Universiteit Leiden
ICT in Business

Blockchain technology

An exploratory case study to identify the underlying principles and to determine the corresponding capabilities.

Name:	Job Bakker
Studentnr:	s1473662
Date:	23/08/2016
1st supervisor:	Dr. Werner Heijstek
2nd supervisor:	Drs. Marvin Meeng

MASTER'S THESIS

Leiden Institute of Advanced Computer Science (LIACS)
Leiden University
Niels Bohrweg 1
2333 CA Leiden
The Netherlands

Acknowledgements

First of all, I would like to thank my company supervisor Prof. dr. Ben van Lier for his excellent guidance throughout this process. The regular discussions that we had during this research were not only valuable but also enjoyable moments that helped me with completing this master thesis.

Secondly, I would also like to thank my first supervisor Dr. Werner Heijstek and second supervisor Drs. Marvin Meeng for their guidance and feedback. Besides that, I would like to thank the participants of the interviews for freeing up their time and sharing their knowledge and experience with me.

Executive summary

Experts estimate that there will be 38.5 billion connected devices to the Internet in 2020 [1]. All these devices will form different groups of autonomous computers that must be able to exchange data in a secure way in order to prevent major incidents. This implies that there is a need for new technologies, which can deal with the growing amount of connected devices and the different types of distributed systems. Therefore, this study has examined if the identified underlying principles of blockchain technolgy have the capabilities to add additional value to the current distributed solutions in multiple domains. This was done by conducting an exploratory case study that consisted of a critical literature review, which was followed by a number of expert-interviews. These were subsequently analysed by applying the grounded theory strategy of Charmaz [2].

The critical literature review showed that the underlying principles of blockchain technology are able to handle a large amount of devices that safely exchange data in complex and distributed environments. Besides safety, these principles also enable the removal of the traditional Trusted Third Party (TTP) and stimulate the self-regulating character of a network. This removes a set of devious and inefficient procedures and improves the integrity of the overall process by empowering nodes to perform tasks in an autonomous way.

By performing the expert-interviews it became clear that the respondents were unfamiliar with some of the identified underlying principles. This indicated that there is a lack of knowledge about multiple core elements of blockchain technology among the experts. This was an unexpected result and it turned out that several interviewees found it difficult to view the technology just as a combination of the underlying principles. Besides that, a selection of related use cases were also discussed during the interviews to test what the respondents considered to be value adding implementations of blockchain technology. It appeared that the smart industry use case in the industrial sector had the most potential. This was due to the amount of laws and regulations to comply with, the fact that there is no flow of money involved and the less risky nature of the data.

This research consisted of an exploratory case study that was conducted to identify the underlying principles of blockchain technology and to determine the corresponding capabilities. The previous paragraphs emphasized that there is a lack of knowledge and that the technological know-how needs to be improved in order to enable the development of blockchain implementations. It shows that this study has been a baseline measurement and indicates that there is much more research needed in order to fully understand blockchain technology.

Contents

List of Acronyms

ARPA Advanced Research Project Agency. 13

ARPANET Advanced Research Project Agency Network. 13

BGP Byzantine Generals Problem. 9, 23–25, 29, 36, 42

CAT Computerized Axial Tomography. 51, 63

CPS Cyber Physical Systems. 15, 17, 18, 36, 60, 62

CPU Central Processing Unit. 31

ICT Information and Communication Technology. 10

IIC Industrial Internet Consortium. 16

IoT Internet of Things. 9, 10, 12, 14, 15, 17–20, 36, 37, 60, 62

IP Internet Protocol. 13

MRI Magnetic Resonance Imaging. 51, 63

NASA National Aeronautics and Space Administration. 21

NCP Network Control Program. 13

NFC Near Field Communication. 15

NIST National Institute of Standards and Technology. 15

RFID Radio Frequency IDentification. 14, 15

SIFT Software Implemented Fault Tolerance. 21, 22, 26

SMR State Machine Replication. 21, 41

TCP Transmission Control Protocol. 13

TTP Trusted Third Party. 2, 10, 36, 39, 40, 48, 58, 62

ULSS Ultra-Large-Scale Systems. 18, 36, 60, 62, 63

WWW World Wide Web. 13, 14

List of Figures

List of Tables

Chapter 1

Introduction

The first few years since its technical origin in 1983, the Internet was nothing more than a technology that enabled daily computer communications [3]. This changed since the introduction of the term Internet of Things (IoT) in 1999 [4]. After that time, it became also possible to connect other physical objects than computers to the Internet. This led to the development of connecting all sorts of devices to the Internet, in fact experts estimate that there will be 38.5 billion connected devices in 2020 [1]. Connecting such an amount of devices to the Internet introduces several security threats. Actually, all these devices form different groups of autonomous computers on the IoT or so called distributed systems [5]. A serious threat to those kind of systems is the Byzantine Generals Problem (BGP). This problem is defined as malfunctioning components that give conflicting information to different parts of a system [6]. When this information is accepted, there can arise various faults and failures in a system that have severe consequences. Therefore, a solution for this threat is needed in order to enable a safe use of the IoT in the future. Blockchain technology is seen as a possible solution for the BGP because it empowers a system to identify components that give conflicting information. In short, it is an encrypted chronological database that is recorded by a network of computers [7]. This database is spread among every device in a network so it can easily be noticed and ignored when a device distributes conflicting information. It implies that the use of a traditional Trusted Third Party (TTP) will be redundant because there is no more need to verify information externally.

1.1 Research questions

The introduction above sketches a problem that can be framed by composing multiple research questions. Answering these questions will help to reach the objective of this study that is described in section 1.2.

Main question: Are the underlying principles of blockchain technology capable to add additional value to the current distributed solutions in multiple domains?

Sub questions:

1. What kind of developments have led to combining more and more objects in networks and what are the related consequences for secure communication?
2. What are the principles underlying blockchain technology and which role do they play in achieving secure communication?
3. Are experts in different domains aware of these underlying principles and what do they consider to be value adding implementations of blockchain technology?

1.2 Research objective

At this point in time, blockchain technology is considered to be the next disruptive technology in the Information and Communication Technology (ICT) sector. This may seem a bit exaggerated to some but the idea to remove the TTP in all sort of situations is revolutionary. Therefore, the objective of this research is to determine to which extent the principles behind blockchain technology can add additional value in multiple domains. In order to do so, it is important to understand why this technology has emerged and what the underlying principles are. This contributes to the current body of literature and helps to unravel the mysteries around this new technology. By verifying and expanding this information through the performance of various expert-interviews, this study will also provide a baseline measurement. This measurement is useful to make predictions about the use of this technology in multiple domains. These will enable companies to identify blockchain related opportunities and helps them to develop their own value adding implementations. Overall, this study will expand the current knowledge about blockchain technology and gives future directions for potential solutions.

To find out to which extent the underlying principles of blockchain technology can add additional value, there are a number of steps that need to be taken. First of all, it must be clear what opportunities and consequences there arise through the developments in the IoT field. When these are identified, it is time to decide if blockchain technology could be an appropriate solution for these consequences. In order to do this, the underlying principles of this technology will be examined and a number of practical examples are reviewed that make use of those principles. With this information it is possible to evaluate if the principles can actually be used in multiple domains.

The next phase is to conduct several expert-interviews to define if others also recognize the potential of the underlying principles behind blockchain technology.

This will be done by reviewing these principles and discussing a number of use cases in which they are used. After these interviews, it is possible to make an inventory of the available knowledge and different developments around the technology in multiple domains. Combining these outcomes, enables the possibility to estimate to which degree the principles behind blockchain technology will be used in the future.

1.3 Research method

The chosen research method for this assignment is an exploratory case study that uses a theory-building structure. At this time, blockchain technology is a new and complex topic that needs exploration in order to determine the key principles behind this technology. This will be done through conducting a critical literature review which is followed by a number of expert-interviews. This critical review helps with understanding a field, its key theories, the concepts, ideas, major issues and debates [8]. After that, the expert-interviews will be performed to gain additional knowledge in a qualitative way. These interviews have a semi-structured nature, which means that there are a few themes and key questions that need to be covered. The remaining portion of the interview can vary per case because the experts will have different knowledge levels.

After the interviews have been conducted it is time to analyse the data with a methodological approach called grounded theory. This term refers to a theory that is grounded in or developed inductively from a set of data [8]. The key element of this methodology is coding and consists of two principal stages according to Charmaz' strategy [2]. It starts with the initial coding stage, which involves the disaggregation of the collected data into conceptual units that are provided with a label (see appendix B). The second stage is focused coding and includes the re-analysing of data to test which of the initial codes can be used to categorise larger units of data. When new data is collected during the coding process it is important to apply the principle of constant comparison. This is needed to check for similarities and differences, to promote consistency and to aid in the overall process of analysis. Using the constant comparison technique stimulates abduction, which helps to gain additional insights in order to create new conceptual possibilities that can be further examined. It also promotes a higher level of analytical coding because the researcher moves between inductive and deductive thinking [8]. The complete overview of the research method is depicted in figure 1.1.

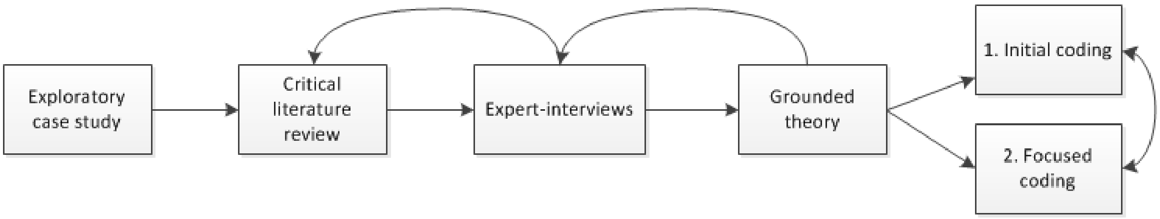

Figure 1.1: Research method overview

During and after these data collection and analysis techniques the exploratory case study will be formed using a theory-building structure. With this structure it is possible to gradually reveal new parts of the theoretical argument that is being made [9]. In this way, the final outcome will be a report which not only gives future research directions but also contains a theoretic overview of the key principles behind blockchain technology.

In order to ensure construct validity in this study it is important to take several tactics into account [9]. The first one is the use of multiple sources of evidence, which creates a form of triangulation that increases the reliability of the data. The second tactic implies the establishment of a chain of evidence, this enables an external observer to follow the derivation of evidence from the initial research questions to the final conclusions. The last one includes the reviewing of the case study report drafts by key informants, which improves the overall quality of the report.

1.4 Structure

The next chapter starts with a brief historical overview to describe the events that led to the Internet in its current form. These will explain the rise of the IoT and show the related opportunities and consequences of this network. After that, chapter 3 will describe the developments that led to the emergence of the field blockchain technology. By examining these developments, it will become clear what the underlying principles of the technology are and which functionalities it can offer. The results of these findings will be discussed in the next chapter, which contains the answers of the expert-interviews that have been conducted in multiple domains. These will help to verify the earlier identified underlying principles and functionalities of blockchain technology. Based on these outcomes, it is time to answer the research questions of this study. Together with the corresponding implications and the need for future work this will be done in chapter 5.

Chapter 2

History of the Internet of Things

The aim of this chapter is to give a brief historical overview of the developments that led to the Internet in its current form. These developments will explain the emergence of the Internet of Things and show the related opportunities and consequences of this network.

2.1 Brief historical overview of the Internet

In August 1962 Licklider envisioned a globally interconnected set of computers through which everyone could quickly access data and programs from any site [3]. Later that year, he became the first director of the Advanced Research Project Agency (ARPA) that funded the groundwork for the technical foundation of the Internet. This research project was an initiative of the United States Department of Defense and led to the development of a pioneering packet switching network in 1967. It was called the Advanced Research Project Agency Network (ARPANET) and made use of the Network Control Program (NCP), a Host-to-Host protocol which enabled users to begin with the development of all sorts of applications [3]. One of those applications was electronic mail and initially fulfilled the need for an easy coordination mechanism between the developers of the ARPANET. It became the largest network application for over a decade until the network implemented the Transmission Control Protocol (TCP) and the Internet Protocol (IP) suite in 1983, also known as TCP/IP. This suite enabled the network to meet the needs of an open-architecture network environment, which resulted in a split of multiple networks [3]. Since then, the Internet is considered as a well-established technology that supported different communities in their daily computer communications.

This development became accelerated when Berners-Lee started the World Wide Web (WWW) initiative in 1989, which was a project designed to bring a global information universe into existence by using available technology [10].

It enabled the dream of extending the human intellect by making collective knowledge available to each individual by using machines. The WWW model consists of a combination of the hypertext, information retrieval and wide area networking technique. It introduced an easy way to make information publicly available for everyone who owned a computer. The model resulted in a significant raise of the total amount of information that was available on the web and transformed the Internet into a mature medium.

2.2 Moving towards the Internet of Things

The term ubiquitous computing initiated the first shift towards a different kind of Internet and was introduced by Weiser in 1988. It can be defined as the goal to have non-intrusive availability of computers throughout the physical environment, virtually, if not effectively, invisible to the user [11]. Eventually, Weiser predicted that computers will weave themselves into the fabric of everyday live until they are indistinguishable from it. By the development of products such as scratchpads, live boards and page-sized pads in the early 90's his ideas slowly evolved into practical applications [12]. Another factor that played a role in the successful realization of ubiquitous computing was the growing distribution of home computers among households. Take for example the Netherlands, in the period of 1985 until 1995 the adoption of home computers in this country grew from 7 to 39% [13]. This indicates that computers were becoming part of people their everyday lives in the mid 90's and that Weisers' vision was correct.

2.3 The Internet of Things

The second shift that changed the Internet started when Ashton introduced the term IoT in 1999. It can be defined as a world where physical objects are seamlessly integrated into the information network, in which they will become active participants in business processes. Services are available to interact with these 'smart objects' over the Internet, query their state and any information associated with them, taking into account security and privacy issues [14]. The unique characteristic of the IoT is that the Internet is used to interconnect physical objects that communicate with each other or with humans. This created new opportunities for users, manufactures and companies to fulfill their needs and support them in their everyday activities [15].

This development started with the use of Radio Frequency IDentification (RFID) tags on physical objects. When these tags are mounted to a physical object it enables the ability to transfer data to a RFID reader in a wireless way.

The first tags were passive, so they did not have on-board power supplies and took the required energy from the query signal transmitted by a RFID reader nearby [16]. Later on, they also used semi-passive and active tags who got there power supply through batteries. The RFID tags are considered to be one of the cornerstones of the IoT because it initiated the evolution towards more advanced tagging technologies such as Near Field Communication (NFC).

Important to understand is how technologies such as RFID and NFC have contributed to the development of the IoT in its current form. Initially, physical objects were separately tagged to enable the ability to transfer data towards a reader. This exchange of information was so useful that they started with integrating these tagging technologies into physical objects. By doing so, this ability became standardized and indicated a next phase in the emergence of the IoT.

2.3.1 Cyber Physical Systems

According to the National Institute of Standards and Technology (NIST), Cyber Physical Systems (CPS) are smart systems that encompass computational (i.e. hard- and software) and physical components, seamlessly integrated and closely interacting to sense the changing state of the real world [17]. The distinctive feature of CPS compared to traditional embedded systems is that the former ones are able to interpret the physical world and can subsequently use this information in performing their tasks. This enables several applications for CPS in areas such as avionics, transportation, factory automation, electronic healthcare and smart grid systems [18].

In order to function properly, a CPS is dependent on the interaction between computational and physical components. This introduced a new phase in the development of the IoT because at first, physical objects were able to function properly without computational components. Around that time, there already existed objects with integrated components but these only provided additional functionalities. If they stopped working the object could still perform its intended tasks but this became impossible within a CPS. This demonstrates that the advent of the Internet and the evolution into the IoT has set a revolutionary transformation of systems in motion.

2.3.2 The Industrial Internet

The last part of the IoT definition states that physical objects can become active participants in business processes. This is exactly what the industrial sector tries to achieve by connecting those objects to the Internet.

The so called Industrial Internet is an internet of things, machines, computers and people, enabling intelligent industrial operations using advanced data analytics for transformational business outcomes [19]. It is a phenomenon that unlocks several new business models for companies because it empowers them to make use of technologies like remote access and data analytics. In this way, the Industrial Internet can offer new insights into what functionalities work and which do not. Based on these outcomes, companies can not only focus on optimizing their products but are also able to produce conform their customer needs [20].

In order to help industries with connecting their systems to the Industrial Internet the Industrial Internet Consortium (IIC) has specified a reference architecture framework. The four viewpoints of this framework are illustrated in figure 2.1.

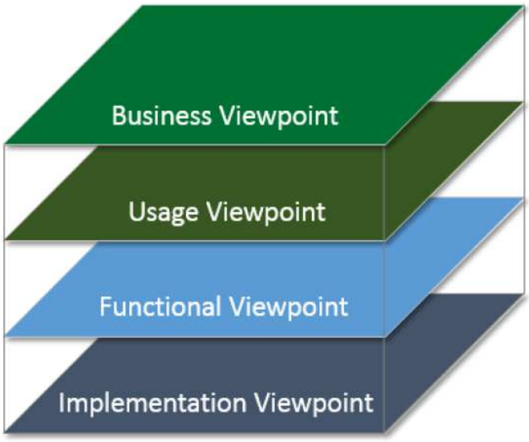

Figure 2.1: Industrial Internet architecture viewpoints [19]

Each viewpoint deals with another area of concerns and gives industries a general overview of the things to keep in mind. The business viewpoint is related to the concerns of the identification of stakeholders and their business vision, values and objectives. One view deeper is the usage viewpoint which addresses the concerns of expected system usage and involves sequences of activities involving human or logical users. The functional viewpoint focuses on the functional components to support the usages and activities of the overall system. Finally, the implementation viewpoint deals with the technologies needed to implement functional components, their communication schemes and their lifecycle procedures [19].

Besides the specified framework, the reference architecture document of the IIC also covers a set of other important topics related to connecting systems to the Industrial Internet.

With the help of such documents multiple countries have set up their own initiatives to contribute to the Industrial Internet. Examples of these initiatives are advanced manufacturing (United States of America), industrie 4.0 (Germany) and smart industry (The Netherlands). These countries demonstrate that terms like the IoT, CPS and the Industrial Internet do not only appeal to the imagination but also provide real business value. The overall goal of these initiatives is to determine new ways in which technology can improve existing products or processes.

2.4 The consequences of these developments

Since the introduction of the IoT the number of connected devices to the Internet has grown significantly. At this point, the number of connected devices to the Internet is 13.4 billion and already exceeds the world population over two times. Experts estimate that these growth will continue exponentially and that the amount of connected devices to the Internet will be 38.5 billion in 2020 [1]. By connecting such an amount of devices to the Internet it is not hard to imagine that there are besides the described opportunities also some consequences.

