

# Poster: Supply Chain Object Discovery with Semantic-enhanced Blockchain

Michele Ruta, Floriano Scioscia, Saverio Ieva, Giovanna Capurso, Eugenio Di Sciascio  
Department of Electrical and Information Engineering, Polytechnic University of Bari

Bari, Italy

{michele.ruta, floriano.scioscia, saverio.leva, giovanna.capurso, eugenio.disciascio}@poliba.it

## ABSTRACT

Supply chains can be seen as cyber-physical networks grounded on object identification and tracking. Conventional trust models featuring centralized information management architectures and simplistic things classification lend two of the most relevant limitations to current solutions. *Blockchain* introduces novel and a valuable trust approaches while semantic technologies better permit a things description. This paper introduces a semantic-enhanced blockchain platform allowing a flexible object discovery. It is based on validation by consensus of smart contracts and adopt a semantic matchmaking between queries and object annotations expressed w.r.t. ontology models. Early experiments assess the good behaviour of the proposed framework.

## CCS CONCEPTS

• **Networks** → *Cyber-physical networks*; • **Computer systems organization** → *Peer-to-peer architectures*; • **Information systems** → *Web Ontology Language (OWL)*;

## KEYWORDS

supply chains, blockchain, semantic matchmaking, object discovery

### ACM Reference format:

Michele Ruta, Floriano Scioscia, Saverio Ieva, Giovanna Capurso, Eugenio Di Sciascio. 2017. Poster: Supply Chain Object Discovery with Semantic-enhanced Blockchain. In *Proceedings of 15th ACM Conference on Embedded Networked Sensor Systems, Delft, The Netherlands, November 2017 (SenSys 2017)*, 2 pages.

[https://doi.org/10.475/123\\_4](https://doi.org/10.475/123_4)

## 1 INTRODUCTION AND MOTIVATION

Supply chains are complex socio-technical systems for planning, executing and monitoring the transfer of products from suppliers to final customers in an efficient, dependable and traceable way. Information sharing across supply chains is based on linking unique IDs of objects –tagged by means of Radio Frequency Identification (RFID) transponders or barcodes– with records in back-end databases; Electronic Product Code Information Services (EPCIS) [6] is the most relevant industry standard. Unfortunately, EPCIS can manage only elementary data-oriented object attributes.

---

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

*SenSys 2017, November 2017, Delft, The Netherlands*

© 2017 Copyright held by the owner/author(s).

ACM ISBN 123-4567-24-567/08/06...\$15.00

[https://doi.org/10.475/123\\_4](https://doi.org/10.475/123_4)

By exploiting knowledge representation technologies borrowed from the Semantic Web, richer product descriptions, endowed with machine-understandable semantics, can be associated to tags [2]. This enables more advanced and meaningful business intelligence analysis, by helping the various actors in a supply chain track and retrieve specific things or groups of things, based on their properties and not only on detection events.

EPCIS and similar standards adopt Public-Key Infrastructures (PKI) for data security and integrity. As trust in a supply chain cannot be granted beyond “friends of a friend” [4], a central authority is designated which is trusted by all partners and owns the most critical elements of the information sharing network. This makes the central authority both a bottleneck and a single point of failure for the whole supply chain, in case of overload, technical errors or malicious behaviors. A radically different approach is becoming popular in latest years, based on *blockchain* technology [1]. In a nutshell, a blockchain is a distributed database, which records transactions occurred in a given time span in blocks, chained by means of cryptographic hashes. The reliability of such a structure comes from the fact that every transaction is approved by *consensus* of the majority of nodes acting in the system. This prevents any single node or small group of nodes from tampering with data and ensures all and only valid transactions are recorded. Blockchain systems do not require mutual authentication and trust, as the protocol itself guarantees data integrity and non-repudiation. The completely distributed model also improves scalability and eliminates single points of failure, thus making blockchain an effective *trustless* collaboration platform and a natural fit for asset tracking and supply chain applications [8]. The simplest blockchain approaches rely on transactional ledgers for asset transfer. More advanced solutions enable the execution (and validation by consensus) of *Smart Contracts* (SCs), *i.e.*, software stubs which automatically process the terms of an agreement. Blockchains based on SCs enable more flexible systems, supporting general-purpose application logic [1].

