



Bringing privacy, security and performance to the internet of things through usage control and blockchains

Nathanaël Denis, Maryline Laurent, Sophie Chabridon

► To cite this version:

Nathanaël Denis, Maryline Laurent, Sophie Chabridon. Bringing privacy, security and performance to the internet of things through usage control and blockchains. Privacy and Identity Management. Between Data Protection and Security, 644, Springer International Publishing, pp.57-72, 2022, IFIP Advances in Information and Communication Technology, 10.1007/978-3-030-99100-5_6 . hal-03754051

HAL Id: hal-03754051

<https://hal.science/hal-03754051>

Submitted on 21 Feb 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Bringing Privacy, Security and Performance to the Internet of Things through Usage Control and Blockchains

Nathanaël Denis, Sophie Chabridon, and Maryline Laurent

SAMOVAR, Télécom SudParis, Institut Polytechnique de Paris, France
`firstname.surname@telecom-sudparis.eu`

Abstract. The Internet of Things (IoT) is bringing new ways to collect and analyse data to develop applications answering or anticipating users' needs. These data may be privacy-sensitive, requiring efficient privacy-preserving mechanisms. The IoT is a distributed system of unprecedented scale, creating challenges for performance and security. Classic blockchains could be a solution by providing decentralisation and strong security guarantees. However they are not efficient and scalable enough for large scale IoT systems, and available tools designed for preserving privacy in blockchains, e.g. coin mixing, have a limited effect due to transaction cost and rate.

This article provides a framework based on several technologies to address the requirements of privacy, security and performance of the Internet of Things. The basis of the framework is the IOTA technology, a derivative of blockchains relying on a direct acyclic graph to create transactions instead of a linear chain. IOTA unlocks distributed ledgers performance by increasing throughput as more users join the network, making the network scalable. However, IOTA is not designed for privacy protection and is complemented in this paper by privacy-preserving mechanisms: merge avoidance and decentralised mixing. Finally, privacy is reinforced by introducing usage control mechanisms for users to monitor the use and the dissemination of their data.

Keywords: IoT · Privacy · Blockchain · IOTA · PET · Usage Control.

1 Introduction

The Internet of Things (IoT) is a ubiquitous network where connected devices exchange data between each other as well as with users [6]. The devices collect data about their environment and usually transfer them to centralised cloud service providers, also known as CSPs. The CSPs process the data in order to provide a real-time and customised service to customers. Due to the amount of devices concerned, their heterogeneity and the personal nature of the data gathered, privacy and security are at risk in IoT systems, thus resulting in the need for new privacy-preserving solutions, well-tailored for the Internet of Things.

Currently, the most common model centralised around CSPs is troublesome for the IoT both for privacy and security reasons. Indeed, cloud service providers

must not be automatically trusted and may snoop on users' data [16]. Besides, they can be vulnerable to internal attacks, from malicious employees as well as accidental disclosures or external attackers [16]. Availability can be a matter of concern too, as physical infrastructure can be damaged, e.g. because of a fire [19]. Furthermore, centralisation hinders performance, specifically by increasing the cost of deployment and maintenance [22], which limits scalability.

Blockchain has been drawing attention as a solution to security issues, because of its properties regarding decentralisation and the removal of intermediate third-parties (cf. Section 2.1). However, conventional blockchains are not suitable for IoT systems, as they are computationally expensive, not scalable enough and introduce memory and bandwidth overhead [6]. Besides, while conventional blockchains address security issues, they provide no more than pseudonymity. Privacy in blockchains is a specific topic, different from security, that must be addressed using dedicated tools.

This article is structured as follows: Section 2 summarises the current state of the art about blockchains, usage control and privacy in the Internet of Things. Section 3 describes the car sharing use case over which both system and threat models are elaborated. Our framework for supporting privacy, security and performance in the IoT, is explained in Section 4 before concluding in Section 6.

2 Related work

Considering the need for decentralisation, security and privacy in the Internet of Things, this section identifies blockchain (2.1) and usage control (2.2) technologies as candidate solutions and discusses their current limitations and state of the art. We eventually discuss privacy of blockchain transactions in Section 2.3.

2.1 Blockchain

A blockchain is a "distributed and immutable ledger made out of unalterable sequence of blocks" [22]. This technology provides several properties of interest for the Internet of Things [5]: 1) decentralisation; 2) ability to audit the data; 3) disintermediation; 4) transparency. Decentralisation and disintermediation are particularly relevant for large scale deployments and for security, as they limit the extent of data leaks and prevent potential misbehaviour from CSPs.