2.4.1 Dependability

Connecting all these various types of physical objects to the Internet increases the dependability of the IoT. This dependability is twofold, it is not only related with the amount of information but also with the final use of it. The growing amount of information contributes to the dependability on one end because it enhances the use of the IoT to search for specific information. Apart from that, the number of physical objects that need a particular part of this information to function properly is also rising. When objects are unable to reach this information it also affects the people who try to use these physical objects, which introduces the other end of dependability. Through the developments described earlier on, this erosion of boundaries between people and systems will continue [21]. Eventually, people will become elements of the system and are dependable on the IoT to perform their daily activities.

2.4.2 Smart industry

Throughout the previous sections, it became clear that developments like CPS and the Industrial Internet are a result of the numerous existing techniques to connect physical objects with each other. The positive side of this trend is that an initiative such as smart industry in the Netherlands is beginning to take shape.

In 2014, there were already 320.000 companies who joined it ranging from the agro-food domain to companies in the high-tech industry [22]. They are currently working on a number of action points to accelerate the digitalization and to enhance the competitive strength of the Dutch industry. These include the capturing of existing knowledge, acceleration through field labs and strengthening the overall foundation.

2.4.3 Ultra-Large-Scale Systems

In a few years' time, the combination of such developments will lead to a new phenomenon called Ultra-Large-Scale Systems (ULSS). Which are systems that push far beyond the size of today's systems by every measure: number of lines of code; number of people employing the system for different purposes; amount of data stored, accessed, manipulated, and refined; number of connections and interdependencies among software components; and number of hardware elements [21]. Some important characteristics of a ULSS are that it is decentralized, heterogeneous, inherently conflicting and continuously evolving. This places unknown demands on aspects like software acquisition, production, deployment, management, documentation, usage, and evolution practices.

2.4.4 Pushing the boundaries

Gradually, the IoT has been pushing the boundaries of the Internet and this implies that there are several aspects that need to be solved in different ways. If this does not happen major incidents are unavoidable as is depicted in figure 2.2. A well-known example is the discovery of the computer worm Stuxnet in 2010 that infected the software of multiple industrial sites in Iran, including a uranium enrichment plant.

The figure shows that the majority of these incidents are caused by substandard security measures. Therefore, it is time to seriously reconsider the security measures that are implemented in existing systems. A dangerous environment arises when concepts as CPS, the Industrial Internet and ULSS are combined with the trend to blindly connect all sorts of physical objects to the Internet. One logical but important step is to question if it is always necessary to connect a physical object to the Internet. This mentality alone can prevent a large share of the minor incidents that occur on a daily basis in the IoT environment. Besides that, it is critical that the parties who are developing software or objects that will be connected to the IoT become aware of the possible threats. There are multiple elements in the IoT network identified that show patterns of vulnerabilities, examples are boundary protection, information flow enforcement, remote access and physical access control [23].

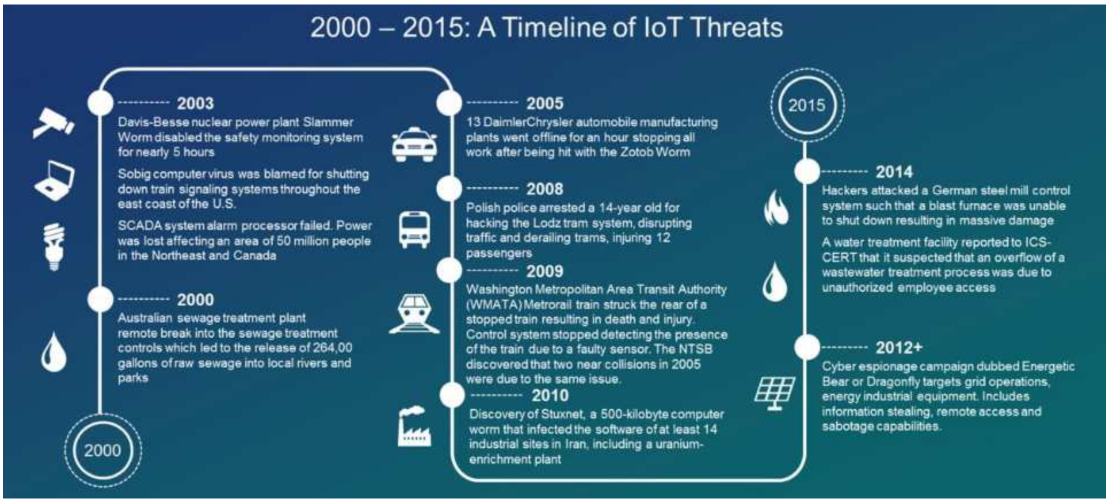

2000 – 2015: A Timeline of IoT Threats

2000

Australian sewage treatment plant remote break into the sewage treatment controls which led to the release of 264,00 gallons of raw sewage into local rivers and parks

2003

Davis-Besse nuclear power plant Slammer Worm disabled the safety monitoring system for nearly 5 hours

Sobig computer virus was blamed for shutting down train signaling systems throughout the east coast of the U.S.

SCADA system alarm processor failed. Power was lost affecting an area of 50 million people in the Northeast and Canada

2005

13 DaimlerChrysler automobile manufacturing plants went offline for an hour stopping all work after being hit with the Zotob Worm

2008

Polish police arrested a 14-year old for hacking the Lodz tram system, disrupting traffic and derailing trams, injuring 12 passengers

2009

Washington Metropolitan Area Transit Authority (WMATA) Metrorail train struck the rear of a stopped train resulting in death and injury. Control system stopped detecting the presence of the train due to a faulty sensor. The NTSB discovered that two near collisions in 2005 were due to the same issue.

2010

Discovery of Stuxnet, a 500-kilobyte computer worm that infected the software of at least 14 industrial sites in Iran, including a uranium-enrichment plant

2012+

Cyber espionage campaign dubbed Energetic Bear or Dragonfly targets grid operations, energy industrial equipment. Includes information stealing, remote access and sabotage capabilities.

2014

Hackers attacked a German steel mill control system such that a blast furnace was unable to shut down resulting in massive damage

A water treatment facility reported to ICS-CERT that it suspected that an overflow of a wastewater treatment process was due to unauthorized employee access

Figure 2.2: A timeline of IoT threats [23]

Focusing more on those areas and taking the appropriate measures to improve security are important steps to reduce the number of major incidents. A combination of awareness and technical know-how will be able to make the IoT a safer place but this requires time and effort.

2.5 Recap

This chapter gave a brief historical overview of the Internet to describe the main developments that led to the IoT in its current form. It is important to realize that this technological evolution not only creates new opportunities but also has some consequences. When taking these consequences into account the IoT can enable those new opportunities in a safely and responsible manner. In the next chapter it is time to zoom in on a pervasive technology that can improve the security in the IoT environment. This technology has the ability to safely store data in a distributed database that subsequently can be exchanged between the different nodes in a network. With the prediction that the number of connected devices to the Internet and the amount of information will only grow, this is a logical development.

Chapter 3

The emergence of blockchain technology

In chapter 2 it became clear that there were a number of important developments that led to the IoT in its current form. This is also applicable to the emergence of the field blockchain technology and these developments will be described in this chapter. After this description, it will be apparent what the underlying principles of the technology are and which functionalities it can offer. By mapping the principles behind blockchain technology it is easier to determine which underlying concepts need to be used in future implementations.

3.1 Generic approach

The introduction of chapter 1 made it clear that a blockchain is in essence an encrypted chronological distributed database. This implies that the underlying principles of blockchain technology have their origin in the distributed computing field. Therefore, there will be several distributed computing developments discussed in this chapter to unravel the combination of elements that form a blockchain. An outline of these elements is illustrated in figure 3.1.

3.2 Brief introduction distributed systems

The IoT is an environment where billions of devices are connected with each other as section 2.4 showed. All these devices form distributed systems because they consist of a collection of distinct processes which are spatially separated, and which communicate with one another by exchanging messages [24]. In order to explain how a distributed system can function in a proper manner there are various terms that need to be explained in the following sections. After that, the role of blockchain technology in the context of distributed systems will also be clear.

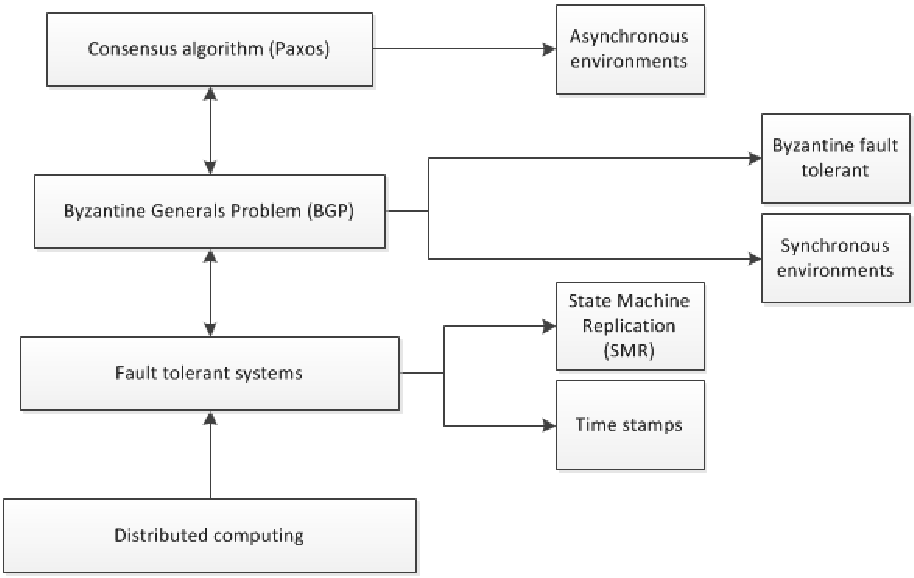
Consensus algorithm (Paxos)
Asynchronous environments
Byzantine fault tolerant
Byzantine Generals Problem (BGP)
Synchronous environments
State Machine Replication (SMR)
Fault tolerant systems
Time stamps
Distributed computing

Figure 3.1: Distributed computing elements

3.3 Fault tolerant

An essential characteristic of any system, including a distributed system is that it needs to be fault tolerant. This is defined as the ability of a system to continue to perform its specified tasks after the occurrence of faults [25]. In figure 3.2 are the different causes for faults defined and these are categorized by the techniques to improve or maintain a system's normal performance.

One of the first fault tolerant systems was developed in 1978 by the National Aeronautics and Space Administration (NASA), which was a computer for critical aircraft control applications. It achieved fault tolerance by the replication of tasks among processing units. This enabled the system to dynamically reconfigure itself in order to bypass faulty units when the software detected and analyzed errors [26]. There are some other important concepts involved which made this system fault tolerant and those will be treated in the next subsections.

3.3.1 State Machine Replication

A distributed system makes use of a technique called State Machine Replication (SMR) to implement fault tolerance. By replicating states of a system among multiple processes, system malfunctions through software and hardware faults can be prevented as the Software Implemented Fault Tolerance (SIFT) computer of the NASA showed [26].

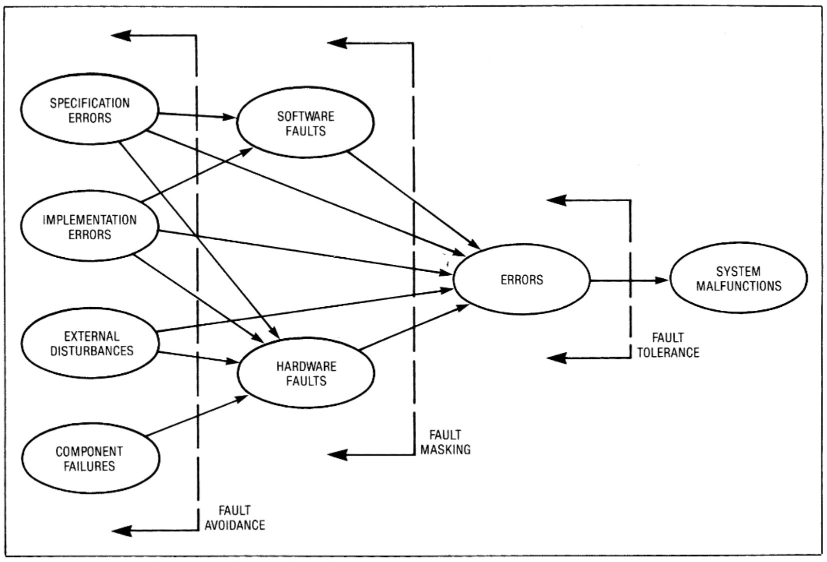

Figure 3.2: An overview of causes for system malfunctions [25]

Another challenge where SIFT and every other distributed system has to deal with is that each individual process consists of a sequence of events. In order to function properly, the system has to be capable of determine the order of these different events.

3.3.2 Time stamps

To determine this order, Leslie Lamport introduced the concept of logical clocks in a distributed system. With these logical clocks there was no more need to use physical time because a system could use the so called "happened before" relation [24]. This relation bypassed the accuracy problems with real clocks and is capable of keeping track of the order of events through a mathematical algorithm.

Until this point, it was possible to make a distributed system fault tolerant for software and hardware faults but this is only valid under one specific assumption. This assumption describes that all the components within a distributed system function according to their specification. Unfortunately, there can arise situations in which components function voluntarily or involuntarily maliciously and thereby causing system faults. To handle this additional type of faults, distributed systems also have to be Byzantine fault tolerant.

3.4 Byzantine fault tolerant

The problem related to the ability of a distributed system to be Byzantine fault tolerant is called the Byzantine Generals Problem (BGP). It is defined as malfunctioning components that give conflicting information to different parts of a system [6]. To describe the problem, Lamport uses the following metaphor of a group of generals of the Byzantine army camped with their troops around an enemy city: Communicating only by messenger, the generals must agree upon a common battle plan. However, one or more of them may be traitors who will try to confuse the others. The problem is to find an algorithm to ensure that the loyal generals will reach agreement. By introducing malfunctioning components that give conflicting information in a distributed system, there needs to be a way for the system to reach agreement. Without this agreement, it becomes impossible for a distributed system to function properly as is demonstrated by the following BGP examples.

3.4.1 Impossibility results

The impossibility results describe the situations in which it is not possible to reach agreement with oral messages when there is a single traitor among three generals. Important to note is that an oral message is one whose contents are completely under the control of the sender, so a traitorous sender can transmit any possible message. For simplicity, the only possible decisions in these examples are “attack” or “retreat”. In figure 3.3 is the situation depicted in which lieutenant 2 is the traitor and the other two generals are loyal.

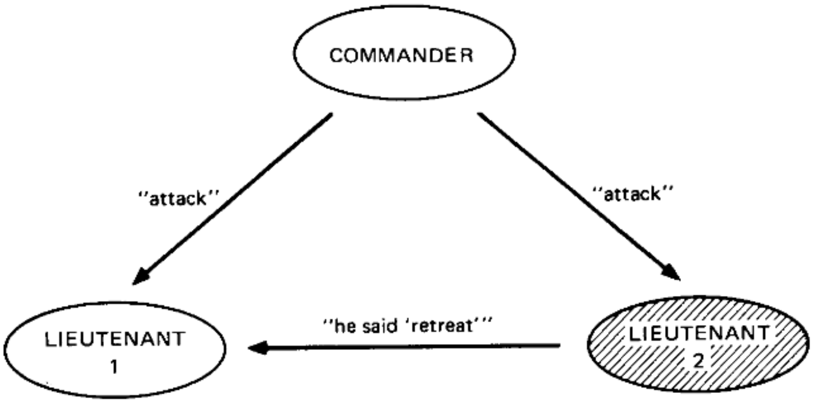

Figure 3.3: When the lieutenant is a traitor [6]

The situation in which the commander is the traitor and the other two generals are loyal is illustrated in figure 3.4. In both situations, lieutenant 1 is not able to determine who the traitor is and therefore it becomes impossible to reach agreement.

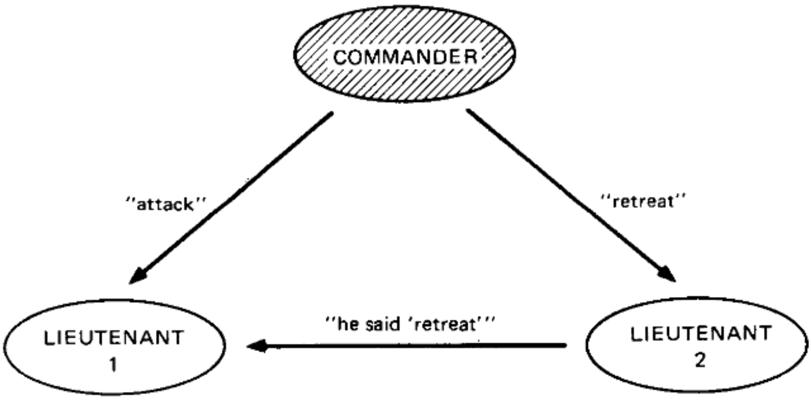

Figure 3.4: When the commander is a traitor [6]

3.4.2 Solutions

There are two types of solutions for the BGP and each of these uses a different kind of message. In the examples shown above, oral messages could not provide a solution for reaching agreement but perhaps the other type could offer an outcome.

Signed messages

The ability of a traitor to transmit any possible message makes the BGP very difficult to solve. By restricting this ability in the form of unforgeable signed messages, it becomes easier to solve the problem. In this case, every message contains a signature so that each lieutenant is able to determine who the sender of the message is. Another modification is that every lieutenant sends his received message from the commander to the other lieutenants, including the signature of the commander and itself. This is shown in figure 3.5, where the commander is the traitor and the other two generals are loyal. With these signed messages the lieutenants are able to identify that the commander is a traitor because his signature appears on two different orders. Therefore, the lieutenants can reach agreement about the fact that the commander is not loyal and this solves the impossibility results among three generals. The impossibility results remains intact for two generals because it is still impossible to reach agreement when there is a single traitor in this situation.

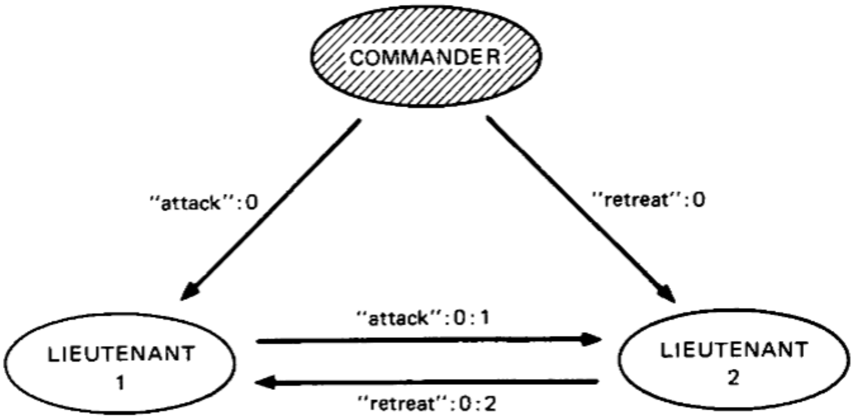

Figure 3.5: A solution with signed messages [6]

Oral messages

Another variant is to use a solution with oral messages, which will require more communication before the traitor is identified because the messages have no signature. This implies that every lieutenant sends his received message from the commander to the other lieutenants. In figure 3.6 is the situation depicted in which the commander is a traitor and sends arbitrary values to the other three lieutenants, who are loyal. By comparing their set of orders with each other, the lieutenants are able to find out that the commander is spreading different orders and can identify that he is a traitor. After comparing their messages, there is no need to return those also to the commander because they already know who the traitor is. Unfortunately, this method is less effective because it cannot solve the earlier described impossibility results. This means that there are at least four generals required to reach agreement about the fact that there is a traitor among them.

