Blockchain-assisted Information Distribution for the Internet of Things

George C. Polyzos and Nikos Fotiou

Mobile Multimedia Laboratory, Department of Informatics
School of Information Sciences and Technology
Athens University of Economics and Business
Evelpidon 47A, 11362 Athens, Greece,

{polyzos,fotiou}@aueb.gr

Abstract

The Internet of Things (IoT) is envisioned to include billions of pervasive and mission-critical sensors and actuators connected to the (public) Internet. This network of smart devices is expected to generate and have access to vast amounts of information, creating unique opportunities for novel applications but, at the same time raising significant privacy and security concerns that impede its further adoption and development. In this paper, we explore the potential of a blockchain-assisted information distribution system for the IoT. We identify key security requirements of such a system and we discuss how they can be satisfied using blockchains and smart contracts. Furthermore, we present a preliminary design of the system and we identify enabling technologies.

Keywords: access control, authentication, identity management, provenance verification

1. Introduction

In 2010, the number of devices connected to the Internet was bigger than the population of the earth [1]. We live in an era where the manufacturing cost of devices with processing and connectivity capabilities is significantly low, hence, we expect that the number of devices connected to the Internet will continue to increase steadily. Furthermore, the notion of "connected device" is not used anymore solely for referring to traditional computing devices—such as personal computers, or smart phones—but it is also used to describe everyday devices—including refrigerators, scales, TVs—which, thanks to the recent technological advances, can connect to the Internet and provide novel services and enhanced end-user experience. This new trend of connected devices supported by the continuously decreasing manufacturing cost of sensors and actuators, fuels the vision of the Internet of Things (IoT). With the IoT, various devices will autonomously exchange (meaningful) information, targeting to improve our daily life, making at the same time the boundaries between the cyber and the physical worlds even more blurred.

The IoT creates new challenges which cannot be overcome simply by using technologies designed for the "traditional" Internet [2]. Nevertheless, overcoming these challenges is a decisive factor that may determine whether the IoT will eventually prevail or not and to what degree.

The security and privacy challenge. Even if we consider that the security solutions currently used in the Internet are successful and viable (an opinion that is questionable and highly debatable), applying them directly to the IoT will not yield the desirable results for the following reasons:

- Things often do not have the necessary computational power to perform complex cryptographic operations.
- In many IoT application scenarios, Things are (physically) exposed to malicious users.
- It is not always feasible to (remotely) connect to a Thing. For example, a Thing may be mobile and unreachable at that time, or it may be in "sleep mode" for conserving energy.

This challenge becomes even more important if we consider that in many cases Things can collect sensitive and personal information, and may control critical aspects of our daily life (for example power or energy distribution, home security, road safety, etc.). All these highlight the need for new, robust, and resilient security solutions that will not depend on the capabilities and properties of the Things. We also believe that open solutions are an advantage, in the end, in this domain.

The sustainability challenge. While users tend to change, or upgrade their "traditional" computing devices, this is not the case with devices such as their refrigerator, their oven, or even their car (in some parts of the world). However, this creates concerns about how such devices will withstand in a connected world. For example, will it be possible to upgrade their operating system throughout their life time, or will we end-up with a fragmented network that will include old, insecure devices?

Moreover, many IoT scenarios concern cases where a Thing is part of a bigger infrastructure, or of an extremely isolated system, making Thing replacement difficult, costly, or even impossible. For example, temperature or earthquake sensors and fire alarms can be installed during the construction of a building or a bridge, pollution detection sensors can be installed on the bottom of the sea, or bio-signal detection devices can be put inside the body of a patient or on the back of a wild animal.

The trust model challenge. Probably, the biggest breakthrough of the IoT is the interaction between the cyber and the physical worlds. Indeed, the IoT is envisioned to include devices of daily use, as well as devices that greatly affect our life. It is evident that we need a (new) trust model that will enable the successful and effective interaction of all these devices with little human intervention, or even none at all. This model should include actions taking place in the physical world (e.g., actuations) with real and often significant impact, it should enable transactions, and it should facilitate novel compensation and accountability mechanisms.

In the following Section 2 we argue that these challenges can be overcome with the help of blockchains and smart contracts. Then, in Section 3, we present the design of a Blockchain-assisted information distribution system for the IoT. We conclude this paper with a discussion and a roadmap for further related research.