Blockchain topology Blockchains can be of three types: public, private or consortiums [22]. *Public blockchains* do not control access and are called permissionless, while private and consortiums do have a control layer and are called permissioned blockchains. Public blockchains are distributed and tamper-proof ledgers which are not controlled by a single entity and are open to anyone. New entries can be appended to the ledgers as long as the network participants agree. To this end, the participants use a consensus method in order to determine who can add a new block to the chain. Conversely, *private blockchains* restrict access to the public. Access to the network and involvement in the consensus protocol

rely on authorisations, and require a third-party. Therefore, private blockchains are not completely decentralised, which has several consequences: 1) they are not as secured as public blockchains because they can not provide the same level of computational power; 2) being partly centralised, scalability and security are decreased; 3) private blockchains logically raise privacy levels as the data are restricted; 4) apart from large scale deployments, network response time is better and computational requirements are reduced. Private blockchains are consequently appropriate for some IoT use cases, in particular when high scalability is not needed. Finally, *consortium blockchains* are partially private blockchains, shared between several institutions instead of a single one. All these institutions are directly involved in the consensus protocol. The only concrete difference between consortium and private blockchains is the number of governing institutions. As a consequence, they will be considered as private blockchains in this paper.

Consensus methods for the Internet of Things Blockchains implement consensus methods to agree on which data can be appended to the ledger. Consensus methods are paramount in blockchains as they enable decentralisation. Moreover, the blockchain network is as secure as its consensus method is robust. Therefore, modifying the consensus method allows to trade security for performance, and the parameters of the blockchain network are deeply impacted by the selected consensus method. Performance of blockchains can be qualified as follows [22]:

- *throughput*, generally measured in transactions per second (TPS);
- *latency*, also referred to as block time, the time between the creation of two blocks on the blockchain;
- *network overhead*;
- *storage overhead*;
- *scalability*, to be understood as scalability in terms of the number of participants [27]. Scalability in terms of transaction processing capacity is directly linked to throughput.

Conventional blockchains heavily rely on *proof of work* (PoW) mechanisms, which are computationally expensive and not suitable for resource-constrained devices of the Internet of Things. The main alternative to proof of work in mainstream blockchains is the *proof of stake* (PoS), where the node responsible for block creation is chosen at random based on its proportional stake in the network. While this removes the resource-hungry computational race, it still introduces new issues. It is based on a monetary concept, the stake, which excludes many IoT use cases, including sensors, that do not require the use of currencies. Proof of stake gives the power to the most important holders, partially centralising the blockchain network. Finally, *Proof of Elapsed Time* (PoET) is another IoT-friendly consensus method. While miners still have to solve the computation puzzle, the winner is chosen based on a random wait time. The next block is created by the miner whose timer expires first, and miners are not competing.




Fig. 1. Tangle transaction graph compared to traditional blockchains

However, the verification of the right timer execution is done with a *Trusted Environment Execution* provided by Intel. Consequently, this consensus depends upon Intel which goes against the decentralisation property.

To make a blockchain network suitable for large scale IoT deployments, all these properties must be achieved simultaneously. To this end, the current literature is looking for specific consensus method for the Internet of Things. Raghav *et al.* [17] propose a lightweight consensus mechanism for blockchain in the IoT. This consensus method is called *Proof of Elapsed Work And Luck* (PoEWAL). Its performance, energy consumption and latency are compared to those of several consensus methods, including Proof of Work and Proof of Stake. It turns out its performance is overall better than Proof of Stake considering different parameters, without introducing monetary concepts, making it suitable for the IoT. Another line of research focuses on the use of artificial intelligence to integrate IoT with blockchains, especially to improve the consensus method. Salimitari *et al.* [21] propose a framework for consensus in blockchain-based IoT systems with the support of machine learning. Actually, their solution consists in a 2-step consensus protocol, first detecting anomalies with machine learning, then using the *Proof of Byzantine Fault Tolerance* (PBFT) consensus. PBFT consensus allows a distributed system to reach a consensus even though a small amount of nodes demonstrate malicious behaviour.

Direct acyclic graph is the most well-known alternative to blockchains

It is used by the IOTA technology [14] to build the Tangle, IOTA's graph of transactions. To issue a new transaction, a node of the IOTA network has to validate two pending transactions known as tips (cf. Figure ??). In blockchains, blocks can be composed of several transactions. However, in the Tangle, each block is composed of only one transaction. A transaction is pending until confirmed by another transaction.

Then, the node processes a light proof of work to prevent spam. This unique system ensures scalability, as more users means faster tip validations, whereas

common blockchains tend to saturate when the number of users increases. IOTA does not require a computationally expensive proof of work for strong security, but uses a proof of work affordable for IoT devices to protect from spam. Moreover, there are no rewards for the proof of work which implies the transactions are free, thus making micropayments possible, a boon for many IoT use cases. Storing a potentially huge ledger on devices is another issue to consider. For nodes with insufficient storage capacity, local snapshots can be created, removing old transactions and strongly reducing the size of the Tangle. Yet, the IOTA technology has a major flaw, because it is at the moment partly centralised. Indeed, IOTA relies on a *coordinator node* run by the IOTA Foundation, i.e. the foundation who created and has been developing IOTA, whose mission is to directly or indirectly validate transactions [25]. It does not completely centralise the network as all the nodes verify that the coordinator node does not break the consensus rules, yet it can freeze funds, ignore transactions and is a single point of failure, i.e. if the coordinator stops, after an attack or by purpose, transactions are no longer validated. In order to solve the coordinator issue, the IOTA Foundation is planning to release IOTA 2.0 by the end of 2021, after launching the test network in June 2021 [1]. The removal of the Coordinator is likely to be achieved by introducing new components, particularly a new consensus method called *Fast Probabilistic Consensus* (FPC) and a node accountability system to protect against basic attacks [15].