3.4.3 Implications

The solutions described in subsection 3.4.2 show the advantage of making a distributed system Byzantine fault tolerant. By doing so, a distributed system is able to deal with malfunctioning components and is no longer vulnerable for the BGP. Overcoming the BGP was a breakthrough in distributed computing but the provided solutions were inherently expensive due to the high number of messages that is required. Therefore, the provided solutions are mainly applicable in situations where extremely high reliability is a prerequisite [6]. Another limitation was that the solutions only focused on distributed systems that made use of a synchronous environment.

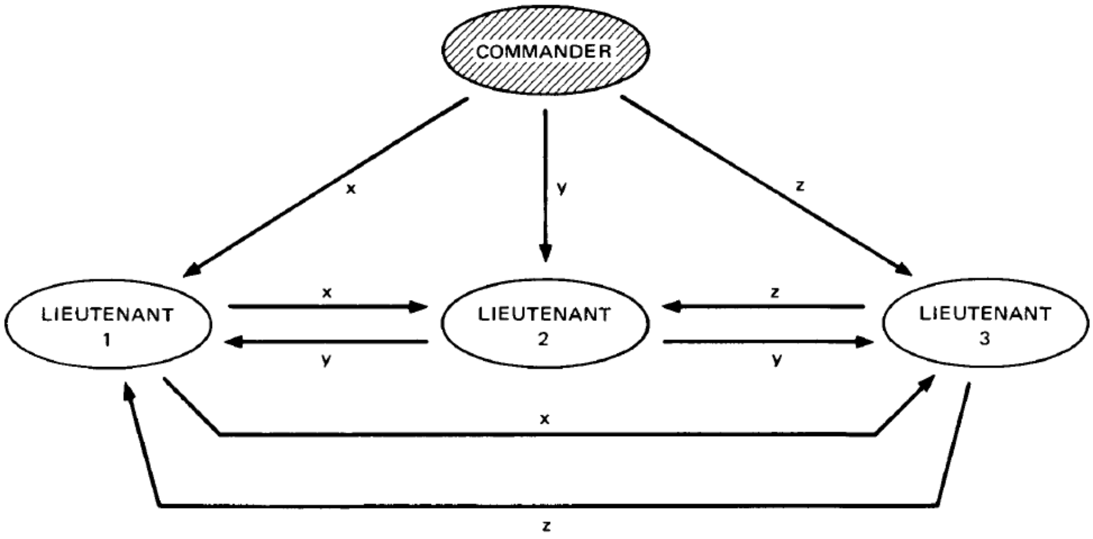

Figure 3.6: A solution with oral messages [6]

This indicated, that there was a demand for a consensus algorithm that could be implemented in asynchronous environments were a lower degree of reliability is required.

3.5 Paxos

The Paxos algorithm was developed in 1990 by Leslie Lamport and defines a number of steps to reach consensus. It is another way to implement fault tolerance in a distributed system then the solutions that were previously disclosed. The algorithm is suitable for asynchronous environments, which are distributed systems where there is no global clock to keep track of time [27]. A few years earlier in 1985, consensus in such environments seemed impossible if there was even one faulty process [28]. The difference with synchronized distributed systems like SIFT was that these required a higher level of reliability. Therefore, each processor had its own clock that was periodically resynchronized to ensure fault tolerance in this system [26]. However, there were also situations in which systems did not need such a level of reliability and for these asynchronous distributed systems Lamport developed the Paxos algorithm.

3.5.1 The solution

In order to describe the solution for reaching consensus in asynchronous distributed systems it is necessary to define the problem first.

The problem is defined as a collection of processes that can propose values. In order to ensure that a single value among the proposed values is chosen there needs to be a consensus algorithm. This consensus algorithm arranges that if no value is proposed then no value should be chosen. If a value has been chosen, the processes should be able to learn the chosen value [29].

The Paxos algorithm consists of two different phases and three classes of agents (proposers, acceptors and learners) who each have their own role. A proposer sends a proposal to a set of acceptors, which the acceptors may accept or decline, when accepted by a majority of acceptors this proposed value is learned by the learners. An instance of the actual protocol starts with phase 1(a) in which the proposer selects a proposal number and sends a prepare request with that number to a majority of acceptors. The request is send to a majority to guarantee that only a single value is chosen by each individual acceptor. In a practical situation this request of a proposer is initiated by a client that wants to issue a command to a central server for example. In the first two steps of figure 3.7 the client sends a request to the proposer who prepares a request with proposal number 1 and sends this towards a majority of acceptors. When an acceptor receives a prepare request it checks if the number of that request is higher than any of the prepare requests to which it has earlier responded. Without this check, an acceptor can choose the same prepare request multiple times and this is not desirable. If the prepare request satisfies this check, the acceptor responds to this request in a standardized way. The acceptor promises that it does not accept any more proposals that have a lower number and it responds with the highest numbered proposal that it has accepted so far (if any). This completes phase 1(b) and is depicted in the third step of figure 3.7 where 1 is the highest number of the received prepare request and Va, Vb and Vc are the individual highest numbered proposals of each acceptor.

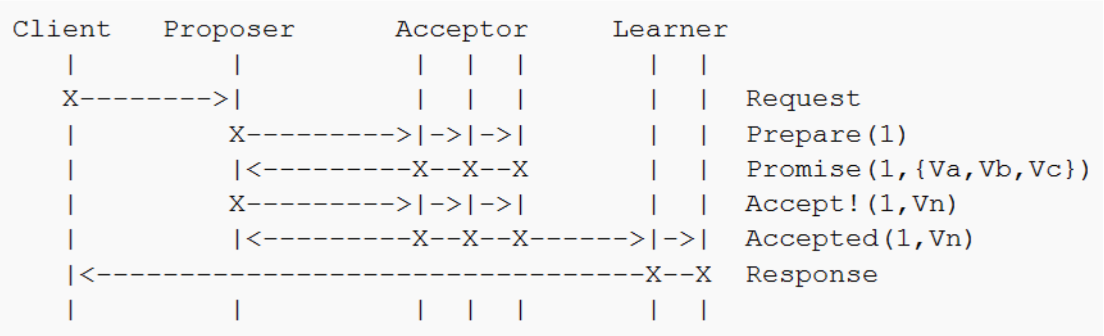

```
Client   Proposer      Acceptor     Learner
   |         |          |  |  |       |  |
   X-------->|          |  |  |       |  |  Request
   |         X--------->|->|->|       |  |  Prepare(1)
   |         |<---------X--X--X       |  |  Promise(1,{Va,Vb,Vc})
   |         X--------->|->|->|       |  |  Accept!(1,Vn)
   |         |<---------X--X--X------>|->|  Accepted(1,Vn)
   |<---------------------------------X--X  Response
   |         |          |  |  |       |  |
```

Figure 3.7: Basic flow of Paxos [30]

Phase 2(a) of the algorithm starts when the proposer receives a response to its prepare requests from a majority of acceptors. Subsequently, it sends an accept request to each of those acceptors for a proposal with the selected proposal number and a specific value. This value is either the value of the highest numbered proposal among the responses or is any value if the responses from the acceptors reported no proposals. Step four of figure 3.7 shows the accept request that is send by the proposer and contains the selected proposal number and the value of the highest numbered proposal among the responses. When an acceptor receives an accept request from the proposer for a promised proposal it needs to check one final condition in phase 2(b). This condition ensures that an acceptor can only accept the proposal if it has not already responded to a prepare request with a higher number. In step five of figure 3.7 is the situation illustrated in which the proposal is accepted by the acceptors. The acceptance includes sending the accepted proposal back to the proposer and to all the learners, otherwise they are not able to learn that a value has been chosen. After that, it is time for the final step in figure 3.7 where the learners send a response to the client to demonstrate that everyone is aware of the latest proposal.

3.5.2 Paxos variants

Since the introduction of Paxos in 1990 there have been several alternative uses and optimizations of the algorithm, which led to a number of variants. One of the first alternative uses of the algorithm was to use it not only in processors but also in disks and was called Disk Paxos. In this way, it was possible to create a reliable distributed system with a network of both processors and disks. The first benefit of this approach is that community disks are cheaper than computers, so it is more economical to use redundant disks for fault tolerance than using redundant computers. Another benefit is that disks do not run programs on application-level, which means that they are less likely to crash than computers [31].

The initial design of the Paxos algorithm was only able to make a distributed system fault tolerant and not Byzantine fault tolerant. This changed since the publication of Fast Byzantine Paxos, which demonstrated a variant of the algorithm that was not only Byzantine fault tolerant but also faster. It required only two communication steps to achieve asynchronous Byzantine consensus in the common case and did not make use of expensive digital signatures. To achieve this consensus in only two steps, Fast Byzantine Paxos needed a higher number of acceptors than in other Byzantine consensus protocols [32].

Another optimization of the protocol is associated with the leader-based characteristic of the Paxos algorithm.

In fact, this characteristic creates a situation in which more work is performed by the leader replica than by the non-leader replicas. When the number of replicas or the load on the system increases the leader replica quickly reaches the limits of one of its resources, which negatively influences the scalability. Therefore it is important that the workload is evenly distributed among all the replicas, to ensure that the leader has only a minimal additional workload. This is achieved by distributing the handling of client communication, disseminating client requests among all replicas and by executing the ordering protocol on id's. These measurements enable S-Paxos to reach a significantly higher throughput rate than the standard Paxos algorithm for any given number of replicas. It creates a situation where there is no more need to make a trade-off between fault tolerance and performance [33].

Besides the use of Paxos in processors and disks the NetPaxos variant introduced the possibility of implementing the protocol also in network devices. Moving the protocol into network devices would create considerable performance benefits for distributed applications. It not only significantly increases the throughput rate of switches but also reduces the latency. A small drawback of this approach is that it requires changes to the underlying switch firmware but these changes are feasible in existing hardware. Overall, it would have a great impact on both the services built with Paxos and the applications that make use of those services [34].

The selection of Paxos variants described above outlines the general evolution of the algorithm over the years. In this period, there have been introduced several other variants of the algorithm but these contained small optimizations or are comparable with the previously treated variants. This brief overview of variants highlighted that the algorithm has enabled multiple opportunities for reaching consensus in asynchronous distributed systems.

3.5.3 Implications

The introduction of the Paxos algorithm changed the field of distributed computing because it solved the problem of reaching consensus in asynchronous distributed systems [27]. This was an important step, because the consensus solutions till that time were inherently expensive to use in situations where a lower degree of reliability was required. The problem of reaching consensus and the BGP were no longer relevant in distributed computing environments, which stimulated the development of large distributed systems. In section 3.6 a selection of those systems will be treated to show in which ways the described techniques are applied in practice.

3.6 Combining these techniques

The main developments of distributed computing have been covered in the previous sections to explain the individual elements in this domain that led to blockchain technology. These elements itself are not new but it is the combination of these elements that makes blockchain technology a groundbreaking development. At this time, there is no standardized definition for the technology because it is such an emerging field. Another problem is that most of the existing definitions are related to the Bitcoin block chain instead of the overall concept of blockchain technology. Therefore, I wrote the following definition of this technology: Blockchain technology is a distributed Byzantine fault tolerant transaction database that contains a chain of data blocks which need to be verified in a standardized way to reach and maintain consensus. The best known example of this technology is the Bitcoin, which will be treated in the next section.

3.7 Bitcoin

Satoshi Nakamoto, which is the inventor of Bitcoin defines it as a peer-to-peer electronic cash system [35]. This would allow online payments to be sent directly from one party to another without going through a financial institution. By empowering the peer-to-peer network to verify transactions itself, there is no more need for the traditional trusted third parties. It is a system that is based on cryptographic proof instead of relying on trust to make payments over a communication channel like the Internet. In order to achieve this, Bitcoin uses a number of techniques that are based on the distributed computing developments described in the previous sections.

3.7.1 Underlying techniques

Bitcoin consists of a combination of techniques that were developed in the 70's, 80's and 90's. First of all, the electronic coins in the system are chains of digital signatures which make use of the public and private key concept to sign the transactions [36]. Those chains must be verified in each new transaction to check if the history of ownership is correct. The problem with this technique is that cases in which the coins are double-spend cannot be verified. Therefore, there must be a way for the payee to know that the previous owners did not sign any earlier transactions [35]. To identify these potential double-spend transactions it is necessary to be aware of all the transactions. Without a trusted third party, there needs to be a list of transactions in the order in which they were received.

To ensure that the transactions are ordered in the right way, an important component of Bitcoin is a timestamp server.

This server takes a hash of a block of items to be timestamped and publishes this hash to every node in the network. The timestamp proves that the hashed data existed at the time the stamp contains, otherwise it could not get into the generated hash [35]. In order to implement this timestamp server subsequently on a peer-to-peer basis the use of a proof-of-work system is required.

This system involves scanning for a hash that begins with a specific number of zero bits. When this hash is found, a node can send the solution together with the block to other nodes so it can be verified and added to the chain. An advantage of this proof-of-work algorithm is that the solution can be verified by executing a single hash, while finding the solution can take a lot of time. This is related to the number of zero bits that is required for a solution, when this amount increases it exponentially affects the average time that is needed to find the solution [35]. Once enough Central Processing Unit (CPU) effort has been expanded to satisfy the proof-of-work the block cannot be changed without redoing the work. The work to change a block becomes greater when later blocks are chained after it because this includes redoing all the work of those blocks. This distributed computation system or proof-of-work algorithm is considered to be the key innovation of Bitcoin [37]. It not only enables the Bitcoin network to arrive at consensus about the state of transactions but also offers protection against double-spending attacks.

3.7.2 The network flow

With the combination of a timestamp server and the proof-of-work techniques the network is able to run. The enumeration below describes this general flow of the Bitcoin network [35]:

1. New transactions are broadcast to all nodes.
2. Each node collects new transactions into a block.
3. Each node works on finding a proof-of-work solution for its block.
4. When a node finds a proof-of-work, it broadcasts the block to all nodes.
5. Nodes accept the block only if all transactions in it are valid.
6. Nodes express their acceptance of the block by working on the next one.

By repeating this process the Bitcoin block chain is formed and contains an overview of confirmed transactions in the network.

3.7.3 Comparison

When comparing the underlying techniques of Bitcoin with the principles underlying blockchain technology there is quite some overlap.

Unfortunately, Nakamoto's specification of the network mainly focuses on the verification components used in Bitcoin. The fault tolerance and consensus characteristics are hardly described, while these also play an important role in the proper functioning of the network. Maybe, the emphasis lies on the verification aspect because Bitcoin is based on the self-interest principle of individual users. This explains why the system needs such an extensive mechanism like the proof-of-work algorithm to ensure the correctness of each transaction. Despite its brief specification, Bitcoin remains the first example that has implemented the main principles behind blockchain technology.

A remarkable finding in Bitcoin is the emphasis on verification while I believe that in most situations consensus is a much stronger principle to use in actual blockchain implementations. To give an example, the energy consumption that is necessary to find the solution for a new block in Bitcoin is not desirable in most situations. This high amount of energy is caused through the extensive verification mechanism that is used. For virtual currencies there are several arguments to support the choice for such a strong verification mechanism but this is only one specific purpose of blockchain technology. When looking at other purposes of the technology, a clearly defined consensus mechanism can also achieve the right amount of security to guarantee safe blockchains. This is related to a number of changed variables in many implementations like the overall openness, the required scalability and the sensitivity of the recorded data.

3.8 Different examples

One of the misunderstandings about Bitcoin is that it is the only possible representation of blockchain technology. This assumption leads to the idea that every other blockchain implementation must use exactly the same techniques as Bitcoin. Therefore, this section will describe some other examples to make it clear that there are also other implementations possible with the principles underlying blockchain technology.

3.8.1 Google Megastore

Google Megastore is a storage system developed to meet the requirements of today's interactive online services [38]. It handles more than three billion write and twenty billion read transactions daily and stores nearly a petabyte of primary data across several global datacenters. For such a large scale distributed system it is important that the data is consistent and easily available. To ensure the consistency of the data, Megastore has implemented the Paxos algorithm as a way to reach consensus.

The introduction of a few optimizations and innovations made the algorithm suitable for Google's system. One of the differences from the description in section 3.5 is the introduction of a service called the coordinator. This is a server that tracks a set of entity groups for which its associated replica has observed all Paxos writes [38]. For each log position a distinguished replica is chosen as leader alongside the preceding log position's consensus value. By making use of these coordinators the availability of the data improves because it allows fast local reads from any datacenter. This is possible with the help of a write algorithm, which makes sure that the state of the coordinator remains conservative. It improves the overall stability of the coordinators, which results in a lower unavailability percentage in case of replica failures.

Comparison

When comparing the system behind Google's Megastore with blockchain technology there are a number of similarities. In the specification of this distributed system is extensively described how consensus is reached in a fault tolerant way. These important principles of blockchain technology are fulfilled by using an implementation of Paxos, which is a widely known consensus algorithm. The terms transactions and blocks are also mentioned a few times but unfortunately not in detail.

Verification wise, there is in Google Megastore no need to use the extensive proof-of-work algorithm that is implemented by Bitcoin. This is due to the difference in openness between the two networks, which requires another verification form to provide the necessary safety. On the other hand, the lighter verification aspect in Megastore increases the need for an extensive consensus mechanism. In Bitcoin this is mainly achieved by the strict verification procedure of mining blocks. The differences in these implementations show that there are multiple combinations possible with the principles underlying blockchain technology. It will depend on the specific situation in which ratio these principles are related to each other.

3.8.2 Theoretical examples

The examples thus far have demonstrated mainly the practical side of blockchain technology but there are also several other examples with a more theoretical character. According to Melanie Swan, Bitcoin is a blockchain 1.0 implementation that includes currency and the deployment of cryptocurrencies in applications related to cash [39]. Besides this category, she also identified two other categories that are consecutively named blockchain 2.0 and 3.0.

Blockchain 2.0 refers to contracts and implies the entire slate of economic, market, and financial applications on the blockchain that are more extensive than simple cash transactions. Currently, there are several developments in this category with smart contracts on a popular blockchain platform called Ethereum. The last category is blockchain 3.0 and this entails blockchain applications beyond currency, finance, and markets [39]. Think of areas like government, health, science, literacy, culture, and art, where multiple startups are working on useful concepts for blockchain technology.