This paper investigates the adoption of Semantic Web technologies for item description and discovery in blockchain infrastructures, in order to reap the combined benefits for next-generation supply chains. Some proposals already exist in this direction [3, 7], but the one presented here is the first complete integration of a standard blockchain framework (*Iroha* - <https://www.hyperledger.org/projects/iroha>) with ontology-based object/product discovery. Semantic annotations of objects are disseminated on the distributed ledger, without the need for centralized storage or coordination. A matchmaking process grounded on non-standard inferences in [9] allows logic-based ranking of items in the supply chain w.r.t. a given query and formal explanation of discovery outcomes, in order to improve object retrieval and business intelligence.

A short technical description of the proposal, early results and future perspectives are presented hereafter.

## 2 PROPOSED APPROACH

The proposed framework defines a decentralized on-line object discovery layer built upon a basic blockchain. Item tracking operations have been revisited as SCs in order to comply with an opportunistic and distributed execution validated by consensus. Basic functions are reported in what follows.

**A. Object registration.** Supply chain nodes register as semantic-enabled agents in the blockchain. Tagged objects are registered as *assets* in the blockchain-based stream storage. In order to describe objects without depending on a backend, attached tags are annotated w.r.t. a domain ontology, modeling high-level qualitative information. Data-oriented attributes (*e.g.*, latitude and longitude of a location, production and expiration date) are also included. Every ontology is uniquely identified by a Uniform Resource Identifier (URI) to cope with the co-existence of several product domains in the same blockchain. Each asset is uniquely identified by its Electronic Product Code (EPC) and is available for discovery and usage as long as its tag remains in the operating range of a node (*owner*).

**B. Object discovery.** In order to search for a (set of) item(s), a node sends a multicast request containing: (i) the annotated request in Web Ontology Language (OWL)<sup>1</sup>, (ii) the ontology URL, (iii) a minimum semantic relevance threshold and (iv) a reference area as pair center-radius. The proposed approach adopts a *gossiping* (*a.k.a.*, *epidemic*) dissemination [5] for discovery requests and to aggregate results. The requester node randomly selects  $n$  nodes applying a location-based pre-filtering: nodes outside the query area are discarded. Each receiving node preselects among its assets only the ones managing the same reference ontology as the request. The embedded reasoning engine Mini-ME [9] computes the semantic distance between request and descriptions of selected assets through a semantic matchmaking. Hence, a relevance-ranked list of tagged objects is calculated, combining the semantic-based metric with geographical distances. Further  $n$  nodes are randomly selected in order to propagate the request. Nodes receiving the forwarded requests behave in the same way, up to a search depth threshold  $m$ . Each queried node returns results to the sender, which propagates them back to the original requester following the same routes of the direct flow.

**C. Explanation.** This is an optional step in a typical discovery process, invoked when a requester needs a justification of the matchmaking outcome for a specific object. It can be useful in a request refinement process.

**D. Object selection.** In order to obtain details of retrieved objects, the requester sends unicast messages to their owners indicating the object's EPC. Semantic descriptions along with additional data-oriented attributes are so returned.

All the above transactions are recorded on the blockchain for robustness, traceability and accountability purposes. In this way blockchain offers not only a decentralized collaboration infrastructure, but also a ledger allowing a full traceability up to sales preventing product counterfeiting.

<sup>1</sup>OWL 2 Web Ontology Language Document Overview (2nd Ed.), W3C Recommendation, 2012-12-11, <https://www.w3.org/TR/owl2-overview/>

## 3 EARLY RESULTS AND PERSPECTIVES

In order to assess effectiveness and scalability, early implementation and performance evaluation were carried out by extending *Iroha*. The SCs described in Section 2 were implemented in Java. Simulations exploiting *Docker* containers, each representing a supply chain node, were conducted in scenarios with 10, 50 and 150 nodes respectively. The average time for accomplishing requests and memory usage per node were evaluated in the above scenarios with increasing response timeout (*e.g.*, 2s, 6s, 10s). Furthermore the average hit ratio per node (*i.e.*, the percentage of requests replied with at least one object satisfying the constraints) was considered. All results reported hereafter are referred to the worst case in the test evaluations.