2.2 Usage control

Usage control, as an extension of access control, monitors how the data can be used after initial access. It was first proposed by Sandhu and Park as the UCON model [13]. The UCON model extends traditional access control by introducing attribute mutability as well as new decision factors, namely obligations and conditions. Obligations are requirements to be fulfilled by the subject to be granted access to. Conditions are subject-independent environmental requirements for allowing access. Since attributes are mutable, authorisations and obligations can be done before or during the access. They are referred as pre-authorisations and ongoing-authorisations, or respectively pre-obligations and on-going obligations. Improving user's control over the data is crucial to achieve privacy in IoT systems [4], and UCON provides the technical basis to enable this control.

Modern Usage Control Systems (UCS) integrate Data Flow Control (DFC) to complement UCON To actually control the usage, another concept was introduced to complement UCON: Data Flow Control [7], [12]. Data Flow Control (DFC) aims at controlling the flow of information, and ensuring the data are not disseminated to irrelevant actors. Therefore, DFC trackers are components of modern data usage control systems (UCS), whose purpose is to improve their behaviour, especially when multiple copies of the data are distributed over numerous devices.

The integration of usage control with blockchains is a recent topic of research Khan *et al.* [10] propose to integrate UCON in blockchains relying on the Hyperledger Fabric, a permissioned blockchain. For the authors, the purpose of introducing UCON is to monitor assets continuously to cover all possible access control models. Rizos *et al.* [18] suggest to extend UCON to distributed systems in order to strengthen the IoT security. More precisely, they adapt UCON to the MQTT and CoAP protocols. Finally, Kelbert and Pretschner [9] developed a fully decentralised usage control for distributed systems, including data flow tracking. In several situations, their decentralised policy enforcement outperforms a centralised one.

2.3 Transaction privacy

While blockchain transactions are thought to be anonymous, the reality is more nuanced. Public blockchains do not require identifying information to make a transaction worldwide. Yet, transactions are publicly broadcast. The transaction content, as well as the operation itself disclose information about the individuals involved. Interested third parties automatically collect and analyse this information, for several purposes including law enforcement [11]. By default, public blockchains only provide pseudonymity, and anonymity provided the linkage between the pseudonym and the real identity is not possible. Yet, two behaviours facilitate significantly the re-identification analysis: address reuse and super-clusters with high centrality. Using address clustering, i.e. partitioning the addresses into subsets likely controlled by the same entity, combined with address tagging and graph analysis, it is possible to re-identify more than 69% of wallets stored by Bitcoin lightweight clients [11].

Privacy-preserving techniques have been designed to mitigate the effectiveness of de-anonymisation The most well-known tools for enforcing privacy in transactions are mixing services and merge avoidance, which can theoretically be added on top of any blockchain [23]. Both aim at obfuscating the transactions by adding new fictional ones. In merge avoidance, a single transaction between two users is split into numerous transactions for both users, who create several addresses. Mixing relies on the same principle, but spurious transactions are created by the mixing service users, possibly multiple times before being sent to the actual target (cf. Figure 2). Note that the service mixes the coins of several users to remove the linkage between the sender and the receiver of a transaction. Besides, some cryptocurrencies have been specifically designed to enforce privacy in their transactions, such as Zcash (ZEC) [3] and Monero (XMR) [20], obfuscating the transactions with several Privacy-Enhancing Technologies (PETs) and cryptographic tools.

Privacy in the IOTA technology Apart from using a Direct Acyclic Graph instead of a blockchain, IOTA has several features that change the concerns related to privacy. Its main asset is the free transaction cost, making merge