Exchanges

A quite concrete example of developments in the blockchain 2.0 area is related to the blockchain initiatives of different exchanges worldwide. Different exchanges are exploring the opportunities of blockchain technology but so far the NASDAQ and Deutsche Börse Group are the ones that already developed a platform and prototypes. Unfortunately, because these developments are so recent there is not a lot of information available besides some news articles. NASDAQ's Linq is considered to be the first platform from an established financial services firm to demonstrate how asset trading could be managed digitally by using blockchain technology. The biggest benefit of this platform is that it has the ability to remove the need for pen-and-paper or spreadsheet-based record-keeping in the private shares trading market. It provides a form of immutable recordkeeping and also a chain of custody for users [40]. The prototypes of Deutsche Börse Group focuses more on a process called corporate proxy voting. This entails the participation of shareholders in a company's annual shareholders meeting who can exercise their vote on the matters to be considered, without physically attending the meeting. It is expected that this will bring significant improvements in the areas of transparency, accuracy and consistency [41].

Industry

The examples in the blockchain 3.0 area have a more abstract character because there are several ongoing developments but these mainly take place on a conceptual level. This is understandable, because blockchain technology just entered the 2.0 phase and the first concrete examples besides cryptocurrencies are starting to take shape. Therefore, the following example describes a hypothetical situation in the industry sector.

Imagine a company that has multiple factories, and each factory has their own collection of machines that individually generate a set of data. Besides monitoring the production process, this data is used by the company to get an overview of broken parts or spare ones that need to be ordered. Once in a while, this results in a set of orders for the different part suppliers.

The problem is that this process is time consuming and causes delays because ordered parts are out of stock on a regular basis. A possible solution for this problem is a blockchain implementation that operates in a private environment with the factories of the company and a number of selected suppliers. In this case, the ledger will be a collection of the agreed transactions from machines that indicate which parts are broken or spare ones that need to be ordered. To guarantee availability of this data, each factory site constantly synchronizes all the transactions of the other factories to ensure multiple complete copies of the ledger. Correctness of the data is achieved by making the implementation fault and Byzantine fault tolerant, which enables the ability to deal with faulty and conflicting components. In order to make sure that all nodes agree with the transactions in the ledger, there is also an algorithm in place that defines the steps to reach consensus. The final element is that the transactions in the ledger must be encrypted to assure that competing suppliers cannot see the orders directed to each other. This would be a distortion of competition and gives some suppliers valuable information, which leads to undesirable situations.

Healthcare

The situation outlined above describes a potential implementation of blockchain technology in the industry sector. It is now time to determine if a similar hypothetical example can be applied in the healthcare sector.

Imagine an elderly person that is capable of living on its own with the help of several medical related devices. Each of these devices generates a set of data which is used by different caregivers to monitor the health of the client. Currently, clients are not capable to easily gain insight in the data that is collected by these devices and which parties make use of it. This can be changed by introducing a blockchain implementation that operates in a private context with an involved client, a number of devices and a selection of caregivers. There are multiple parties that have an instance of the ledger in this context, which is a collection of the agreed transactions from the medical related devices and those from the caregivers. To guarantee the availability of this data, each caregiver and a number of selected devices are constantly synchronizing transactions to ensure multiple complete copies of the ledger. The distinction is that the shared ledger on the selected devices is a subset of the shared ledger that the caregivers maintain. This is related to the limited storage capacity of the devices, therefore those ledgers only contain an overview of operations performed on client related transactions. Correctness of the data is achieved by making the implementation fault and Byzantine fault tolerant, which enables the ability to deal with faulty and conflicting components. In order to make sure that the two groups of nodes agree with the transactions in the ledger, there is a need for different consensus levels.

The first level is related to the devices that have to reach consensus with each other before they can write transactions in the ledger. Secondly, the caregivers have to reach consensus with each other before any of them can perform operations on the transactions in the ledger. The final element is that the transactions in the ledger must be encrypted to ensure that anyone outside the described context is unable to make sense of the data. If not, it will be an invasion of the client its privacy and this is in conflict with the applicable legislation.

3.9 The need for blockchain technology

There are two general trends that contributed to the development of blockchain technology. One trend is related to the growing number of devices, connecting those to the IoT network and creating CPS and eventually ULSS which contain vast amounts of information. To guarantee liveness and safety in the IoT environment, the field of distributed computing had to come up with several solutions to deal with this trend. This were techniques to make a system fault tolerant, solving the BGP and reaching consensus with Paxos. The combination of these two trends ensured that the developments in both fields continued but that the total amount of information in distributed systems grew significantly over the years. Therefore, parts of the information in these systems needs to be recorded and reviewed by a so called TTP in order to be useful for different actors. The use of such a TTP is devious and inefficient in several situations and this is where blockchain technology can offer a solution. It has the capabilities to replace a TTP because it offers the same functionalities and removes the human factor, which increases the integrity of the overall process.

3.10 Recap

This chapter started with describing the developments that led to the emergence of the field blockchain technology. Defining these made the principles clear that formed this technology and the functionalities it can offer. The chapter ended with a number of practical examples, which helped to visualize the different implementations that are possible with this technology. In the next chapter it is time to examine if the identified principles behind blockchain technology are also known and supported in the field. This will be done by performing a number of expert-interviews in multiple domains in which these principles and several related use cases are discussed.

Chapter 4

Results

In chapter 2 and 3 the emergence of the IoT and the underlying principles of blockchain technology were discussed. Based on these findings, several expert-interviews in multiple domains have been conducted to verify the identified underlying principles and to discuss several related use cases. The results of these interviews will be treated in this chapter.

4.1 Expert-interviews

The main goal of the interviews was to test if the experts are familiar with the underlying principles of blockchain technology that were found during the critical literature review, see section 3.1. By discussing this set of elements, it will be possible to identify the general knowledge level of the respondents. This will form a key indicator for determining the current state of blockchain technology and gives insight in potential problem areas.

Every expert-interview had a predefined sequence of questions that started with a few introductory ones about the general concepts of Bitcoin and blockchain technology. When the distinction between these concepts was clear, the awareness of the underlying principles was tested in the second category. After that, the interviewees were asked to choose the most potential use case out of a small selection. The last category consisted of some future related questions about the development of blockchain technology. An overview of the interview questions can be found in appendix A.

4.1.1 Prerequisites

In order to perform a proper baseline measurement (section 1.2) of the general concepts, the underlying principles and related use cases it was important to select experts from different domains.

The complete list of interviewees is depicted in table 4.1 and contains experts from the healthcare, logistics and financial sector. Two research institutions, a ministry and a university served as representatives for the public sector. This list was prepared with the help of my company coach Ben van Lier and the interviewees were either verbally or electronically approached. Another selection criteria for the participants was their common search for alternative implementations of blockchain technology. This is demonstrated by the fact that they have performed experiments in a blockchain lab, set up a research group or joined a consortium.

Name	**Function & Company**
Oskar van Deventer	Senior Scientist Media Networking at TNO
Patrick van Beers	Senior Director Digital Platform Solutions at Philips Research
Mark Hennessy	Cloud Services Architect at Philips Research
Dr. Sunil Choenni	Head of Statistical Data and Policy Analysis at the Ministry of Security and Justice
Henk-Jan Vink	Director Networked Information at TNO
Dr. Eric Pauwels	Group Leader Intelligent Systems at Centrum Wiskunde en Informatica (CWI)
Martijn Siebrand	Program Manager Supply Chain Finance at TKI Dinalog
Prof. dr. Ron Berndsen	Head of Market Infrastructures Policy at De Nederlansche Bank (DNB)
Mark Buitenhek	Global Head Transaction Services at ING
Prof. dr. Jaap van den Herik	Director Leiden Centre of Data Science at Leiden University
Prof. dr. Bas Edixhoven	Professor of Geometry at Leiden University
Raymond van Bommel	PhD candidate at Leiden University

Table 4.1: List of interviewees

An essential part of the interviews was to test if the experts were aware and knowledgeable of the identified underlying principles of blockchain technology. By transforming the general concepts, the underlying principles and the related use cases into figures it became easier to explain these components during the interviews. This enabled the possibility to identify the different levels of expertise in a fast and concrete manner.

Because of the exploratory and inductive nature of this study, it is advised to use semi-structured interviews in order to stimulate the development of grounded theory [8]. This explains the choice for multiple overarching categories in the questionnaire, which were treated consecutively to stimulate a generic order during each interview. Those categories also enhanced the coding process because it provided a basic outline to start with the categorization process. On top of that, it prevented a lack of coherence to make sure that a consistent story could be derived after performing and comparing the various interviews. The results of the interviews will be treated according to these overarching categories.

4.1.2 Usefulness

In total there are nine interviews conducted during this study of which seven were suitable to apply the described coding strategy in section 1.3. By applying the initial and focused coding stages it became possible to make a number of categories and group multiple codes into these categories. Subsequently, these categories were used to derive results from the interviewees their perception about the general concepts, the underlying principles and a couple of related use cases. The other two interviews are more considered to be general discussions about these topics than actual interviews. Nevertheless, they had a valuable contribution in the determination of the overall perception around blockchain technology in several domains.

4.2 The general concepts

The interviews started with three introductory questions about the familiarity of the interviewees with Bitcoin and blockchain technology. These questions helped to identify if the knowledge levels of the interviewees were sufficient enough to discuss the underlying principles and related use cases. As expected, all respondents were aware of these two concepts and understood the difference between them. The concept of the TTP, earlier described in section 3.9, was further clarified with the help of figure 4.1. In this situation the traditional TTP contains the ledger, which is a collection of all the actions that the nodes performed in a network.

An important constraint is that a node can only perform actions in such a network with the approval of the TTP. Therefore, the TTP creates a single point of failure in the network because it is the only entity that is responsible for making all the decisions.

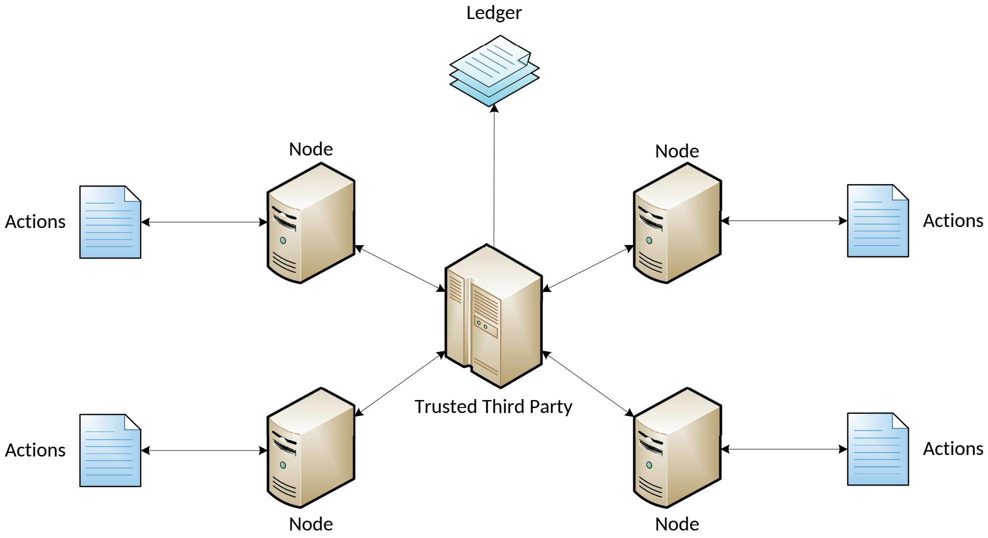

Figure 4.1: Trusted third party overview

4.2.1 Blockchain technology

When comparing the TTP situation with blockchain technology, there are some significant distinctions. First of all, there is a removal of the TTP in blockchain technology implementations in order to create a self-regulating network. Such a network is nothing more than a collection of nodes that are autonomously capable of performing actions and making decisions. This implies that the nodes in a self-regulating network need to contain a shared ledger, which is more than just a collection of all the actions that the nodes performed in a network. It also includes all the decisions that the nodes have accepted after consensus is reached. In this way, each node is capable to perform its individual actions and make decisions about the actions of other nodes without the interference of a TTP as shown in figure 4.2. To make sure that a self-regulating network is able to operate safely there are a number of underlying principles used, which will be addressed in the next section. A final remark is that a node in a self-regulating network does not necessarily have to be a desktop computer or a server. It can also be a laptop, mobile phone or tablet to name just a few examples.

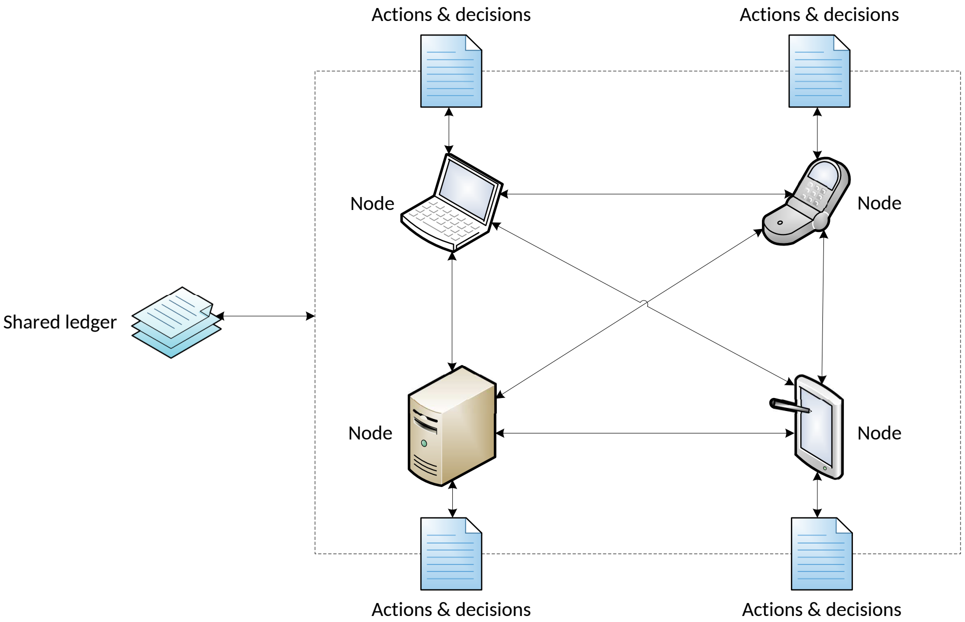

Figure 4.2: Blockchain technology overview

4.3 The underlying principles

In the sections 3.3, 3.4 and 3.5 a theoretical explanation of the underlying principles of blockchain technology was given. During the interviews, the awareness and knowledge levels of the experts were tested in another way. By using various figures that represented those theoretical explanations it became easier to identify those different levels. In order to measure the knowledge levels of the interviewees, each answer was assigned to one of the following three categories: unknown, partially known and known. The answer was identified as unknown if one of the underlying principles was a new term for a respondent. It was considered to be partially known if the expert was slightly aware of an underlying principle but could not explain it in detail. An answer was classified as known if the interviewee was fully aware of the underlying principle and could explain it in detail.

4.3.1 Fault tolerance

The first underlying principle that was tested is illustrated in figure 4.3 and is related to the term fault tolerance, as described in section 3.3. A network is considered to be fault tolerant if it is still able to perform its tasks when one or more components fail. This can be achieved by using techniques such as SMR (subsection 3.3.1) and time stamps (subsection 3.3.2), which enable a network to replicate the states of nodes and determine the order of events.

The respondents were all aware of the underlying principle fault tolerance and could explain the purpose of it, which is shown in table 4.2. The unidentified column represents the percentage of interviews that was unsuitable for the earlier described coding strategy. To conclude, the main task of fault tolerance is dealing with hardware related failures but there is also a form that is able to deal with software related failures. This form is called Byzantine fault tolerant and offers a solution for the BGP, see section 3.4.

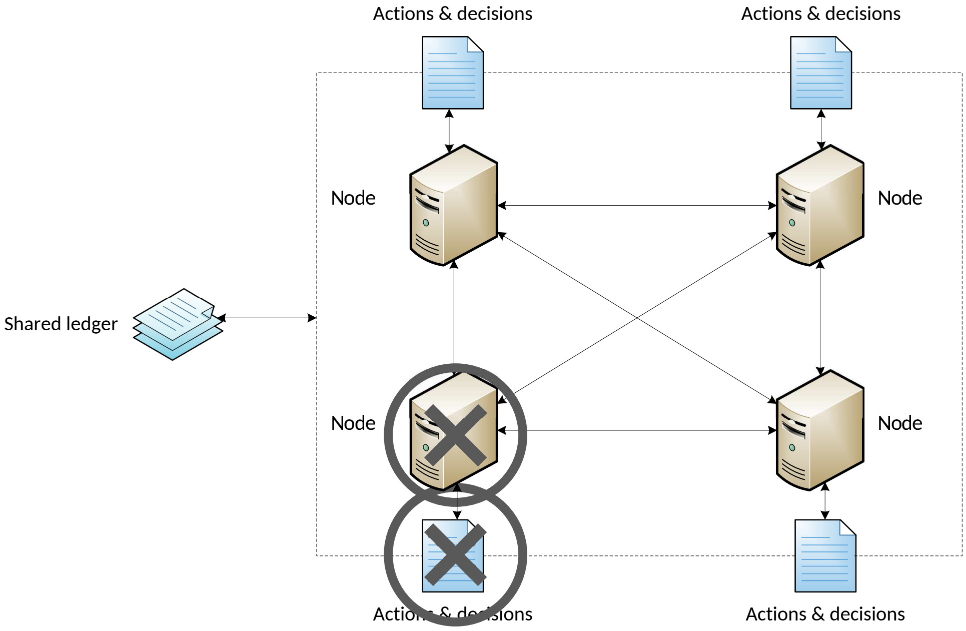

Figure 4.3: Underlying principles - Fault tolerance

Underlying principle	Unknown	Partially known	Known	Unidentified
Fault tolerance	0%	0%	77,8%	22,2%

Table 4.2: Awareness of fault tolerance

4.3.2 Byzantine fault tolerant

Byzantine fault tolerant is the second underlying principle that was tested during the interviews. A network is considered to be Byzantine fault tolerant if it is still able to perform its tasks when malicious components give conflicting values, as stated in section 3.4. This can be done through the exchange of signed messages (subsection 3.4.2) between the different nodes in the network.

When each node compares its messages with every other node in the network they can collectively determine which node is sending conflicting values, this process is depicted in figure 4.4. It starts with the malicious node that sends the conflicting values X, Y and Z to the other nodes in the network (1, X for example). To test if there is a malicious node in the network the other nodes send their received value with the corresponding sender to the other participants (2, X:1 for example). When the nodes compare these signed messages with each other they can identify that node 1 is malicious because its signature appears on conflicting values. In this way, the network is not only able to deal with hardware related failures but also with software related failures.

The software related variant Byzantine fault tolerant was a new term for four of the nine respondents. Three interviewees had heard of the term before and only two of them were familiar with this underlying principle as table 4.3 illustrates. The unidentified column represents the percentage of interviews that was unsuitable for the earlier described coding strategy.

