

Building a Knowledge Graph of Distributed Ledger Technologies

Lukas König ^{a,*} and Sebastian Neumaier ^a

^a *Department of Computer Science & Security, St. Pölten University of Applied Sciences, Austria*
E-mail: lukas.koenig@fhstp.ac.at

Abstract. Distributed ledger systems have become more prominent and successful in recent years, with a focus on blockchains and cryptocurrency. This has led to various misunderstandings about both the technology itself and its capabilities, as in many cases blockchain and cryptocurrency is used synonymously and other applications are often overlooked. Therefore, as a whole, the view of distributed ledger technology beyond blockchains and cryptocurrencies is very limited. Existing vocabularies and ontologies often focus on single aspects of the technology, or in some cases even just on one product. This potentially leads to other types of distributed ledgers and their possible use cases being neglected. In this paper, we present a knowledge graph and an ontology for distributed ledger technologies, which includes security considerations to model aspects such as threats and vulnerabilities, application domains, as well as relevant standards and regulations. Such a knowledge graph improves the overall understanding of distributed ledgers, reveals their strengths, and supports the work of security personnel, i.e. analysts and system architects. We discuss potential uses and follow semantic web best practices to evaluate and publish the ontology and knowledge graph.

URI: <https://w3id.org/DLTOntology>

DOI: [10.5281/zenodo.6497619](https://doi.org/10.5281/zenodo.6497619)

Keywords: Distributed Ledger Technology, Blockchain Security, Ontology, Knowledge Graph

1. Introduction

While the success of blockchains and especially cryptocurrencies has helped to spread the word about distributed ledger technology (DLT), it has led to misunderstandings about what the technology actually consists of and what it is capable of. In [1] this is highlighted very well with a comparison to facial tissues and the brand product Kleenex, which often get used interchangeably, even though the latter is just one implemented variant of the former. Such a misunderstanding of the technology can hinder its adoption. Each variant of distributed ledger technology operates in a slightly different way with regard to distinct requirements, which therefore means that there are also different usage scenarios and implementations that can be realized with distributed ledgers. Limiting the broad field of the technology to just one single variant of it leads to a distorted over-representation and use of that one variant, because the valid alternatives are simply not emphasized in the same way. There are observable attempts of reverting this trend by stating the technicalities and comparing different types of distributed ledger systems to one another, with new research and reviews stating up to 5 different types that are regularly seen as equally valid implementations of a distributed ledger system [2, 3].^{1 2}

* Corresponding author. E-mail: lukas.koenig@fhstp.ac.at.

¹<https://i.pinimg.com/originals/9d/d1/1b/9dd11bbaf5f025f3bd345e6816a4fc16.png>

²https://medium.com/@support_61820/different-types-of-dlts-and-how-they-work-cfd4eb218431

One way to unify the attempts of providing clarity on distributed ledgers is to use a knowledge graph to capture different technical capabilities and the wider ecosystem. In recent years, the concept of knowledge graphs has become increasingly popular, the idea being to use a graph-based data model to collect and convey knowledge about the real world [4]. To define the semantics of the terms used in the graph, an ontology, i.e. a formal representation of the relationships and classes in the graph, is used [5] As we point out in detail in Section 5, an all-inclusive approach towards a distributed ledger ontology is currently still lacking. Existing ontologies and vocabularies typically focus on a limited sub-group of distributed ledger systems, or build an ontology for a single distributed ledger product or application. Blockchains and blockchain-based systems are often the focus of such works, which leaves a large part of distributed ledgers unaccounted for.

This work therefore establishes a modular ontology, which encompasses and relates the following concepts and areas: Distributed ledger technology with a broader perspective than just a limited focus on blockchains, broader security considerations and technical implications of such systems, as well as organizational boundaries and real world applicability. Additionally, this ontology and the respective knowledge graph aims to clarify the distinctions between different types and systems of distributed ledgers, as it includes a broader spectrum than merely the view on blockchain as the paragon of the entire technology.

The model and subsequent knowledge graph created in this work can be used in a multitude of ways: Besides the already mentioned perspective of the entire ecosystem, it also allows for a more distinct observation of the included subjects. For instance, it can be used for threat analysis and risk assessments by identifying attacks and vulnerabilities of a system. In particular, security analysts will benefit from such a model, as it enables them to store and retrieve critical information about systems they might use, and to highlight further measures where they are needed to secure a system.

The main contributions of this work can be summarized as follows:

- a survey of literature and information on distributed ledger technologies, use-cases, fields of application, standardizations and legal issues;
- a set of competency questions that cover the main aspects of the collected information;
- an RDFS-based ontology that allows to model the surveyed aspects;
- a knowledge graph of relevant entities and relations, extracted from the literature;
- an evaluation of the model via SPARQL queries; an online interface to execute the queries.