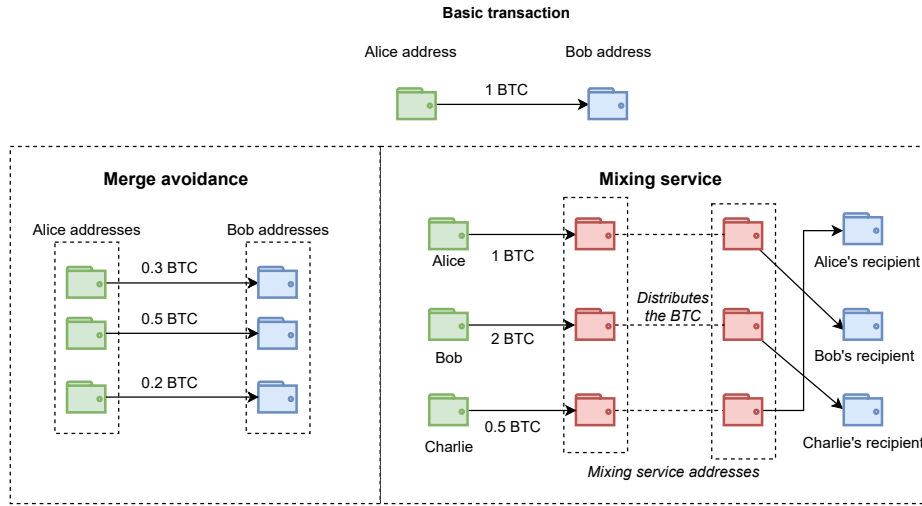


Fig. 2. Obfuscation with merge avoidance and mixing on the Bitcoin blockchain

avoidance particularly relevant as transactions can be virtually split into infinite sub-transactions. Decentralised mixing is then relevant as the network does not rely on financial motivation. To this end, Sarfraz [23] designed a decentralised mixing service for the IOTA network, which requires no mixing fees. Mixing consists in joining coins from different senders before swapping their receivers, in order to remove linkage. Conversely, IOTA has some properties harmful to privacy. Indeed, the removal of the mining process prevents from creating tokens without taint. A token is considered as tainted if it belonged to at least one identifying address on the IOTA ledger. All IOTA tokens were created in the first *genesis* transaction. Only iotas that have never been linked to any identifiable address, i.e. an address belonging to someone who has been re-identified, can be considered as untainted [28].

3 System model

To identify the needs in terms of performance, security and privacy for large scale deployments of IoT systems, a car club illustrating scenario is first proposed in Section 3.1. Section 3.2 then highlights the security and privacy threats based on this scenario.

3.1 Scenario

Car clubs (UK) or car sharing (US) is a model of car rental where people rent cars for a short period of time, often by the hour. They differ from classic rental models in that the owners of the cars are individuals themselves, instead of an

agency. The context is highly dynamic, as many users may enter the car club or leave it on the same day. In order for the users to interact with the system, an application is responsible for registration and asking or granting access to the vehicles.

Mainly for security reasons, the car owners have the right to watch over their cars and know where they are, almost in real-time. The position of the cars as well as their navigation produce data about the car renters which are sent to the car owners.

The agents of the system can therefore be summarised as follows:

- the car owners, who propose their vehicles on the renting market;
- the car renters, who pay for renting the vehicles;
- the car itself, which sends data to the owners such as location, and whose access must be monitored;
- the Access Server (AS), which is responsible for managing the access to the cars;
- the Usage Control System (UCS), which monitors the data generated by other agents;
- the mixing service, responsible for obfuscating the transactions to preserve privacy.

Actually, both the Access Server and the Usage Control System control access, respectively to a physical object - the car - and to the data. They will be referred to as *controllers*. The UCS also prevents the dissemination of the data to irrelevant actors.

3.2 Privacy and security threat model

Considering the system agents, the threat model identifies four attackers:

1. the car owners, who have legitimate access to some sensitive data of the car renters. The processing and dissemination of these data are restricted by the UCS;
2. car owners colluding with each other to gather big sets of data;
3. the mixing service, which can get the addresses of senders and receivers in order to provide its service, if used in a centralised fashion;
4. external attackers, who wish to disable the UCS to help car owners disseminate data to other agents.

The car owners as well as the mixing service are considered honest-but-curious, which means they will fulfil their mission, but will snoop on the data of the users requesting their services. Honest-but-curious attackers will not consider the network interactions but only the transactions. For instance, they will not use the IP addresses to re-identify the users. External attackers are conversely malicious and may try actively to neutralise the UCS to enable car owners to disseminate their data. The main motivation of honest-but-curious attackers is to gather as much data as possible. Depending on the data obtained by the attackers, the *risks to privacy* are the following ones:

- re-identification: the attacker obtains information that enables linking the pseudonym to the real identity of the car renter;
- inference: “This category covers attacks where the attacker has used existing knowledge to aid the attack” [8]. An inference attack occurs when an attacker is able to infer valuable information from trivial information. For example, in our scenario, an attacker could infer working hours by gathering transactions timestamps.

Concurrently, there are *risks to security* because an agent of the system - namely the UCS - is responsible for the data protection. External attackers can be interested in neutralising the UCS, e.g. by disabling or modifying the UCS, to enable car owners to collude.

4 Proposed framework

Regarding the different challenges for large scale deployments of IoT systems, as illustrated by the car sharing scenario (cf. Section 3.1), a set of complementary tools is needed to match privacy, security and performance requirements simultaneously. To this end, the originality of this article is to design a framework with the following aspects (cf. Figure 3):

1. IOTA technology, as the most promising solution matching IoT performance requirements;
2. IOTA Access, an open-source framework used to control access to IoT devices. It is developed by the IOTA Foundation to complement the IOTA technology;
3. a Usage Control System, for car renters to monitor the usage of the data they produce;
4. a decentralised mixing service coupled with merge avoidance, to obfuscate the transactions and improve users’ privacy.