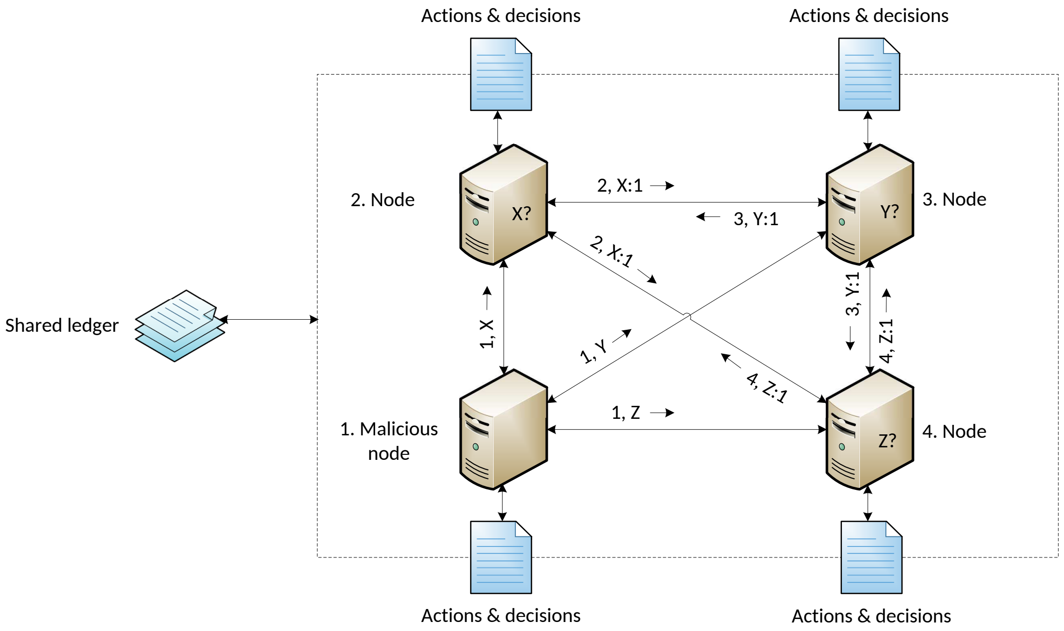

Figure 4.4: Underlying principles - Byzantine fault tolerant

4.3.3 Consensus algorithm

The third underlying principle that was tested during the interviews is a consensus algorithm like Paxos, which is actually a voting protocol.

Underlying principle	Unknown	Partially known	Known	Unidentified
Byzantine fault tolerant	44,4%	11,1%	22,2%	22,2%

Table 4.3: Awareness of Byzantine fault tolerant

This algorithm is developed to reach consensus in the stream of values that nodes sent to each other in a network, as described in section 3.5. It is important that for each iteration of the consensus flow only a single uniform value is chosen. Otherwise, there can arise a situation in which multiple values are chosen during a single iteration of the consensus flow. This will lead to different values in the individual ledgers of certain nodes, which harms the integrity of the overall network. An example of a consensus flow where only a single value is introduced to the overall network is shown in figure 4.5. This figure is based on the consensus flow in subsection 3.5.1 and uses a question mark to propose a value and an exclamation mark to promise or accept a request. It starts with the proposer, who proposes the value X with proposal number 1 to the acceptors (X, 1?). In the next step, the acceptors respond to the proposer with a promise to accept this request (X, 1?!). After that, the proposer sends out an accept request to the acceptors (X, 1?!!). The acceptors respond with an accepted message (X, 1), which contains the values of the original request from the proposer. This accepted message (X, 1) is also send to the learner who did not participate as an acceptor in the consensus procedure. Finally, the learner sends this accepted message (X, 1) as a response to the proposer to notify him that he is also aware of the latest proposal. For clarification purposes, the different steps of the Paxos flow are once more described in table 4.4.

Step	Actor	Action
1.	Proposer	Proposes a value with proposal number to the acceptors.
2.	Acceptor	Responds with a promise to accept this request.
3.	Proposer	Sends out an accept request to the acceptors.
4.	Acceptor	Responds with an accepted message.
5.	Acceptor	Sends this accepted message to the learner.
6.	Learner	Sends the received accepted message to the proposer.

Table 4.4: The flow of Paxos

As with the previous principle, a consensus algorithm was a new term for four of the nine respondents. Besides that, there were three experts who globally knew the meaning of such an algorithm.

In contrast to the earlier results there was not a single interviewee who could explain this underlying principle in detail, which is depicted in table 4.5.

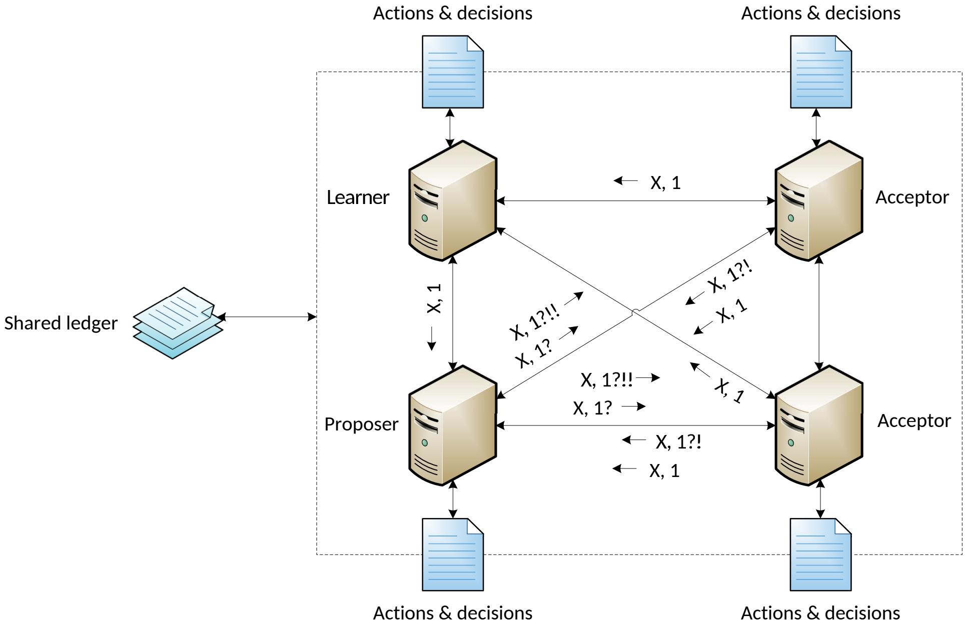

Figure 4.5: Underlying principles - Consensus algorithm

Underlying principle	Unknown	Partially known	Known	Unidentified
Consensus algorithm	44,4%	33,3%	0%	22,2%

Table 4.5: Awareness of consensus algorithm

4.3.4 The missing link

With the underlying principles thus far it is possible to create a network that is fault tolerant, Byzantine fault tolerant, and is able to reach consensus. Nevertheless, there is also a mechanism needed that specifies the conditions under which the proposed values in a network are legit. Without this mechanism it is impossible to define the values that nodes will accept after an iteration of the consensus flow. This means that each proposed value in the network will be accepted and that the ledger becomes incredibly large and inconsistent. Therefore, there is an additional principle needed that is comparable with the preliminary protocol described in [27]. This protocol contained a set of specific constraints that guaranteed consistency and allowed progress in the network.

In other words, it specified the conditions under which the proposed values were legit. According to my view, this underlying principle can be seen as an "overlaying protocol" because it is an additional element that regulates the flow of values in a distributed network. Where the first three underlying principles are generally applicable, the implementation of the overlaying protocol will be context dependent. This is related to the unique value conditions that apply to every context. In figure 4.6 is a situation illustrated in which a node tries to propose a number instead of a letter to the other nodes in the network (1, 5). In this specific context, a number is specified as a non-legitimate transaction and needs to be discarded. The nodes are able to detect this constraint by checking the overlaying protocol (5?) and will therefore refuse the proposed value. This prevents the start of an unnecessary consensus flow in the network to reach agreement about a non-legitimate transaction.

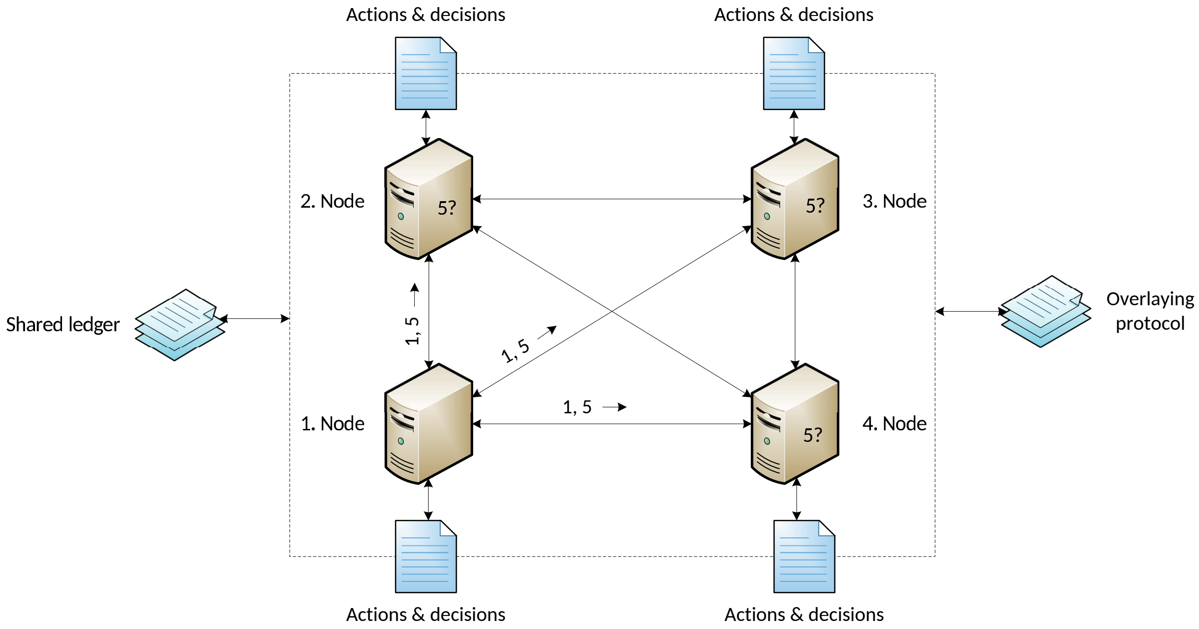

Figure 4.6: Underlying principles - Overlaying protocol

The last underlying principle is not extensively described in the literature, which is demonstrated by the fact that it was a new term for five of the nine respondents. There were two interviewees who recognized the role of the overlaying protocol and referred to other examples in which a similar protocol is used. On the other hand, none of the experts considered the overlaying protocol as a fundamental element of blockchain technology before the interviews. An overview of the awareness with the overlaying protocol is depicted in table 4.6.

Underlying principle	Unknown	Partially known	Known	Unidentified
Overlaying protocol	55,5%	22,2%	0%	22,2%

Table 4.6: Awareness of overlaying protocol

4.3.5 Overview

When combining the four underlying principles of the previous subsections there is a self-regulating network created that has several properties. Fault tolerance and Byzantine fault tolerant are the first two and guarantee that a network is able to deal with failing and malicious components. Besides that, there needs to be a consensus algorithm in place that enables the network to reach consensus. The overlaying protocol is the last element and defines under which conditions a value is legit. For the sake of completeness, the overview of blockchain technology with the four underlying principles is shown in figure 4.7.

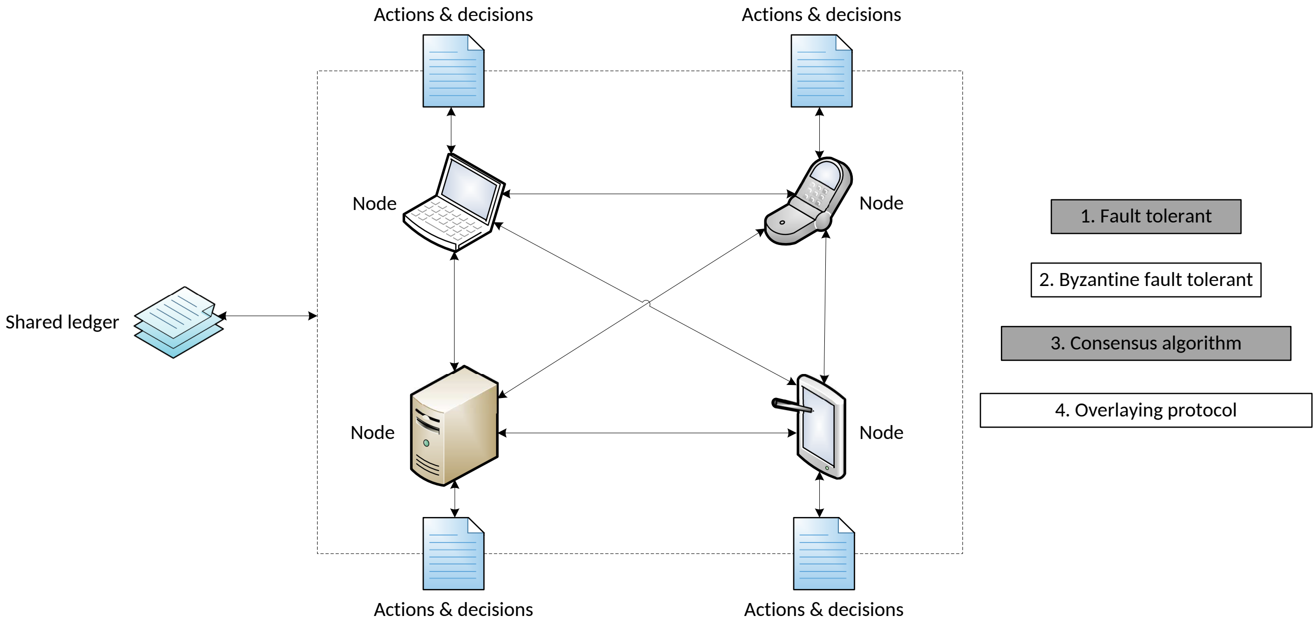

Figure 4.7: Underlying principles - Overview

Of these four underlying principles, only fault tolerance needs to be largely implemented in a hardware related way to deal with the failure of components. The three other principles are purely implemented in a piece of software that consists of algorithms and protocols, which forms the core of a blockchain.

Whereby the overlaying protocol is the differentiating factor because it specifies the conditions under which the proposed values are legit. The remaining principles are concerned with creating the foundation that is necessary to let a blockchain function properly. Together, I believe that these four underlying principles represent the fundamentals of blockchain technology.

Lack of knowledge

In table 4.7 is the awareness of the underlying principles illustrated. The unidentified column represents the percentage of interviews that was unsuitable for the earlier described coding strategy. Remarkable is the fact that only fault tolerance is known to all respondents that were interviewed. This indicates that there is a lack of knowledge about multiple core elements of blockchain technology among the experts. It is important to keep this gap in mind, because it could influence the interviewees their opinion about the potential of the related use cases in the next section.

Underlying principle	**Unknown**	**Partially known**	**Known**	**Unidentified**
Fault tolerance	0%	0%	77,8%	22,2%
Byzantine fault tolerant	44,4%	11,1%	22,2%	22,2%
Consensus algorithm	44,4%	33,3%	0%	22,2%
Overlaying protocol	55,5%	22,2%	0%	22,2%

Table 4.7: Awareness of the underlying principles

4.4 Related use cases

After reviewing the general concepts and explaining the underlying principles it was time to discuss a number of related use cases during the interviews. These contain descriptions of concepts in which the earlier described underlying principles could have additional value. In section 2.5 it became clear that blockchain technology has the ability to safely store and exchange data between nodes in a distributed network. This characteristic enables the capability to connect large numbers of devices with each other in several domains without the need for a TTP. Therefore, the use cases in this section describe hypothetical situations in which new collections of devices are connected with each other. Once connected, these devices will be able to store and exchange data based on the underlying principles of blockchain technology. This will not only provide more transparency but it also improves the efficiency and effectiveness of various processes.

By asking the respondents subsequently which of those use cases had the most potential, it became possible to estimate in which direction blockchain technology is headed. There are four use cases in total, which are related to the logistics, healthcare and industrial sector. The financial sector was deliberately not part of the sample in order to determine the potential of blockchain technology in other sectors. By introducing these use cases, the interviewees are also able to see the practical side of this new technology and the possibilities that it can offer.

4.4.1 Smart container

The first use case that was discussed during the interviews is related to the concept of a smart container. This will be an autonomous container that is able to keep track of its goods and flow of money. By storing this information in a shared ledger that is accessible for different stakeholders in the transportation process, the overall transparency will be significantly improved. An outline for the smart container use case is depicted in figure 4.8. There are two types of context in this use case, an internal and an external one. The external context is displayed in the top rectangle and consists of multiple smart containers that are transported by a smart ship or truck. These entities write a limited amount of data in the shared ledger after they reach consensus. The internal context is represented by the bottom rectangle and is a collection of transportation companies. They collect and use the data that is generated by the containers and the transportation vehicles. In this way, the shared ledger becomes an overview of the data that the containers and transportation vehicles share with the transportation companies throughout the entire process. Both of these contexts are fault and Byzantine fault tolerant and are part of the same self-regulating network. The conditions under which these proposed values are legit will be specified in the overlaying protocol. As not all entities in the external context are equipped with enough capacity to store the entire ledger, they only have a subset that contains their own actions and decisions. Capacity is not an issue in the internal context, which means that these nodes are capable of storing the whole shared ledger. These instances of the ledger can be consulted by the smart containers, smart vehicles and transport companies in the network. To guarantee privacy, each entity can only read the subset of the ledger that is relevant for performing its tasks. This can be achieved by implementing a layer of encryption, which prevents the invalid use of data.

For four of the nine respondents, the smart container did not have the most potential of the different use cases. Despite this result, there were three of the experts who believed that the smart container had the most potential. Two of them mentioned that they considered the logistics sector as the most appropriate one to start experimenting in. This is due to the lower related risks, which reduces the threshold to start experimenting in this sector compared to the other use cases.