2. Blockchains and smart contracts

A blockchain is a *distributed ledger* of transactions maintained by a network of *untrusted* nodes. Each block of the blockchain contains a list of transactions organized in a Merkle tree; new blocks are added to the blockchain by the *miners*. Blockchains are often referred to as a *democratic* way of maintaining transactions as they rely on consensus for confirming transactions and require no central authority. We distinguish two types of blockchains: open and closed.

Open blockchains. Open blockchains are blockchains where anybody can become a miner. In these blockchains, the addition of a new block involves the computation of a solution to a computationally intensive puzzle. The miner that successfully solves the puzzle floods the block in the network: if this block becomes accepted by > 50% of the miners, then it is added in the blockchain. Miners have incentives (usually monetary) to calculate a valid block. Two well known implementations of this type of blockchains are Bitcoin [3] and Ethereum [4].

Closed blockchains. In this type of blockchains the number and the identity of each miner is predefined and it cannot be changed. The protocols used by these chains to achieve consensus are simpler and are based on solutions to the Byzantine Generals' Problem [5]. A well known implementation of this type of blockchain is IBM's Hyperledger [6].

All blockchain implementations are based on a virtual *coin*. The basic types of records maintained by the distributed ledger are transactions related to the creation of new coins (this can be done only by the miners), as well as the transfer of a coin form one user to another. In addition to these simple transactions, many blockchains allow distributed ledgers to store *smart contracts.* A smart contract is an autonomous application with pre-defined inputs and outputs that can be executed by a miner in a deterministic way. Any user can invoke a smart contract, the outcome of which is recorded as a transaction in the distributed ledger.

We now revisit the challenges discussed in the previous section and we discuss how blockchains and smart contracts can help us to overcome them.

Blockchains enable novel security mechanisms. By not relying on a centralized trusted entity, blockchain technology offers significant security advantages: it does not suffer from single points of failure, it prevents censorship, and it contributes to the scalability of the overlay architecture. Blockchains provide user authentication through public keys and digital signatures. Moreover, blockchains allow two entities that do not have established any trust relationship to securely communicate with each other. Finally, since all transactions are public and the distributed ledger can only be appended, blockchains contribute to a system's transparency and facilitate the realization of accountability mechanisms.

Blockchains contribute to the sustainability of a system. Blockchain implementations, such as Bitcoin, handle thousands of transactions per day (in the case of Bitcoin this is translated into millions of dollars) and have been proven to be resistant against numerous cyber attacks. Moreover, many research teams around the world tend to agree that the underlay technologies of these blockchains are secure. Therefore, it is expected that an architecture where (i) (meta-)information is stored in a distributed ledger, (ii) all critical operations are implemented using simple transactions or smart contracts, and (iii) security mechanisms—including identification, access control, and secure channel establishment—are built using smart contracts, can lead to sustainable systems. This mainly happens because most critical information storage and exchange is *delegated* to the blockchain, while endpoint devices can be "dumb" and untrusted, with very little maintenance requirements.

Blockchains enable new trust models. Blockchains are built around transactions: blockchains allow two or more entities to perform transactions using a digital asset (the coin), the mapping of which to the physical world depends on each specific application (for example, a coin can be translated into real money, into a domain name, or even to an actuation such as the transfer of energy from one device to another. Moreover, smart contracts allow two or more entities to establish a trust relationship without relying on a commonly trusted entity: the blockchain can reliably and deterministically enforce this relationship (i.e., the smart contract) when needed.

