t}}$ was successful or not; it is set to *true* in the affirmative case, to *false* in the negative one, and to `NULL` if $Tr_{jq_{k_t}}$ is still in progress; *(vii)* $content_{jq_{k_t}}$ indicates the content "exchanged" from ι_{j_k} to ι_{q_k} during $Tr_{jq_{k_t}}$.

In its turn, $content_{jq_{k_t}}$ presents the following structure:

$$content_{jq_{k_t}} = \langle format_{jq_{k_t}}, fileName_{jq_{k_t}}, size_{jq_{k_t}}, topics_{jq_{k_t}} \rangle$$

Here: *(i)* $format_{jq_{k_t}}$ indicates the format of the content exchanged during $Tr_{jq_{k_t}}$; the possible values are: "audio", "video", "image" and "text"; *(ii)* $fileName_{jq_{k_t}}$ denotes the name of the transmitted file; *(iii)* $size_{jq_{k_t}}$ indicates the size in bytes of this content; *(iv)* $topics_{jq_{k_t}}$ denotes the set of the content topics; it consists of a set of keywords representing the subjects exchanged during $Tr_{jq_{k_t}}$. It can be formalized as: $topics_{jq_{k_t}} = \{kw^1_{jq_{k_t}}, kw^2_{jq_{k_t}}, \ldots, kw^w_{jq_{k_t}}\}$.

Now, we can define the set $tranSet_{j_k}$ of the transactions performed by ι_{j_k} in $\mathcal{I}_k$. Specifically, let $Inst_k$ be the set of the instances of $\mathcal{I}_k$. Then:

$$tranSet_{j_k} = \bigcup_{\iota_{q_k} \in Inst_k, \iota_{q_k} \neq \iota_{j_k}} tranSet_{jq_k}$$

In other words, the set $tranSet_{j_k}$ of the transactions performed by an instance ι_{j_k} is given by the union of the sets of the transactions from ι_{j_k} to all the other instances of $\mathcal{I}_k$.

4.2 Definition of a thing's profile

In this section, we present our definitions of instance and object profiles. They represent a preliminary knowledge, mandatory to fully understand the rest of our approach. To introduce them, we need to present the following operators:

- $\uplus$: it receives a set $\{entitySet_1, entitySet_2, \cdots, entitySet_t\}$ of entity sets and performs their union, not eliminating duplicates but reporting the number of their occurrences. Therefore, this operator returns a set of pairs $\{(entity_1, ne_1), (entity_2, ne_2), \cdots, (entity_w, ne_w)\}$ in which the pair $(entity_r, ne_r)$ indicates the r^{th} entity and the number of its occurrences. In counting the number of occurrences, $\uplus$ takes the presence of synonymies and homonymies into account. These properties can be computed (for terms, images, etc.) by applying the classical approaches proposed in past literature [13, 20].

- $avgFileSize$: it receives a set of files and computes their average size.

We are now able to define the profile $\mathcal{P}_{jq_k}$ of the relationship existing between two instances ι_{j_k} and ι_{q_k}, which performed a set $tranSet_{jq_k} = \{Tr_{jq_{k_1}}, Tr_{jq_{k_2}}, \cdots, Tr_{jq_{k_v}}\}$ of transactions. Specifically:

$$\mathcal{P}_{jq_k} = \langle reasonSet_{jq_k}, sourceSet_{jq_k}, destSet_{jq_k}, avgSzAudio_{jq_k}, avgSzVideo_{jq_k}, \\ avgSzImage_{jq_k}, avgSzText_{jq_k}, successFraction_{jq_k}, topicSet_{jq_k} \rangle$$

where:

- $reasonSet_{jq_k} = \biguplus_{t=1..v}(reason_{jq_{k_t}})$;
- $sourceSet_{jq_k} = \biguplus_{t=1..v}(source_{jq_{k_t}})$;
- $destSet_{jq_k} = \biguplus_{t=1..v}(dest_{jq_{k_t}})$;
- $avgSzAudio_{jq_k} = AvgFileSize_{t=1..v}\{fileName_{jq_{k_t}} | format_{jq_{k_t}} = \text{“}audio\text{”}\}$;
- $avgSzVideo_{jq_k} = AvgFileSize_{t=1..v}\{fileName_{jq_{k_t}} | format_{jq_{k_t}} = \text{“}video\text{”}\}$;
- $avgSzImage_{jq_k} = AvgFileSize_{t=1..v}\{fileName_{jq_{k_t}} | format_{jq_{k_t}} = \text{“}image\text{”}\}$;
- $avgSzText_{jq_k} = AvgFileSize_{t=1..v}\{fileName_{jq_{k_t}} | format_{jq_{k_t}} = \text{“}text\text{”}\}$;
- $successFraction_{jq_k} = \frac{|\{Tr_{jq_{k_t}} | Tr_{jq_{k_t}} \in tranSet_{jq_k}, success_{jq_{k_t}} = true\}|}{v}$;
- $topicSet_{jq_k} = \biguplus_{t=1..v}(topics_{jq_{k_t}})$.

If we introduce the operator $\bigsqcup$, which compactly represents the set of the operations described above, needed to obtain a profile of a pair of instances $\mathcal{P}_{jq_k}$ starting from the corresponding transactions, we can formalize the previous tasks with only one operation as:

$$\mathcal{P}_{jq_k} = \bigsqcup_{t=1..v} Tr_{jq_{k_t}}$$

Furthermore, let ι_{j_k} be the instance of the object o_j in the IoT $\mathcal{I}_k$. Let $Inst_{j_k}$ be the set of the instances of $\mathcal{I}_k$ with which ι_{j_k} performed at least one transaction in the past. In this case, we can define the profile $\mathcal{P}_{j_k}$ of ι_{j_k} as :

$$\mathcal{P}_{j_k} = \bigsqcup_{\iota_{q_k} \in Inst_{j_k}} \mathcal{P}_{jq_k}$$

Finally, let o_j be an object and let $\{\mathcal{I}_1, \mathcal{I}_2, \cdots, \mathcal{I}_l\}$ be the set of the IoTs it participates to. Let $ObjInst_j$ be the instances of o_j in the IoTs of the MIoT. We can define the profile $\mathcal{P}_j$ of o_j as:

$$\mathcal{P}_j = \bigsqcup_{\iota_{j_k} \in ObjInst_j} \mathcal{P}_{j_k}$$

Everything we have seen so far regards the profile of an instance from a “content-based” perspective (i.e., taking its past behavior into account). Beside this perspective, another one can be considered, i.e., the “collaborative filtering” perspective (i.e., based on the similarity of the behaviors of the instance neighbors). However, it is out of the scope of this paper.

4.3 Example

Consider a smart shopping center as the one described in Section 2.2. It consists of three buildings, one for each store. These last are a supermarket, an electronics store and a clothing store (see Figure 1). We can associate a MIoT $\mathcal{M}$ with this center. $\mathcal{M}$ consists of three IoTs:

$$\mathcal{M} = \{\mathcal{I}_1, \mathcal{I}_2, \mathcal{I}_3\}$$

$\mathcal{I}_1$ (resp., $\mathcal{I}_2$, $\mathcal{I}_3$) connects all the instances of the smart objects of people accessing the supermarket (resp., electronics store, clothing store).

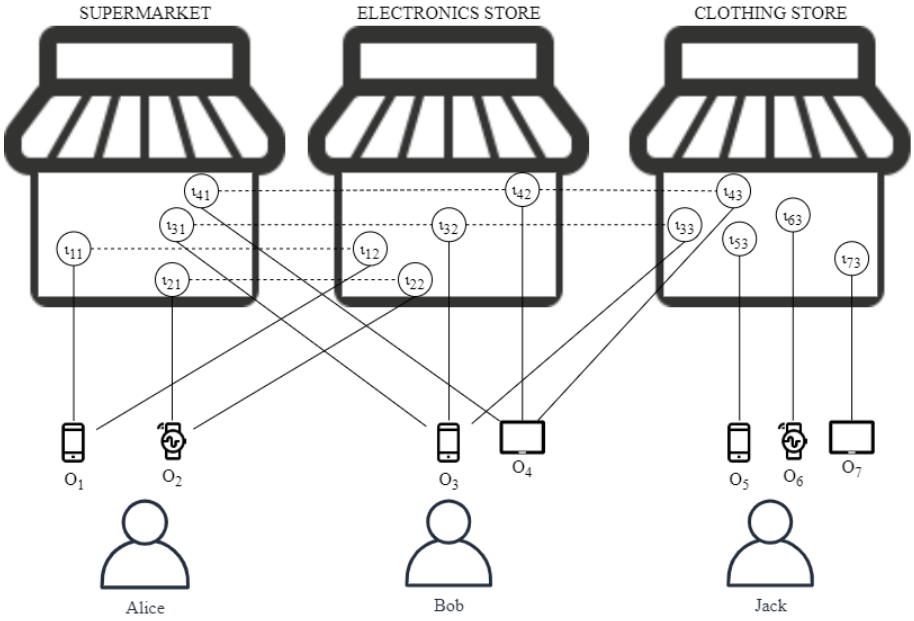

Figure 1: An example of a MIoT associated with a smart shopping center

Consider three customers: *(i)* Alice, who owns a smartphone o_1 and a smartwatch o_2; *(ii)* Bob, who owns a smartphone o_3 and a tablet o_4; *(iii)* Jack, who owns a smartphone o_5, a smartwatch o_6 and a tablet o_7.

Alice visits the supermarket and the electronics store. ι_{1_1} (resp., ι_{1_2}) represents the instance of the Alice's smartphone in $\mathcal{I}_1$ (resp., $\mathcal{I}_2$); instead, ι_{2_1} (resp., ι_{2_2}) denotes the instance of the Alice's smartwatch in $\mathcal{I}_1$ (resp., $\mathcal{I}_2$).

Analogously, Bob visits the supermarket, the electronics store and the clothing store. ι_{3_1}, ι_{3_2} and ι_{3_3} (resp., ι_{4_1}, ι_{4_2} and ι_{4_3}) denote the instances of the smartphone (resp., tablet) of Bob in the three stores.

Finally, Jack visits only the clothing store. He has a smartphone o_5, a smartwatch o_6 and a tablet o_7; ι_{5_3}, ι_{6_3} and ι_{7_3} represent the instances of these smart objects in $\mathcal{I}_3$.

In Figure 1, the dashed line between ι_{1_1} and ι_{1_2} indicates that they are two instances of the same object o_1. An analogous semantics regards the dashed lines between ι_{2_1} and ι_{2_2}, ι_{3_1} and ι_{3_2}, ι_{4_1} and ι_{4_2}, ι_{3_2} and ι_{3_3} and, finally, ι_{4_2} and ι_{4_3}.

5 Proposed approach

5.1 Architecture

Figure 2 shows the MIoT architecture that we designed for supporting our approach. Each colored circle represents a distinct IoT of the MIoT. The stack composed by "Transaction metadata", "Instance metadata" and "Object metadata" handles the transaction, instance and object metadata described in Section 4.1.

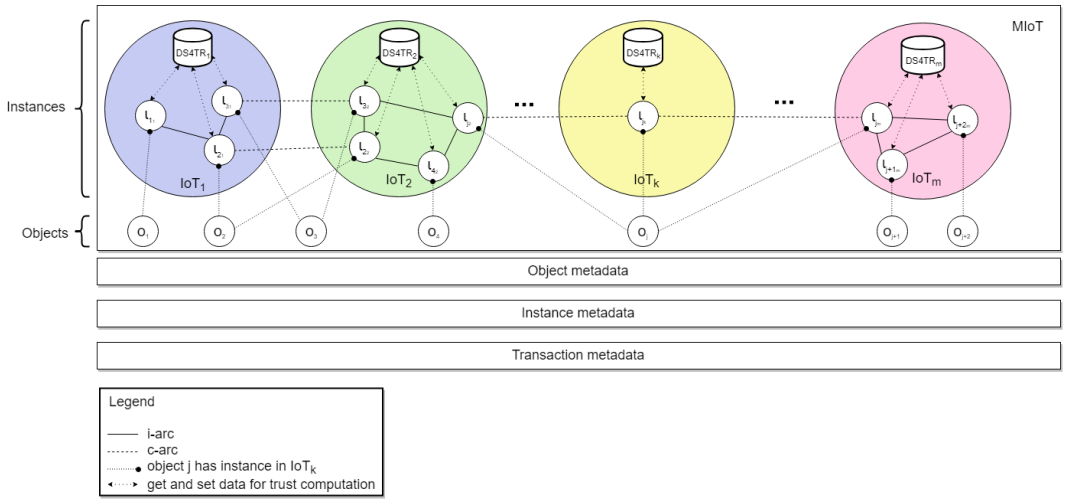

Figure 2: Schematic representation of the proposed MIoT architecture

The aim of this architecture is to provide a scalable way to manage both the storage of support data and the tasks necessary for the computation of trust and reputation. We propose to add a repository, called *DS4TR* (Data Storage for Trust and Reputation) in each IoT, which stores the data necessary for the computation of trust and reputation. From a logical viewpoint, *DS4TR* is separated from the other objects of the corresponding IoT. Actually, from an implementation point of view, *DS4TR* could be either deployed through a cloud service or embedded in one of the objects operating in the IoT. As for security issues involving *DS4TR*, the following reasonings hold: As previous highlighted, the overall scenario consists of two different cases, one in which *DS4TR* is provided by a cloud service and another in which it is part of an IoT. In the former case, *DS4TR* represents a trusted-third party. For the sake of space, and since this issue does not present the core of this paper, we do not describe this case in detail. We only refer to relevant techniques to protect trusted-third parties in a cloud environment [40, 48] already proposed in past literature.

