Future Generation Computer Systems 49 (2015) 113–124

Contents lists available at ScienceDirect

Future Generation Computer Systems

journal homepage: www.elsevier.com/locate/fgcs

Dynamic and flexible selection of a reputation mechanism for heterogeneous environments

Ginés Dólera Tormo [b,*,1], Félix Gómez Mármol [a], Gregorio Martínez Pérez [b]

[a] NEC Europe Ltd., Kurfürsten-Anlage 36, 69115 Heidelberg, Germany

[b] Department of Information and Communications Engineering, University of Murcia, Murcia, 30100, Spain

HIGHLIGHTS

- Design a flexible mechanism to select the most suitable trust and reputation model to apply on-the-fly.
- Selection of the reputation model based on current system conditions and expected performance measurements.
- The approach guarantees a smooth and automatic transition between the reputation computation engines.
- Experiments conducted to validate the feasibility of the approach.
- Focus on heterogeneous Internet of Things environments.

ARTICLE INFO

Article history:
Received 4 April 2014
Received in revised form
3 June 2014
Accepted 14 June 2014
Available online 23 June 2014

Keywords:
Trust and reputation management
Internet of things
Interchangeable computation engine

ABSTRACT

Current trust and reputation management approaches usually offer rigid and inflexible mechanisms to compute reputation scores, which hinder their dynamic adaptation to the current circumstances in the system where they are deployed. At most, they provide certain parameters which are configurable or tunable. Yet, this is not enough for such heterogeneous and dynamic environments as the ones introduced by Internet of Things (IoT). In this paper we propose a rupture with this old philosophy and have therefore designed and prototyped a flexible mechanism to select the most suitable trust and reputation model to apply on-the-fly, amongst a pool of predefined ones, considering both the current system conditions and the selected performance measurements, which, to the best of our knowledge, is missing nowadays. Additionally, this mechanism guarantees a smooth transition between different computation engines avoiding abrupt changes in the computed reputation scores. Conducted experiments prove that our solution is able to identify and start up the most suitable trust and reputation model depending on the current system conditions (number of users, allocated resources, etc.) and expected performance measurements (accuracy, scalability, robustness, etc.).

1. Introduction

Trust and reputation management systems are widely spread and used today in the Internet. We find them in a myriad of service provisioning scenarios, ranging from pure e-Commerce ones to blogs, social networks, video streaming services, etc. [1–3]. Furthermore, an extensive amount of research work has been performed as well in applying trust and reputation management techniques to P2P networks [4], wireless sensor networks [5], vehicular ad hoc networks [6], Cloud Computing [7], Identity Management systems [8], collaborative intrusion detection networks [9,10], etc.

Nevertheless, though this large variety of systems and scenarios constitute a proof of the applicability and feasibility of trust and reputation management solutions, they also lead to a so far neglected problem raised by widely dynamic environments as the ones introduced by Internet of Things (IoT). In IoT environments, many heterogeneous devices (i.e. widely having dissimilar elements, features or behaviors) define dynamic, complex and distributed frameworks [11–13]. Obviously, each system/scenario has different and specific requirements and particularities in terms of infrastructure design, participating entities, communication capabilities, exchanged data, etc. Even

* Corresponding author. Tel.: +34 868 887646; fax: +34 868 884151.
E-mail addresses: ginesdt@um.es (G. Dólera Tormo), felix.gomez-marmol@neclab.eu (F. Gómez Mármol), gregorio@um.es (G. Martínez Pérez).

[1] The work presented in this paper was performed while working at NEC Laboratories Europe.

http://dx.doi.org/10.1016/j.future.2014.06.006
0167-739X/

more, the scenarios might be dynamically and continuously changing (their topology, their committed resources, etc.).

Current trust and reputation models usually provide certain configuration parameters aimed to tune the behavior of the deployed mechanism. However, this settings feature is quite often not able to offer the high dynamicity required to adapt the model to different circumstances that may happen in IoT environments. As shown in [14] each trust and reputation model has its advantages and shortcomings, and most of the times their drawbacks are related to the current conditions of the system [15] (number of participants, number of feedbacks, feedbacks storage capabilities, computational capabilities, etc.). Moreover, the expected performance measurements (accuracy, robustness, resilience against attacks, etc.) also affect on the selection of the most appropriate reputation models, and they usually depend on the requirements of the scenario.

Therefore, a flexible mechanism to select the most suitable trust and reputation model to apply on-the-fly considering both the current system conditions and the expected performance measurements, to the best of our knowledge, is missing today. Notwithstanding such fact, we should not ignore the interoperability between models, since every model might not be applicable to all scenarios. However, reputation scores are usually calculated through a module within the model, known in the literature as *reputation computation engine*, which implements different algorithms to aggregate recommendations depending on the scenario.