IOTA and its Tangle are introduced along with IOTA Access, the framework developed by the IOTA Foundation to control the access to devices. IOTA Access is meant for any device, ranging from sensors to vehicles. The Usage Control System, which controls the data and how they are disseminated, is embedded into IOTA Access. The mixing service is external to the Tangle and they interact with one another. Merge avoidance can be programmed directly by the user, when sending the transactions to the mixing service.

The IOTA Access framework is composed of three main components: a policy database to store the access control policies, a client so that the car owners can define their policies and can grant access to their cars, and finally a server monitoring the access and interacting with the Tangle. As the Access Server (AS) already controls the access to vehicles, the UCS is embedded into the AS even if the controlled objects differ. Indeed, the AS controls access to a physical device, the car, while the UCS monitors access to the data and prevents dissemination.

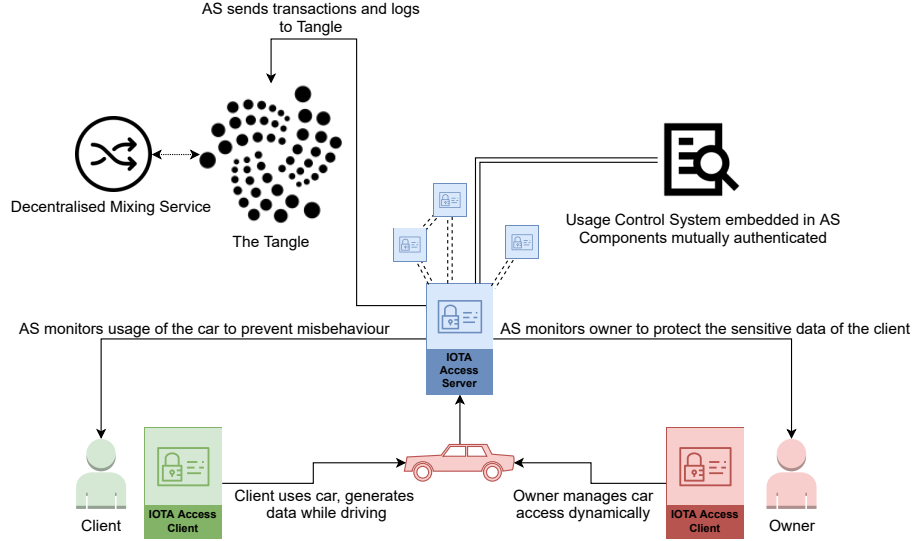


Fig. 3. Framework and relationships between agents

Decentralising the framework First, we emphasise that the IOTA Access server is already decentralised, as it can be deployed by anyone. In our use case, the most suitable solution is to pick one external trustworthy server to connect to, which is realistic as a list of trustworthy IOTA nodes is maintained by the community¹. The same principle could be extended to IOTA Access servers, with a list of the top public ones.

Merge avoidance and mixing are used jointly to increase the effect of obfuscation. Effectiveness of merge avoidance is increased due to free transactions on the IOTA network. For the same reason, mixing is more efficient as the nodes involved in the mixing service do not have to pay for the transactions. Indeed, if IOTA nodes were encouraged to participate for money, decentralised mixing services would become vulnerable to edge insertion attacks [26] where nodes can claim undue rewards. Therefore, our framework uses a decentralised mixer to remove the threat of linkage and re-identification, and without introducing the edge insertion issue.

The Usage Control System must be decentralised as well in order to benefit from the IOTA 2.0 (without the coordinator) and to be resilient to some attacks like denial of service. Kelbert and Pretschner [9] implemented a decentralised usage control system. It is achieved by distributing the components of the UCS responsible for the policy enforcement, and it addresses both the UCON and the data flow tracking aspects. Additionally, decentralising the UCS reduces the communication and performance overhead compared to a centralised policy enforcement. In our framework, the Usage Control Systems are deployed along

