Increasing protection and autonomy in the IoT through a two-tier blockchain framework

(Discussion Paper)

Enrico Corradini[1], Serena Nicolazzo[2], Antonino Nocera[3], Domenico Ursino[1] and Luca Virgili[1]

[1]*Department of Information Engineering, Polytechnic University of Marche*
[2]*Daisy Lab, Politechnic University of Marche*
[3]*Department of Electrical, Computer and Biomedical Engineering, University of Pavia*

Abstract

In this paper, we propose an approach that uses a two-tier blockchain framework and a trust-based protection mechanism to increase the security and autonomy of smart objects in the IoT. The proposed approach groups the involved smart objects into suitable communities. The two blockchains perform different, but complementary, tasks. Indeed, the first-tier blockchain is local and records probing transactions performed to evaluate the trust of one smart object in another. Periodically, after a time window, the probing transactions are aggregated to determine the reputation of each smart object within its community and the trust of one community in each of the others. These values are stored in the second-tier blockchain. This paper describes the proposed approach, the underlying framework, the behavior, the security model and a test carried out to evaluate its performance.

Keywords

Internet of Things, Blockchain, Protection, Autonomy, Reliability, Trust, Reputation,

1. Introduction

In recent years, the Internet of Things (IoT) paradigm has become increasingly successful and pervasive. However, at the same time, it has posed new challenges. Indeed, the IoT involves the presence of a large number of smart objects cooperating with each other. These objects are often characterized by constraints on storage, computational capability, criticality and sensitivity of used services and applications. At the same time, they show a great dynamism. In this scenario, the protection of smart objects and the possibility of granting them autonomy represent two challenges that must be faced simultaneously.

As for *protection*, [1] presents an approach to address this issue when it comes to privacy. It partially hides object features, but allows for their full usage to support inter-object communication. However, this approach does not provide a scalable, reliable and secure framework

SEBD 2021: The 29th Italian Symposium on Advanced Database Systems, September 5-9, 2021, Pizzo Calabro (VV), Italy
e.corradini@pm.univpm.it (E. Corradini); serena.nicolazzo.sn@gmail.com (S. Nicolazzo); antonino.nocera@unipv.it (A. Nocera); d.ursino@staff.univpm.it (D. Ursino); l.virgili@pm.univpm.it (L. Virgili)
https://kmitd.github.io/ilaria/ (S. Nicolazzo)
0000-0002-1140-4209 (E. Corradini); 0000-0003-2719-9526 (S. Nicolazzo); 0000-0003-2120-2341 (A. Nocera); 0000-0003-1360-8499 (D. Ursino); 0000-0003-1509-783X (L. Virgili)

for the IoT devices. As for *autonomy*, in order to make objects independent from each other during their interactions, it is necessary to include a solution allowing objects to add/remove contacts, and to identify which features/services are provided by other objects [2]. Moreover, it is crucial to design a mechanism assessing the ability of an object to correctly provide the necessary features/services. The previous reasoning shows that protection and autonomy are two strongly correlated aspects whose effective management requires the definition of trust and reputation mechanisms. However, due to the peculiarities of the IoT scenario discussed above, existing solutions for sensors or P2P networks are not directly applicable [3].

To achieve a fully distributed solution in this context, each smart object should be able to build a representation of the behavior of other objects in the IoT. To this end, it should be capable of unambiguously knowing the sequence of actions of its peers. To achieve this goal, the blockchain technology is unanimously recognized as one of the most effective strategy [4]. However, even the only monitoring of the public ledger is a heavy and expensive task for smart objects with a low computational capability in presence of a high volume of transactions. To overcome this problem, some authors have proposed to adopt approaches based on the use of a validity window and the aggregation of historical data within it. However, if the volume of transactions is big, this approach may be too expensive and unfeasible for the IoT [5].

This paper aims at providing a contribution in this setting. In fact, it proposes a two-tier blockchain framework to increase the protection and autonomy of smart objects in the IoT. Following the intuition proposed in [1], our approach organizes smart objects into communities. Moreover, it uses the local first tier of the framework to manage the reputation of each smart object within the community it belongs to. It also uses a validity window, coupled with a lightweight blockchain, to face the high transaction volume. The organization of objects in communities allows our approach to control the size of the blockchain, thus avoiding excessive loads for smart objects. Finally, it uses the second global tier to record aggregate data related to communities, as well as the trust values that each community assigns to the others. To implement the tasks of our approach, we leverage the blockchain's smart contract technology. It has already been successfully used in the context of the IoT, e.g., to implement single and multi-party authentication for an IoT device (see [6] for an example).