The remainder of this paper is structured as follows: Section 2 provides an overview about distributed ledger technologies, as well as a short introduction to Ontologies, Knowledge Graphs and their terminologies. Section 3 details the methods, i.e., the information collection, the competency questions and evaluation of the ontology. Section 4 describes the ontology in detail, and it explains the elements and contents of the knowledge graph. In Section 5 we discuss related work and conclude in Section 6.

2. Background

2.1. DLT Data Structures

Blockchain. The most well-known example of a distributed ledger technology is a blockchain system. Especially the specific blockchain application Bitcoin sparked a hype around cryptocurrencies that lead to an often misunderstood terminology where Bitcoin and Blockchain are used interchangeably [6, 7], which is not true however. There are many more blockchain applications and cryptocurrencies are merely one aspect of the broader picture. Other examples of blockchain technologies are Ethereum and its renewed introduction of smart contracts, which have become increasingly popular over the last years, as well as the Hyperledger project by the Linux Foundation and IBM [8].

At its technological core, a blockchain is a data structure that (cryptographically) links blocks of data into a list of blocks. Blockchains are usually further separated into either public, private or federated blockchains, according to the degree of access that is implemented where a public blockchain is accessible to anyone and a private blockchain is restricted and usually used internally [2, 9].

1 *Directed Acyclic Graph.* Directed Acyclic Graphs (DAG) as a data structure used for distributed ledger systems 1
2 started to emerge with the applications NXT, and IOTA with its *Tangle*, which is especially focused on IoT, micro 2
3 transactions and devices with lower computational power [2, 3]. Instead of having data blocks hashed into one 3
4 single line, DAGs are used where connections between transactions in the graph show where a valid transaction 4
5 is present. Since each transaction is processed individually and does not require the formation of blocks as with a 5
6 regular blockchain, there is technically no transaction limit for DAGs and therefore they are very scalable. 6

7
8 *Hashgraph.* The most prominent product for Hashgraphs is *Hedera Hashgraph* as a next generation DLT that 8
9 focuses on fairness and transaction speed with regard to the order of events as they unfold. It is very efficient to use 9
10 as it uses virtual voting for reaching consensus and information about transactions are shared with a gossip protocol. 10
11 Technically, the data structure used for Hashgraphs is similar to a DAG, which is bound to the sequence of time 11
12 and the nodes participating in the network. By that they fundamentally differentiate themselves from regular DAGs 12
13 [2, 3]. 13

14
15 *Holochain.* Contrary to the previous types of DLTs, the major change with *Holochain* is the shift from a focus on 14
16 data itself towards an agent-centric structure which some consider the purest form of distributed ledger, as every 16
17 node virtually is responsible for its own ledger. Validation is done by the introduction of a ruleset to the Holochain 17
18 network. This ruleset is referred to as DNA and can be used to verify a node and to spot malicious actors. The 18
19 changes introduced with Holochain make mining obsolete, which makes it a very energy efficient alternative [3]. 19

20 2.2. Security Challenges, Opportunities and Use Cases for Blockchains & Distributed Ledger Systems 20

21
22
23 As with regular ICT systems, distributed ledger technologies are not immune to threats. While the introduction 23
24 of distributed ledgers by itself can mitigate certain risks already, it also introduces new dangers that need to be 24
25 considered. Especially since there are a number of myths revolving around the technology that make it seem almost 25
26 invincible [6, 8]. 26

27
28 *Organizational Challenges.* One of the challenges of using blockchains for supply chains or as a mean of ensuring 27
29 transparency and auditability is that if manual input is required, the people involved are still prone to errors and 28
30 false input [10]. Additionally, the knowledge about and familiarity with distributed ledgers is still limited generally 29
31 speaking. This lack of knowledge also hampers the willingness to become an early adopter of the technology [6, 30
32 7, 10]. Another challenge would be a lack of standardization for distributed ledgers or the integration into existing 32
33 standardization frameworks or questions about governance, legal requirements, and regulations, especially for cyber 33
34 security [11, 12]. This challenge is only aggravated by the fact that there is no unified terminology or vocabulary as 34
35 well [13]. 35

36
37 *Technical Challenges.* So far there has been intensive research on the vulnerabilities and attacks on distributed 36
38 ledger systems. A major part of these focuses on blockchain systems where *Bitcoin* and *Ethereum* are often times 37
39 the center of attention. Examples for blockchain attacks are 51%-Attacks, Fork Attacks, and Selfish Mining [14]. 38
40 Additionally, blockchains are often affected by scalability issues and low transaction rates due to mining and consen- 39
41 sus mechanisms [15, 16]. For other forms of distributed ledger, like a DAG, there is a strong focus on IoT and IIoT 40
42 [17], where security features that are present in a blockchain structure are sacrificed for an increase in performance. 42