In the latter case, which is a scenario typical of MIoT, each object is a peer of a P2P architecture and *DS4TR* is part of an IoT. In order to obtain a stable and reliable IoT, it is fundamental to define a good strategy to build *DS4TR*. For instance, we can select some reliable nodes from the IoT. To assure node reliability, we can exploit our trust and reputation model to compute a ranking of nodes, ordered by their reputation values. Then, we can choose the first τ ones to compose a *DS4TR*. Recall that, as a typical situation in a P2P scenario, each selected object stores only a part of the whole data repository. However, in order to maintain a certain level of fault-tolerance, some parts can be replicated on multiple objects. The parameter τ represents a tradeoff between reliability and

performance of the network. The greater τ, the more trustworthy the network, but the lower its speed. Of course, the setting of τ depends on the context in which the overall model is developed. However, once a node is selected to be part of a *DS4TR*, it has to save a portion of the trust and reputation repository of its IoT.

Another interesting aspect to consider is how to check whether the nodes contributing to a *DS4TR* are properly working or not. Each transaction made by a node is part of our model, so that we can compute the trust and reputation values of these objects. To maintain a high level of reliability, we can set a threshold, say th_{rep}, which represents the minimum reputation value that a node must have to be part of a *DS4TR*. If a node of *DS4TR* obtains a reputation value lower than th_{rep}, it leaves the repository and another node is chosen (by following an approach similar to the one presented above) to replace it.

Beside *DS4TR*, cross-nodes play an important role in the computation of trust and reputation. Indeed, they allow every node of an IoT to ask for trust data regarding participants to other IoTs of the MIoT. Finally, after instances completed a transaction, each of them has to add a feedback about the other part. Obviously, this feedback has to be added in the *DS4TR* node(s) corresponding to the transaction participants.

This architecture can face scalability issues because each IoT has its own repository to save data. In this way, the problem of bottlenecks in the network is highly mitigated. A careful reader could point out that requests coming from different IoTs could overwhelm a *DS4TR* node. However, in the intrinsic architecture of a MIoT, there are much less transactions between two different IoTs than within an IoT.

5.2 Definition of trust and reputation

In this section, we provide our definitions of the different levels of trust and reputation in a MIoT.

5.2.1 Trust of an instance in another one of the same IoT

Let ι_{j_k} and ι_{q_k} be two instances of an IoT $\mathcal{I}_k$. We want to define the trust T_{jq_k} of ι_{j_k} in ι_{q_k}. Actually, this trust is not unique, because it depends on both the topic and the format of the data exchanged during the corresponding transactions. As a consequence, T_{jq_k} is a matrix and has a value for each possible combination of topics and formats. Since the possible formats are 4 (i.e., "audio", "video", "image", and "text"), T_{jq_k} is a $|topicSet_{jq_k}| \times 4$ matrix. The element $T_{jq_k}[u, v]$ of this matrix indicates the trust of ι_{j_k} in ι_{q_k} regarding the topic u delivered in the format v. This trust depends on several factors, namely: *(i)* the fraction of successful transactions; *(ii)* the overall number of transactions, which is an indicator of the robustness of the result; *(iv)* the size of exchanged files; *(v)* the timestamp of the last transaction, which is an indicator of the possible obsolescence of the relationship between ι_{j_k} and ι_{q_k}.

In order to define $T_{jq_k}[u, v]$, we must introduce some notions. Specifically:

- $tranSet_{jq_k}[u, v]$ is the subset of the transactions $tranSet_{jq_k}$ whose content presents the topic u in the format v at least once;
- $OKTranSet_{jq_k}[u, v]$ is the fraction of successful transactions in $tranSet_{jq_k}[u, v]$;

- $maxNumTranSet_k[u, v]$ is the maximum number of transactions, concerning the topic u in the format v, performed between two given instances of $\mathcal{I}_k$. It is defined as:

$$maxNumTranSet_k[u, v] = max_{\iota_{j_k} \in Inst_k, \iota_{q_k} \in Inst_k, \iota_{j_k} \neq \iota_{q_k}} |tranSet_{jq_k}[u, v]|$$

- $size_{jq_k}[u, v]$ is the size of contents concerning the topic u in the format v exchanged by ι_{j_k} and ι_{q_k}.

- $maxSize_k[u, v]$ is the maximum size of the overall contents concerning the topic u in the format v, exchanged between two instances of $\mathcal{I}_k$. It is defined as:

$$maxSize_k[u, v] = max_{\iota_{j_k} \in Inst_k, \iota_{q_k} \in Inst_k, \iota_{j_k} \neq \iota_{q_k}} size_{jq_k}[u, v]$$

We are now able to define $T_{jq_k}[u, v]$. It consists of a pair $(V_{jq_k}[u, v], LTS_{jq_k}[u, v])$. $V_{jq_k}[u, v]$ can be computed as a weighted mean of the parameters introduced above. Specifically:

$$V_{jq_k}[u, v] = \frac{\alpha \cdot OKTranSet_{jq_k}[u, v] + \beta \cdot \frac{|tranSet_{jq_k}[u,v]|}{maxNumTranSet_k[u,v]} + \rho \cdot \frac{size_{jq_k}[u,v]}{maxSize_k[u,v]}}{\alpha + \beta + \rho}$$

Here, α, β and ρ denote the weights of the three components of the mean. We have experimentally set $\alpha = 0.55$, $\beta = 0.35$ and $\rho = 0.10$ (see Section 6.1 for all details).

$LTS_{jq_k}[u, v]$ is the last ending timestamp concerning a transaction of $tranSet_{jq_k}[u, v]$.

Example (...cnt'd)

Consider the smart shopping center described in Section 2.2 and assume that Alice and Bob enter the electronics store. Assume, also, that Alice needs information about smart home products. In this case, the smartphone of Alice asks the other smart objects of the customers in the store if they have information about smart home products sold there (for example, in order to know the current promotions of the store). Assume that Bob had already visited the store several times in the last days for searching information and buying a smart home product. The smartphone of Bob (o_3) can answer the smartphone of Alice (o_1) and the corresponding instances ι_{1_2} and ι_{3_2} can start their interaction. Here:

- $tranSet_{13_2}$ denotes the transactions exchanged between o_1 and o_3 in the electronics store (i.e., between ι_{1_2} and ι_{3_2}). Assume that $tranSet_{13_2} = 150$.

- $tranSet_{13_2}[u, v]$ represents the transactions of $tranSet_{13_2}$ about the topic u (in our case, smart home products) in the format v (for instance, videos). Assume that $tranSet_{13_2}[u, v] = 90$.

- $OKTranSet_{13_2}[u, v]$ is the fraction of successful transactions of $tranSet_{13_2}[u, v]$. Suppose that some transactions of $tranSet_{13_2}[u, v]$ failed because the smartphone of Bob had connection problems with the WiFi of the store. Assume that the successful transactions of $tranSet_{13_2} = 85$ so that $OKtranSet_{13_2}[u, v] = \frac{85}{90} = 0.94$.

- $maxNumTranSet[u, v]$ is the maximum number of videos about smart home products exchanged between two smart objects in the electronics store. Assume that $maxNumTranSet[u, v] = 110$.

- $size_{13_2}[u, v]$ is the overall size of the videos concerning smart home products exchanged between the smartphones of Alice and Bob. Assume that $size_{13_2}[u, v] = 10\ MB$.
- $maxSize[u, v]$ is the maximum overall size of the videos about smart home products exchanged between two smart objects in the electronics store. Assume that $maxSize[u, v] = 12\ MB$.
- $V_{13_2}[u, v]$ is the value of the trust about videos on smart home products that the smartphone of Alice has in the smartphone of Bob in the electronics store. It is equal to:

$$V_{13_2}[u, v] = \frac{0.55 \cdot 0.94 + 0.35 \cdot \frac{90}{110} + 0.10 \frac{10}{12}}{0.55 + 0.35 + 0.10} = 0.89$$

- $LTS_{13_2}[u, v]$ is the timestamp of the last video about smart home products that the smartphone of Bob sent to the smartphone of Alice in the electronics store.

5.2.2 Trust of an object in another one of the MIoT

Let o_j and o_q be two objects of $\mathcal{M}$. Let $\mathcal{M}_{jq} = \{\mathcal{I}_1, \cdots, \mathcal{I}_l\}$ be the subset of the IoTs of $\mathcal{M}$ that simultaneously contain one instance of o_j and one instance of o_q. In this case, it is possible to define the trust $T_{jq}[u, v]$ of o_j in o_q regarding the topic u delivered in the format v. Also in this case, $T_{jq}[u, v]$ consists of a pair $(V_{jq}[u, v], LTS_{jq}[u, v])$. Here:

- $V_{jq}[u, v]$ is set to the average of the trusts that the instances of o_j have in the instances of o_q in the IoTs of $\mathcal{M}_{jq}$:

$$V_{jq}[u, v] = \frac{\sum_{k=1..l} V_{jq_k}[u, v]}{l}$$

- $LTS_{jq}[u, v]$ is set to the maximum ending timestamp of any transaction simultaneously involving one instance of o_j and one instance of o_q:

$$LTS_{jq}[u, v] = max_{k=1..l} LTS_{jq_k}[u, v]$$

Example (...cnt'd)
Consider the smart shopping center described in Section 2.2. Assume that Alice and Bob first enter the supermarket and then the electronics store. Consider the smartphone of Alice (o_1) and the one of Bob (o_3) and assume that they interact in both stores to help Alice find the smart home products she desires.

Assume that the value of the trust about videos on smart home products that the smartphone of Alice has in the smartphone of Bob in the supermarket is equal to 0.93[1]. In the example in Section 5.2.1, we have seen that the trust about videos on smart home products that the smartphone of Alice has in the smartphone of Bob in the electronics store was equal to 0.89.

As a consequence, the value of the overall trust about videos on smart home products that the smartphone of Alice has in the smartphone of Bob in the whole smart shopping center is:

[1]This value is obtained by proceeding in the same way as we did for the electronics store in Section 5.2.1.

$$V_{13}[u,v] = \frac{0.93+0.89}{2} = 0.91$$

Instead, $LTS_{13}[u,v]$, i.e. the ending timestamp of the last video on smart home products that the smartphone of Bob sent to the smartphone of Alice coincides with the value of $LTS_{13_2}[u,v]$ computed in the example of Section 5.2.1, because Alice ad Bob entered first the supermarket and then the electronics store.

5.2.3 Reputation of an instance in an IoT

Let $\mathcal{I}_k$ be an IoT and let ι_{j_k} be an instance of $\mathcal{I}_k$. The reputation $R_{j_k}[u,v]$ of ι_{j_k}, regarding the topic u in the format v, depends on the following factors: *(i)* the number of instances from which ι_{j_k} received transactions in the past; *(ii)* the trust that these instances have in ι_{j_k}; *(iii)* the reputation of these instances in $\mathcal{I}_k$; *(iv)* their oldness.

To formalize this type of dependencies, the classical approach involves the usage of the PageRank formula. To proceed in this direction, it is necessary to introduce the following preliminary definitions:

- $Age_{q_k}[u,v]$ is the number of days spent from the first transaction, regarding the topic u delivered in the format v, performed by the object o_q in the IoT $\mathcal{I}_k$.
- $Age_k^{max}[u,v]$ is the maximum number of days spent from the first transactions, regarding the topic u delivered in the format v, performed by an object in the IoT $\mathcal{I}_k$.
- $R_k^{max}[u,v]$ is the maximum reputation of an instance of $\mathcal{I}_k$, regarding the topic u delivered in the format v.