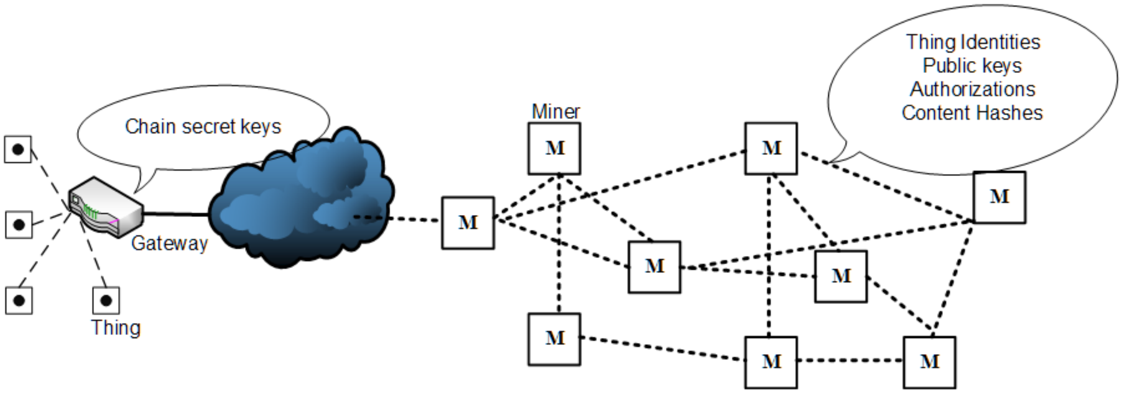

Figure 1: System Design

3. Blockchain-assisted information distribution

Blockchain transactions require public-key encryption operations (such as digital signatures). However, not all Things can support this computationally intensive task. For this reason, our design adopts a gateway-oriented approach, where all blockchain-related operations are offloaded to a gateway, which in return provides an appropriate API for the Things to invoke. This approach is compatible with the Ethereum client side architecture. We assume that all Things that act as information providers and optionally the Things that act as information consumers, are identified by a globally unique identifier. This identifier, as well as its mapping to the Thing's network address have been announced in the blockchain. This provides us with a secure way to identify and locate Things. Moreover, Thing identifiers are associated with one or more blockchain-specific public keys. Finally, it is assumed that there exist a service that allows end-users to learn the identities of the Things that provide the desired information or service. Figure 1 gives an overview of our design.

3.1. Identification and trust management

The predominant trust mechanism in the Internet is the public-key infrastructure (PKI). PKI maps an identity (i.e., a domain name) to a cryptographic primitive (a public key) using a security certificate. Relying on PKI for the IoT is not optimal, since PKI is fragile. The validity of a PKI certificate is attested through a chain of trust composed of certificate authorities (CAs). CAs are entities that vouch for the validity of a security certificate by digitally signing it. A number of CA certificates are pre-configured in user equipment and these CA certificates are de-facto considered trusted. Any CA certificate that has been signed by a trusted CA is also considered trusted creating this way the chain of trust. A recent study [8] found that such a chain is currently composed of 1832 certificates, belonging to 683 organizations. Each of these CA certificates can be used to verify a security certificate for any identity, therefore this chain of trust is as weak as its weakest link (i.e., CA). A worrying observation is that a malfunctioning CA may affect an entity with which it has no direct relationship whatsoever. Moreover, such behavior cannot be easily detected by end-users or end-devices since in most cases software will happily accept any valid certificate.

Blockchains can greatly improve this situation. Using blockchain implementations, such as the Namecoin [7], identity owners can "announce" their identities, as well as related cryptographic primitives. We presented such an approach in [9], where identity owners announce in the blockchain all information required to implement identity based encryption.

3.2. Provenance verification and information tracking

Every Thing that generates an information item may announce its hash to the blockchain, which can act as a distributed and secure timestamping service. This announcement can be used to resolve information ownership conflicts, or even to detect counterfeit products (e.g., by embedding this hash to a real-world object through a printable QR-code). In addition to this announcement, a Thing may create a smart contract that will provide "authorization" to access information items. This smart contract will accept as input a user identity and an amount of "coins" and will output a payment receipt. Furthermore, all information distribution operations can be recorded as transactions to the blockchain. With this, information can be "tracked", i.e., it will possible to know at any given time the identities of the users that hold an information item.

3.3. Authentication and Access control

Access control and endpoint authentication in the IoT is a challenging problem. In [10] we designed and implemented a lightweight solution that solves this problem by allowing the delegation of security operations to a third party, referred to as the Access Control Provider (ACP). The main idea of this solution is that IoT service providers store access control policies in ACPs and in return ACPs generate secret keys which are stored in Things. These keys are generated, during a setup phase, using a secure hash with input the Thing identifier. Additionally, Things are configured with *pointers* (e.g., a URL that points to an ACP) to the access control policies that protect sensitive information. Every time a client requests access to a protected information item the Thing uses a secure hash function to generate a session key. The secret key used by that function is the key generated by the ACP and the inputs of the hash functions are: (a) the pointer to the policy that protects the item and (b) a random nonce. The Thing transmits the nonce and the pointer to the client, which in return requests authorization from the appropriate

ACP (over a secure channel). The ACP has all the necessary information required to calculate the session key: *if the client is authorized*, the ACP calculates the session key and transmits it back to the client. Providing that: (i) the Thing has not lied about its identity and (ii) the messages exchanged between the client and the Thing have not been modified, the Thing and the client end up sharing a secret key. This key can be used for securing subsequent communications (e.g., by using DTLS).

