ARTICLE IN PRESS

Future Generation Computer Systems ▮ (▮▮▮▮) ▮▮▮–▮▮▮

Contents lists available at ScienceDirect

Future Generation Computer Systems

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

Trust management for distributed heterogeneous systems by using linguistic term sets and hierarchies, aggregation operators and mechanism design

Christian Esposito [a], Aniello Castiglione [a,*], Francesco Palmieri [a], Massimo Ficco [b]

[a] Department of Computer Science, University of Salerno, Via Giovanni Paolo II, 132 I-84084 Fisciano (SA), Italy

[b] Department of Industrial and Information Engineering, Second University of Naples, Via Roma 29, I-81031 Aversa (CE), Italy

HIGHLIGHTS

- We have modeled the problem of trust computation by means of fuzzy theory with crisp and fuzzy values.
- We have proposed to deal with heterogeneity by applying Linguistic Hierarchies and transformations among different hierarchies.
- We have applied game theory so as to promote truth-telling behaviors in reputation dissemination.
- We presented simulations of our solution and demonstrated the benefits it can provide.

ARTICLE INFO

Article history:
Received 27 July 2015
Received in revised form
17 November 2015
Accepted 3 December 2015
Available online xxxx

Keywords:
Heterogeneity
Trust management
Linguistic fuzzy term sets
Mechanism design

ABSTRACT

The current trend in system development is integrating already-existing systems in order to realize large-scale infrastructures. Several of these infrastructures exhibit stringent security requirements that must be handled by properly managing the trust relationships within both the different systems involved, and all the external entities interacting with the integrated infrastructure. Trust management is made complex by the intrinsic heterogeneity characterizing the integrated systems. To handle such a heterogeneity, we propose the application of fuzzy logic, combining it with proper means to deal with heterogeneous fuzzy sets. We present a technique to combine qualitative and quantitative specifications of trust scores aiming at periodically computing a new trust degree, by also considering reputation scores collected from other systems within the infrastructure. Last, we applied the game theory in order to promote truth-telling behavior during the process of reputation information collection. We finally present some experimental results showing the effectiveness of our approach in heterogeneous distributed environments.

1. Introduction

The current trend in building complex information systems is not to realize them ex-novo, but to take advantage of existing solutions as well as integrating and engineering them in order to obtain a larger hardware/software infrastructure able to perform more complex operations than the ones provided by the individual components. In current industrial practices, we can find several concrete examples of this trend, such as the EU-funded project Single European Sky ATM Research (SESAR) [1] aiming at federating the existing Air Traffic Management (ATM) centers in order to obtain a seamless pan-European ATM infrastructure, or several projects started around the world for the implementation of integrated Health Information Systems (HIS) [2] whose main goal is providing Electronic Health Records (EHR) [3] or Electronic Medical Records (EMR) [4] to citizens by integrating and federating the information hold by public and/or private health centers within a country or among countries in order to support internal and external patient mobility [5].

In most of the cases, these integration projects assume the existing component systems as black boxes, *i.e.*, characterized by specific interfaces and without any knowledge of their internal implementation, and introducing any changes to their internal operations and/or organization is quite difficult for both economic or convenience reasons. This is likely to cause several problems, spanning from interoperability (which can occur at technological,

* Corresponding author. Tel.: +39 089969594; fax: +39 089969821.
E-mail addresses: christian.esposito@dia.unisa.it (C. Esposito), castiglione@acm.org, castiglione@ieee.org (A. Castiglione), fpalmieri@unisa.it (F. Palmieri), massimo.ficco@unina2.it (M. Ficco).

http://dx.doi.org/10.1016/j.future.2015.12.004
0167-739X/

syntactical and semantical levels) to security. In this work, we are interested to some security issues characterizing large-scale distributed infrastructures made of off-the-shelf systems, and more specifically to the problem of trust management in heterogeneous distributed operating scenarios. It is reasonable to assume that each individual system involved in the integration process, especially within the context of critical domains such as the ones of the above-mentioned examples, allows only trusted users and cooperating systems to successfully request certain operations and/or access certain resources. Such authorization strategy is mostly driven by the concept of trust and by the perception of reputation that the overall system has of the requesting entities. In the current academic literature and industrial practice, we can find different trust models and management solutions. There is an on-going debate on standardizing the language/abstraction to be used to express trust policies, such as the concept of a security token in WS-Trust [6] or the trust assertions expressed with the Security Assertion Markup Language (SAML) [7], or alternatively, or alternatively using ontologies to define a common vocabulary and semantics for trust management [8–10]. However, rather than expressing trust by means of a proper formalism, we are more interested in building a trust degree based on direct observations, propagating trust considerations by means of reputations, and dealing with selfish and/or malicious systems/users that hide or manipulate reputations in order to force the assignment of a good trust degree to malicious entities or a bad trust degree to honest ones.