¹ <https://trinity.iota.org/nodes>

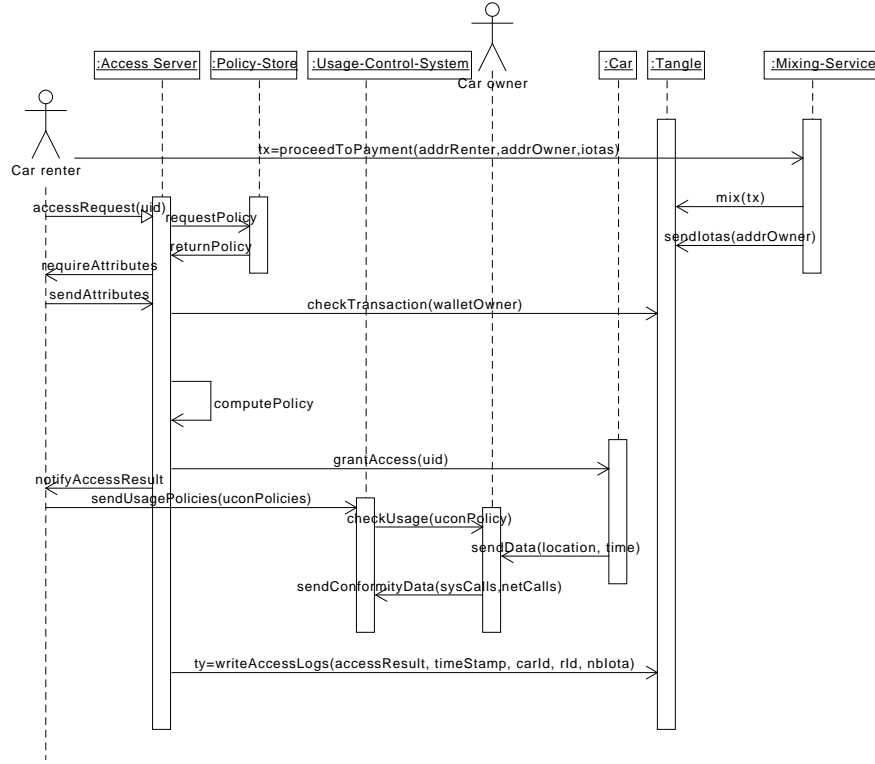


Fig. 4. The sequence diagram related to our framework

the IOTA Access servers which are decentralised as well, enabling the integration of Kelbert and Pretschner’s solution.

Sequence of interactions Figure 4 details the sequence of interactions in our framework through our use case (cf. Section 3.1)

These interactions unfold as follows. First, the car renter requires access to the vehicle. Before the access can be granted, the car renter must send iotas to the car owner’s wallet. To avoid re-identification, the car renter sends the transaction to the mixing service, which removes the linkage between the car renter and owner while creating false intermediate transactions to obfuscate the transactions. The access request is then sent to the Access server, which then requests the car renter attributes to be able to take an access decision. Then, it fetches the car access policies from the policy store, before evaluating them and returning the access result to the car. If the evaluation is positive, the car is unlocked and the car renter can get inside. Afterwards, the car renter sends its data usage policies to the UCS, before using the car. These data usage policies are

composed of the authorisations, the obligations to be fulfilled by the car owner and finally the conditions on the system. For example, a pre-obligation of the car owner is to accept to send the logs for the UCS, e.g. by reading instructions and ticking a box to actually agree. When driving, the car renter generates navigation data, relayed by the cars to the car owners. To comply with the car renter policies, the UCS monitors the data usage of the car owner, who sends in return the mandatory data to enable the monitoring. These mandatory data are composed of system calls and some networks interactions, to ensure the car owner does not process the data in a forbidden way, e.g. store or disseminate the data. When the navigation is over, the Access server writes logs on the Tangle, detailing the result of the access policies evaluation, the hours, the amount of iotas spent and finally the pseudonym of the car to simplify car management when a car owner has several vehicles. The addresses of the car renters i.e. their pseudonyms on the Tangle, are written in the logs as well when they request access to a car. However, to prevent attackers from analysing a user behaviour with the logs, it is paramount to avoid address reuse at this point. This can be done by generating a new address for each payment directly in the IOTA Access client. Note that creating a new address for each transaction is an increasingly default behaviour, implemented on Trinity, the IOTA official ledger.

5 Privacy and security analysis

This section analyses the privacy and security risks in the system. It considers two different aspects: the privacy risk of inference from the collected car renters data, and the security risks to the Usage Control System.

5.1 Inference and re-identification attacks

First, Table 1 describes each combination of attackers, the data type they have access to, the place those data are stored in the system, the risk to privacy associated with each profile of attacker and finally an example of a privacy leakage associated to this risk.

Table 1: Inference and re-identification attacks according to the attackers' profile

Attacker type	Data type	Data storage	Risk	Example
Honest-but-curious	Transaction	Tangle	Inference	Purpose of payment
Car owner (alone)	Location	Owner's device	Inference	Renter's job
Car owners (colluding)	Location	Owners' devices	Further inferences	Renter's job
Ext. attacker & car owners	Location	Owners' devices	Simplified inferences	Data sets on renters

Any user has access to the transactions on the Tangle, which are public and contain privacy-sensitive timestamps, users' addresses and values, i.e. how many iotas are sent to a car owner. Based on these elements, any honest-but-curious attacker, namely the mixing service or a car owner, can attempt to:

- use the blockchain transactions to re-identify the users. In case the blockchain protocol is not used properly, an attacker may re-identify the car renters, but also the car owners they interact with [11]. To mitigate this threat, the mixing service is used to obfuscate the address of the car owner and remove the linkage between car owners and car renters;
- use the blockchain transactions for inference attack, e.g. use the amount of tokens in the transactions to infer for which purpose the payment is done. Note that the merge avoidance mechanism integrated in our framework can help reduce the risk of inference by splitting the transactions into several smaller ones, thus making harder to guess the purpose of the transactions.