We are now able to define the formula for the computation of $R_{j_k}[u,v]$. In particular:

$$R_{j_k}[u,v] = \gamma + (1-\gamma) \cdot \frac{\sum_{\iota_{q_k} \in nbh^{in}(\iota_{j_k})} T_{qj_k}[u,v] \cdot R_{q_k}[u,v] \cdot \frac{Age_{q_k}[u,v]}{Age_k^{max}[u,v]}}{|nbh^{in}(\iota_{j_k})|}$$

In this formula, γ is the damping factor generally adopted in the PageRank. It determines the minimum absolute reputation assigned to an instance of $\mathcal{I}_k$. From a more abstract viewpoint, it allows us to tune the fraction of the absolute reputation that ι_{j_k} transmits to ι_{q_k}.

$R_{j_k}[u,v]$ belongs to the real interval $[\gamma, +\infty)$. In order to obtain a reputation value belonging to the interval $[0,1]$ and, at the same time, to normalize the reputations of the instances of the IoTs of the MIoT, we define the relative reputation $\widehat{R_{j_k}}[u,v]$ of ι_{j_k} in $\mathcal{I}_k$ as follows:

$$\widehat{R_{j_k}}[u,v] = \frac{R_{j_k}[u,v]}{R_k^{max}[u,v]}$$

Example (...cnt'd)
Consider the smart shopping center described in Section 2.2. We want to evaluate the reputation of the smartphone of Bob in the electronics store, i.e. the reputation of ι_{3_2} in $\mathcal{I}_2$. Here:

- $Age_{1_2}[u,v]$ (resp., $Age_{2_2}[u,v]$, $Age_{4_2}[u,v]$) is the number of days since the first transmission of a video on smart home products performed by the smartphone of Alice (resp., the smartwatch of Alice, the smartwatch of Bob) in the electronics store. Assume that $Age_{1_2}[u,v] = 20$, (resp., $Age_{2_2}[u,v] = 20$, $Age_{4_2}[u,v] = 70$).

- $Age_2^{max}[u, v]$ is the maximum number of days since the transmission of a video on smart home products performed by a smart object in the electronics store. Assume that $Age_2^{max}[u, v] = 75$.
- Assume that all the objects currently present in the electronics store are totally connected to each other. As a consequence, $nbh^{in}(\iota_{3_2}) = \{\iota_{1_2}, \iota_{2_2}, \iota_{4_2}\}$.
- $R_2^{max}[u, v]$ is the maximum reputation of a smart object transmitting a video on smart home products in the electronics store. Assume that $R_2^{max}[u, v] = 0.68$.
- γ is the minimum absolute reputation assigned to an object in the electronics store. Assume $\gamma = 0.30$.
- $T_{13_2}[u, v]$ (resp., $T_{23_2}[u, v]$, $T_{43_2}[u, v]$) is the value of the trust that the smartphone of Alice (resp., the smartwatch of Alice, the smartwatch of Bob) has in the smartphone of Bob, regarding videos on smart home products sent in the electronics store. Assume that $T_{13_2}[u, v] = 0.95$, $T_{23_2}[u, v] = 0.90$ and $T_{43_2}[u, v] = 1$.
- $R_{1_2}[u, v]$ (resp., $R_{2_2}[u, v]$, $R_{4_2}[u, v]$) is the value of the reputation of the smartphone of Alice (resp., the smartwatch of Alice, the smartwatch of Bob), regarding videos on smart home products sent in the electronics store. Assume that $R_{1_2} = 0.98$, $R_{2_2} = 0.93$ and $R_{3_2} = 0.96$.
- The reputation of the smartphone of Bob, regarding videos on smart home products sent in the electronics store, is obtained as:

$$R_{3_2}[u, v] = 0.30 + (1 - 0.30) \cdot \frac{0.95 \cdot 0.98 \cdot \frac{20}{75} + 0.90 \cdot 0.93 \cdot \frac{20}{75} + 1 \cdot 0.96 \cdot \frac{70}{75}}{3} = 0.62$$

- The normalized reputation $\widehat{R_{3_2}}[u, v]$ of the smartphone of Bob, regarding videos on smart home products sent in the electronics store, is obtained as:

$$\widehat{R_{3_2}}[u, v] = \frac{0.62}{0.68} = 0.91$$

5.2.4 Reputation of an object in a MIoT

Let o_j be an object of $\mathcal{M}$. Let $\mathcal{M}_j = \{\mathcal{I}_1, \cdots, \mathcal{I}_l\}$ be the subset of the IoTs of $\mathcal{M}$ containing one instance of o_j.

The reputation of o_j depends on both the trust that its instances receive in each IoT of $\mathcal{M}_j$ and the reputation of the object, which the instance providing this trust refers to. To formalize this concept, we can say that the reputation of o_j, regarding the topic u delivered in the format v, is defined as follows:

$$R_j[u, v] = \delta + (1 - \delta) \cdot \frac{\sum_{k=1..l} \sum_{\iota_{q_k} \in nbh^{in}(\iota_{j_k})} V_{qj}[u, v] \cdot R_q[u, v]}{l \cdot |nbh^{in}(\iota_{j_k})|}$$

As in the previous case, this formula is similar to the PageRank one. δ is the damping factor and its semantics is analogous to the one of γ seen in Section 5.2.3.

At this point, it is necessary to proceed with the normalization of $R_j[u, v]$. This task is performed in a way analogous to the one defined for the instance reputation in the previous section:

$$\widehat{R_j}[u,v] = \frac{R_j[u,v]}{R^{max}[u,v]}$$

Example (...cnt'd)
Consider the smart shopping center described in Section 2.2. The reputation $\widehat{R_3}[u,v]$ of the smartphone of Bob, when it sends videos on smart home products in the whole smart shopping center, can be computed in a way analogous to the computation of the reputation $\widehat{R_{3_2}}[u,v]$ of the smartphone of Bob in the electronics store, illustrated in the example of Section 5.2.3. For this reason, and due to space limitations, we do not report all details of the computation of $\widehat{R_3}[u,v]$ below.

5.2.5 Reputation of an IoT in a MIoT

The reputation of an IoT $\mathcal{I}_k$ in $\mathcal{M}$, regarding the topic u delivered in the format v, is given by the average of the reputations of the objects of $\mathcal{M}$ having one instance in $\mathcal{I}_k$.

If we introduce the set Obj_k of the objects having one instance in $\mathcal{I}_k$, the reputation $\widehat{\mathcal{R}^k}[u,v]$ of $\mathcal{I}_k$ in $\mathcal{M}$ can be formalized as follows:

$$\widehat{\mathcal{R}^k}[u,v] = \frac{\sum_{j \in Obj_k} \widehat{R}_{j_k}[u,v]}{|Obj_k|}$$

Example (...cnt'd)
Consider the smart shopping center described in Section 2.2. The reputation $\widehat{\mathcal{R}^2}[u,v]$ of the IoT associated with the electronics store, when the objects present therein send videos on smart home products in the smart shopping center, can be computed in a way analogous to the computation of the trust $T_{13}[u,v]$ of the smartphone of Alice in the smartphone of Bob, when this last sends video on smart home products in the smart shopping center, as illustrated in the example of Section 5.2.1. For this reason, due to space limitations, we do not report all details of the computation of $\widehat{\mathcal{R}^2}[u,v]$ below.

5.2.6 Trust of an IoT in another IoT

The trust $\mathcal{T}^{hk}[u,v]$ of an IoT $\mathcal{I}_h$ in an IoT $\mathcal{I}_k$, regarding the topic u delivered in the format v, consists of a pair $(\mathcal{V}^{hk}[u,v], \mathcal{LTS}^{hk}[u,v])$.

$\mathcal{V}^{hk}[u,v]$ is defined as the average of the trust values of any object of $\mathcal{I}_h$ in any object of $\mathcal{I}_k$, with which it performed at least one transaction.

To formally define $\mathcal{V}^{hk}[u,v]$, we must introduce the set $tranSet_{j_k}[u,v]$ of the transactions that any instance ι_{j_h} of $\mathcal{I}_h$ carried out with any instance of $\mathcal{I}_k$ and having in their content the topic u delivered in the format v. After having introduced $tranSet_{j_k}[u,v]$, we can define $\mathcal{V}^{hk}[u,v]$ as follows:

$$\mathcal{V}^{hk} = \frac{\sum_{j \in Obj_h} \sum_{q \in tranSet_{j_k}[u,v]} V_{jq}[u,v]}{\sum_{j \in Obj_h} |tranSet_{j_k}[u,v]|}$$

$\mathcal{LTS}^{hk}[u,v]$ is the last ending timestamp that can be found in a transaction involving any instance of $\mathcal{I}_h$ with any instance of $\mathcal{I}_k$.

Example (...cnt'd)
Consider the smart shopping center described in Section 2.2. The trust $\mathcal{T}^{21}[u, v]$ of the IoT associated with the electronics store in the IoT associated with the supermarket, when the objects in this last network send videos on smart products, can be computed in a way analogous to the computation of the trust $T_{13_2}[u, v]$ of the smartphone of Alice in the smartphone of Bob, when it sends videos on smart home products in the electronics store, illustrated in the example of Section 5.2.1. For this reason, due to space limitations, we do not report all details of the computation of $\mathcal{T}^{21}[u, v]$ below.

5.2.7 Trust of an object in an IoT

Let o_j be an object of $\mathcal{M}$ and let $\mathcal{I}_k$ be an IoT of $\mathcal{M}$. Again, the trust $\mathcal{T}_j^k[u, v]$ of o_j in $\mathcal{I}_k$, regarding the topic u delivered in the format v, consists of a pair $(\mathcal{V}_j^k[u, v], \mathcal{LTS}_j^k[u, v])$. In the computation of $\mathcal{T}_j^k[u, v]$ we must distinguish two cases, namely:

- o_j has one instance ι_{j_k} in $\mathcal{I}_k$. In this case, let $Inst_{j_k}[u, v]$ be the set of the instances of $\mathcal{I}_k$ with which ι_{j_k} carried out at least one transaction involving the topic u delivered in the format v in the past. $\mathcal{V}^{hk}[u, v]$ is defined as the average of the trusts of ι_{j_k} in all the instances of $Inst_{j_k}[u, v]$. More formally:

$$\mathcal{V}_j^k[u, v] = \frac{\sum_{q \in Inst_{j_k}[u,v]} V_{jq_k}[u, v]}{|Inst_{j_k}[u, v]|}$$

  $\mathcal{LTS}_j^k[u, v]$ is the last ending timestamp that can be found in a transaction involving ι_{j_k} and any instance of $Inst_{j_k}[u, v]$.

- o_j has no instance in $\mathcal{I}_k$. In this case, the trust of o_j in $\mathcal{I}_k$ is equal to the sum of the trusts of the instances of o_j in the IoTs it belongs to, weighted by the trust of the corresponding IoT in $\mathcal{I}_k$. More formally, let $\mathcal{M}_j = \{\mathcal{I}_1, \cdots, \mathcal{I}_l\}$ be the set of the IoTs of $\mathcal{M}$ containing one instance of o_j. In this case:

$$\mathcal{V}_j^k[u, v] = \frac{\sum_{h=1..l} \mathcal{V}_j^h[u, v] \cdot \mathcal{V}^{hk}}{l}$$

  $\mathcal{LTS}_j^k[u, v]$ is the maximum LTS among the ones associated with $\mathcal{T}_j^h[u, v]$, $1 \leq h \leq l$. Formally speaking:

$$\mathcal{LTS}_j^k[u, v] = max_{h=1..l} \mathcal{LTS}_j^h[u, v]$$

Example (...cnt'd)
Consider the smart shopping center described in Section 2.2. The trust $\mathcal{T}_3^2[u, v]$ that the smartphone of Bob has in the IoT associated with the electronics store, when the objects in this last network send videos on smart home products, can be computed in a way analogous to the computation of the trust $T_{13_2}[u, v]$ of the smartphone of Alice in the smartphone of Bob, when it sends videos on smart home products in the smart shopping center. We have illustrated the computation of $T_{13_2}[u, v]$ in the example of Section 5.2.1. For this reason, due to space limitations, we do not report all details of the computation of $\mathcal{T}_3^2[u, v]$ below.

6 Experiments

In this section, we present the set of experiments that we carried out to evaluate the performance of our approach from several viewpoints. First of all, in Subsection 6.1, we describe how we experimentally set the weight of α, β and ρ in the computation of the trust of an instance into another one of the same IoT, in order to give an example of how we set the weights in our approach. Then, we describe our testbed in Subsection 6.2, whereas, in Subsections 6.3, 6.4, and 6.5, we illustrate our tests, along with the underlying motivations and the results obtained. Finally, in Subsection 6.6, we present an experiment to evaluate the accuracy of our approach.