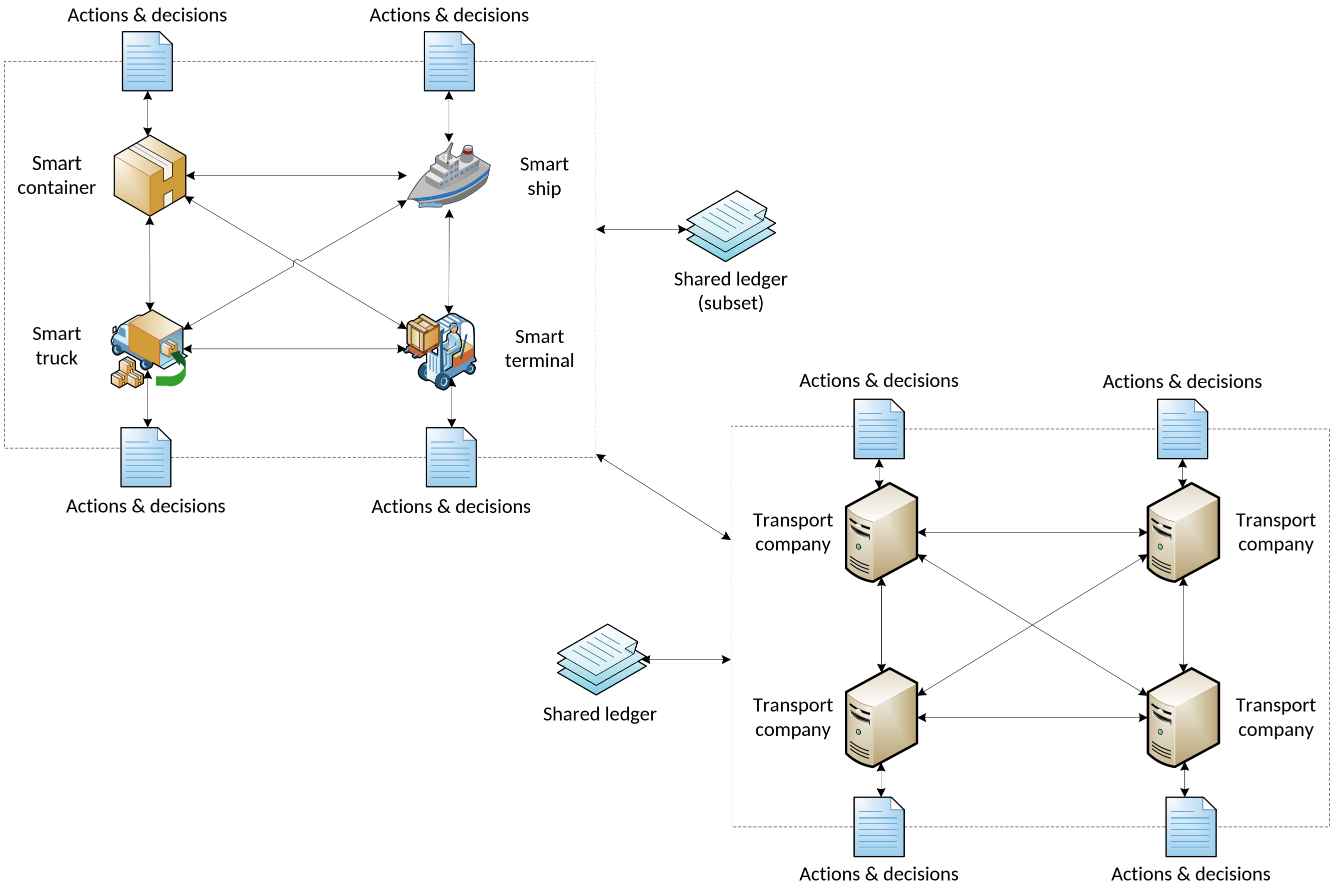

Figure 4.8: Related use cases - Smart container

Another argument was that the smart container use case could potentially generate the most business value. An overview of the smart container its potential is shown in table 4.8. The unidentified column represents the percentage of interviews that was unsuitable for the earlier described coding strategy.

Use case	Less potential	Most potential	Unidentified
Smart container (Logistics sector)	44,4%	33,3%	22,2%

Table 4.8: Potential of the smart container use case

4.4.2 Cure

The second use case in the interviews was associated to a concept in the healthcare sector. With this concept, the aim is to combine the various data collections of hospitals in order to make these easier accessible for doctors. This could be done by connecting devices such as Magnetic Resonance Imaging (MRI) or Computerized Axial Tomography (CAT) scanners with each other. Together, these devices will form a reference framework based on previous scans that enables doctors to quickly determine the right diagnose and to improve the overall success rate of treatments. The solution for this use case is illustrated in figure 4.9 and is comparable with the smart container outline in the previous section. In the top rectangle is the external context displayed, which consists of multiple MRI and CAT scanners that write a limited amount of data in the shared ledger after they reach consensus. As not all entities in this context are equipped with enough capacity to store the entire ledger, they only have a subset that contains their own actions and decisions. The internal context is represented by the bottom rectangle and is a collection of hospitals that store the whole shared ledger and make use of this data. With the help of this collected data, these hospitals are able to perform actions and take decisions. In this way, the shared ledger will become an overview of the decisions that the hospitals have taken based on MRI and CAT scanner related data. By making this easily accessible in the self-regulating network, the doctors have an extra tool that could help them with improving the overall success rate of treatments. The conditions in the overlaying protocol will determine which proposed values are legit to store. Again, it is important that the context is fault and Byzantine fault tolerant in order to be able to deal with faulty and conflicting components. Besides that, privacy is also an aspect that needs to be guaranteed in the healthcare sector. Each entity can therefore only read the subset of the ledger that is relevant for performing its tasks. This can be achieved by implementing a layer of encryption, which prevents the invalid use of data.

For six of the nine respondents, the cure use case did not have the most potential of the four use cases. The difficulty to comply with the large number of laws and regulations in the healthcare sector was the main argument for this opinion. In addition, there is no room for errors because of the sensitive character of the data compared to other sectors. Only one interviewee believed that the cure use case had indeed the most potential. This was related to some comparable existing solutions in the healthcare sector that could have additional value by making use of blockchain technology. An overview of the cure use case its potential is depicted in table 4.9. The unidentified column represents the percentage of interviews that was unsuitable for the earlier described coding strategy.

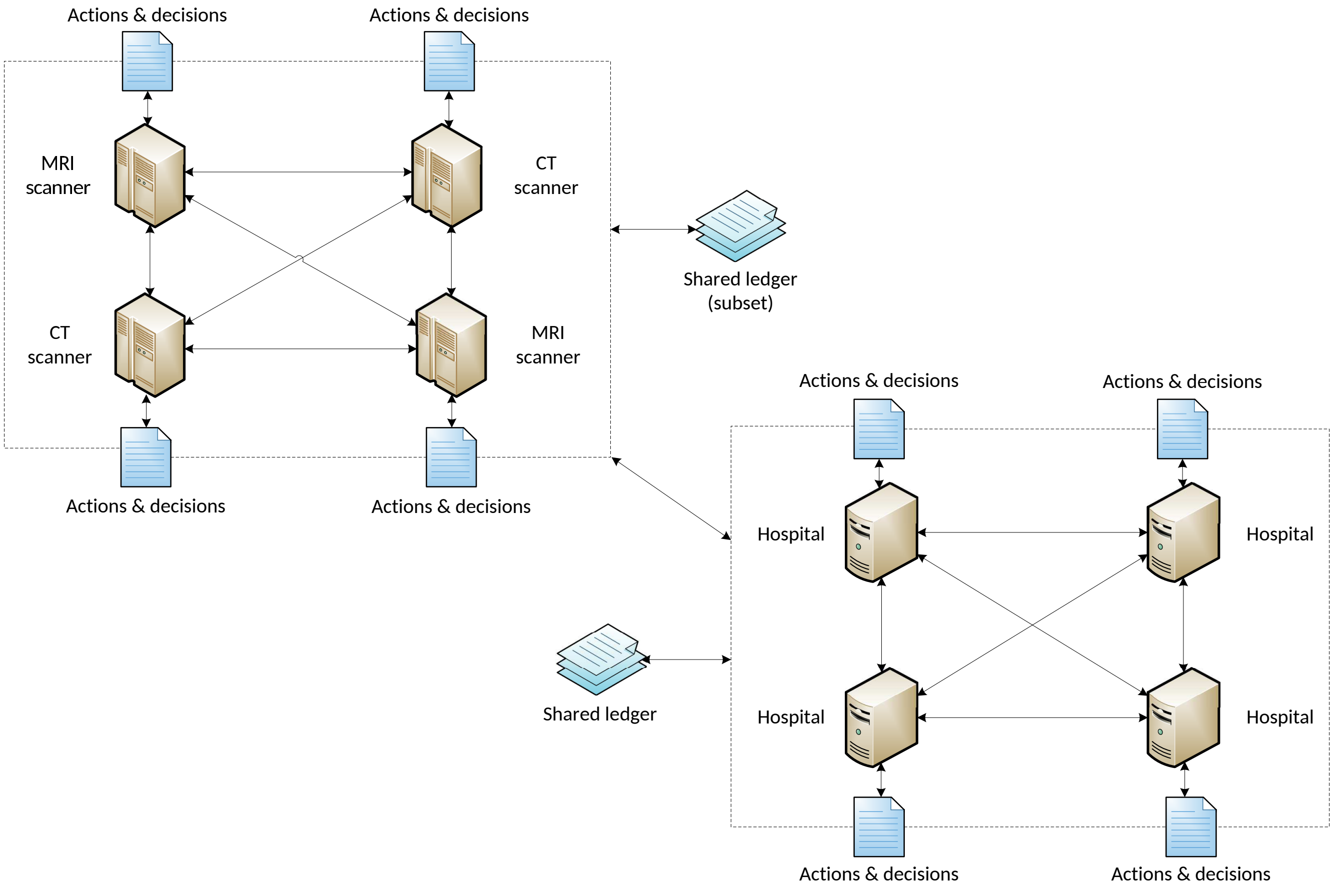

Figure 4.9: Related use cases - Cure

Use case	Less potential	Most potential	Unidentified
Cure (Healthcare sector)	66,6%	11,1%	22,2%

Table 4.9: Potential of the cure use case

4.4.3 Smart industry & Care

The other two use cases that were discussed during the interviews are related to the industrial and the healthcare sector. In subsection 3.8.2 these concepts are described in detail, so in this section they will only be treated by explaining the corresponding figures. The goal of the smart industry concept was to make the collection of machinery related data in smart factories partially public for suppliers. This can be realized by connecting the different machines in such factories with each other and store portions of their data in a shared ledger.

By doing so, part suppliers will be immediately aware when new parts are needed and this can increase the overall efficiency in the supply chain. The outline for this use case is shown in figure 4.10 and consists of two types of context. In the top rectangle is the external context represented, which contains several smart factories that each write a limited amount of data in the shared ledger after they reach consensus. The internal context is displayed in the bottom rectangle and consists of various suppliers that collect and use the data that is generated in the external context. In this way, the shared ledger becomes an overview of the data that the machines in the smart factories share with the suppliers during the production process. The conditions under which these proposed values are legit will be specified in the overlaying protocol.

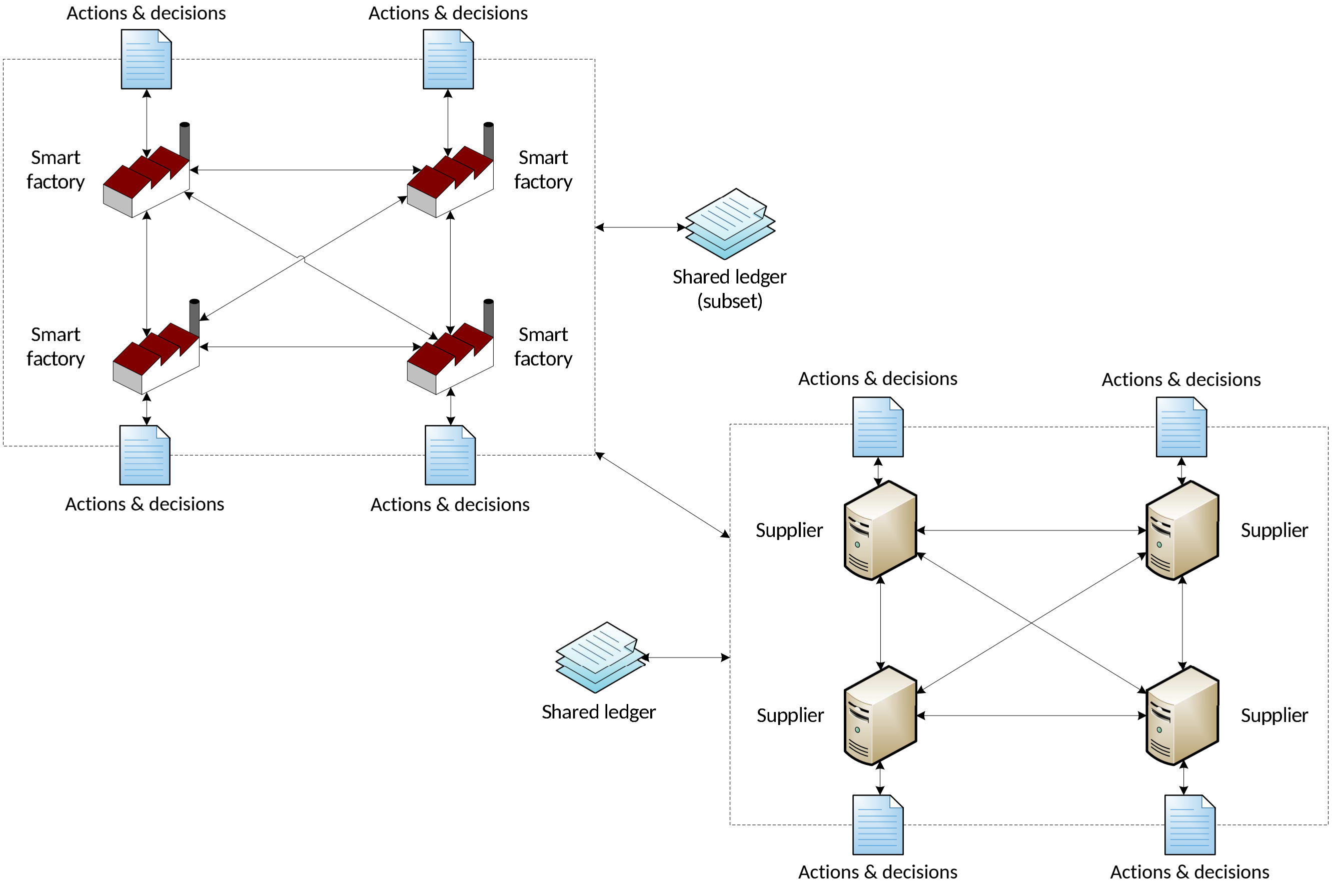

Figure 4.10: Related use cases - Smart industry

For four of the nine respondents, the smart industry use case had the most potential of the different use cases.

It is considered to be the most feasible one because there are not that many laws and regulations to comply with as in the healthcare sector. Besides that, there is no flow of money involved in contrast to the smart container concept, which will simplify the creation of an actual prototype. Finally, the nature of the data that is exchanged in the smart industry use case is less risky than in the other sectors. Despite these arguments, there were three of the experts who had another view and thought that one of the other concepts had more potential. An overview of the smart industry use case its potential is illustrated in table 4.10.

Use case	**Less potential**	**Most potential**	**Unidentified**
Smart industry (Industrial sector)	33,3%	44,4%	22,2%

Table 4.10: Potential of the smart industry use case

Care

In the care use case was the objective to give clients insight in the data that medical related devices collect and which parties make use of it. This can be achieved by connecting the medical related devices in a house with each other and store their data in a shared ledger. As a result, the transparency in the healthcare sector will be improved and the client gets more control over its own data. The outline for the care concept is depicted in figure 4.11 and is broadly comparable with the cure use case. In the top rectangle is the external context displayed, which consists of a combination of devices and smart sensors in the house of an elderly person. These devices write their collected data in the shared ledger after they reach consensus. The internal context is represented by the bottom rectangle and contains a collection of caregivers that store the whole ledger and make use of this data. In this way, the shared ledger will become an overview of the data that the medical related devices share and the actions that the caregivers conducted with this data. The conditions under which these proposed values are legit will be specified in the overlaying protocol.

For six of the nine respondents, the care use case did not have the most potential of the four use cases. The arguments for this viewpoint were once again the large amount of legislation in the healthcare sector and the fact that there is no room for errors with such sensitive data. Once more, only one interviewee believed that the cure use case had indeed the most potential. This is because there are no existing solutions for the care concept yet, which implies that blockchain technology could have additional value in this area.

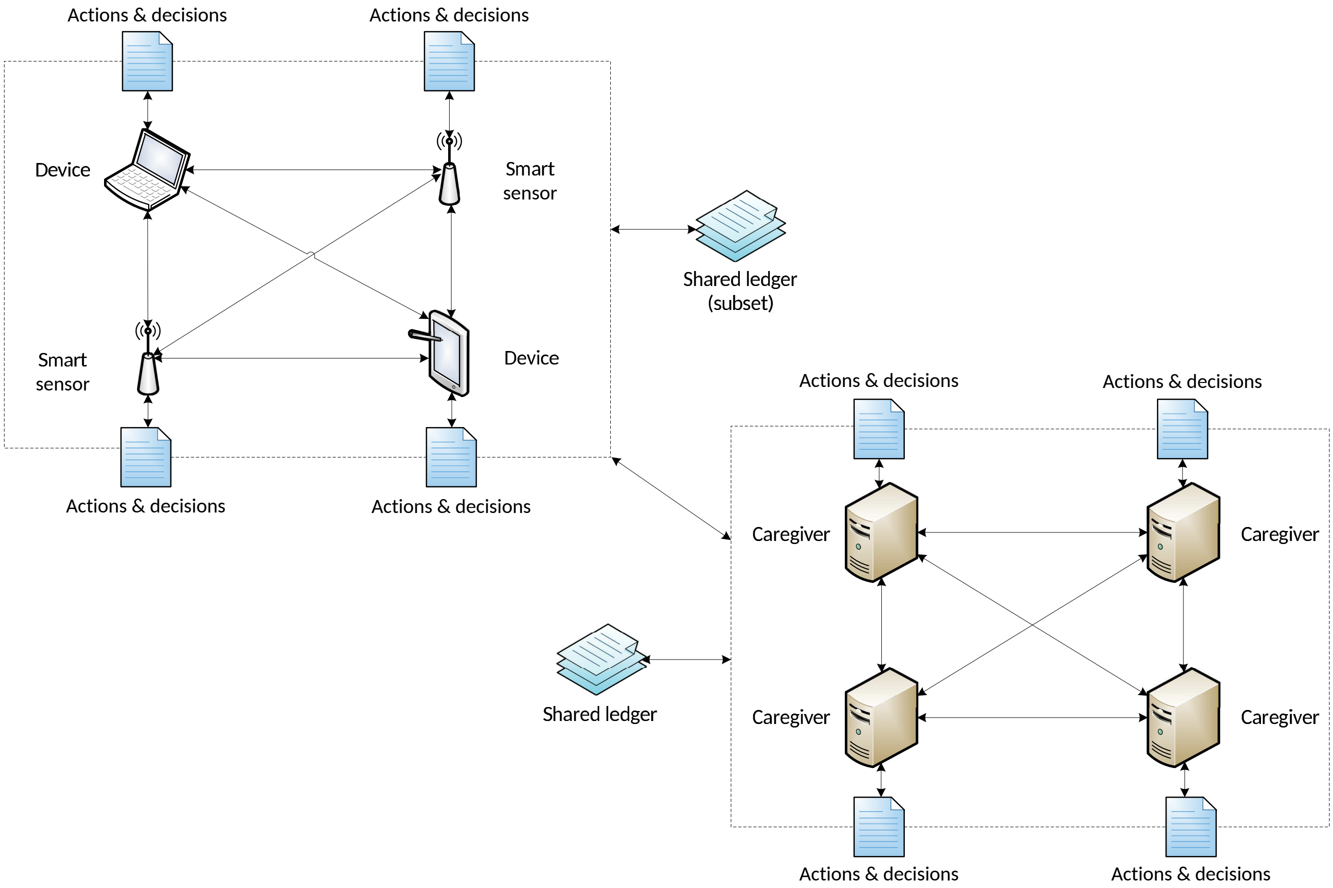

Figure 4.11: Related use cases - Care

Use case	Less potential	Most potential	Unidentified
Care (Healthcare sector)	66,6%	11,1%	22,2%

Table 4.11: Potential of the care use case

4.4.4 Overview

When comparing the potential of the four use cases that were described in the previous subsections it is important to keep the knowledge gap about the underlying principles (subsection 4.3.5) in mind. The implications of this knowledge gap will be described in the next chapter, see subsection 5.2.1. As shown in table 4.12, the smart industry use case in the industrial sector has the most potential according to the respondents.