The outline of this paper is as follows: Section 2 describes the proposed framework. Section 3 illustrates our security model. Section 4 presents an experimental comparison with two other approaches. Finally, Section 5 draws conclusions and takes a look at some possible future developments.

2. Technical description of our approach

In our model, the main actor is the smart object. It has a profile characterized by: *(i)* an identifier; *(ii)* a set of features regarding it; *(iii)* a set of services it offers; *(iv)* the information that other smart objects need for communicating with it (e.g., its MAC and IP addresses, etc.). The smart objects of our framework can be partitioned into communities, based on some rules. Each smart object belongs to exactly one community. A source smart object can communicate with a target one through suitable transactions. These last can be classified in: *(i) ordinary*, if the source requests a service/feauture to the target; *(ii) probing*, if the source (called trustor) wants to test

what the target (called trustee) declared to assess its reliability. Transactions can also be *intra-community* (resp., *inter-community*) if they involve smart objects from the same community (resp., from different communities).

Each community has an associated local blockchain that records information about transactions having one of its smart objects as trustor. The overall IoT has associated a global blockchain, which records aggregate information produced periodically from probing transactions recorded in the local blockchains. Specifically, the global blockchain stores: *(i)* the list of smart objects belonging to each community and, for each of them, the corresponding reputation scores; *(ii)* the trust of each community in the others of the IoT.

The interaction mechanisms between smart objects allows each of them to understand which features/services can be provided by another one with which it is in contact [7]. To allow a smart object to asses whether another one is reliable in providing the features/services it advertises, we adopt the approach of [8] based on probing transactions. To certify them, we use a blockchain-based solution. Adopting blockchains in the IoT context poses important challenges concerning the large number of nodes involved, the large amount of data generated, and the low computational power of many smart objects. As specified in the Introduction, to address these challenges, our approach leverages a two-tier blockchain framework. It groups smart objects into appropriate communities, based on certain criteria. Within these communities, smart objects adopt control mechanisms to identify anomalous behaviors and make interactions as secure as possible. Our approach is independent of the way communities are built. It only requires that smart objects in a community should have some level of redundancy in the features/services offered.

Thus, the first layer of the framework is a blockchain underlying a community; it represents a local public ledger storing all the probing transactions performed within the community. There are several approaches to implement lightweight blockchains for the IoT context [9, 10] that could be adopted to create this layer. In our case, we could use any of them, for instance IOTA (www.iota.org), which supports smart contracts via the QUBIC protocol [11]. The second layer of the framework is a global blockchain that involves the whole IoT and only stores the *aggregate* information of the different communities. It could be implemented with any blockchain, such as Ethereum (www.ethereum.org) or HyperLedger (www.hyperledger.org). In Figure 1, we report the general architecture of our approach.

In order to limit the volume of transactions to be analyzed, communications between devices take place within time windows, where the devices perform ordinary and probing transactions. These last are randomly generated; each smart object can decide to test another one belonging to its community with a certain probability, while considering the features and services offered by it. The reliability of the tested smart object can be verified thanks to the support of other smart objects belonging to the same community of the tester and providing the same feature/service. These tests compute the trust scores associated with smart objects in their communities, and all of them are stored in their corresponding local blockchain. Figure 2 shows a summary representation of our intra-community probing scheme and the computation of the trust of a trustor tr_i in a trustee te_j.

After a defined time window, the reputation of each smart object in its community is derived, thanks to the aggregation of the results of its probing transactions. A smart contract of the blockchain is responsible for the reputation computation. The results of this task, which consists