43
44 *Opportunities and Use Cases* There are plenty of other fields of use for distributed ledger systems outside of the 43
45 world of cryptocurrencies. Since that is where the hype started, the introduction of other financial applications that 44
46 use the technology came to no surprise. This includes not only payment services but also the financial infrastructure 45
47 and international finance for example [9, 18]. However, distributed ledgers are also already extensively used in 46
48 supply chains around the globe as well. Instances of that would be shipping and logistics, ethical supply chains, or 47
49 food safety [6, 10]. Another opportunity for distributed ledgers is their use for government services, be it for voting, 48
50 record-keeping, transparency or combating fraudulent activities [9, 18]. Other recommended and proposed fields 49
51 for the use of distributed ledger technology are for example healthcare, identity management, the internet of things, 50
notary and insurance services, and (intellectual) property ownership [6, 19]. 51

3. Methodology

Methodologically, our ontology creation process follows the recommendations by Noy and McGuinness [5]. The modelling and creation of the ontology followed an iterative process and was done using Protégé. The applied steps can be summarised as follows:

1. In the first step we determine the *domain and scope* of the ontology.
2. We gather literature about distributed ledger technologies, their components and existing standards, application domains, security vulnerabilities and threats (cf. Sec. 3.2).
3. Second, we define a set of questions “that a knowledge base based on the ontology should be able to answer.” [5] We derive use cases and create *competency questions* based on the collected information (cf. Section 3.3).
4. We create an ontology to express the collected information (cf. Section 4).
5. We evaluate the ontology by providing SPARQL queries to answer the collected competency questions.
6. Finally, we instantiate the developed ontology by with named entities and relations extracted from the literature (cf. Section 4.1).

3.1. Ontology Scope

The scope of the ontology is a model for a holistic, modular approach on the field of distributed ledger technology that does not focus on a single or a subset of technologies, but rather the wider distributed ledger ecosystem. The goal is to include factors which affect the implementation and operation of distributed ledgers, including existing standards and its usage across the industries.

As this ontology does not serve as an overview of one single product or technology stack, the technological composition of a system is merely a starting point for a broader observation that includes its strengths and weaknesses, as well as the general applicability of a system for specific use cases. It therefore serves as an application-neutral knowledge base for the broader field of distributed ledger technology.

However, to achieve this general applicability it was necessary to focus on the broader picture, which means that detailed technical implementations and processes like message flows and triggers, and involved accounts are out of scope.

3.2. Information Collection

The core concept of this ontology encompasses the three major areas *technology, business & market use, and legal & standardization*.

In Section 2 we already discussed major structural differences between distributed ledger systems, including challenges and threats. The research in this mentioned section forms the basis for the technical part of the ontology, with additional information from further comparisons³. Additionally, technological differences between variants of blockchain/distributed ledger systems or smart contract platforms have been discussed and compared in [20, 21], from where we drew additional input for the ontology. Further information on the technological composition of distributed ledgers and possible technology stacks can be found in [22, 23] and from technical reviews⁴.

For use cases and the real-world application of blockchains and distributed ledgers, we build upon existing assessments and research on the possibilities and likelihood of distributed ledgers being used as a tool in different business sectors, which can be found in e.g. [6, 9]^{5,6}. On top of that, these business sectors can further be split up into specific use cases, with a majority of these being mentioned in Section 2.2 as well.

When it comes to organizational standardization of distributed ledger technology, finished work that is entirely focused on it is scarce. Reports like [7] offer a prediction and assessment of when and what to expect regarding standardization areas, others see it as a hurdle which is responsible for slow adoption rates of the technology, as

³<https://i.pinimg.com/originals/9d/d1/1b/9dd11bbaf5f025f3bd345e6816a4fc16.png>

⁴<https://101blockchains.com/web-3-0-blockchain-technology-stack/>

⁵<https://101blockchains.com/blockchain-digital-transformation/>

⁶<https://101blockchains.com/wp-content/uploads/2019/01/Applications-of-Distributed-Ledger-Technology-DLT.png>

mentioned in 2.2. However, there are many standards in the works and several standardization organizations have set up focus groups or committees to create new normative standardization reference material specifically catering towards distributed ledger technology and blockchain [11].