The smart container in the logistics sector came in second place and the two healthcare related uses cases cure and care finished last. Summarized, there were two main reasons why these other use cases did not have enough potential according to the experts. This was due to the risky nature of the data that would be stored in a shared ledger and the amount of laws and regulations to comply with.

Use case	**Less potential**	**Most potential**	**Unidentified**
Smart container (Logistics sector)	44,4%	33,3%	22,2%
Cure (Healthcare sector)	66,6%	11,1%	22,2%
Smart industry (Industrial sector)	33,3%	44,4%	22,2%
Care (Healthcare sector)	66,6%	11,1%	22,2%

Table 4.12: Potential of the related use cases

4.5 Additional results

The three key elements of the interviews were the discussion about the general concepts, the explanation of the underlying principles and the evaluation of the related use cases. This has led to several results that have been described in the previous sections of this chapter. Nevertheless, there were also a number of additional results found after analysing the interviews.

4.5.1 Identified problems

The first category of additional results contains various identified problems that are related to blockchain technology. In the previous section, the related use cases showed that the data in the shared ledger needs to be encrypted in order to guarantee privacy. This believe is represented by a majority of the experts, who mentioned that privacy is an important aspect to take into account when implementing blockchain technology. A number of other identified problems were related to the management side such as compliance, governance and standardization. Some of the respondents emphasized the critical role of these elements because they can either stimulate or hold back the development of the technology.

Besides that, there were also multiple problems associated to implementation choices like the amount of data to store, enabling data removal, the handling of forking and selecting a proper incentive to mine. A portion of the interviewees was aware that making the wrong choices in this area will have a negative effect on the success rate of any blockchain implementation. The last two identified problems were linked to specific Bitcoin issues as processing capacity and scalability. This reveals that some of the experts found it difficult to separate the Bitcoin and blockchain technology concepts from each other. Even though, there was a variety of respondents who underlined that alternative blockchain solutions must be better capable of dealing with these characteristics in order to be broadly applied. An overview of the identified problems and the number of times that they are grounded is illustrated in table 4.13.

Identified problem	**Grounded**
Amount of data	3 times
Compliance	3 times
Data removal	2 times
Forking	2 times
Governance	2 times
Incentive to mine	3 times
Privacy	6 times
Processing capacity	2 times
Scalability	3 times
Standardization	2 times

Table 4.13: Identified problems

4.5.2 Theory related statements

The next category of additional results includes a selection of theory related statements about the general concepts and underlying principles. In table 4.14 is a subset of these different remarks depicted, including the number of times that they are grounded. The statements showed that several experts agreed with the underlying principles and that these helped them to gain more insight in blockchain technology. Besides that, there were also two interviewees who identified the distributed systems origin of the principles.

Overall, multiple respondents considered the theory in the interviews as logical findings that formed a clarifying overview in which the difference between the general concepts became clear. Finally, there was a portion of the experts that highlighted the interesting character and fundamental nature of this research.

Theory related statement	Grounded
Agrees with the underlying principles	4 times
Clarifying overview	2 times
Different concepts are clear	3 times
Distributed systems principles	2 times
Fundamental research	3 times
Insightful principles	3 times
Interesting work	3 times
Logical findings	2 times

Table 4.14: Theory related statements

4.5.3 Useful considerations

The last category consists of several useful considerations that were derived from the interviews. A majority of the respondents acknowledged that blockchain technology has several disruptive characteristics but that it is important to keep searching for the added value. This can be realized by comparing the benefits of this new technology to existing solutions in order to prevent that it is cut off too early. Currently, there is a risk that blockchain becomes a technology that is looking for a problem instead of offering solutions. Another point that was made by some of the interviewees is the present lack of knowledge, which implies that additional research is needed. The last useful consideration is that society will not be ready for blockchain technology because it consists of the revolutionary idea that there is no more TTP needed. This is contradictory with the current situation and it will take time before people are used to this new concept. In table 4.15 is a selection of these considerations shown and the number of times that they are grounded.

4.6 Recap

This chapter started with explaining the structure of the expert-interviews that were conducted during this research. After that, the results associated to these interviews were discussed based on a number of figures and tables.

Useful consideration	**Grounded**
Additional research is needed	1 time
Compare benefits to existing solutions	3 times
Cutting it off too early	1 time
Disruptive characteristics	4 times
Lack of knowledge	2 times
Search for the added value	3 times
Society is not ready	1 time
Technology looking for a problem	1 time

Table 4.15: Useful considerations

This led to the general concepts, the underlying principles, related use cases and additional results sections of this chapter. In the next chapter it is time to answer the research questions and to evaluate this study. This evaluation will contain a number of implications and the need for future work.

Chapter 5

Discussion

With the help of the previous chapters it is possible to answer the sub questions of this study. Based on this information, the main question of this study will be addressed together with the corresponding implications and the need for future work.

5.1 Answering the research questions

In order to address the main question of this study there are several sub questions that need to be answered first. The first sub question was related to the introduction of chapter 1, which showed that the number of connected devices will only grow in the next few years. This is due to developments like the IoT, CPS and the Industrial Internet that were described in section 2.3. Eventually, this will result in smart industries and ULSS (subsection 2.4.3) in which secure communication is crucial to prevent incidents with massive amounts of data. Therefore, there is a technology needed that has the characteristics to deal with these complex and distributed environments. Blockchain technology is considered to be such a technology and by answering the second sub question it became clear which capabilities it can offer.

5.1.1 The role of the underlying principles

By defining the role of the underlying principles it was possible to determine in a step-by-step manner how these contributed to achieving secure communication in distributed systems. First of all, the technology needs to be fault tolerant (section 3.3) to keep performing its specified tasks after the occurrence of faults. This guarantees that sent communication messages will be replicated in multiple processing units so they do not disappear when one or multiple components in a network fail.

Secondly, section 3.4 described that malfunctioning components that give conflicting information to different nodes in a distributed network must also be detected by making it Byzantine fault tolerant. By doing so, it becomes impossible to reach agreement about conflicting messages that are transmitted throughout the network and this will ensure that a blockchain implementation continues to function properly. In the third place, there is a mechanism required such as Paxos (section 3.5) that enables the network to reach consensus about the messages that are communicated between nodes. This is to ensure that a blockchain network performs its actions based on one or more proposed messages that are accepted by a majority of the nodes. Finally, as described in subsection 4.3.4 there is an overlaying protocol needed that specifies the conditions under which the proposed messages in the network are legit. In this way, only the proposed messages that fulfil these specified conditions will be taken into consideration by initiating a consensus flow in the network. Experts need to be aware of these underlying principles in order to give an informed opinion about value adding implementations other than virtual currencies. This was tested by conducting the expert-interviews and provided enough information to answer the third sub question.

5.1.2 Expert awareness & useful implementations

The awareness levels of the underlying principles (table 4.7) illustrated that only fault tolerance is known to all the respondents that were interviewed. Two of the interviewees were familiar with the term Byzantine fault tolerant and none of the experts was well aware of a consensus algorithm like Paxos or the overlaying protocol. This indicated that there is a lack of knowledge about multiple core elements of blockchain technology. Based on this limited knowledge, the respondents gave their opinion about the potential of some related use cases. As depicted in table 4.12, the smart industry use case has the most potential according to four of the nine interviewees. This is related to the smaller amount of laws and regulations to comply with, there is no flow of money involved and the nature of the data is not that risky. Three experts supported the smart container concept that finished in second because of its relatively low threshold, limited amount of related risks and potential business value. In both the cure and the care use case only one respondent believed that these concepts had the most potential due to the additional value that they could offer. This indicates that the interviewees had a preference for the implementations in the industrial and logistics sector. In the next subsection it is time to combine the different answers from the three sub questions in order to address the main question of this study.

5.1.3 Additional value of the technology

In the next few years, new technologies are needed that can deal with the growing amount of connected devices in different types of distributed systems. Another effect of earlier described developments like the IoT, CPS, smart industries and ULSS is that the amount of information will raise significantly in the future. This implies that the massive amounts of data in such distributed systems need to be exchanged in a safely manner to prevent major incidents, see subsection 2.4.4. By analysing the underlying principles of blockchain technology, it became clear that this technology has the characteristics to handle a large amount of devices that exchange information in complex and distributed environments. The fact that these principles not only guarantee safety but also enable the removal of the traditional TTP stimulates the self-regulating character of a network. Compared to the existing solutions in distributed environments, this is a significant change because it removes the devious and inefficient procedures that the TTP currently performs. With the help of the underlying principles the nodes are able to perform these tasks in an autonomous way, which improves the efficiency and integrity of the overall process. In theory, this additional value can be offered in any domain but there are circumstances that make it harder to actually implement blockchain technology. Think of the risky nature of the data and the amount of laws and regulations to comply with that were mentioned during the expert-interviews. Nevertheless, this does not change the fact that the technology still offers additional value when these aspects are taken into account. The answers on the research questions of this study identified some points of interest, these will be further discussed in the next section.

5.2 Implications

The first identified point of interest is the lack of knowledge about certain underlying principles of blockchain technology among the experts. This was an unexpected result and illustrates that some of the basic elements in the distributed computing field are unknown to the interviewees. On top of that, several respondents found it difficult to unleash the Bitcoin concept and view the technology just as a combination of the underlying principles. This explains the fact that some experts are sceptical about implementations without the proof-of-work algorithm, see subsection 3.7.1. The benefit of applying this form of abstraction is that the proof-of-work algorithm is nothing more than an implementation of the overlaying protocol in a specific context. When the context changes and has a more closed nature there is no more need for such an extensive mechanism to verify the correctness of transactions. By taking the context of a blockchain implementation into account it is possible to estimate to which extent each element of the technology must be used.

In general, when blockchain technology is regarded as a set of principles instead of only the Bitcoin implementation there will emerge much more opportunities. This can reduce the knowledge gap because there needs to be more awareness and understanding of the underlying principles in order to develop various implementations.

5.2.1 Use case perception

As described in subsection 5.1.2, the respondents had several arguments to choose the industrial and logistics sector as the ones with the most potential. The question is of their opinions represent the reality because they were based on limited knowledge about the underlying principles. Maybe if there were more interviewees who viewed blockchain technology only as a set of principles some use cases could have ended up higher than they did. An example would be the cure concept, because this use case is broadly comparable with the smart factory case where machines in different factories are connected with each other. The data that the MRI and CAT scanners in hospitals collect is not directly patient related, which makes it less risky to store in a shared ledger. This also implies that the amount of laws and regulations to comply with will decline. Such an example shows that estimating the potential of the different use cases is mainly a matter of perception. More important is that the experts realize that the related use cases were merely some examples of blockchain technology and that it can be applied in several other situations. One optimization is for example to not only connect the same sort of devices in different domains with each other but to use a variety of devices. This combination will create completely new distributed environments and could lead to the first ULSS.

5.3 Validation

There have been multiple validation moments in this study to verify the correctness of the described theory in the previous chapters. First of all, my company supervisor Prof. dr. Ben van Lier validated the outcomes of the conducted critical literature review (section 1.3) and provided feedback to improve the overall quality. After that, we also organised a meeting with a group of external stakeholders to discuss their blockchain related knowledge and experience. This meeting was a suitable moment to validate the underlying principles once more before starting with the actual expert-interviews. These interviews were considered to be the third validation moment during this research and provided an additional data set to analyse. This enabled the possibility to combine the outcomes of the critical literature review with the data set from the expert-interviews.

The integration of the concepts in these two data sources led to the results of this study, which were described in chapter 4. These results will be validated again by organising a second meeting with the group of external stakeholders after this research. Besides that, the feasibility of the underlying principles and related uses cases can be validated in the future by conducting experiments.

5.4 Future work

This research consisted of an exploratory case study that was conducted to identify the underlying principles of blockchain technology, which were subsequently tested among multiple experts. The previous sections made it clear that there is a lack of knowledge and that the technological know-how needs to be improved in order to enable the development of alternative implementations. Actually, there are two approaches for solving this problem and to determine which one to choose it is necessary to take the knowledge level into consideration. When there is hardly any knowledge about the underlying principles it is wise to start with additional research to better understand these principles and the capabilities that they can offer. In environments where there is more than a basic understanding of the underlying principles, a strategy could be to start with conducting some experiments. When it turns out that some aspects need extra attention during these experiments it will be necessary to perform additional research.

At Centric, which is the company where I conducted this study, the approach is to start with some experiments in order to develop their own blockchain based on the underlying principles. The goal is to test the feasibility of these principles by connecting four nodes (subsection 3.4.2) in a network that perform actions and make decisions, which are stored in a shared ledger. If this succeeds, the number of connected devices can gradually be increased to simulate the amount of nodes in larger networks. When this does not cause any problems, the next step could be to implement such a blockchain solution in the industrial or logistics sector, see table 4.12. A bonus is that I will remain closely involved in this entire process and can actually find out if the described theory works in practice.

5.5 Reflection

In the beginning of this study I was unfamiliar with the whole concept of blockchain technology and its underlying principles. Along the way, I acquired more knowledge about the different components of the technology and how it can provide additional value in several situations. The more people I spoke, the more I became aware that there is a serious lack of knowledge about the fundamentals of blockchain technology.

This became evident when some of the people I met considered me as the “expert”, which was a nice but unexpected experience. Looking back, at this moment in time I have more questions about the technology than when I started with this study. It shows that this research is only a minor step and that there are much more steps that need to be taken in order to fully understand blockchain technology.
Therefore, I would like to end with the following quote from the Greek philosopher Socrates (469 - 399 BC) that represents the process of the last few months:

The only true wisdom is in knowing you know nothing.

Bibliography

[1] Juniper Research, *Internet of Things: Consumer, industrial & public services 2015-2020*, 2015.

[2] K. Charmaz, *Constructing Grounded Theory.* Sage Publications, 2006.

[3] B. M. Leiner, V. G. Cerf, D. D. Clark, R. E. Kahn, L. Kleinrock, D. C. Lynch, J. Postel, L. G. Roberts, and S. Wolff, “A brief history of the internet,” *ACM SIGCOMM Computer Communication Review*, vol. 35, no. 5, pp. 22–31, 2009.

[4] I. Bojanova, G. Hurlburt, and J. Voas, “Imagineering an internet of anything,” *Computer*, no. 6, pp. 72–77, 2014.

[5] W. Emmerich, “Distributed system principles,” University College London, 1997.

[6] L. Lamport, R. Shostak, and M. Pease, “The byzantine generals problem,” *ACM Transactions on Programming Languages and Systems (TOPLAS)*, vol. 4, no. 3, pp. 382–401, 1982.

[7] A. Wright and P. De Filippi, “Decentralized blockchain technology and the rise of lex cryptographia,” *Available at SSRN 2580664*, 2015.

[8] M. Saunders, P. Lewis, and A. Thornhill, *Research methods for business students.* Pearson Education India, 2011.

[9] R. Yin, *Case study research: Design and methods.* Sage Publications, 2003.

[10] T. Bemers-Lee, R. Cailliau, J.-F. Groff, and B. Pollermann, “World-wide web: the information universe,” *Electronic Networking: Research, Applications and Policy*, vol. 2, no. 1, pp. 52–58, 1992.

[11] M. Weiser, “Ubiquitous computing,” *Computer*, vol. 26, no. 10, pp. 71–72, 1993.

[12] M. Weiser, “The computer for the 21st century,” *Scientific American*, vol. 265, no. 3, pp. 94–104, 1991.

[13] G. Alberts and R. Oldenziel, *Hacking Europe.* Springer, 2014.

[14] S. Haller, S. Karnouskos, and C. Schroth, "The internet of things in an enterprise context," *Lecture Notes in Computer Science*, vol. 5468, pp. 14–28, 2009.

[15] D. Miorandi, S. Sicari, F. De Pellegrini, and I. Chlamtac, "Internet of things: Vision, applications and research challenges," *Ad Hoc Networks*, vol. 10, no. 7, pp. 1497–1516, 2012.

[16] L. Atzori, A. Iera, and G. Morabito, "The internet of things: A survey," *Computer networks*, vol. 54, no. 15, pp. 2787–2805, 2010.

[17] National Institute of Standards and Technology, *Foundations for innovation in cyber-physical systems*, Energetics Incorporated, Columbia, 2013.

[18] K.-J. Park, R. Zheng, and X. Liu, "Cyber-physical systems: Milestones and research challenges," *Computer Communications*, vol. 36, no. 1, pp. 1–7, 2012.

[19] Industrial Internet Consortium, *Industrial internet reference architecture*, 2015.

[20] Bosch Software Innovations GmbH, *Industrial Internet: Putting the vision into practice*, 2015.

[21] L. Northrop, P. Feiler, R. P. Gabriel, J. Goodenough, R. Linger, T. Longstaff, R. Kazman, M. Klein, D. Schmidt, K. Sullivan *et al.*, *Ultra-large-scale systems: The software challenge of the future*, Software Engineering Institute Carnegie Mellon, 2006.

[22] Smart Industry, *Actieagenda smart industry*, 2014.

[23] Cisco, *IoT threat environment*, 2015.

[24] L. Lamport, "Time, clocks, and the ordering of events in a distributed system," *Communications of the ACM*, vol. 21, no. 7, pp. 558–565, 1978.

[25] B. W. Johnson, "Fault-tolerant microprocessor-based systems," *IEEE Micro*, vol. 4, no. 6, pp. 6–21, 1984.

[26] J. H. Wensley, L. Lamport, J. Goldberg, M. W. Green, K. N. Levitt, P. M. Melliar-Smith, R. E. Shostak, and C. B. Weinstock, "Sift: Design and analysis of a fault-tolerant computer for aircraft control," *Proceedings of the IEEE*, vol. 66, no. 10, pp. 1240–1255, 1978.

[27] L. Lamport, "The part-time parliament," *ACM Transactions on Computer Systems (TOCS)*, vol. 16, no. 2, pp. 133–169, 1998.