The blockchain technology can further improve this solution, resulting in more secure and sustainable systems. Firstly, the mapping between information item identifiers and pointers to policies can be stored in a distributed ledger. With this, (i) an end-user can be sure that this is a valid mapping, (ii) for any update related to this mapping, no modification has to be transmitted to the Things. Furthermore, our originally proposed solution depends on end-users relaying communications. With blockchain technologies such as Catenis [11], messages can be securely exchanged using a distributed ledger. This has the following advantages: (i) the protocol becomes robust against faulty end-user protocol implementations (note that major single sign-on systems have been breached in the past due to poor implementations—see for example [12]), and (ii) storing ACP decisions in a public ledger is a countermeasure against malicious Things that do not respect ACP verdicts.

3.4 Accountability

One of the most significant properties of blockchains and public ledgers is transparency: all transactions are recorded in the ledger and this information cannot be removed or modified. This property facilitates the development of efficient accountability mechanisms: by making sure that all information distribution transactions are recorded in the blockchain, malicious activities—such as transmission of malicious information, distribution of DRM protected content—can be traced back to the users that performed them. Moreover, all major blockchain implementations make sure that only valid transactions are recorded in the ledger. This property provides non-repudiation to our system, i.e., it is not possible for users to claim that they did not approve a transaction.

4. Conclusion

Many researchers and companies around the world advocate that blockchain technology will contribute to the security and, eventually, to the deployment of the IoT. Nevertheless, little work has been done on how this vision can be realized: this paper is a step in this direction. Firstly, we identified key security and trust related challenges and we discussed how blockchains can be used to overcome them. Secondly, we presented the design of a blockchain-assisted information distribution system for the IoT and we analyzed how key security mechanisms can be built by leveraging blockchain technology. We argue that the enabling technologies of our system are (almost) available, nevertheless there are still open issues. Our design relies on a gateway in which Things delegate blockchain operations: securing this gateway requires further research. Future work in this area should therefore be concentrated on how our system (and other similar systems for the IoT that rely on blockchains) can be secured in the presence of untrusted—or even malicious gateways. Furthermore, many argue that public distributed ledgers add many privacy threats. To this end, private preserving blockchain technologies—such as z-cash [13]—should be studied. Also, efficiency and scalability are very significant problems at present. We believe that indirection and careful choice of what needs to be on the blockchain(s) can lead to significant advances. Finally, we can mention that systems relying on multiple different blockchains can be made to interoperate (see e.g. the Interledger effort [14]), allowing independent evolution of systems, exploitation of different properties and characteristics, and truly open systems and diverse business models.

References

[1] D. Evans, “The Internet of Things – How the Next Evolution of the Internet is Changing Everything,” Cisco, April 2011.

[2] IBM Institute for Business Value, “Device democracy: Saving the future of the Internet of Things,” 2014.

[3] Bitcoin home page, available at: https://www.bitcoin.com/ [accessed 18 May 2017].

[4] Ethereum home page, available at: https://www.ethereum.org/ [accessed 18 May 2017].

[5] L. Lamport, R. Shostak, and M. Pease, “The Byzantine generals problem,” *ACM Transactions on Programming Languages and Systems*, Vol. 4, No. 3, pp. 382-401, 1982.

[6] IBM blockhain home page, available at: https://www.ibm.com/blockchain/ [accessed 18 May 2017].

[7] Namecoin home page, available: https://namecoin.org/ [accessed 18 May 2016].

[8] Z. Durumeric, J. Kasten, M. Bailey, and J. A. Halderman, “Analysis of the https certificate Ecosystem,” Proc. ACM IMC, Barcelona, Spain, pp. 291-304, October 2013.

[9] N. Fotiou and G.C. Polyzos, “Decentralized name-based security for content distribution using blockchains,” Proc. IEEE INFOCOM Workshops, San Francisco, CA, pp. 415-420, April 2016.

[10] N. Fotiou, T. Kotsonis, G. F. Marias, G.C. Polyzos, “Access Control for the Internet of Things,” Proc. International Workshop on Secure Internet of Things (SIoT), Heraklion, Greece, pp. 29-38, September 2016.

[11] Catenis home page, available at: http://blockchainofthings.com/ [accessed 18 May 2017].

[12] R. Wang, S. Chen, and X. Wang. “Signing me onto your accounts through Facebook and Google: a traffic-guided security study of commercially deployed single-sign-on web services,” Proc. IEEE Symposium on Security and Privacy, San Francisco, CA, pp. 365-379, May 2012.

[13] Z-cash home page, available at: https://z.cash/ [accessed 18 May 2017].

[14] S. Thomas, E. Schwartz, A. Hope-Bailie, “The Interledger Protocol,” Internet-Draft, July 2016.