Legal issues on the other hand are still mostly unresolved and require further considerations [24]. Especially when it comes to governance and (shared) responsibility or liability when using blockchains. The introduction of specified laws and a legal basis, in particular for governance and adoption of the technology, is seen as a crucial factor by the authors of [25]. With ever rising numbers in cyber crime and based on the existing catalogues of laws on technology, it can be expected that distributed ledgers will receive their own set of laws in the future.

3.3. Evaluation

The competency questions in Table 1 are based on the collected information (see section 3.2) and grouped into three main categories:

- *Technology and Security* involves questions regarding the components of a DLT system, as well as technical threats and vulnerabilities of systems and components. From this category we derive relations between vulnerabilities of the respective systems and components, and attacks that exploit these vulnerabilities.
- *Industry and Application* involves questions about applications, business sectors, and use cases of DLT systems. Relevant relations in this category are the applications of DLT systems in specific use cases, as well as the mapping of use cases to a business sector.
- *Standardization and Regulation* involves questions about standards, technical controls, standardization organizations, and relevant laws. From this category we derive relations regarding compliance to standards, as well as the technical control of specific DLT components.

Table 1
Competency Questions used to evaluate and validate the ontology.

ID	Competency Question
Technology and Security	
T1	Which components are part of the distributed ledger system?
T2	Which technical threats have to be considered regarding the consensus algorithm of the system?
T3	What are known smart contract vulnerabilities?
T4	What are the data structures of the system?
T5	Which types of DLT attacks could be used against a system or its components?
Industry and Application	
I1	Which use cases can be realised with distributed ledger technology?
I2	Which industries could use distributed ledger systems?
I3	Which distributed ledger systems are used for public transportation and smart cars?
I4	Which types of record keeping could be realized with distributed ledger technology?
Standardization and Regulation	
S1	Which standardization organizations are active in regards to distributed ledger technology?
S2	Which normative references do exist for distributed ledger systems?
S3	Which industry standards do exist?
S4	What are relevant laws in regards to distributed ledger systems?
S5	What are organizational controls and mitigations for a distributed ledger system and/or component?
S6	Is there an industry initiative that directs and regulates the used distributed ledger system?

In Listing 1, 2, and 3 we give three SPARQL translations of the competency questions. In 1 we ask for existing technical threats that potentially threaten the consensus algorithm of a system; in 2 we translate question I4, which asks for specific use cases of a DLT system; the query in Listing 3 lists standardization organizations that actively publish DLT standards.

```

1  PREFIX : <https://w3id.org/DLTOntology#>
2
3  SELECT ?threat WHERE {
4    ?threat :threatens [ a :ConsensusAlgorithm ]
5  }

```

Listing 1 T2 - Which technical threats have to be considered regarding the consensus algorithm of the system?

```

9  SELECT ?dltsystem ?usecase WHERE {
10   ?dltsystem :isSpecializedFor ?usecase .
11   ?usecase a :RecordKeeping .
12 }

```

Listing 2 I4 - Which types of record keeping could be realized with distributed ledger technology?

```

15 SELECT ?stdorga ?standard WHERE {
16   ?stdorga a :StandardizationOrganization ;
17           :creates ?standard .
18   ?dltsystem :compliantTo ?standard .
19 }

```

Listing 3 S1 - Which standardization organizations are active in regards to distributed ledger technology?

The complete list of questions as SPARQL queries for all competency questions, including the respective classes and properties, can be found in the online documentation: <https://w3id.org/DLTOntology>.

4. DLT Ontology and Knowledge Graph

Based on the collected information and competency questions we developed *DLT Ontology*. The ontology covers the technical setup and components of a DLT system, security aspects, as well as its applications and use cases. It covers 115 classes and 15 properties, and consists of a total of 571 triples.

The DLT ontology provides classes and properties which are used to describe various aspects of a distributed ledger ecosystem. The core concepts are displayed in Figure 1: *DLTSystem* and *DLTComponent* describe the essential parts of a DLT system; *Vulnerability* and *Attack* connect potential security aspects of the system; *UseCase*, and *BusinessSector* link existing applications of the described systems.

The relations between the included classes are described in detail in Table 2, where a property-centric view of the core concepts, including domains and ranges of the properties of the ontology is provided.

4.1. Knowledge Graph

The knowledge graph uses the above introduced DLT Ontology, and is based on the information collection described in Section 3.2. It consists of three parts: (i) standards and legal authorities [11], (ii) technical details, vulnerabilities and security aspects [14], and (iii) use cases and business sectors [26]. Additionally, we reviewed white papers for further information about the technical collocation of DLT systems [27].⁷

In total, the knowledge graph consists of 746 triples; Table 3 lists the number of entities for the core classes.