6.1 Setting of weights

In this experiment, we aimed at determining the values of α, β and ρ in the computation of the value of the trust of an instance in another one of the same IoT (see Section 5.2.1). First of all, we observe that, roughly speaking, α represents the weight of the fraction of the correct transactions between ι_{j_k} and ι_{q_k}, β denotes the significance of the number of transactions existing between ι_{j_k} and ι_{q_k}, whereas ρ indicates the weight of the size of the content exchanged between ι_{j_k} and ι_{q_k}.

To perform this experiment, we initially selected 100 smart objects and we connected them to form an IoT. Then, we let them to perform 100,000 transactions through which they exchanged data. At the end of this task, we computed the values of trust for each pair of objects by applying the formulas reported in Section 5.2.1.

Afterwards, we forced some fictitious wrong behaviors in the transactions between the smart objects of the network. The entity of the error was varying; it was evaluated by some domain experts as *null*, *small*, *medium*, *high* and *very high*. In particular, we made sure that 20% of the transactions had a *null* (resp., *small*, *medium*, *high*, *very high*) error.

After this, for each pair of smart objects, we asked the human experts to evaluate the overall perturbation caused in their transactions by the wrong behavior induced in the experiment. The possible evaluations were *negligible*, *small*, *medium*, *high* and *very high*.

At this point, for each pair of smart objects, we recomputed the value of trust with the perturbed transactions. Then, we compared the size of the change of the trust values against the values themselves. We considered as *negligible* (resp., *small*, *medium*, *high* and *very high*) the perturbation caused in a trust value if its change was less than 20% (resp., between 20% and 40%, between 40% and 60%, between 60% and 80%, more than 80%) of the original value.

We repeated this last part of the experiment with different combinations of values of α, β and ρ. In particular, the adopted combinations are the ones reported in the first three columns of Table 1.

Finally, for each weight combination, we computed the percentage of times the evaluation of trust perturbations performed by our approach and the one carried out by human experts coincided. The obtained values are reported in the fourth column of Table 1. From the analysis of this table, we can see that the optimal combination of values is $\alpha = 0.55$, $\beta = 0.35$ and $\rho = 0.10$. We can also observe that the combinations slightly differing from the previous one produce a small number of errors. On the other side, as long as the combinations differ from the optimal ones, the errors increase. This witnesses the optimal resilience of our approach that, however, is (correctly) sensitive to weight errors when these become high.

Interestingly, in this experiment, we were guided only by the semantics of the three measures weighted by α, β and ρ, and not on the nature of the scenario where it was performed. As a conse-

α	β	ρ	*Percentage of errors*
0.35	0.25	0.05	35.67 %
0.35	0.25	0.10	30.34 %
0.35	0.25	0.15	25.86 %
0.35	0.35	0.05	25.56 %
0.35	0.35	0.10	20.87 %
0.35	0.35	0.15	15.12 %
0.35	0.45	0.05	15.76 %
0.35	0.45	0.10	10.95 %
0.35	0.45	0.15	5.56 %
0.45	0.25	0.05	25.38 %
0.45	0.25	0.10	20.54 %
0.45	0.25	0.15	15.37 %
0.45	0.35	0.05	15.54 %
0.45	0.35	0.10	10.48 %
0.45	0.35	0.15	5.83 %
0.45	0.45	0.05	5.69 %
0.45	0.45	0.10	5.25 %
0.45	0.45	0.15	5.58 %
0.55	0.25	0.05	15.28 %
0.55	0.25	0.10	10.94 %
0.55	0.25	0.15	5.59 %
0.55	0.35	0.05	5.93 %
0.55	0.35	0.10	0.50 %
0.55	0.35	0.15	5.73 %
0.55	0.45	0.05	5.28 %
0.55	0.45	0.10	10.62 %
0.55	0.45	0.15	15.28 %
0.65	0.25	0.05	5.74 %
0.65	0.25	0.10	5.34 %
0.65	0.25	0.15	5.79 %
0.65	0.35	0.05	5.93 %
0.65	0.35	0.10	10.48 %
0.65	0.35	0.15	15.52 %
0.65	0.45	0.05	15.28 %
0.65	0.45	0.10	20.58 %
0.65	0.45	0.15	25.92 %
0.75	0.25	0.05	5.36 %
0.75	0.25	0.10	10.83%
0.75	0.25	0.15	15.28 %
0.75	0.35	0.05	15.27 %
0.75	0.35	0.10	20.74 %
0.75	0.35	0.15	25.94 %
0.75	0.45	0.05	25.38 %
0.75	0.45	0.10	30.19 %
0.75	0.45	0.15	35.18 %

Table 1: Setting of the weights α, β and ρ in the computation of the trust of an instance in another one of the same IoT

quence, even if this experiment could be repeated each time we want to determine the values of α, β and ρ in the most disparate scenarios, we are confident that the values obtained are general and do not depend on the application environment in which our approach is operating.

6.2 Adopted Testbed

In order to perform the next experiments, we had the necessity to create several MIoTs with different sizes, ranging from hundreds to thousands of nodes. Since, currently, real MIoTs with the size and the variety handled by our model do not exist yet, we constructed a MIoT simulator. This tool starts from real data and returns simulated MIoTs with certain characteristics specified by the user.

The MIoTs created by our simulator follow the paradigm described in Section 4.1. Our MIoT simulator is also provided with a suitable interface allowing a user to "personalize" the MIoT to construct by specifying the desired values for several parameters, such as the number of nodes, the maximum number of instances of an object, and so forth.

To make "concrete" and "plausible" the created MIoTs, our simulator leverages a real dataset. It regards the taxi routes in the city of Porto from July 1^{st} 2013 to June 30^{th} 2014. It can be found at the address `http://www.geolink.pt/ecmlpkdd2015-challenge/dataset.html`. Each route con-

tains several Points of Interests corresponding to the GPS coordinates of the vehicle.

We partitioned the city of Porto in six areas and associated a real IoT with each of them. Our simulator associates an object with a given route recorded in the dataset and an object instance with each partition of a route belonging to an area. It creates a MIoT node for each instance and a c-arc for each pair of instances belonging to the same route. Furthermore, it creates an i-arc between two nodes of the same IoT if the length of the time interval between the corresponding routes is less than a certain threshold th_t. The weight of the i-arc indicates the length of this time interval. The value of th_t can be specified through the constructor interface. Clearly, the higher th_t, the more connected the constructed MIoT.

As far as instance profiles are concerned, since there are no available thing profiles, we had to simulate them. However, we aimed at making them as real as possible. For this purpose, we performed a sentiment analysis task for each of the six areas in which we partitioned the city of Porto and for each day which the dataset refers to. To carry out this task, we leveraged IBM Watson on the social media and blogs available in it. Having this data at disposal, our simulator assigns to each instance the most common topics (along with the corresponding occurrences) discussed in that area in the day on which the corresponding route took place. The constructed MIoTs are returned in a format that can be directly processed by the cypher-shell of Neo4J.

The interested reader can find the MIoTs adopted in the experiments described in Subsections 6.3 and 6.4 at the address `http://daisy.dii.univpm.it/miot/datasets/trustReputation`.

We carried out all the tests presented in this section on a server equipped with an Intel I7 Quad Core 7700 HQ processor and 16 GB of RAM with Ubuntu 16.04 operating system. To implement our approach, we adopted: *(i)* Python, as programming language; *(ii)* Neo4J (Version 3.4.5), as underlying DBMS.

6.3 Computation time

Our first test is devoted to evaluate the computation time of our approach. Indeed, since it could operate in large MIoTs, whose IoTs could consist of even hundreds of nodes, it is necessary to verify if, in these real cases, the time it needs to return a result is still acceptable.

6.3.1 Trust of an instance in another one of the same IoT and of an object in another one of the MIoT

In this experiment, we considered several MIoTs having a different number of nodes. Given a MIoT $\mathcal{M}$, we considered all its IoTs $\mathcal{I}_1, \cdots, \mathcal{I}_6$ (see Section 6.2). For each IoT, we computed the trust of each of its nodes in the others. At the first iteration, we set the value of the trust of a node in any other of the MIoT to 0.5. In other words, we decided to assume a "neutral" policy in order to not outweigh either positively or negatively on the trust of one node in another.

We performed a total number of 60,000 transactions in the MIoT and we recomputed all trust values every 10 transactions (in the following, we call *epoch* an interval of 10 transactions). We measured the time required by our approach to compute the trust of an instance in another one of the same IoT against the number of epochs for MIoTs having a different number of nodes (ranging from 10 to 1,000). Finally, we averaged the obtained computation times for all the instances of the MIoT. The results obtained are reported in Figure 3.

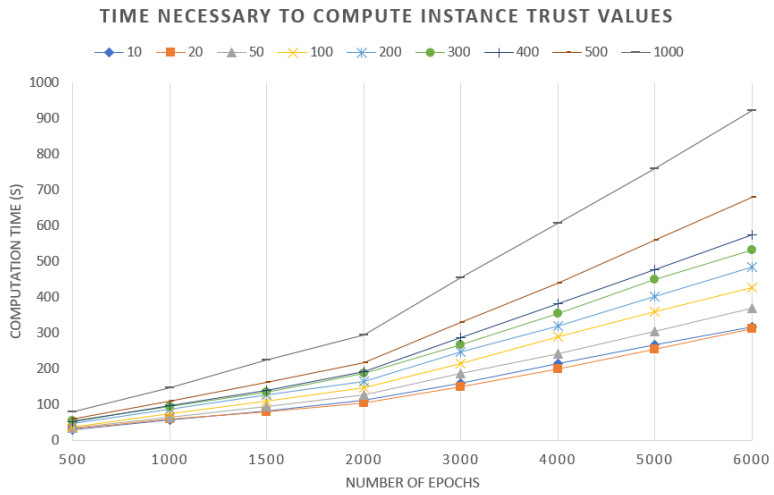

Figure 3: Average time of the computation of the trust of an instance in another one of the same IoT against the number of epochs for MIoTs with different numbers of nodes

From the analysis of this figure we can observe that the computation time is very small when the numbers of epochs is less than 2,000, independently of the size of the MIoT. After 2,000 epochs, it starts to increase more quickly. In this case, if the number of the nodes of the MIoT is lower than 1,000, the computation time and the quickness of its growth are still acceptable. Instead, in presence of a MIoT with more than 1,000 instances, the computation time tends to become excessively high and quickly unacceptable.

As shown in Section 5.2.2, the trust of an object in another one of $\mathcal{M}$ is easily determined by computing the average values of the trust of its instances in the ones of the other object in the IoTs where both of them are present. The additional computations necessary to obtain it, once the trust values of the corresponding instances have been determined, are negligible. As a consequence, all the considerations about the computation time that we made for the trust of an instance in another one of the same IoT can be extended to the trust of an object in another one of the MIoT.

With regard to this result, we observe that the cases in which the computation time begins to become unacceptable regard scenarios that we do not currently find in real cases. In fact, in order to start having computational problems, we should be in presence of a MIoT consisting of more than 1,000 instances. If we consider that, in real cases, the number of IoTs in a MIoT is currently less than 10, we should have more than 100 objects which simultaneously want to interact in all the IoTs of the MIoT. This highly unlikely scenario could be still managed by our approach if the number of epochs used to compute the trust values is less than 2,000. Now, since an epoch corresponds to 10 transactions, this means that our approach starts to present an excessive computation time only in presence of about 100 objects wanting to simultaneously interact in 10 different IoTs of the MIoT and performing at least 20,000 simultaneous transactions.

Actually, the current MIoTs would consist of at least 3-5 IoTs. The number of objects in each IoT that want to simultaneously interact is less than 50. As a consequence, in real cases, a MIoT consists of at most 200-400 instances. Furthermore, not all the objects want to interact with all the other ones. In fact, the number of pairs of objects wanting to interact with each other is very limited and do

not exceed 400. In addition, in real cases, the overall number of transactions necessary to compute a stable value of trust for each pair of interacting instances does not exceed 20. With all the hypotheses above, our approach would need at most 8,000 transactions to determine stable values of trust for each pair of interacting objects. This number is much smaller than the limit value of 20,000 transactions. Clearly, we think that, in the future, the size and the density of MIoTs will increase; however, the computing power available in servers should increase too. In any case, if this increase would be not sufficient, we could adopt two countermeasures to make the computation time still reasonable. Indeed: *(i)* we could increase the number of transactions associated with an epoch (for instance, from 10 to 100 or to 1,000); *(ii)* we could use distributed and parallel processing to perform trust evaluation.