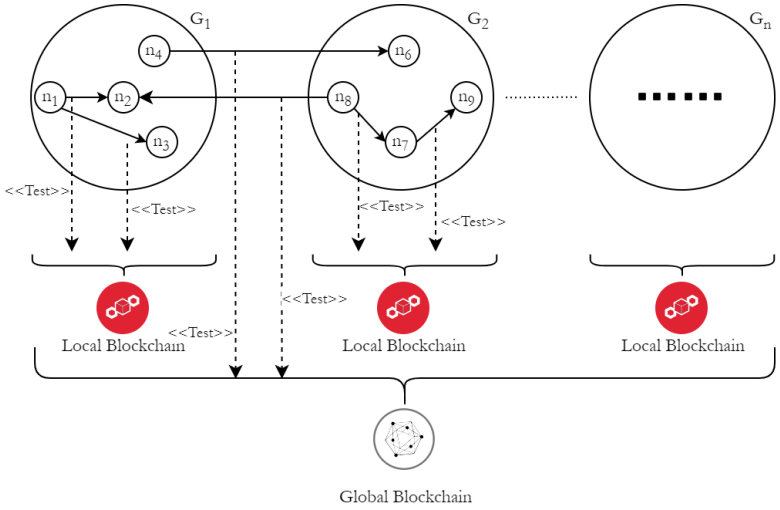

Figure 1: General architecture of our approach

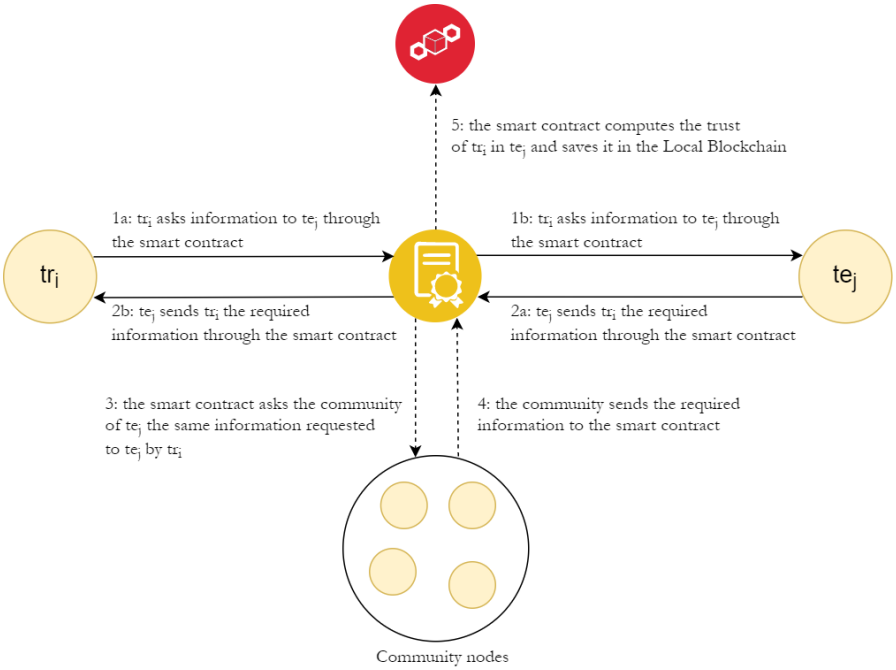

Figure 2: Computation of the trust of a trustor tr_i in a trustee te_j

of the list of community members and the corresponding reputation scores, is published in the global blockchain (see Figure 3). Smart objects having a reputation below a certain threshold are automatically removed from their community.

Our approach also provides protection during the interactions between objects from different communities. Specifically, when two objects from different communities contact each other,

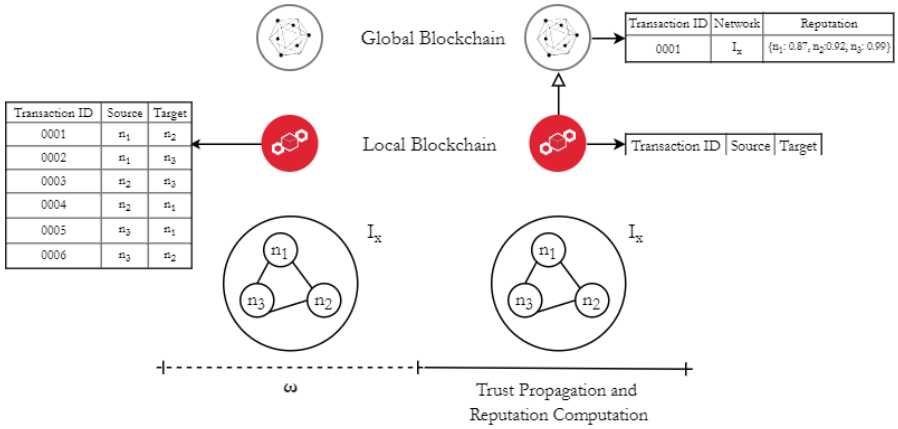

Figure 3: Transaction aggregation and computation of the reputation of the smart objects of a community

one of them may undergo a test with a certain probability. To perform this test, the tester object requests a feature/service among those provided by the tested object. This request is also sent to other objects in the same community of the tested object. The test result is stored in the local blockchain of the community of the tester object. After a certain time window, the results of the tests performed by the smart objects belonging to different communities are aggregated. Following this reasoning, it is possible to obtain a trust value of a community towards other communities with which at least one transaction between the corresponding objects has taken place. These trust values are also saved in the global blockchain.