4.2. Availability of the Resources

DLT Ontology and Knowledge Graph. The knowledge graph and respective documentation are available online at <https://w3id.org/DLTOntology> under the CC-by-4.0 license; a DOI is provided via the Zenodo repository.⁸ The

⁷<https://www.r3.com/reports/corda-technical-whitepaper/>, <https://ethereum.org/en/whitepaper/>

⁸<https://doi.org/10.5281/zenodo.6497619>

Table 2
Main Properties of the DLT Ontology.

Property	Description
canExploit <i>Domain:</i> Attack <i>Range:</i> Vulnerability	The <i>canExploit</i> property links an attack on a distributed ledger system or component to the respective vulnerability, which could be used by an attacker to enable further malicious activity in an effort to manipulate or gain control of the system.
threatens <i>Domain:</i> Attack <i>Range:</i> DLTSytem/Component	This property indicated that there is an existing attack on a distributed ledger system or one of its components. Attacks can manipulate specific components and data, or generally weaken the structural integrity of the entire system.
hasVulnerability <i>Domain:</i> DLTComponent <i>Range:</i> Vulnerability	This property indicates that one or multiple of the components used in a distributed ledger system carry with them one or an array of vulnerabilities. Having an overview of existing vulnerabilities is crucial for securing the system.
mitigates <i>Domain:</i> Technical/Org.Mitigation <i>Range:</i> DLTSytem/Component	Once disaster strikes, it is important to have a series of mitigations in place, to reduce the impact of an attack, malicious event or generally in case of the system breaking down. This property links a technical or organizational mitigation measure to a distributed ledger system or specific component.
hasComponent <i>Domain:</i> DLTSytem <i>Range:</i> DLTComponent	It links an individual component to the DLT system. It is used to show what a system is comprised of and to get a better understanding of the elements and building blocks. Such an overview allows for an analysis or the implementation of measures for a single component specifically and to tailor solutions to certain requirements.
isCompliantTo <i>Domain:</i> DLTSytem/Component <i>Range:</i> Standard	When it comes to using distributed ledger systems in an enterprise setting, national and/or international standardization can become a necessity. This property expresses compliance to existing standardization or normative references on a system or component level.
controls <i>Domain:</i> Technical/Org.Control <i>Range:</i> DLTComponent/System	Implemented controls can secure components and the entire system itself. This property links a technical or organizational control measure to a distributed ledger system or individual component.
createsStandard <i>Domain:</i> Stand.Organization <i>Range:</i> Standard	Behind every normative reference material and standardization there is an organization with the needed authority to publish such material, which can be either a national institute or an international organization. This property links a standard to its organization.
isRegulatedBy <i>Domain:</i> DLTSytem <i>Range:</i> Law	As with regular ICT-systems, a distributed ledger system will have to comply with the legal status of the respective country it is used in by any organization. This property connects existing legislation for ICT and DLT systems.
isSpecializedFor <i>Domain:</i> DLTSytem <i>Range:</i> UseCase	Not all implementations of distributed ledger systems are the same or similar in their nature and complexion. This property links existing DLT systems which are tailored for the delivery of a specific use case.
hasUseCase <i>Domain:</i> BusinessSector <i>Range:</i> UseCase	This property is used to show which use cases could be used in a selected industry. Not all forms of usage of distributed ledger systems are fit for all types of business.
isUsedFor <i>Domain:</i> DLTSytem <i>Range:</i> BusinessSector	There are certain distributed ledger systems that are specialized for the use in one specific field or industry. This property is used to highlight this exact relation.
comesFrom <i>Domain:</i> IndustryInitiatives <i>Range:</i> BusinessSector	It is entirely possible for several organizations of an industry form an initiative to steer and develop the future of distributed ledger technology within that industry in a common effort. This property shows the origins of an industry initiative with regards to the industry of involved parties.
hasBusinessSector <i>Domain:</i> UseCase <i>Range:</i> BusinessSector	Distributed ledger technology can be used to realize a variety of different use cases. However, not all of them are a fit for any field or industry. This property highlights the connection between a use case and the specific business sector/industry it is used in.