6.3.2 Reputation of an instance in an IoT and of an object in the MIoT

In this experiment, we considered the same MIoTs adopted in the previous one and, for each instance of an object, we computed its reputation in the corresponding IoT. Also in this case, at the first iteration, we set the initial reputation of each instance to 0.5. In this case, we performed a total number of 600,000 transactions.

Reputation is intrinsically much more static than trust. As a consequence, it appears more reasonable to assume epochs of 100 transactions, instead of 10. We measured the time required by our approach for computing the reputation of each instance in its IoT against the number of epochs for the MIoTs adopted in the previous experiment. Then, we averaged these values for all the instances of the MIoT. The results obtained are reported in Figure 4.

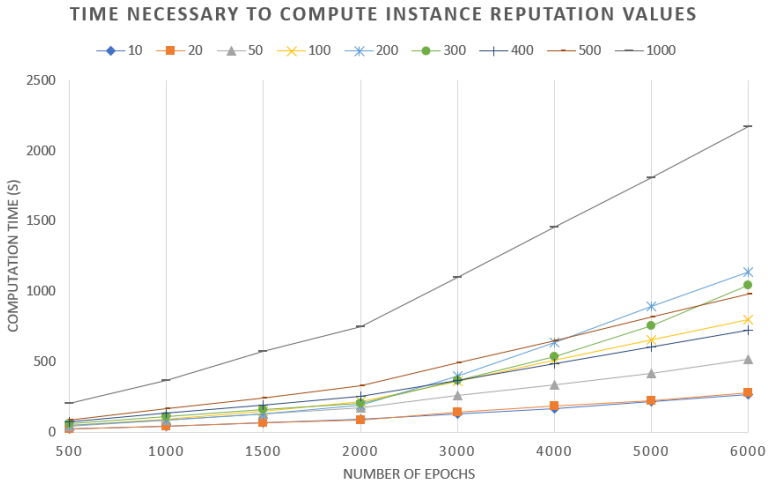

Figure 4: Average time of the computation of the reputation of an instance in its IoT against the number of epochs for MIoTs with different numbers of nodes

From the analysis of this figure, we can observe that the time necessary to compute the values of instance reputation is always low when the number of epochs is lower than, or equal to, 2,000 and the number of nodes is lower than, or equal to, 500. When the number of nodes is higher, the computation time increases, even if it is still acceptable for a number of epochs lower than, or equal to, 2,000. When the number of epochs is higher than 2,000, the computation time starts to rapidly increase. It tends to become unacceptable when the number of nodes is higher than 500 and the number of epochs is higher

than 2,000. With regard to this result, we observe that all the reasonings about the computation of trust in real cases, which we have presented at the end of Section 6.3.1, can be extended here to the computation of reputation in real cases.

As illustrated in Section 5.2.4, the definition of the reputation of an object in a MIoT is structurally similar to the definition of the reputation of an instance in an IoT. As a consequence, all the considerations about the computation time that we have illustrated above can be easily extended to this last case.

6.3.3 Reputation of an IoT in the MIoT

As shown in Section 5.2.5, the reputation of an IoT in the MIoT is obtained by averaging the reputation of the objects having one instance in it. The computation of the average is negligible after that the reputation of the corresponding objects has been determined. Therefore, all the considerations about the computation time, which we made for the reputation of an object in the MIoT, can be extended to the reputation of an IoT in the MIoT.

6.4 Values and Distributions

In this section, we investigate the trends of the trust and reputation values for instances and objects. Furthermore, we analyze the distribution of these values in order to describe how the corresponding shape changes over specific epochs.

6.4.1 Trust of an instance in another one of the same IoT

In order to investigate the features characterizing the trust of an instance in another one of the same IoT, we applied the guidelines described in Section 6.3.1 even if, this time, we focused on values and not on computation time. In Figure 5, we report the average values of this trust against the number of epochs for the same MIoTs we introduced in the previous section. From the analysis of this figure, we can observe that, initially, as the number of transactions increases, the trust values increase too. This fact is justified by considering that, as the number of correct transactions increases, the instances "have more confidence" in each other. This increase is particularly evident until to 1,000 epochs (i.e., 10,000 transactions). When the number of epochs ranges between 1,000 and 2,000, the value of the trust still slightly grows, even if the speed of growth is much lower than before. Finally, after 4,000 epochs, the average trust reaches an approximately fixed value in some cases, whereas, in other ones, it grows very slowly.

Observe that, in Figure 5, the performance of our approach depends on the number of epochs and the number of instances. Clearly, the increase of the number of epochs always leads to an increase of the trust between instances or objects. On the other side, it requires a higher number of transactions and, ultimately, a higher computational and time cost. Clearly, a tradeoff is necessary between these two exigencies. In particular, in cases where computational costs are more important than accuracy, it is better to choose a number of epochs lower than 2,000. By contrast, whenever accuracy is extremely important and computational costs can be partially sacrificed, it is better to choose a number of epochs higher than 2,000. As for the number of instances, it depends on the scenario on which the MIoT is operating, and cannot be tuned by the operator.

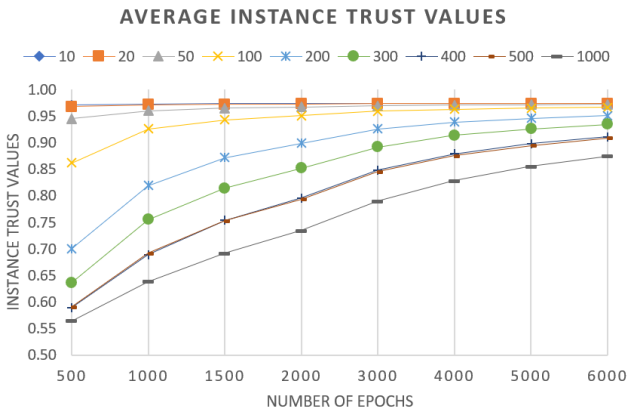

Figure 5: Average values of the trust of an instance in another one of the same IoT against the number of epochs for MIoTs with a different number of nodes

We also computed the distributions of the trust values after 1,000, 2,000 and 3,000 epochs. We performed this computation for the MIoT with 300 instances adopted in the previous experiments. The results obtained are reported in Figure 6. From the analysis of this figure we can observe that the distribution shape moves to the right. This phenomenon is very evident when passing from 1,000 to 2,000 epochs, but it is still significant also when passing from 2,000 to 3,000 epochs.

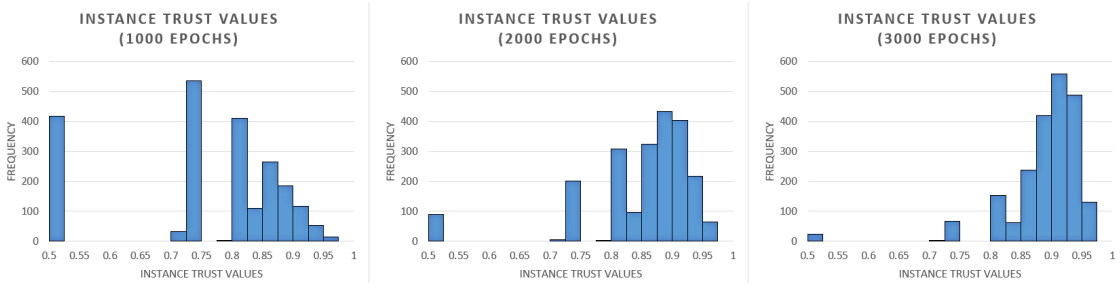

Figure 6: Distribution of the trust of an instance in another one of the same IoT after 1,000, 2,000 and 3,000 epochs for the MIoT with 300 instances

A final parameter that we computed is the standard deviation after 1,000, 2,000 and 3,000 epochs. The values we obtained were 0.1385, 0.0934, and 0.0626, respectively. This result is extremely interesting; as a matter of fact, already after 1,000 epochs, the values of the standard deviation are acceptable. Furthermore, when passing from 1,000 to 2,000 and from 2,000 to 3,000 epochs, we can observe a quick decrease of the corresponding values. This denotes a high stability of the overall instance trusts that can be already observed after only 1,000 epochs.

6.4.2 Trust of an object in another one of the MIoT

In this experiment, we applied the guidelines described in Section 6.3.1, but we focused on trust values and not on computation time. The average values of the trust of an object in another against the number of epochs for the MIoTs introduced previously is reported in Figure 7. From the analysis of this

figure, we can observe that the trend of the trust values for objects is analogous to the corresponding one for instances, discussed in Section 6.4.1. Actually, by carefully examining Figures 5 and 7, we can observe an "extremization" of some phenomena. For instance, in small MIoTs, the object trust is constantly equal to about 1. Furthermore, when the MIoTs are medium or large, the trust values increase very quickly until to 1,000 epochs. This increase is still significant from 1,000 to 3,000 epochs. Finally, it becomes very small after 3,000 epochs.

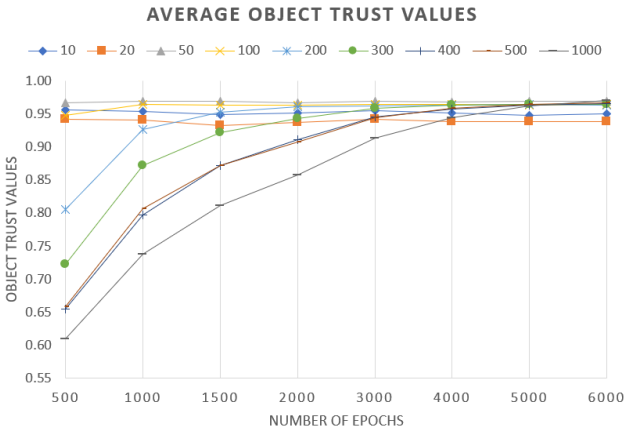

Figure 7: Average values of the trust of an object in another one of the MIoT against the number of epochs for MIoTs with a different number of nodes

This conclusion can be also extended to the distribution of the object trust values, reported in Figure 8, for the usual MIoT with 300 instances. As for the analysis of the standard deviation, we obtained that, after 1,000, 2,000 and 3,000 epochs, its values are 0.1911, 0.1093, and 0.0716, respectively. We can observe a rapid decrease when passing from 1,000 to 2,000 and from 2,000 to 3,000 epochs. This is an indicator of the stability of the obtained values for object trust.

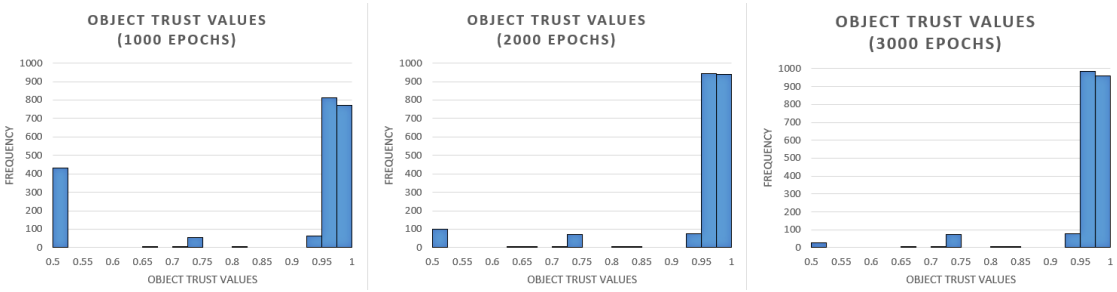

Figure 8: Distribution of the trust of an object in another one of the same IoT after 1,000, 2,000 and 3,000 epochs for the MIoT with 300 instances

Observe that, analogously to Figure 5, also in Figure 7 the performance of our approach depends on the number of epochs and the number of instances. With regards to these two parameters, the same reasonings we have proposed for Figure 5 can be applied to this figure.

6.4.3 Reputation of an instance in an IoT, of an object in the MIoT and of an IoT in the MIoT

In order to investigate the variation of the reputation of an instance in an IoT against the number of epochs we applied the guidelines described in Section 6.3.2. The corresponding results are reported in Figure 9.