**Time:** in small and medium scenarios, mean turnaround time of object discovery was 0.23 s and 0.30 s, respectively. With 150 nodes, instead, it reached 7 s, due to the needed consensus among a larger number of components. **Memory:** RAM consumption is mainly related to Java Virtual Machine requirements for *Iroha*. Memory usage peak per node was always below 180 MB. **Hit ratio:** it was closely related to the number of nodes. The best results were obtained in the 50 nodes scenario, with 99.7%. In the 10 nodes scenario it was 83.3%, likely due to fewer objects available in the whole network. Finally, in the 150 nodes scenario hit ratio was only 35.7% with request timeout set to 2 s: *Docker* scalability issues led to increased probability of timeout expiration. This was confirmed by an improved value of 79.4% with 10 s timeout.

Early results supported the feasibility of the proposal: performance was satisfactory up to medium scale supply chains. Nevertheless, further optimization is needed, in order to improve scalability. Testbed migration toward a *Docker Swarm* deployment in a cluster computing environment will allow larger-scale simulations. Finally, further supply chain SCs for advanced product management and comparisons with state-of-the-art approaches are under investigation.

## REFERENCES

- [1] Konstantinos Christidis and Michael Devetsikiotis. 2016. Blockchains and Smart Contracts for the Internet of Things. *IEEE Access* 4 (2016), 2292–2303.
- [2] Roberto De Virgilio, Eugenio Di Sciascio, Michele Ruta, Floriano Scioscia, and Riccardo Torlone. 2011. Semantic-based RFID Data Management. In *Unique Radio Innovation for the 21st Century: Building Scalable and Global RFID Networks*, D.C. Ranasinghe, Q.Z. Sheng, and S. Zeadally (Eds.). Springer.
- [3] Matthew English, Sören Auer, and John Domingue. 2016. Block Chain Technologies & The Semantic Web: A Framework for Symbiotic Development. In *Computer Science Conference for University of Bonn Students*, Jens Lehmann, Harsh Thakkar, Lavdim Halilaj, and Rohan Asmat (Eds.). 47–61.
- [4] Joseph Galaskiewicz. 2011. Studying supply chains from a social network perspective. *Journal of Supply Chain Management* 47, 1 (2011), 4–8.
- [5] Márk Jelasity, Alberto Montresor, and Ozalp Babaoglu. 2005. Gossip-based aggregation in large dynamic networks. *ACM Transactions on Computer Systems (TOCS)* 23, 3 (2005), 219–252.
- [6] Andrew Kennedy, Michele Southall, Gena Morgan, Ken Traub, et al. 2014. *EPC Information Services (EPCIS) Specification*. Technical Report. GS1. Available: <http://www.gs1.org/epcis/epcis/1-1>.
- [7] Henry M. Kim and Marek Laskowski. 2016. Towards an Ontology-Driven Blockchain Design for Supply Chain Provenance. (2016). <https://ssrn.com/abstract=2828369>.
- [8] Kari Korpela, Jukka Hallikas, and Tomi Dahlberg. 2017. Digital Supply Chain Transformation toward Blockchain Integration. In *Proceedings of the 50th Hawaii International Conference on System Sciences*. 4182–4191.
- [9] Floriano Scioscia, Michele Ruta, Giuseppe Loseto, Filippo Gramegna, Saverio Ieva, Agnese Pinto, and Eugenio Di Sciascio. 2014. A Mobile Matchmaker for the Ubiquitous Semantic Web. *International Journal on Semantic Web and Information Systems (IJSWIS)* 10, 4 (2014), 77–100.