In order to tackle the aforementioned problem, we have designed and prototyped a system [16] which is able to dynamically select the most appropriate reputation computation engine according to the current conditions of the system (in terms of number of users, number of feedbacks, available bandwidth, available storage capacity, etc.), as well as to the desired performance measurements (i.e., accuracy, user satisfaction, adaptability, resilience to certain attacks, etc.), aimed to deal with the dynamicity introduced by IoT. The suitability of each reputation computation engine is computed based on predefined inference rules, which relate system conditions to performance measurements. These rules are defined using fuzzy sets [17], in order to improve the flexibility in the rule definition process, usually performed by system administrators. Additionally, our approach guarantees a smooth transition between different computation engines, avoiding abrupt changes in the computed reputation scores.

Applying trust and reputation to IoT is no longer a matter of designing and developing tunable trust and reputation models, but to provide a pool of them, analyzing their intrinsic characteristics in order to determine under which conditions they provide the best outcomes for each of the desired performance measurements. In this way, knowing which are the performance metrics required by each scenario, and monitoring the system conditions, the system is able to dynamically select the most suitable model according to the parameters of such a scenario. The models are automatically swapped in order to have the most suitable reputation model working at each moment.

Furthermore, how to measure the suitability could depend on the scenario, since the performance metrics to be optimized might vary from one environment to another. For instance, a target of a reputation model within a IoT-based sensors environment could be to minimize the consumption of required resources [18], whereas for a IoT cloud ecosystem the target could be to improve the accuracy of the applied trust and reputation mechanisms no matter how much computation is required [19].

The remainder of this article is structured as follows. Section 2 presents a description of the problem being addressed in this work. Then, in Section 3 we introduce our solution for dynamically selecting the most appropriate reputation model on-the-fly, while in Section 4 a mechanism to avoid abrupt changes when reputation computation engines are swapped is described. Section 5 shows some conducted experiments in order to validate the feasibility of the proposed system. Later, Section 6 provides the main references and related works, while in Section 7 the main conclusions and lines of future work are outlined.

2. Problem statement

IoT is predicted to revolutionize the way organizations implement their information systems and applications [20]. Expecting tons of devices seamlessly integrated into information networks, IoT enables unlimited scalability and greater flexibility all at a contained cost. It introduces several new business models based on unlimited application scenarios and new smart services. On the other hand, IoT raises new challenges, and companies and organizations supporting this concept have to face a number of burdens and barriers to its deployment. Trust, at the core of these concerns, is identified as a critical component to allow the IoT to reach its greatest potential [12,21].

In such a dynamic and distributed environment, static agreements are generally hard to apply for managing trust relationships between different entities, such as those based on Service-Level Agreements (SLA), where rigid (and usually complex) contracts are applied. Reputation management systems have emerged in the last years as a solution for this kind of scenarios, where the trust of a given entity is dynamically acquired from analyzing its past and recent behavior.

Reputation management systems propose mechanisms to allow an entity to somehow determine if another entity can be taken as reliable or not, in order to get some services or exchange some information between them. Firstly, these systems try to collect as much information as possible about the behavior of a given entity, which usually comes from recommendations based on past experiences. Secondly, all gathered recommendations are aggregated in order to calculate a reputation score for such an entity [14].

The computed score will be used to decide the level of trustworthiness that a particular entity has. If such entity has enough reputation, the communication process is therefore triggered between the entities. Finally, the service consumer provides the satisfaction with regards to the received service which, in turn, will be used as a recommendation for calculating future reputation scores.

For instance, we can think of an IoT scenario in the eHealth context, where information about patients is collected through heterogeneous sensors, in order to track their status. These sensors might be deployed through complex network structures, where the trust information could be hardly managed in a single point. As an alternative, reputation management systems could collect recommendations from different sources, to decide whether a given entity (e.g. a sensor, a set of them, an intermediary node, or even a network path [22]) is trustworthy enough to be in charge of such a sensitive data.

Reputation management systems have been proposed in different contexts and they have been applied to many different scenarios [1]. Since the behavior of the reputation system mainly depends on the requirements of the scenario where such system is deployed, a lot of mechanisms to accomplish the aforementioned steps have been defined. For instance, depending on the computation constraints of the devices which aggregate and compute the reputation values, a different reputation computation engine might be used, whereas a different way of collecting recommendations may be used depending on their network capabilities.