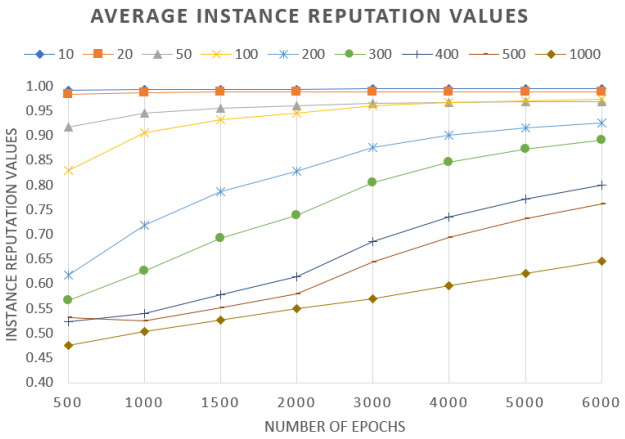

Figure 9: Average values of the reputation of an instance in its IoT against the number of epochs for MIoTs with a different number of nodes

From the analysis of this figure, we can observe that the trend of the reputation values shows a continuous (even if slow) increase against the number of epochs. This can be explained by observing that, analogously to what happens to communities of people, as time passes and the number of transactions increases, objects tend to trust each other. As a consequence, the number of failed transactions decreases, which leads to an increase of the reputation of the objects performing them. Observe that the reputation value is higher for smaller networks. This reflects a general trend also observed in social networks of humans and, more in general, in communities of people. In fact, in a small community, the corresponding members tend to trust each other more.

The distribution of the corresponding values, for the usual MIoT with 300 instances, is reported in Figure 10. This figure represents a further confirmation of what we observed in Figure 9. Indeed, we can note that the shape of the distribution does not significantly change over time, but, as the number of epochs increases, the distribution values move to the right. This phenomenon is much more evident when passing from 1,000 to 2,000 epochs than when passing from 2,000 to 3,000 ones.

Finally, the values of the standard deviation of the instance reputation after 1,000, 2,000 and 3,000 epochs are 0.0950, 0.0682, and 0.0535, respectively. This extremely low and quite constant values evidence that the results obtained are acceptable and stable over time.

A similar procedure can be applied to object reputation, whose values against the number of epochs for the usual MIoTs are reported in Figure 11, and whose value distributions for the MIoT with 300 instances are shown in Figure 12.

From the analysis of Figure 11, we can observe that the values of object reputation are often smaller than the corresponding ones of instance reputation, even if they are still acceptable. This is explained by the fact that object reputations refer to the whole MIoT and not to a single IoT, i.e., to a larger and more variegate scenario than the one characterizing the evaluation of instance reputations.

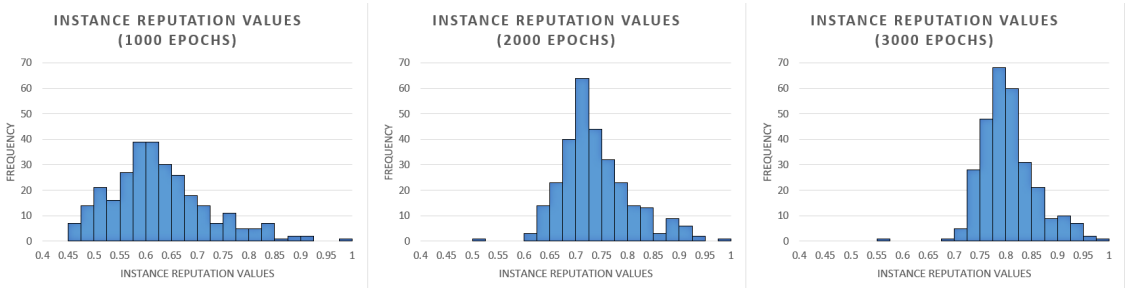

Figure 10: Distribution of the reputation of an instance in its IoT after 1,000, 2,000 and 3,000 epochs for the MIoT with 300 instances

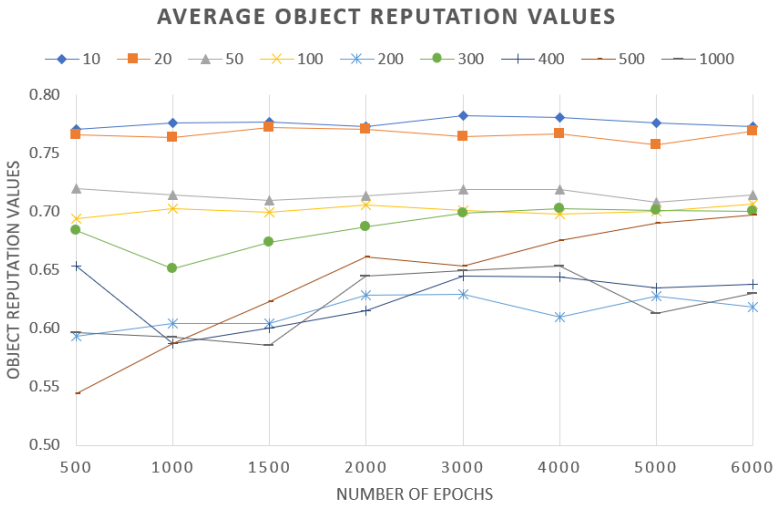

Figure 11: Average values of the reputation of an object in its MIoT against the number of epochs for MIoTs with a different number of nodes

In this context, it is clearly more difficult for an object to acquire and maintain trustworthiness.

We point out that the same observations and conclusions which we have drawn for Figures 5 and 7 can be extended to Figures 9 and 11.

The distributions of Figure 12, performed for the usual MIoT with 300 instances, confirm these observations. Indeed, in this case, we can observe that the distribution shape is roughly the same after 1,000, 2,000 and 3,000 epochs, but, differently from what happens for instance reputation values, it moves to the right only very slightly as the number of epochs increases. In other words, in this case, object reputation values show only a very small increase over time. The reasons are the same as the ones reported for Figure 11.

The value of the standard deviation after 1,000, 2,000 and 3,000 epochs are 0.1269, 0.1369 and 0.1403, respectively. These values are higher than the ones characterizing the standard deviation of instance reputation, even if they are still acceptable and quite constant over time. This evidences that the object reputation scenario is certainly more difficult to handle than the instance reputation one, even if it can be still maintained under control.

Finally, the reputation of an IoT in the MIoT is obtained by averaging the reputations of the

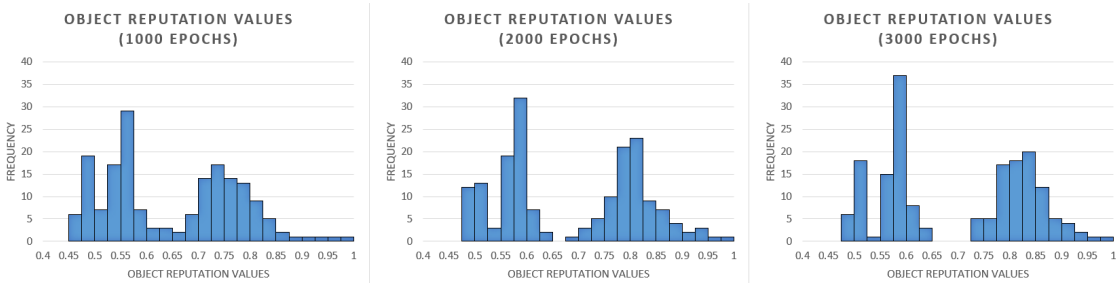

Figure 12: Distribution of the reputation of an object in the MIoT with 300 instances after 1,000, 2,000 and 3,000 epochs

objects having one instance in it. As a consequence, the corresponding values and distributions are very similar to the ones illustrated for the objects in the MIoT. Therefore, due to space limitations, we do not report them here.

6.5 Resilience

This experiment aimed at evaluating the robustness of our approach against the possible anomalies of the trust values assigned by an instance to another. We conducted it on the usual MIoT with 300 instances adopted for the previous experiment. In particular, we assumed the average value of the trust of an instance in another of the same IoT against the number of epochs for the MIoT with 300 instances (shown in Figure 5) as the "ground truth", i.e., as the case with no anomalies. After this, we considered two possible extreme anomalies. The former assumed that a fraction $X\%$ of instances constantly assigns a trust equal to 1 to all the other instances, independently of exchanged transactions, and all the transactions regarding these instances are reported as successful, independently of their real result (we call them "positive anomalies" in the following). The latter assumed an opposite behavior; therefore, it assumed that, independently of exchanged transactions, a fraction $Y\%$ of instances associates a value of 0 with the trust in all the other instances and all the transactions concerning these instances are reported as failed, independently of their real results (we call them "negative anomalies" in the following). We computed the average values of trust against the number of epochs for several fractions of positive or negative anomalies (namely, 5%, 10%, 15%, 20%, 30%). The results obtained are reported in Figures 13 and 14.

First, let us consider Figure 13, which refers to positive anomalies. From the analysis of this figure, we can observe that our approach is very resilient to this kind of anomaly. For example, the presence of 20% of positive anomalies leads to an increase of the trust values ranging from 10.24% at 500 epochs to 1.29% at 6,000 epochs.

After having examined positive anomalies, we analyze negative ones. They are reported in Figure 14. From the analysis of this figure, we can observe that our approach is sensitive to them. For instance, the presence of 20% of negative anomalies leads to a decrease of the trust values ranging from 22.06% at 500 epochs to 25.75% at 6,000 epochs, which is much higher than the corresponding one seen for positive anomalies. Even more interesting, when the fraction of negative anomalies reaches 30% of the MIoT instances, we can observe a strong fall of the trust values. In fact, its decrease ranges from 46.16% at 500 epochs to 52.87% at 6,000 epochs, which implies that the behavior of our approach is no longer acceptable.

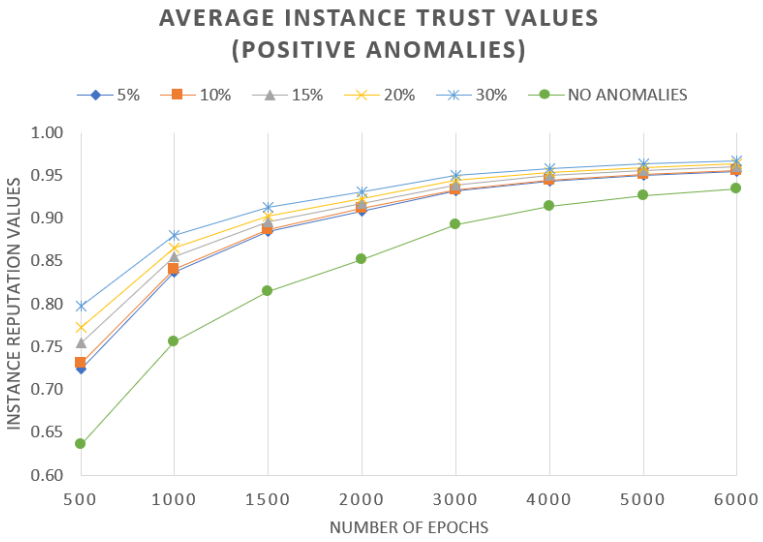

Figure 13: Average values of the trust of an instance against the increase of positive anomalies for the MIoT with 300 instances

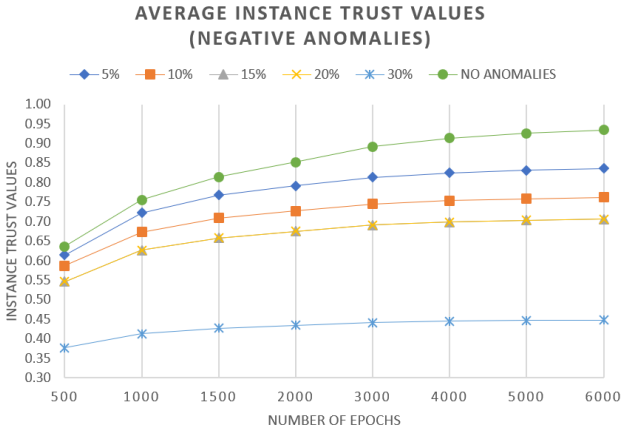

Figure 14: Average values of the trust of an instance against the increase of negative anomalies for the MIoT with 300 instances

The overall analysis of positive and negative anomalies allows us to conclude that our approach is very resilient to positive anomalies; perhaps, it is excessively resilient to them when they become high. An opposite behavior can be observed for negative anomalies. Our approach allows users to find them very easily; however, it is excessively sensitive to them when they are few.

An analogous reasoning can be drawn for the resilience of our approach to compute the reputation of an instance in an IoT. Analogously to what we have done for trust, we considered the average values of the reputation of an instance in its IoT against the number of epochs for the MIoT with 300 instances (shown in Figure 9) as the "ground truth", i.e., as the case with no anomalies. After this, we operated in the same way as we had operated for trust. In this case, the variation of the average reputation value against the number of epochs in presence of positive (resp., negative) anomalies is

reported in Figure 15 (resp., 16). We computed it for several fractions of positive (resp., negative) anomalies (namely, 5%, 10%, 15%, 20%, 30%).