Trust degrees can be expressed by using a numeric representation [11], for example by means of a non negative real number between 0 or 1, which can be interpretable as the probability to trust a certain entity, or of a real number between -1 and 1, where at the two extremes we have respectively the indication of untrusted or trusted entities while in the middle we can find the representation of various shades of trustability. Typically, the systems being integrated are manually configured, by a human operator acting as administrator of the infrastructure, with a starting value of trust for certain users/systems and/or groups of them. For this reason, a more human-friendly approach to trust expression is needed, empowered by the use of linguistic labels, such as given adjectives like LOW or HIGH related to trust towards a given entity. Such a linguistic approach to trust representation is not new, since some solutions in this sense are available in the literature, as presented in [12,13]. However, such solutions fail in dealing with the intrinsic vagueness of humans and do not assume uncertain linguistic information, so that the operator cannot decide on a single linguistic label, but, assuming an ordered linguistic term sets, it can pick two linguistic labels with the first one ranked below the second one. As a practical example, if we assume the linguistic term set as composed of LOW, SCARCE, MIDDLE, GOOD, HIGH, an operator can select one of these terms but also the interval (SCARCE, GOOD) in the case of uncertainty on the trust degree related to a given entity.

Moreover, since multiple heterogeneous systems have to be integrated, it is possible that not all of them agree on the same linguistic term sets, but in the overall infrastructure, we can find a multi-granularity of linguistic information, *i.e.*, some systems operate with a linguistic term set composed of five terms, as the previous example, and other ones with a set of four terms, or more than five terms. Having a linguistic representation of trust, a problem to be addressed is how to translate a given set of observations of past interactions with an entity into a proper trust degree. Moreover, not only the trust degree must be expressed by using linguistic terms, but also the various reputations of a given entity collected from other systems. Such collected reputations and the trust degree determined from direct observation must be properly aggregated in order to reach a final trust consideration related to a specific entity. Last, but not least, we cannot assume that the collected reputations may be all honest, but it is possible that an external system, or a group of them, may cheat on the sent reputations by acting in a rational/selfish and/or malicious manner. Such behavior can result in neglecting to pass over the reputation information or providing false reputations.

In the current work, we address the aforementioned trust and reputation management issues by means of three different strategies, constituting its main contribution directions:

1. we employ fuzzy set theory to deal with uncertain linguistic information, quantitative determination of trust degree and multi-granularity in linguistic term sets;
2. we make use of linguistic aggregation operators to composite linguistic terms in order to reach the final determination of trust degree based on collected reputations and direct observation;
3. we adopt game theory to promote collaboration and truth saying among systems when spreading their reputations about certain entities.

In order to evaluate the effectiveness of the proposed approach, we have performed several simulations of a large-scale distributed infrastructure, with several interacting users located at different portions of the system, by using the OMNET++ event-driven simulation environment. We evaluate the latency associated with computing the trust degree of a requesting entity, and the number of rational entities operating within the system. In particular, experimental results show that the latency is kept minimal, as well as adopted approach is able to reduce the number of rational players by making the proposed strategy less appealing.

The paper is structured as follows. Section 2 provides the needed background to support the description of our proposal and an analysis of related work on the topic of trust management in distributed systems. Section 3 presents the approach and details its fundamental contributions, which have been experimentally verified, as discussed in Section 4, in order to test the quality and performance of the proposed trust management solution. We conclude with Section 5, which contains some final remarks and plans for future work.

2. Background and related work

Trust is commonly defined as the degree of belief to be established in order to assume that the right user/entity is accessing the right resource in order to conduct the allowed operation. In secure solutions built from the aggregation and cooperation of multiple heterogeneous systems and entities, trust is used as a measure to decide the granting or denying of a given request received from an authenticated user/entity, *i.e.*, for which the identity is known at the time of the decision and whose veridicality is certified by a trusted third party. Trust management [14] is the overall process based on collecting data or a-priori configuration information in order to establish a trust relationship between the involved system and an authenticated user/entity. Sometimes, such a relationship may have a dynamic nature, where the interactions between the two parties are continuously monitored in order to adjust the trust degree behind the relationship itself. Basically, we can have two broad classes that describe the available solutions for trust management: the first one is based on policies [15], expressed in a given standard format and language such as eXtensible Access Control Markup Language (XACML) [16], and the other is based on reputations [17,18]. In the first case, the trust management is limited to verifying the credentials provided by requesting entities and the satisfaction of one or more policies in order to grant access and establish a trust relationship. As long as one of the available policies is satisfied, the requesting entity is considered trustworthy and requests are granted. In the second case, trust relationships are established based on reputations, which are observations on the

Please cite this article in press as: C. Esposito, et al., Trust management for distributed heterogeneous systems by using linguistic term sets and hierarchies, aggregation operators and mechanism design, Future Generation Computer Systems (2015), http://dx.doi.org/10.1016/j.future.2015.12.004