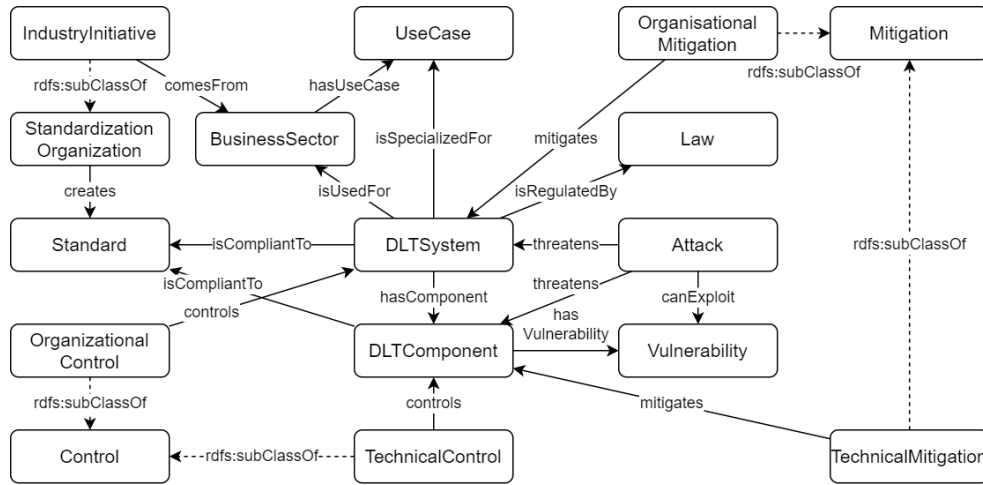


Fig. 1. Overview of the core concepts of the DLT Ontology.

Table 3
Number of triples and entities in the knowledge graph.

Triples	Standard	Std.Orga.	DLTComponent	UseCase	BusinessSector	Vuln.	Attack
746	18	8	9	55	9	7	11

online documentation consists of a comprehensive overview, examples of how to use the ontology, and SPARQL queries for the respective competency questions.

SPARQL Interface. We provide a SPARQL query interface at <https://w3id.org/DLTOntology> which allows to execute the competency questions provided in Section 3.3. The interface is based on the Comunica SPARQL Widget [28].

5. Related Work

Glossaries and Vocabularies for Distributed Ledgers. There have been made various attempts and proposals on the creation of terminologies, glossaries and vocabularies regarding blockchains and distributed ledgers. Most importantly, there is the ISO working group TC 307 [29], the ITU-T *Focus Group on Application of Distributed Ledger Technology* [30], and the German Institute for Standardization (DIN) [31]. On top of that there is a multitude of private companies and blockchain enthusiasts that offer their own glossaries for blockchains and distributed ledgers like Blockchainhub Berlin⁹, 101 Blockchains¹⁰, and ConsenSys [32].

Ontologies for Distributed Ledgers. EthOn (the Ethereum Ontology) [33] is an approach to model the major concepts of the Ethereum blockchain, e.g., Blocks, Accounts, Transactions, etc. The ontology itself is based on OWL [34] and there exist extensions that can be used to describe ERC20 compliant tokens¹¹ and smart contracts.¹² While EthOn focuses on modelling the technical details of a particular blockchain, we propose an approach that allows to model DLT systems, their specific use cases, and their potential security threats, independent of the concrete implementation.

⁹<https://blockchainhub.net/blockchain-glossary/>

¹⁰<https://101blockchains.com/blockchain-definitions/>

¹¹<https://github.com/ConsenSys/EthOn/tree/master/ERC20>

¹²<https://github.com/ConsenSys/EthOn/tree/master/Contracts>

Also related to our efforts is the BLONDiE ontology (Blockchain Ontology with Dynamic Extensibility) [35]. BLONDiE provides classes and properties to describe the structure and related information of the three most prominent blockchain projects – Bitcoin, Ethereum and Hyperledger. However, while our goal is to cover the overall ecosystem and landscape of existing DLT technologies, BLONDiE describes concepts that specifically relate to the implementation of the three technologies (e.g., “Ethereum Payload”, and “Hyperledger Transaction”).

6. Conclusion

Although distributed ledger technology has seen an enormous increase in usage and spread over the last decade, there are still many misconceptions and misunderstandings revolving around it. First and foremost is the equalization of distributed ledger technology with blockchains, where the latter is simply one of many manifestations of the former. Furthermore, there is a lack of holistic vocabularies and conceptualization of the technology as a whole, rather than just one of its manifestations or even a single product. To tackle and overcome present issues, we have contributed the following:

- We developed an ontology for distributed ledgers. This ontology considers various aspects of DLT systems, including threats, vulnerabilities, and the legal situation and standardization of the technology; an overview of the market situation and real-world applications is also included.
- We have demonstrated the use of the ontology by building a knowledge graph containing entities and relations for all the core classes of the ontology, i.e. (i) standards and legal authorities, (ii) components and technical details of DLT systems, as well as corresponding vulnerabilities and (iii) exemplary use cases.
- We have validated the ontology based on a pre-defined set of competency questions. On top of that, a set of SPARQL queries is provided to evaluate the competency questions and thus the capabilities of the knowledge graph.
- The ontology and the knowledge graph are available for public use under an open license; the documentation and downloads can be found at <https://w3id.org/DLTOntology>.