Additionally, car owners having access to some private data can infer sensitive data. For instance, the location of the car renters is very privacy sensitive as it might reveal the car renters driving habits, their jobs, their religion, their hobbies, partially their social graph. Besides, when colluding, car owners can 1) merge their data about a given user to increase the quality of the inference; 2) increase the size of their user base thus improving its value. If a colluding external attacker successfully neutralises the UCS entities, as reported in Section 5.2, car owners can freely share user data through the system and can disseminate their data to a shared database for processing.

As the mixing service is decentralised and according to Sarfraz [23], a node involved in the mixing process is not able to make links between any input or output addresses. Consequently, the mixing service gets into the general category of honest-but-curious attackers (cf. Table 1) as mixing nodes do not benefit - as attackers - from their involvement in the mixing service.

5.2 UCS neutralisation

The UCS is a paramount actor for controlling usage control and data flow transfers between the agents. It is consequently an attractive target, vulnerable to specific attacks which can be partially mitigated [9]:

- *denial of service*: the external attacker can disable temporarily the UCS, threatening the availability of the system and disabling the usage control mechanisms. Modern denial of service attacks are hard to mitigate, but the decentralisation of the UCS, as designed in our framework, alleviates the risk, as well as mutual authentication of all the infrastructure components, e.g. using certificates;
- *privilege escalation*: the external attackers can leverage vulnerabilities as illustrated in Babil *et al.* article [2] to bypass the UCS restrictions. These attacks are very diverse and implementation dependent, therefore out of the scope of this paper.

6 Conclusion

In this paper, we devise a framework to guarantee simultaneously the requirements in terms of performance and security of large scale IoT systems, as well

as the privacy of the users. To do so, we rely on several technologies. IOTA guarantees high performance and scalability with a balanced security fitting IoT needs. Usage control empowers the users with a tool to monitor how their data are used, while coin mixing and merge avoidance introduce obfuscation on the network to protect from re-identification and inference. Using a car sharing scenario, we highlight the threats faced by the agents in using the system, and we analyse the security of our solution.

As soon as the version of IOTA 2.0 without the coordinator is available and the framework feasible, other perspectives will be opened for privacy with cross-chain transactions [24]. As IOTA mixing brings two properties of interest - its free transactions and the support for decentralisation - there is a significant interest to integrate the IOTA mixing concept into other blockchains, to avoid payments of centralised mixing fees on their own networks.

Finally, our car rental use case focuses on access to physical objects. The concept could be taken further and applied to data-centric use cases, closer to the UCON philosophy of controlling access to data and not only objects. Besides, our framework can be applied to any IoT use cases involving large scale deployments, decentralisation and a demand for high processing capacities, all requirements taken together or separately. For instance, a widespread network of vending machines could benefit from this framework, especially for its zero fee transactions.

Acknowledgments

This paper is supported by the Future & Ruptures program of Fondation Mines-Télécom and the Institut Mines-Télécom VP-IP Chair on Values and Policies of Personal Information.

References

1. The Era of IOTA’s Decentralization starts here (2021), <https://blog.iota.org/iotav2devnet/>
2. Babil, G.S., Mehani, O., Boreli, R., Kaafar, M.: On the Effectiveness of Dynamic Taint Analysis for Protecting against Private Information Leaks on Android-based Devices. In: 2013 Int. Conference on Security and Cryptography (SECRYPT). pp. 1–8 (2013)
3. Bowe, H.S., Hornby, T., Wilcox, N.: Zcash Protocol Specification (2016)
4. Cha, S., Hsu, T., Xiang, Y., Yeh, K.: Privacy Enhancing Technologies in the Internet of Things: Perspectives and Challenges. *IEEE Internet of Things Journal* **6**(2), 2159–2187 (2019)
5. Christidis, K., Devetsikiotis, M.: Blockchains and Smart Contracts for the Internet of Things. *IEEE Access* **4**, 2292–2303 (2016). <https://doi.org/10.1109/ACCESS.2016.2566339>
6. Dorri, A.: A Scalable Lightweight Blockchain-based Framework for IoT Security and Anonymity (2020)