[28] M. J. Fischer, N. A. Lynch, and M. S. Paterson, “Impossibility of distributed consensus with one faulty process,” *Journal of the ACM (JACM)*, vol. 32, no. 2, pp. 374–382, 1985.

[29] L. Lamport, “Paxos made simple,” *ACM Sigact News*, vol. 32, no. 4, pp. 18–25, 2001.

[30] G. Solmaz, “Paxos algorithm,” University of Central Florida, 2012.

[31] E. Gafni and L. Lamport, “Disk paxos,” *Distributed Computing*, vol. 16, no. 1, pp. 1–20, 2003.

[32] J.-P. Martin and L. Alvisi, “Fast byzantine paxos,” in *Proceedings of the International Conference on Dependable Systems and Networks (DSN'05)*, 2004, pp. 402–411.

[33] M. Biely, Z. Milosevic, N. Santos, and A. Schiper, “S-paxos: Offloading the leader for high throughput state machine replication,” in *31st Symposium on Reliable Distributed Systems (SRDS)*, 2012, pp. 111–120.

[34] H. T. Dang, D. Sciascia, M. Canini, F. Pedone, and R. Soulé, “Netpaxos: Consensus at network speed,” in *Symposium on Software Defined Networks (SOSR)*, 2015.

[35] S. Nakamoto, “Bitcoin: A peer-to-peer electronic cash system,” *Consulted*, vol. 1, no. 2012, p. 28, 2008.

[36] R. L. Rivest, A. Shamir, and L. Adleman, “A method for obtaining digital signatures and public-key cryptosystems,” *Communications of the ACM*, vol. 21, no. 2, pp. 120–126, 1978.

[37] A. M. Antonopoulos, *Mastering Bitcoin: Unlocking digital cryptocurrencies.* O'Reilly Media, Inc., 2014.

[38] J. Baker, C. Bond, J. C. Corbett, J. Furman, A. Khorlin, J. Larson, J.-M. Leon, Y. Li, A. Lloyd, and V. Yushprakh, “Megastore: Providing scalable, highly available storage for interactive services,” in *Conference on Innovative Data Systems Research (CIDR)*, vol. 11, 2011, pp. 223–234.

[39] M. Swan, *Blockchain: Blueprint for a New Economy.* O'Reilly Media, Inc., 2015.

[40] CoinDesk, *Hands On With Linq, Nasdaq's Private Markets Blockchain Project*, 2015.

[41] Deutsche Börse Group, *Blockchain Technology*, 2015.

Appendix A

Interview questions

Introductory questions

- Can you tell me something about your education and career?
- At the moment there is a hype around Bitcoin, what do you know about this concept?
- Besides Bitcoin there is also much rumour about the underlying concept blockchain technology, what do you know about this technology?

Review of underlying principles blockchain technology

- The concept blockchain technology consists of a number of principles that are considered to be the building blocks of this technology, could you name a few of them?
  - If I mention the term fault tolerance, what comes to mind?
  - Besides fault tolerance the Byzantine Generals Problem (BGP) also plays a role, do you have any idea what this means?
  - The third element is related to the use of a consensus algorithm, what is the function of such an algorithm according to you?
- The overlaying protocol can be considered to be the fourth element that forms a part of blockchain technology, what is your view on this?
- What is the role of the overlaying protocol in an entire blockchain implementation according to you?

Applications for blockchain technology

- What are in your view promising areas for this new technology based on the principles we mentioned earlier?

- We also identified several application areas and are interested in how much potential these areas have in your eyes?
  - Smart container (logistics sector): An autonomous container that is able to keep track of its flow of goods and flow of money and shares this information with different stakeholders.
  - Cure (healthcare sector): Combining hospital data so it becomes accessible for doctors so they can come up with better diagnoses and treatments in the future.
  - Smart industry (industrial sector): Make the data that is collected in factories partially public with other stakeholders to increase the overall efficiency in the supply chain.
  - Care (healthcare sector): Give clients the opportunity to gain insight in the data that is collected by medical related devices and which parties make use of it.

  A choice can be made from the financial, logistics, healthcare or industrial sector.
- Are there at this moment any factors that hold back the adoption of blockchain technology or that must be taken into account?

Control questions

- Do you agree with our findings and the described characteristics of a blockchain or do you have another opinion about this?
- Is the difference clear between Bitcoin, the most famous implementation of blockchain technology and the underlying concept?

Final questions

- How do you think that blockchain technology will develop itself in the coming years?
- Do you believe that blockchain technology is a so called disruptive technology that is comparable with the Internet, cloud computing and the Internet of Things (IoT)?
- Do you think that we discussed the key developments around blockchain technology or do you have anything to add?

Appendix B

Code list

Code	Grounded	Created
a_dissenting_voice	1	25-05-16 11:08
additional_value_for_blockchain	9	19-04-16 16:42
advantages_over_other_approaches	2	19-04-16 11:17
agrees_with_the_principles	5	11-05-16 10:53
agricultural_related_applications	2	11-05-16 13:04
alternative_consensus_mechanisms	1	30-05-16 14:58
alternative_solutions	1	27-05-16 14:08
amount_of_computing_power	1	23-05-16 13:50
application_layer	2	29-03-16 14:50
applications_above_research	1	13-05-16 16:47
arbitrating_the_use_of_data	1	19-04-16 11:10
auditing	1	05-04-16 10:41
bank_related_tasks	1	10-05-16 16:08
based_on_existing_theory	1	01-06-16 15:39
believes_in_the_concept	1	27-05-16 14:04
BIP32_wallet	1	19-04-16 11:20
Bitcoin_knowledgeable	7	04-04-16 13:39
BitLicense	1	04-04-16 15:29
blockchain_example	1	04-04-16 14:44
both_correct_and_incorrect	1	04-04-16 14:11
bridge_currency	1	30-05-16 10:23
building_knowledge	1	13-05-16 16:28
Byzantine_fault_tolerance_knowledgeable	2	17-05-16 16:08
Byzantine_fault_tolerance_unknown	4	04-04-16 13:56
capturing_blockchain_situations	1	19-04-16 10:22
Central_Security_Depository	1	30-05-16 14:19
central_use_case_theme	1	19-04-16 15:59

chain_control	1	05-04-16 10:40
clarifying_overview	2	11-05-16 13:27
classical_distributed_system_principles	4	19-04-16 10:04
clear_owner	1	04-04-16 15:20
clear_principles	3	04-04-16 14:42
combination_of_technologies	1	17-05-16 14:36
combine_logistics_with_finance	1	31-05-16 10:35
comparable_with_cloud_computing	2	24-05-16 11:36
complex_event_processing	1	19-04-16 15:52
complex_transactions	1	30-05-16 14:19
compliance_not_the_issue	1	01-06-16 11:44
complicated_consensus_mechanisms	3	04-04-16 14:29
complying_with_laws_and_regulations	6	19-04-16 14:14
concrete_developments_are_taking_place	2	26-05-16 13:58
confirmation_layer	2	29-03-16 16:14
connectivity_to_hospitals	1	19-04-16 16:44
consensus_algorithms_knowledgeable	2	19-04-16 11:53
consensus_algorithms_unknown	3	04-04-16 14:01
consensus_once_a_day	1	04-04-16 14:45
context_specific_blockchains	1	19-04-16 15:43
contract_related_applications	1	13-05-16 10:35
cost_structures	2	29-03-16 17:25
created_awareness	1	30-05-16 13:37
critical_systems	1	31-05-16 11:08
cryptographic_hashes	3	17-05-16 15:11
cutting_it_off_too_early	1	20-04-16 11:57
darkweb_payment_method	1	10-05-16 13:45
data_analytics	2	19-04-16 15:52
data_becomes_too_complex	1	19-04-16 16:16
data_characteristics	2	19-04-16 09:58
data_duplication	1	04-04-16 13:53
data_side_is_also_important	1	11-05-16 11:08
data_storage_problem	1	13-05-16 11:51
define_legal_solutions	1	19-04-16 14:30
defined_some_principles	1	30-05-16 11:22
defining_blockchain_ecosystems	3	05-04-16 10:28
depending_on_a_community	1	05-04-16 10:39
developed_different_prototypes	1	30-05-16 10:58
developing_standards	1	29-03-16 13:44
developing_technologies	1	29-03-16 13:55
difference_between_concepts_is_clear	3	11-05-16 11:18
difference_between_concepts_is_unclear	1	25-05-16 11:30

difference_with_cloud_computing	2	26-05-16 15:40
difference_with_distributed_databases	1	30-05-16 15:53
different_cryptographic_algorithms	1	17-05-16 16:31
different_mindset	1	01-06-16 11:46
different_techniques_per_layer	1	04-04-16 14:41
different_variants	2	04-04-16 14:46
difficult_variant	1	04-04-16 16:54
discussion_about_the_principles	1	05-04-16 11:00
disruptive	6	05-04-16 10:55
distributed_databases_principles	1	10-05-16 14:09
distributed_ledger	1	12-05-16 09:55
distributed_trust	2	20-04-16 13:19
distribution_of_computing_capacity	2	13-05-16 11:55
distribution_principle	1	10-05-16 14:05
DNB_Coin_project	1	30-05-16 10:24
domain_specific_language	2	10-05-16 15:17
dynamic_environments	1	30-05-16 14:20
early_stages	1	20-04-16 13:28
easy_to_set_up	1	04-04-16 15:32
economical_incentives	1	29-03-16 17:25
enables_automation	1	23-05-16 11:46
encryption_as_integral_component	1	30-05-16 13:51
encryption_by_design	1	30-05-16 15:01
encryption_mechanism	1	30-05-16 11:41
endless_opportunities	1	24-05-16 13:34
enough_potential_to_invest	1	13-05-16 16:15
Ethereum_model	2	04-04-16 13:40
everybody_is_guessing	1	20-04-16 13:45
everyone_can_mine	1	04-04-16 15:26
exaggerated_expectations	1	11-05-16 12:05
exchange_of_data	1	19-04-16 15:44
existing_care_solutions	1	19-04-16 14:00
existing_IT_solutions_vs_blockchain	9	04-04-16 15:08
existing_smart_industry_solutions	1	19-04-16 16:01
experts_are_cautious	1	27-05-16 14:02
exposing_data	1	19-04-16 15:59
facilitate_new_opportunities	1	25-05-16 16:48
fault_tolerance_knowledgeable	8	04-04-16 13:53
fight_against_poverty	1	24-05-16 14:17
finance_is_difficult	1	31-05-16 11:06
financial_related_applications	5	13-05-16 10:39
focused_use_case_testing	1	20-04-16 13:59

forking_problem	2	04-04-16 16:32
forward_error_correction	1	04-04-16 13:52
foundation_form	1	04-04-16 16:33
four_layer_model	2	04-04-16 14:40
full_transparency	5	10-05-16 15:13
fundamental_research	3	05-04-16 11:03
general_use_cases	3	05-04-16 10:27
governance_is_not_the_problem	1	11-05-16 10:44
governance_problem	13	01-04-16 14:41
government_related_applications	1	13-05-16 10:29
healthcare_could_improve_human_lives	1	26-05-16 14:53
healthcare_has_the_most_potential	2	19-04-16 17:00
healthcare_is_difficult	2	13-05-16 11:28
high_risk	1	19-05-16 11:32
hot_item	1	24-05-16 11:22
how_to_make_the_chain	1	12-05-16 11:17
hype	7	05-04-16 10:50
hype_instead_of_solution	1	05-04-16 10:16
ideal_payment_metod	1	25-05-16 11:11
identifying_potential_use_cases	1	19-04-16 11:13
implementation_can_vary_per_context	1	17-05-16 16:14
implementation_differences	3	19-04-16 09:55
important_component	1	04-04-16 14:08
improve_data_quality	1	10-05-16 17:06
incentive_to_mine	5	19-04-16 11:18
inherently_safer	1	12-05-16 11:52
interesting_use_cases	4	05-04-16 10:58
interesting_work	3	10-05-16 15:29
international_differences	1	19-04-16 13:54
internet_bubble	1	05-04-16 10:56
internet_of_things	1	17-05-16 14:11
internet_of_value	1	17-05-16 14:09
involving_multiple_parties	2	19-04-16 14:05
IoT_related_applications	1	23-05-16 11:21
killer_app_is_necessary	1	13-05-16 16:10
lack_of_knowledge	3	13-05-16 11:58
lack_of_standardisation	3	04-04-16 16:31
law_enforcement_problem	2	04-04-16 15:22
lifecycle_management	1	05-04-16 10:41
limited_functionality	1	19-04-16 16:25
limited_number_of_suppliers	1	19-05-16 11:49
linkage_of_patient_data	1	19-04-16 16:44

list_of_pointers	1	19-04-16 13:36
logical_findings_and_characteristics	2	05-04-16 10:43
logistics_has_the_most_potential	3	13-05-16 11:28
logistics_related_applications	4	10-05-16 16:48
low_general_awareness	1	26-05-16 10:56
low_hanging_fruit	1	23-05-16 11:24
many_participants	2	30-05-16 14:21
market_demand	1	13-05-16 13:35
modifying_data	1	10-05-16 14:59
money_is_not_an_issue	1	05-04-16 10:56
money_laundering	1	10-05-16 13:45
more_research_is_needed	1	20-04-16 13:33
mutual_benefit	1	19-04-16 11:29
necessary_principles	1	24-05-16 11:11
need_for_encryption	1	26-05-16 10:37
need_for_other_models	1	05-04-16 11:05
no_additional_value_for_blockchain	1	19-04-16 16:07
no_clear_direction	1	01-06-16 11:42
no_complete_overview	1	13-05-16 13:24
no_disruptive_technology	2	11-05-16 13:10
no_fitting_solution	1	04-04-16 16:41
no_fixed_leader	1	30-05-16 16:33
no_hype	1	11-05-16 11:37
no_interoperability	1	19-04-16 14:04
no_killer_application	1	19-05-16 11:50
no_liability	1	24-05-16 11:44
no_one_size_fits_all_solution	1	19-04-16 15:42
no_single_point_of_failure	4	26-05-16 11:12
no_trusted_third_party	2	12-05-16 10:01
not_agree_or_disagree	2	04-04-16 14:10
not_fully_decentralized	2	04-04-16 14:06
not_one_new_idea	1	17-05-16 14:38
notary_related_tasks	1	10-05-16 16:09
number_portability	1	04-04-16 14:44
old_data_needs_to_be_relevant	1	30-05-16 14:21
only_one_implementation	1	04-04-16 16:31
open_data	9	30-03-16 13:47
originated_from_a_disruptive_technology	1	01-06-16 14:50
overlaying_protocol_knowledgeable	2	04-04-16 14:07
parallels_with_changes_in_society	1	01-06-16 13:37
partly_Byzantine_fault_tolerant_knowledgeable	1	14-04-16 11:14
Paxos_knowledgeable	1	17-05-16 16:44

Paxos_unknown	4	04-04-16 14:01
permission_mechanism	1	04-04-16 16:58
permission_to_use	5	04-04-16 15:19
permissioned_vs_unpermissioned	1	04-04-16 14:57
plausible_principles	1	30-05-16 11:52
possible_explanation	1	04-04-16 13:49
preforming_experiments	1	13-05-16 16:31
preforming_explorations	2	29-03-16 14:00
privacy	10	05-04-16 10:39
private_context	1	01-06-16 10:31
processing_problem	3	13-05-16 10:22
programmable_form_of_value	2	23-05-16 11:37
promising_public_implementations	1	04-04-16 15:14
promising_variant	1	04-04-16 16:45
proof_of_stake	1	19-04-16 15:14
proof_of_work	1	17-05-16 15:11
psychological_barrier	1	23-05-16 14:15
public_blockchains	1	04-04-16 15:18
public_permissioned_blockchains	1	04-04-16 16:41
public_vs_private	1	04-04-16 14:57
publicly_available	1	04-04-16 15:21
random_attack_resistant	1	30-05-16 16:19
redundancy_principle	1	10-05-16 14:05
reference_database	2	10-05-16 15:25
regular_software_cycle	1	01-06-16 13:54
relation_to_Bitcoin	1	26-05-16 15:38
relatively_low_risk	1	19-05-16 11:22
remote_management	1	04-04-16 16:33
removal_of_data	2	19-04-16 15:12
replacing_proof_of_work	4	19-04-16 10:22
research_context	1	04-04-16 14:43
reviewing_data_in_a_granular_way	2	19-04-16 10:59
sale_of_goods	1	11-05-16 13:39
scalability_problem	6	19-04-16 13:38
scattered_knowledge	1	13-05-16 13:27
search_for_alternative_applications	1	26-05-16 16:59
search_for_the_happy_medium	1	24-05-16 13:08
searching_for_specific_use_cases	2	19-04-16 17:27
security_problem	1	13-05-16 10:25
self_learning_rules	2	12-05-16 10:24
self_regulating_system	2	24-05-16 11:44
semi_trusted_third_parties	1	26-05-16 16:58

set_of_integrated_technologies	1	05-04-16 10:52
share_data_anonymously_between_trusted_parties	1	19-04-16 16:41
sharing_aggregated_data	1	10-05-16 17:05
sidechain_concept	1	19-04-16 13:41
simple_data	1	31-05-16 11:51
smart_contracts	6	29-03-16 14:52
smart_industry_has_the_most_potential	4	10-05-16 16:47
smart_industry_related_applications	2	13-05-16 10:42
social_related_applications	1	23-05-16 14:22
society_is_not_ready	1	17-05-16 10:49
specify_the_right_conditions	1	25-05-16 15:19
storage_of_medical_records	1	10-05-16 16:18
storage_of_permissions	1	19-04-16 10:52
tamper_proof	3	12-05-16 10:22
teaching_Bitcoin	1	04-04-16 13:34
technical_vote_system	1	04-04-16 16:33
technology_is_not_the_problem	3	19-04-16 14:07
technology_looking_for_a_problem	1	05-04-16 10:30
technology_that_is_here_to_stay	2	24-05-16 11:31
the_best_combination_of_rules_and_protocols	1	18-05-16 16:36
the_principles_have_additional_value	1	26-05-16 11:42
third_world_related_applications	1	24-05-16 14:07
too_large_quantities_of_data	3	11-05-16 11:05
too_much_redundancy	2	10-05-16 14:50
toy_examples	1	31-05-16 11:10
translation_into_solutions	2	20-04-16 13:59
trust_related_transactions	1	10-05-16 16:12
trust_vs_privacy	1	10-05-16 16:37
trust_vs_verification	1	19-05-16 12:49
underlying_architecture	1	13-05-16 16:41
unpermissioned_blockchains	4	04-04-16 15:21
unsuccessful_technology	1	17-05-16 14:52
usable_technology	1	19-04-16 09:51
use_very_simple_environments	1	19-04-16 16:20
user_controls_the_data	3	19-04-16 10:51
using_blockchain_as_storage	1	19-04-16 13:35
valuable_model	3	04-04-16 14:11
verification_layer	2	29-03-16 16:07
without_a_blockchain	1	05-04-16 10:24