Thanks to the information stored in the global blockchain, when a smart object o_{i_k} of a community C_k wants to interact with a smart object o_{j_q} of a community C_q, $C_q \neq C_k$, o_{i_k} can compute the reliability of o_{j_q} taking into account the reputation of o_{j_q} within C_q and the trust of C_k in C_q.

3. Security Model

In this section, we present the security model associated with our framework. Preliminarily, we highlight that it is based on the assumption that a sufficient number of nodes are available in such a way as to successfully implement our approach. Therefore, our model does not consider anomalous situations or the startup time, in which the number of nodes available in the framework is less than the minimum required.

In the analysis of security properties, we will consider that our threat model includes the following assumptions: *(i)* At most t smart objects can collude to break the security properties of the protocol; *(ii)* the size of all the pruned support partitions is greater than t; *(iii)* an attacker cannot control a whole group of smart objects; moreover, she cannot own all the smart objects providing a certain service; *(iv)* an attacker has no additional knowledge derived from any direct physical access to smart objects; *(v)* the blockchain technologies adopted to implement both the local and the global tiers are compliant with the standard security requirements already

adopted for common blockchain applications. As for the first assumption, we recall that probing transactions are produced collaboratively by several smart objects in our protocol. Some of them might be corrupted, but, according to [12], we assume the honesty of the majority of them.

In the following, we report the list of the security properties that our framework must assure (due to space limitations, we cannot report here the corresponding security analysis):

- *Resistance to the attacks to local and global blockchains* conceived to find vulnerabilities in them.
- *Resistance to self-promoting attacks*, occurring when a smart object manipulates its own reputation to increase it falsely and promote itself.
- *Resistance to whitewashing or self-serving attacks*, occurring when a malicious smart object, with a compromised reputation, tries to quickly degrade the latter with the goal of being removed from the framework and asking to rejoin it again with a fresh start.
- *Resistance to slandering or bad-mouthing attacks*, occurring when one or more attackers try to manipulate the reputation of other smart objects by reporting false data.
- *Resistance to opportunistic service attacks*, occurring when a malicious smart object can provide good or bad services opportunistically.
- *Resistance to ballot stuffing attacks*, occurring when an attacker tries to boost the reputation of bad objects providing good recommendations for them.
- *Resistance to Denial of Service (DoS) attacks*, occurring when an attacker tries to prevent a reputation system from operating properly by flooding it with an excessive number of transactions.
- *Resistance to orchestrated attacks*, occurring when malicious smart objects orchestrate their actions and leverage several of the previous strategies to perform a coordinated and multi-faced attack, which can change over time.

4. Experimental comparison with other approaches

In this section, we compare our approach with other related strategies proposed in the past literature. The related approaches we selected have many similarities with our own in both the reference scenario and the adopted methodologies; instead, their goal is different.

The first approach we selected concerns an intrusion detection system protecting smart devices in vehicular networks [13]. The authors proposed to group the nodes into "clusters", so that security is achieved through the collaboration of nodes inside the identified protected zones. This approach and ours are very similar in two aspects, even if their goals are different. The former is the definition of a security model operating on smart devices and the IoT, whereas the latter is the usage of groups and clusters of things (corresponding to communities of smart objects in our model). The second approach we considered focuses on the modeling of a scheme for grouping objects in such a way as to preserve their privacy [1]. The modeled scheme guarantees the protection of user's privacy in all the IoT scenarios where the knowledge of the object characteristics may lead to attacks based on the collection of user habits and behaviors.