7. Harvan, M., Pretschner, A.: State-Based Usage Control Enforcement with Data Flow Tracking using System Call Interposition. In: 2009 Third Int. Conference on Network and System Security. pp. 373–380 (2009). <https://doi.org/10.1109/NSS.2009.51>
8. Henriksen-Bulmer, J., Jeary, S.: Re-identification Attacks—A Systematic Literature Review. *Int. Journal of Information Management* **36**(6, Part B), 1184 – 1192 (2016). <https://doi.org/https://doi.org/10.1016/j.ijinfomgt.2016.08.002>, <http://www.sciencedirect.com/science/article/pii/S0268401215301262>
9. Kelbert, F., Pretschner, A.: Data Usage Control for Distributed Systems. *ACM Trans. Priv. Secur.* **21**(3), 12:1–12:32 (Apr 2018)
10. Khan, M., Syed, T., Fariz, M., Moreira, F., Branco, F., Martins, J., Gonçalves, R.: Blocku: Extended usage control in and for blockchain. *Expert Systems* **37** (01 2020). <https://doi.org/10.1111/exsy.12507>
11. Martin, H., Christoph, F.: The Unreasonable Effectiveness of Address Clustering. 2016 Intl IEEE Conferences (UIC/ATC/ScalCom/CBDCoM/IoP/SmartWorld) (2016)
12. Myers, A.C., Liskov, B.: A Decentralized Model for Information Flow Control. In: *Proc. of the Sixteenth ACM Symp. on Operating Systems Principles*. pp. 129–142 (1997)
13. Park, J., Sandhu, R.: The UCON ABC Usage Control Model. *ACM Trans. Inf. Syst. Secur.* **7**(1), 128–174 (Feb 2004)
14. Popov, S.: The Tangle. Tech. rep. (2017), https://iotatoken.com/IOTA_Whitepaper.pdf
15. Popov, S.: The Coordicide. Tech. rep. (2020), https://files.iota.org/papers/Coordicide_WP.pdf
16. Qin, X., Huang, Y., Yang, Z., Li, X.: A Blockchain-based Access Control Scheme with Multiple Attribute Authorities for Secure Cloud Data Sharing. *Journal of Systems Architecture* p. 101854 (2020). <https://doi.org/https://doi.org/10.1016/j.sysarc.2020.101854>, <http://www.sciencedirect.com/science/article/pii/S1383762120301405>
17. Raghav, Andola, N., Venkatesan, S., Verma, S.: PoEWAL: A lightweight consensus mechanism for blockchain in IoT. *Pervasive and Mobile Computing* **69**, 101291 (2020). <https://doi.org/https://doi.org/10.1016/j.pmcj.2020.101291>, <http://www.sciencedirect.com/science/article/pii/S1574119220301279>
18. Rizos, A., Bastos, D., Saracino, A., Martinelli, F.: Distributed UCON in CoAP and MQTT Protocols. In: *Computer Security - ESORICS 2019 Int. Workshops, Cyber-ICPS, SECPRE, SPOSE, and ADIoT, Luxembourg City, Luxembourg, September 26-27, 2019 Revised Selected Papers. Lecture Notes in Computer Science*, vol. 11980, pp. 35–52. Springer (2019). https://doi.org/10.1007/978-3-030-42048-2_3, https://doi.org/10.1007/978-3-030-42048-2_3
19. Rosemain, M.: Millions of websites offline after fire at French cloud services firm (2021), <https://www.reuters.com/article/us-france-ovh-fire-idUSKBN2B20NU>
20. van Saberhagen, N.: Cryptonote Monero Whitepaper (2013), <https://github.com/monero-project/research-lab/blob/master/whitepaper/whitepaper.pdf>
21. Salimitari, M., Joneidi, M., Chatterjee, M.: AI-Enabled Blockchain: An Outlier-Aware Consensus Protocol for Blockchain-Based IoT Networks. In: 2019 IEEE Global Communications Conference (GLOBECOM). pp. 1–6 (2019). <https://doi.org/10.1109/GLOBECOM38437.2019.9013824>

22. Salimitari, M., Chatterjee, M., Fallah, Y.P.: A Survey on Consensus Methods in Blockchain for Resource-constrained IoT Networks. *Internet of Things* **11**, 100212 (2020). <https://doi.org/https://doi.org/10.1016/j.iot.2020.100212>, <http://www.sciencedirect.com/science/article/pii/S2542660520300482>
23. Sarfraz, U., Alam, M., Zeadally, S., Khan, A.: Privacy Aware IOTA Ledger: Decentralized Mixing and Unlinkable IOTA Transactions. *Computer Networks* **148**, 361–372 (2019). <https://doi.org/https://doi.org/10.1016/j.comnet.2018.11.019>, <https://www.sciencedirect.com/science/article/pii/S1389128618306972>
24. Shadab, N., Houshmand, F., Lesani, M.: Cross-chain Transactions. In: 2020 IEEE Int. Conference on Blockchain and Cryptocurrency (ICBC). pp. 1–9 (2020). <https://doi.org/10.1109/ICBC48266.2020.9169477>
25. Silvano, W.F., Marcelino, R.: IOTA Tangle: A Cryptocurrency to Communicate Internet-of-Things Data. *Future Generation Computer Systems* **112**, 307–319 (2020). <https://doi.org/https://doi.org/10.1016/j.future.2020.05.047>, <https://www.sciencedirect.com/science/article/pii/S0167739X19329048>
26. Simoes, J., Ferreira, E., Menasché, D., Campos, C.: Blockchain Privacy Through Merge Avoidance and Mixing Services: a Hardness and an Impossibility Result (01 2021)
27. Steen, M., Chien, A., Eugster, P.: The Difficulty in Scaling Blockchains: A Simple Explanation (03 2021)
28. Tennant, L.: Improving the Anonymity of the IOTA Cryptocurrency. <https://laurencetennant.com/papers/anonymity-iota.pdf>