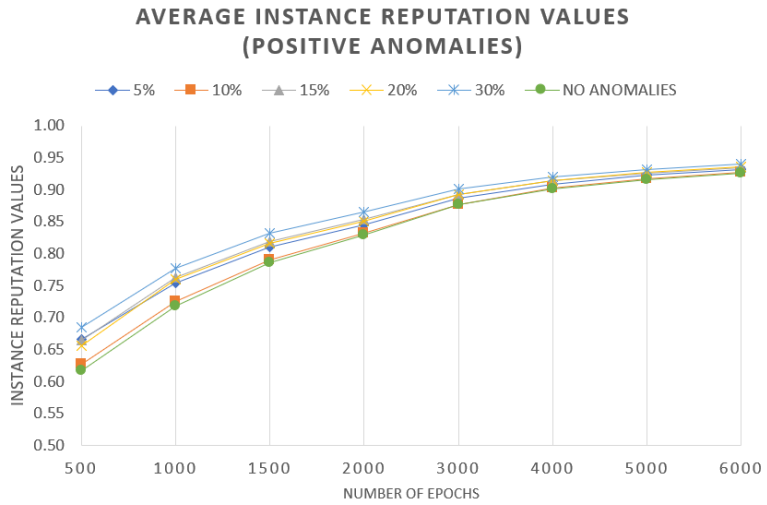

Figure 15: Average values of the reputation of an instance against the increase of positive anomalies for the MIoT with 300 instances

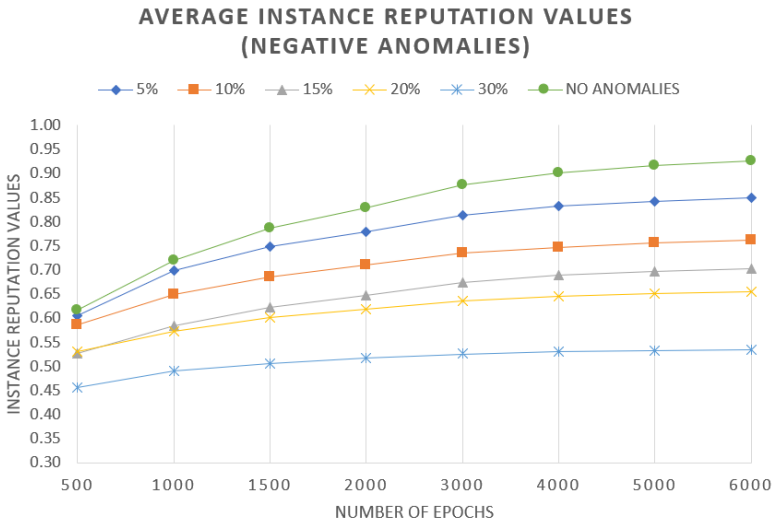

Figure 16: Average values of the reputation of an instance against the increase of negative anomalies for the MIoT with 300 instances

From the comparison of Figures 13 and 15, we can observe that, in presence of positive anomalies, the trend of the resilience for the computation of reputation is very similar to the one regarding the computation of trust; in this case, the increase of the average reputation is lower, as it ranges from 6.24% at 500 epochs to 1.07% at 6,000 epochs.

Analogously, by comparing Figures 14 and 16, we can observe that, in presence of negative anomalies, the trends of the resilience for the computation of reputation is similar (even if more mitigated)

to the one regarding the computation of trust. Indeed, the decrease of the average reputation ranges between 14.17% at 500 epochs and 29.37% at 6,000 epochs. Interestingly, in this case, the fall of the reputation values, when the percentage of negative anomalies passes from 20% to 30%, is less than the corresponding one observed for the trust values.

6.6 Accuracy

In order to measure the accuracy of our approach, we needed a ground truth regarding the trustworthiness of the smart objects involved in the MIoT. Unfortunately, the dataset used in the previous experiments did not have this information. As a consequence, we had to construct a new dataset. This was obtained by drawing inspiration from the smart city scenario described in Section 2.1. In particular, we asked 30 students of our university, 15 males and 15 females, to wear a smartwatch and run in three different parks of our town. The first was near the city center; the second was in a suburb; the third was in a naturalistic area near the sea. In each park we put several smart sensors capable of measuring temperature, humidity and light intensity. During the run of each student in each park, her/his smartwatch communicated with the park sensors to evaluate the environmental quality of the park. Through these communications, the students' smartwatches and the park's smart sensors could interact with each other to evaluate their mutual trust. At the same time, the smart sensors in each park communicated with each other and, thanks to these communications, it was possible to measure the trust of a smart sensor in the other ones of the same park. All these trust values contributed to the computation of the reputation of each sensor of the park.

The interested reader can find this dataset at the address `http://daisy.dii.univpm.it/miot/datasets/trustReputation` clicking on the link regarding this section.

The distribution of the average reputation $\widehat{R_a}$ of all park sensors provided by our approach, against the number of exchanged transactions, is reported in the left part of Figure 17. This figure shows that the average reputation is quite high, and this result was actually not surprising taking the previous experiments into account.

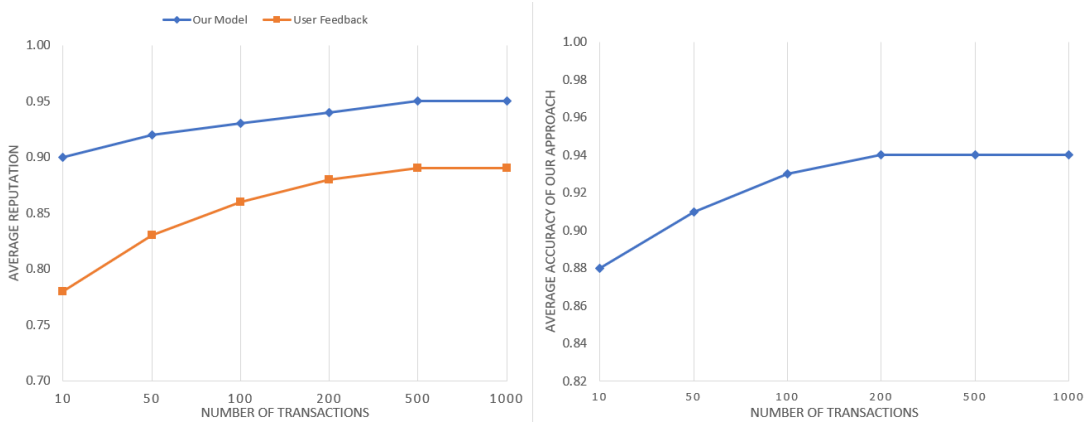

Figure 17: Accuracy of our approach

In order to have the ground truth, we asked each student to provide her/his evaluation of the information provided by each park sensor. To our surprise, we observed that this evaluation was not constant and grew over time. We attributed this to the fact that, as time went by, the runner

gradually adapted better to the environmental conditions of the park where she/he was running and, therefore, was more capable of objectively evaluating the information provided by sensors. The average reputation $\widehat{R_r}$ of all park sensors provided by the runners is reported at the left of Figure 17.

At this point, we were able to compute the average accuracy $\mathcal{A}$ of our approach. Specifically:

$$\mathcal{A} = 1 - |\widehat{R_a} - \widehat{R_r}|$$

The accuracy values obtained by our approach are reported at the right of Figure 17. From the analysis of this figure, we can observe that: *(i)* the accuracy values are always very high; *(ii)* they initially tend to increase over time; *(iii)* after an initial phase, they tend to become very stable and very high. These results allow us to conclude that the accuracy of our approach is certainly very satisfying.

7 Discussion

In this section, we propose some considerations on the experience made by performing the tasks discussed in this paper. In particular: In Subsection 7.1, we summarize the key aspects learned from the experimental validation. In Subsection 7.2, we discuss the possible usage of the extracted knowledge from a practical perspective. In Subsection 7.3, we discuss about how general the results that we found are from a practical point of view. Finally, in Subsection 7.4, we discuss the similarities/differences between communities of people and communities of objects.

7.1 Lesson Learned

The experiments have shown that our approach has an optimal resilience to the setting of weights and thresholds, even if it is (correctly) sensitive to weight errors when these become high (see Section 6.1). The results discussed in Section 6.1 also make us confident that the values obtained for weights are general and valid independently of the application environment in which our approach is operating. The experiments described in Section 6.3 revealed us that the time necessary to compute trust and reputation values is acceptable in all real cases. There are some theoretical situations in which this time could become unacceptable, but these cases are very far from the current real ones. Certainly, in the future, with the enormous development of IoTs, they could become possible, but we are confident that, in the meantime, the computation power of servers will simultaneously increase. In any case, we have specified some countermeasures that could be taken to face this problem, if it will happen in the future.

Section 6.4 revealed that it is possible to define a tradeoff between computation time and accuracy in determining the value of trust and reputation. Specifically, if computation time is the main factor to consider in this last activity, the number of epochs adopted to evaluate trust and reputation should not exceed 2,000 (which is a really huge number in the current real settings). If this number is not exceeded, there is no tradeoff to perform and, therefore, no need to sacrifice accuracy over computation time. By contrast, if the number of epochs is higher than 2,000, and there are more than 500-700 instances in the MIoT, in order to maintain an acceptable computation time, it is necessary to perform some actions that lead to partially sacrificing accuracy over computation time.

Section 6.5 shows that our approach is very resilient to positive anomalies and quite resilient to negative ones.

Last, but not the least, Section 6.6 reveals that: *(i)* the accuracy of our approach is always high; *(ii)* it tends to increase over time; *(iii)* after an initial phase, it tends to become very stable and high.

7.2 Possible usage of the extracted knowledge from a practical perspective

In Section 2, we have described two motivating examples, which illustrate two possible scenarios that could benefit from the approach presented in this paper. The former regards a smart city scenario and describes how people can use the data exchanged by their smart objects and the ones of the city to improve the effectiveness and the efficiency of their activities. The latter concerns a smart shopping center and illustrates how customers can use the data exchanged between their smart objects and the ones of the shopping center to improve the effectiveness and the efficiency of their shopping activities. However, these are only two of the large amount and variety of scenarios that could benefit from our approach. Think, for instance, of the adoption of smart objects to best regulate transports, to improve predictive maintenance in manufacturing, to regulate the patient flow to an hospital during a health emergency, to regulate the visitor flow to an exposition, and so forth.

7.3 Generalization level of results from a practical point of view

Throughout the discussions on the figures presented in Section 6, we have seen that most of the experiments are general (think, for instance, of the ones for setting the values of α, β and ρ presented in Section 6.1) and do not depend on the application environment our approach is operating on.

Furthermore, in Section 3 and in the next ones, we have observed that the limits on the computation time that we could find in our approach regard theoretical cases that are currently very over-dimensioned w.r.t. the real scenarios. In fact, all current real scenarios are fully manageable by our approach, which can be considered general and not limited to some specific scenarios.

The generality of our approach, and its applicability to all real cases, also regard its resilience (as witnessed by the results and the discussion of Section 6.5) and its accuracy (as witnessed by the results and the discussion of Section 6.6).

All the reasonings above make us confident that our approach can be a precious support in a large variety and amount of practical situations, as the ones described in Section 2 and the other ones mentioned above.

7.4 Similarities and differences between communities of people and communities of objects

Sensors and devices are becoming increasingly smart. The amount of data that they can store and the computing power at their disposal are constantly increasing. If, in the past, they were passive entities, without any autonomy, currently they have become increasingly active.

In this scenario, it is not surprising that, for some years, researchers have started to discuss about Social Internet of Things and, if the social relationships investigated by them in the past were extremely simple and elementary, the ones analyzed in the current researches are increasingly rich, complex and variegate.

Smart objects start to have a profile and to show a behavior obtained by implementing artificial intelligence-based algorithms on them. As a consequence, the boundary between what can be done by communities of people and communities of objects becomes increasingly blurred [43].

Clearly, in this discussion, we are considering only the technical viewpoint. However, when we discuss on Social Networking and Social Network Analysis (both if they are applied to humans and if they are applied to smart objects), we must consider that there is also another viewpoint, more related to humanistic and sociological studies. It concerns the investigation of the intrinsic essence distinguishing humans from animals and humans from machines. As for this aspect, we think that the gap between humans and smart objects is still enormous and, in our opinion, it will never be fully filled. But, here, we would open a discussion which is not object of this paper.

8 Conclusion

In this paper, we have proposed a new approach to evaluate trust and reputation in a Multi-IoTs scenario. We have seen that things are becoming increasingly complex, so that it is not out of place to talk about the profile of a thing and to assume for it a behavior similar to the one of a human. In this scenario, it also makes sense to introduce the trust of a thing in another one and the reputation of a thing in an IoT. We have also observed that the number and the variety of available things is leading researchers to model the existing reality as a set of IoTs interacting with each other, instead of a unique IoT. Therefore, we have presented the MIoT paradigm to model this scenario, and we have defined the concept of reputation of a thing or of an IoT in the MIoT.