One limitation of our work in this regard is the lack of existing data that can be integrated in a knowledge graph: While it was possible to gather information on the technical components of DLT systems, existing standardization documents, and potential use cases in different industries, this requires intensive research, extensive manual work, and expert knowledge in the domain.

References

- [1] M.P. Network, Difference Blockchain and DLT, Marco Polo Network Operations (Ireland) Limited, 2018, Accessed: 2021-10-14.
- [2] N.E. Ioini and C. Pahl, A Review of Distributed Ledger Technologies, in: *OTM Confederated International Conferences "On the Move to Meaningful Internet Systems"*, Vol. 11230, H. Panetto, C. Debruyne, H.A. Proper, C.A. Ardagna, D. Roman and R. Meersman, eds, Springer, 2018, pp. 277–288. doi:10.1007/978-3-030-02671-4_16.
- [3] A. Panwar and V. Bhatnagar, Distributed Ledger Technology (DLT): The Beginning of a Technological Revolution for Blockchain, in: *2nd International Conference on Data, Engineering and Applications (IDEA)*, 2020, pp. 1–5.
- [4] A. Hogan, E. Blomqvist, M. Cochez, C. D’amato, G.D. Melo, C. Gutierrez, S. Kirrane, J.E.L. Gayo, R. Navigli, S. Neumaier, A.-C.N. Ngomo, A. Polleres, S.M. Rashid, A. Rula, L. Schmelzeisen, J. Sequeda, S. Staab and A. Zimmermann, Knowledge Graphs, *ACM Comput. Surv.* **54**(4) (2021). doi:10.1145/3447772.
- [5] N. Noy and D. McGuinness, Ontology Development 101: A Guide to Creating Your First Ontology, *Knowledge Systems Laboratory* **32** (2001).
- [6] B. Carson, G. Romanelli, P. Walsh and A. Zhumaev, Blockchain beyond the hype: What is the strategic business value, Digital McKinsey, 2018, pp. 1–13. <https://www.caba.org/wp-content/uploads/2020/04/IS-2018-209.pdf>.
- [7] A. Deshpande, K. Stewart, L. Lepetit and S. Gunashekar, Distributed Ledger Technologies/Blockchain: Challenges, opportunities and the prospects for standards, *Overview report The British Standards Institution (BSI)* **40** (2017), 40. https://www.bsigroup.com/LocalFiles/zh-tw/InfoSec-newsletter/No201706/download/BSI_Blockchain_DLT_Web.pdf.
- [8] M.K. Shrivastava, T.Y. Dean and S.S. Brunda, The Disruptive Blockchain Security Threats and Threat Categorization, in: *2020 First International Conference on Power, Control and Computing Technologies (ICPC2T)*, 2020, pp. 327–338. doi:<https://doi.org/10.1109/ICPC2T48082.2020.9071475>. <https://ieeexplore.ieee.org/document/9071475>.