A way to compare the approaches of [13] and [1] with ours consists of measuring the communication delay introduced by them against the community size. In our approach, we

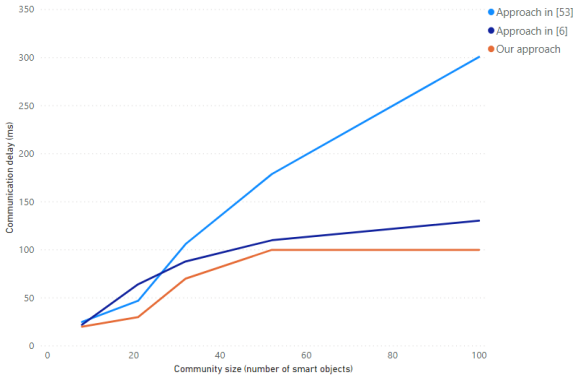

Figure 4: Comparison of the average delay against the community size between our approach and the ones of [13] ad [1]

defined this parameter as the average difference, in terms of delivery time, between a scenario adopting our approach and another not adopting it. Figure 4 shows the results obtained. As we can see, the average delay introduced by our approach ranges from 20 ms to 100 ms. The delay of [13] ranges between 24 ms and 170 ms, whereas the one of [1] ranges from 22 ms to 300 ms. This result tells that our approach is clearly comparable with, and even better performing than, the ones described in [13] and [1]. In conclusion, we can say that our approach returns satisfactory results still keeping low the IoT overhead.

5. Conclusion

In this paper, we proposed an approach adopting a two-tier blockchain framework and a trust-based protection mechanism for increasing the security and autonomy of smart objects in the IoT. The proposed approach and the results obtained are not to be intended as an ending point. By contrast, they represent a starting point for further future activities. For example, we plan to combine our approach with other community-based strategies aiming at ensuring the privacy of smart objects and their owners. The ultimate goal of such a task would be the definition of a single solution handling both privacy and security in the IoT.

References

[1] S. Nicolazzo, A. Nocera, D. Ursino, L. Virgili, A privacy-preserving approach to prevent feature disclosure in an iot scenario, in: Future Generation Computer Systems, volume 105, 2019, pp. 1–8. IEEE.

[2] J. Quevedo, M. Antunes, D. Corujo, D. Gomes, R. Aguiar, On the application of contextual iot service discovery in information centric networks, Computer Communications 89 (2016) 117–127. Elsevier.

[3] F. Buccafurri, G. Lax, S. Nicolazzo, A. Nocera, Accountability-Preserving Anonymous Delivery of Cloud Services, in: Proc. of the International Conference on Trust, Privacy and Security in Digital Business (TRUSTBUS 2015), Springer, 2015, pp. 124–135.

[4] M. Shen, X. Tang, L. Zhu, X. Du, M. Guizani, Privacy-Preserving Support Vector Machine Training over Blockchain-Based Encrypted IoT Data in Smart Cities, IEEE Internet of Things Journal (2019). IEEE.

[5] S. Huh, S. Cho, S. Kim, Managing IoT devices using blockchain platform, in: Proc. of the International Conference on Advanced Communication Technology (ICACT'17), PyeongChang, Korea, 2017, pp. 464–467. IEEE.

[6] M. Khan, K. Salah, IoT security: Review, blockchain solutions, and open challenges, Future Generation Computer Systems 82 (2018) 395–411. Elsevier.

[7] S. Bandyopadhyay, M. Sengupta, S. Maiti, S. Dutta, A survey of middleware for Internet of Things, in: Recent trends in wireless and mobile networks, 2011, pp. 288–296. Springer.

[8] F. Buccafurri, L. Coppolino, S. D'Antonio, A. Garofalo, G. Lax, A. Nocera, L. Romano, Trust-Based Intrusion Tolerant Routing in Wireless Sensor Networks, in: Proc. of the International Conference on Computer Safety, Reliability and Security (SAFECOMP 2014), Springer, Firenze, Italy, 2014, pp. 214–229.

[9] M. Rehman, N. Javaid, M. Awais, M. Imran, N. Naseer, Cloud based secure service providing for IoTs using blockchain, in: Proc. of the IEEE Global Communications Conference (GLOBCOM 2019), Puako, Hawaii, USA, 2019, pp. 1–7.

[10] A. Dwivedi, G. Srivastava, S. Dhar, R. Singh, A decentralized privacy-preserving healthcare blockchain for IoT, Sensors 19 (2019) 326. Multidisciplinary Digital Publishing Institute.

[11] T. Moellers, IOTA-based Business Model Configurations, https://www.alexandria.unisg.ch/257117/ (2018).

[12] P. Fouque, G. Poupard, J. Stern, Sharing decryption in the context of voting or lotteries, in: Proc. of the International Conference on Financial Cryptography (FC'00), Anguilla, Anguilla, 2000, pp. 90–104. Springer.

[13] M. Aloqaily, S. Otoum, I. A. Ridhawi, Y. Jararweh, An intrusion detection system for connected vehicles in smart cities, Ad Hoc Networks 90 (2019) 101842. Elsevier.