Finally, we have presented an experimental campaign showing that our definitions are reasonable, their computation can be performed in an acceptable time, the accuracy they guarantee is high; moreover, they are resilient in presence of marginal errors, whereas they are correctly sensible to greater ones.

In the future, we plan to continue our research efforts in different directions. First, we would like to define a recommendation strategy taking into account the profile, the trust and the reputation of things in an IoT. Furthermore, we would like to exploit trust and reputation to define an approach for detecting malicious behaviors of things. Afterwards, trust and reputation can be used to develop an approach for detecting anomalous objects in a MIoT. Finally, we plan to leverage the concept of profile of a thing to create virtual communities of things having homogeneous interests or being assortative w.r.t. trust and reputation concepts.

Acknowledgments

This work was partially funded by the Department of Information Engineering at the Polytechnic University of Marche under the project "A network-based approach to uniformly extract knowledge and support decision making in heterogeneous application contexts" (RSAB 2018), and by the Marche Region under the project "Human Digital Flexible Factory of the Future Laboratory (HDSFIab) - POR MARCHE FESR 2014-2020 - CUP B16H18000050007".

References

[1] Concise Oxford Dictionary. *https://en.oxforddictionaries.com/*, 2020.

[2] IPSO Alliance. *https://www.ipso-alliance.org/*, 2020.

[3] F. Al-Turjman. *Cognitive Sensors and IoT: Architecture, Deployment, and Data Delivery.* Boca Raton, Florida, USA, 2017. CRC Press.

[4] F. Al-Turjman. Information-centric sensor networks for cognitive IoT: an overview. *Annals of Telecommunications*, 72(1-2):3–18, 2017. Springer.

[5] F. Al-Turjman and S. Alturjman. Context-sensitive access in industrial internet of things (IIoT) healthcare applications. *IEEE Transactions on Industrial Informatics*, 14(6):2736–2744, 2018. IEEE.

[6] M.D. Alshehri and F.K. Hussain. A fuzzy security protocol for trust management in the internet of things (Fuzzy-IoT). *Computing*, 101(7):791–818, 2019. Springer.

[7] A.P. Plageras and K.E. Psannis and C. Stergiou and H. Wang and B.B. Gupta. Efficient iot-based sensor big data collection–processing and analysis in smart buildings. *Future Generation Computer Systems*, 82:349–357, 2018. Elsevier.

[8] D. Artz and Y. Gil. A survey of trust in computer science and the Semantic Web. *Web Semantics: Science, Services and Agents on the World Wide Web*, 5(2):58–71, 2007.

[9] L. Atzori, A. Iera, and G. Morabito. SIoT: Giving a social structure to the Internet of Things. *IEEE Communications Letters*, 15(11):1193–1195, 2011. IEEE.

[10] G. Baldassarre, P. Lo Giudice, L. Musarella, and D. Ursino. A paradigm for the cooperation of objects belonging to different IoTs. In *Proc. of the International Database Engineering & Applications Symposium (IDEAS 2018)*, pages 157–164, Villa San Giovanni, Italy, 2018. ACM.

[11] G. Baldassarre, P. Lo Giudice, L. Musarella, and D. Ursino. The MIoT paradigm: main features and an "ad-hoc" crawler. *Future Generation Computer Systems*, 92:29–42, 2019. Elsevier.

[12] F. Bao, R. Chen, and J. Guo. Scalable, adaptive and survivable trust management for community of interest based internet of things systems. In *Proc. of the International Symposium on Autonomous Decentralized Systems (ISADS'13)*, pages 1–7, Mexico City, Mexico, 2013. IEEE.

[13] P.A. Bernstein, J. Madhavan, and E. Rahm. Generic Schema Matching, Ten Years Later. *Proceedings of the VLDB Endowment*, 4(11):695–701, 2011.

[14] C. Stergiou and K.E. Psannis and B.B. Gupta and Y. Ishibashi. Security, privacy & efficiency of sustainable cloud computing for big data & iot. *Sustainable Computing: Informatics and Systems*, 19:174–184, 2018. Elsevier.

[15] J. Caverlee, L. Liu, and S. Webb. The socialtrust framework for trusted social information management: Architecture and algorithms. *Information Sciences*, 180(1):95–112, 2010.

[16] G. Chen, B.D. Ward, C. Xie, W. Li, Z. Wu, J. Jones, M. Franczak, P. Antuono, and S. Li. Classification of Alzheimer disease, mild cognitive impairment, and normal cognitive status with large-scale network analysis based on resting-state functional MR imaging. *Radiology*, 259(1):213–221, 2011. Radiological Society of North America, Inc.

[17] R. Chen, J. Guo, and F. Bao. Trust management for SOA-based IoT and its application to service composition. *IEEE Transactions on Services Computing*, 9(3):482–495, 2016. IEEE.

[18] H.S. Choi and W.S. Rhee. Social based Trust Management System for Resource Sharing Service. In *Proc. of International Conference on Intelligent Systems, Metaheuristics & Swarm Intelligence (ISMSI'18)*, pages 148–152, Phuket, Thailand, 2018. ACM.

[19] P. De Meo, A. Nocera, D. Rosaci, and D. Ursino. Recommendation of reliable users, social networks and high-quality resources in a Social Internetworking System. *AI Communications*, 24(1):31–50, 2011. IOS Press.

[20] P. De Meo, G. Quattrone, G. Terracina, and D. Ursino. Integration of XML Schemas at various "severity" levels. *Information Systems*, 31(6):397–434, 2006.

[21] P. De Meo, G. Quattrone, and D. Ursino. A query expansion and user profile enrichment approach to improve the performance of recommender systems operating on a folksonomy. *User Modeling and User-Adapted Interaction: The Journal of Personalization Research (UMUAI)*, 20(1):41–86, 2010. Springer.

[22] G. Fortino, W. Russo, C. Savaglio, W. Shen, and M. Zhou. Agent-oriented cooperative smart objects: From IoT system design to implementation. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 48(11):1939–1956, 2017. IEEE.

[23] G. Fortino, W. Russo, C. Savaglio, M. Viroli, and M. Zhou. Modeling Opportunistic IoT Services in Open IoT Ecosystems. In *Proc. of the Workshop "From Objects to Agents" (WOA'17)*, pages 90–95, Scilla (RC), Italy, 2017.

[24] G. Fortino, C. Savaglio, C. E. Palau, J. S. de Puga, M. Ganzha, M. Paprzycki, M. Montesinos, A. Liotta, and M. Llop. Towards multi-layer interoperability of heterogeneous IoT platforms: The INTER-IoT approach. pages 199–232, 2018. Springer.

[25] D. Gambetta. Can we trust trust? In *Trust: Making and breaking cooperative relations*, chapter 13, pages 213–237. 2000.

[26] S. Garruzzo, S. Modafferi, D. Rosaci, and D. Ursino. X-Compass: an XML agent for supporting user navigation on the Web. In *Proc. of the International Conference on Flexible Query Answering Systems (FQAS 2002)*, pages 197–211, Copenhagen, Denmark, 2002. Lecture Notes in Artificial Intelligence, Springer-Verlag.

[27] D. Georgakopoulos and P.P. Jayaraman. Internet of things: from internet scale sensing to smart services. *Computing*, 98(10):1041–1058, 2016. Springer.

[28] A. Jøsang, R. Ismail, and C. Boyd. A survey of trust and reputation systems for online service provision. *Decision support systems*, 43(2):618–644, 2007. Elsevier.

[29] D.H. McKnight and N.L. Chervany. The meanings of trust. *Technical Report MISRC (Management Information Systems Research Center) - Working Paper Series 96-04*, 1996. University of Minnesota.

[30] M. McPherson, L. Smith-Lovin, and J.M. Cook. Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27:415–444, 2001. JSTOR.

[31] M. Nitti, R. Girau, and L. Atzori. Trustworthiness management in the social internet of things. *IEEE Transactions on Knowledge and Data Engineering*, 26(5):1253–1266, 2014. IEEE.

[32] L. Page, S. Brin, R. Motwani, and T. Winograd. The PageRank Citation Ranking: Bringing Order to the Web. In *Proc. of the Seventh International World-Wide Web Conference (WWW 1998)*, pages 161–172, Brisbane, Australia, 1998. Elsevier.

[33] A.P. Plageras, C. Stergiou, G. Kokkonis, K.E. Psannis, Y. Ishibashi, B.G. Kim, and B.B. Gupta. Efficient large-scale medical data (ehealth big data) analytics in internet of things. In *Proc. of the Conference on Business informatics (CBI'17)*, volume 2, pages 21–27, Thessaloniki, Greece, 2017. IEEE.

[34] A.R.G. Ramirez, I. González-Carrasco, G.H. Jasper, A.L. Lopez, J.L. Lopez-Cuadrado, and A. García-Crespo. Towards human smart cities: internet of things for sensory impaired individuals. *Computing*, 99(1):107–126, 2017. Springer.

[35] P. Resnick and R. Zeckhauser. Trust among strangers in Internet transactions: Empirical analysis of eBay's reputation system. In *The Economics of the Internet and E-commerce*, pages 127–157. 2002. Emerald Group Publishing Limited.

[36] J. Sabater and C. Sierra. Review on computational trust and reputation models. *Artificial Intelligence Review*, 24(1):33–60, 2005.

[37] Y.B. Saied, A. Olivereau, D. Zeghlache, and M. Laurent. Trust management system design for the Internet of Things: A context-aware and multi-service approach. *Computers & Security*, 39:351–365, 2013. Elsevier.

[38] C. Savaglio, G. Fortino, and M. Zhou. Towards interoperable, cognitive and autonomic IoT systems: An agent-based approach. In *Proc. of the World Forum on Internet of Things (WF-IoT'16)*, pages 58–63, Reston, VA, USA, 2016. IEEE.

[39] A. Sheikhahmadi and M. Nematbakhsh. Identification of multi-spreader users in social networks for viral marketing. *Journal of Information Science*, 43(3):412–423, 2017. SAGE Publications Sage UK: London, England.

[40] Y. Sun, J. Zhang, Y. Xiong, and G. Zhu. Data security and privacy in cloud computing. *International Journal of Distributed Sensor Networks*, 10(7):190903, 2014. SAGE Publications Sage.

[41] N.B. Truong, T.W. Um, and G.M. Lee. A reputation and knowledge based trust service platform for trustworthy social internet of things. In *Proc. of the International Conference on Innovations in Clouds, Internet and Networks (ICIN '16)*, Paris, France, 2016.

[42] N.B. Truong, T.W. Um, B. Zhou, and G.M. Lee. From personal experience to global reputation for trust evaluation in the social internet of things. In *Proc. of the International IEEE Global Communications Conference (GLOBECOM'17)*, pages 1–7, Singapore, 2017. IEEE.

[43] D. Ursino and L. Virgili. Humanizing IoT: defining the profile and the reliability of a thing in a Multi-IoT scenario. *Towards Social Internet of Things: Enabling Technologies, Architectures and Applications. Studies in Computational Intelligence*, 846:51–76, 2020. Springer Nature.

[44] L.H. Wang, R.C. Bucelli, E. Patrick, D. Rajderkar, E. Alvarez III, M.M. Lim, G. DeBruin, V. Sharma, S. Dahiya, R.E. Schmidt an T.S. Benzinger, B.A. Ward, and B.M. Ances. Role of magnetic resonance imaging, cerebrospinal fluid, and electroencephalogram in diagnosis of sporadic Creutzfeldt-Jakob disease. *Journal of Neurology*, 260(2):498–506, 2013. Springer.

[45] Y. Wang and J. Vassileva. Toward trust and reputation based web service selection: A survey. *International Transactions on Systems Science and Applications*, 3(2):118–132, 2007.

[46] S.R. Yan, X.L. Zheng, Y. Wang, W.W. Song, and W.Y. Zhang. A graph-based comprehensive reputation model: Exploiting the social context of opinions to enhance trust in social commerce. *Information Sciences*, 318:51–72, 2015. Elsevier.

[47] Z. Yan, P. Zhang, and A.V. Vasilakos. A survey on trust management for Internet of Things. *Journal of Network and Computer Applications*, 42:120–134, 2014. Elsevier.

[48] D. Zissis and D. Lekkas. Addressing cloud computing security issues. *Future Generation Computer Systems*, 28(3):583–592, 2012. Elsevier.