- [9] D.B. Rawat, V. Chaudhary and R. Doku, Blockchain: Emerging Applications and Use Cases, *CoRR abs/1904.12247* (2019). <http://arxiv.org/abs/1904.12247>.
- [10] S. Jabbar, H. Lloyd, M. Hammoudeh, B. Adebisi and U. Raza, Blockchain-enabled supply chain: analysis, challenges, and future directions, *Multim. Syst.* **27**(4) (2021), 787–806. doi:10.1007/s00530-020-00687-0.
- [11] L. König, Y. Korobeinikova, S. Tjoa and P. Kieseberg, Comparing Blockchain Standards and Recommendations, *Future Internet* **12**(12) (2020), 222. doi:10.3390/fi12120222.
- [12] V. Gramoli and M. Staples, Blockchain Standard: Can We Reach Consensus?, *IEEE Commun. Stand. Mag.* **2**(3) (2018), 16–21. doi:10.1109/MCOMSTD.2018.1800022.
- [13] J. De Kruijff and H. Weigand, Towards a blockchain ontology, in: *11-th International Workshop on Value Modeling and Business Ontologies, The Luxembourg Institute of Science and Technology (LIST)*, 2017. https://www.list.lu/fileadmin/files/Event/sites/tudor/files/Trainin_g_Center/OTHERS/VMBO2017_paper_5.pdf.
- [14] L. König, S. Unger, P. Kieseberg and S. Tjoa, The Risks of the Blockchain A Review on Current Vulnerabilities and Attacks, *J. Internet Serv. Inf. Secur.* **10**(3) (2020), 110–127. doi:10.22667/JISIS.2020.08.31.110.
- [15] H. Chen, M. Pendleton, L. Njilla and S. Xu, A Survey on Ethereum Systems Security: Vulnerabilities, Attacks, and Defenses, *ACM Comput. Surv.* **53**(3) (2020), 67:1–67:43. doi:10.1145/3391195.
- [16] M. Saad, J. Spaulding, L. Njilla, C.A. Kamhoua, S. Shetty, D. Nyang and D.A. Mohaisen, Exploring the Attack Surface of Blockchain: A Comprehensive Survey, *IEEE Commun. Surv. Tutorials* **22**(3) (2020), 1977–2008. doi:10.1109/COMST.2020.2975999.
- [17] J. Sengupta, S. Ruj and S.D. Bit, A Comprehensive Survey on Attacks, Security Issues and Blockchain Solutions for IoT and IIoT, *J. Netw. Comput. Appl.* **149** (2020). doi:10.1016/j.jnca.2019.102481.
- [18] M. Schlegel, L. Zavolokina and G. Schwabe, Blockchain Technologies from the Consumers’ Perspective: What Is There and Why Should Who Care?, in: *51st Hawaii International Conference on System Sciences, HICSS 2018*, T. Bui, ed., 2018, pp. 1–10. <http://hdl.handle.net/10125/50329>.
- [19] A. Kolan, S. Tjoa and P. Kieseberg, Medical Blockchains and Privacy in Austria - Technical and Legal Aspects, in: *International Conference on Software Security and Assurance, ICSSA 2020, Altoona, PA, USA, October 28-30, 2020*, IEEE, 2020, pp. 1–9. doi:10.1109/ICSSA51305.2020.00009.
- [20] T.T.A. Dinh, R. Liu, M. Zhang, G. Chen, B.C. Ooi and J. Wang, Untangling Blockchain: A Data Processing View of Blockchain Systems, *IEEE Trans. Knowl. Data Eng.* **30**(7) (2018), 1366–1385. doi:10.1109/TKDE.2017.2781227.
- [21] X. Li, Z. Zheng and H. Dai, When services computing meets blockchain: Challenges and opportunities, *J. Parallel Distributed Comput.* **150** (2021), 1–14. doi:10.1016/j.jpdc.2020.12.003.
- [22] Deloitte, Blockchain Technology Stack, Deloitte Touche Tohmatsu India LLP, Accessed: 2021-06-15.
- [23] TE-FOOD, Making business profit by solving social problems - White Paper, Technical Report, TE-FOOD International GmbH., 2018, Accessed: 2021-10-28.
- [24] U.R. Rodrigues, Law and the Blockchain, *Iowa L. Rev.* **104** (2018), 679.
- [25] K. Werbach, Trust, but verify: Why the blockchain needs the law, *Berkeley Tech. LJ* **33** (2018), 487.
- [26] F. ITU, ITU-T Technical Report - Distributed ledger technology use cases, Technical Report, ITU Telecommunication Standardization Sector, 2019.
- [27] D. Hawig, C. Zhou, S. Fuhrhop, A.S. Fialho, N. Ramachandran et al., Designing a distributed ledger technology system for interoperable and general data protection regulation-compliant health data exchange: a use case in blood glucose data, *Journal of medical Internet research* **21**(6) (2019), e13665.
- [28] R. Taelman, J. Van Herwegen, M. Vander Sande and R. Verborgh, Comunica: a Modular SPARQL Query Engine for the Web, in: *Proceedings of the 17th International Semantic Web Conference*, 2018. <https://comunica.github.io/Article-ISWC2018-Resource/>.
- [29] ISO/TC307, Standards by ISO/TC 307 - Blockchain and distributed ledger technologies - Vocabulary, 2020, Accessed: 2021-09-16.
- [30] ITU, Focus Group on Application on Distributed Ledger - D1.1, Distributed ledger technology terms and definitions, Technical Report, International Telecommunication Union, 2019.
- [31] DIN, 16597:2018-02 Terminology for blockchains, 2018.
- [32] ConsenSys, A Blockchain Glossary for Beginners, 2021, Accessed: 2021-09-16.
- [33] J. Pfeffer, A. Beregszazi, C. Detrio, H. Junge, J. Chow, M. Oancea, M. Pietrzak, S. Khatchadourian, S. Bertolo and W.E. Bodell III, EthOn - An Ethereum Ontology, Accessed: 2021-06-14.
- [34] W.O.W. Group, OWL 2 Web Ontology Language, W3C Recommendation, Accessed: 2021-06-28.
- [35] U. Hector and C. Boris, BLONDIE: Blockchain Ontology with Dynamic Extensibility, *CoRR abs/2008.09518* (2020). <https://arxiv.org/abs/2008.09518>.