MSc in Business Administration and Information Systems (E-Business)

Master's Thesis

Blockchain: an exploratory study

Jussi Saha

Supervisors:
Leif Bloch Rasmussen
Ulrik Vilhelm Falktoft

53 pages
98.584 characters
17.1.2017
Copenhagen, 2017

Abstract

Blockchain is the technological innovation behind bitcoin, but it's also much more than that. It's a distributed ledger that allows companies and individuals to collaborate in unprecedented ways without trust or centralization. The goal of this thesis was to identify the main value proposition of blockchain, and to understand how it can be used in different sectors. This was done by analyzing the core concepts of the technology, and how they enable the different use cases. The second goal was to determine the criteria for potential blockchain use cases. The results show that blockchain could potentially disrupt and transform various industries through disintermediation, improvements in efficiency, cost reductions and new business models, but is still at an early stage and has many challenges that have to be overcome. Use of blockchain makes most sense when there are multiple parties present and a lack of trust between them.

Table of contents

Table of figures

1 Introduction

Blockchain is a technology that first emerged from the bitcoin protocol. In the early years, bitcoin was the target of much attention and speculation, but interest has since expanded to the underpinning technology and its potential applications in various sectors. The technology offers a new approach to database management and conducting various transactions between multiple parties without relying on a trusted intermediary. It could potentially have a transformative impact on various industries through disintermediation and disrupting incumbents, and is currently the target of considerable investment and study from all sides, including startups, venture capitalists, market incumbents and even governments. However the technology is still at an early stage and there seem to be as many questions as there are answers about what it could be used for and how. The goal of this thesis is to find answer some of these questions, and to find out if the hype is warranted.

1.1 Conceptual background

Blockchain was initially used to power the bitcoin cryptocurrency. A cryptocurrency can be defined as a decentralized digital payment mechanism, which bitcoin is the first viable example of. The system was first proposed in 2008 by Satoshi Nakamoto, a pseudonym for a person or group of people whose real identity remains unknown. Although the mechanics are described in detail in the original bitcoin whitepaper, Nakamoto doesn't actually use the term "blockchain". As the phenomenon has progressed, its uses have expanded beyond the original bitcoin blockchain. Various other cryptocurrencies were subsequently derived from bitcoin, some with slight modifications to the system and others with major overhauls, for example Litecoin, which proposes faster transaction speeds. Bitcoin has however remained by far the largest and most traded cryptocurrency throughout its history (Coinmarketcap.com 2016). This cryptocurrency use of blockchain is in most cases focused on the same goal: running a virtual currency that enables users to transfer value between one another in a trustless and decentralized manner.

In the last few years, the use of blockchains has expanded to new areas. A key development in the domain was the launch of Ethereum, a programmable general-purpose blockchain. Whereas bitcoin has a single very specific use (virtual currency), Ethereum proposes to take the decentralized and trustless aspects of bitcoin, and enable execution of complex programs and smart contracts, acting as a sort of "world computer". Various applications are being built on the Ethereum platform, and also other blockchain projects.

There are no clear definitions for many of the concepts related to blockchain technology. The term *blockchain* first started to be used with the bitcoin blockchain, although Nakamoto's original paper doesn't explicitly mention it. One view is that the original bitcoin is the only blockchain, and others are mere variations of the same technology. Ethereum founder Vitalik Buterin (2014) has dubbed this approach "bitcoin maximalism" and states: "Bitcoin maximalists often use "network effects" as an argument, and claim that it is futile to fight against them". The bitcoin maximalist approach however ignores the fact that bitcoin is a specific tool for a specific job, and in other cases some other type of implementation may work better.

Another, broader view often includes more than just bitcoin. Deloitte (2016) include in their definition all blockchains built on the technology, and specifically state that they don't have to be built on the bitcoin architecture. Their definition is also based the blockchain being distributed: in a centralized model information exists in one place, and in a replicated one in many, but it originates from one central authority and the others are merely copies. In a blockchain, information is distributed so that everyone can consume but also produce information. This scope of read-write access will be covered in more detail in a later chapter.

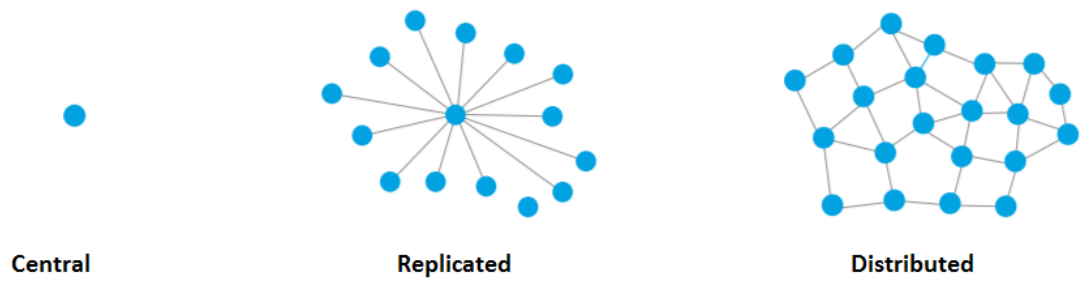

Figure 1: blockchain is a distributed architecture

The term *distributed ledger* has started to be used alongside the term blockchain, especially in the financial sector. One of the reasons for this difference in terminology is that when financial institutions first started exploring uses for the technology, they wished to distance themselves from the negative reputation of bitcoin that came from the fact that besides its many legitimate uses, it's sometimes also used by criminals. Strictly speaking, it could be said the while blockchains are usually distributed ledgers, not all distributed ledgers have to use all aspects of the blockchain approach: The terms are however often used interchangeably, since there are yet no set conventions in the industry.

Because blockchain is so new, one of the key considerations is what exactly even qualifies as a blockchain. Gideon Greenspan (2015) argues that the exact definition doesn't matter, because while there are many different models, they all share a sufficient number of technical similarities, and have different but useful applications. But besides from a definition perspective, the question is important in order to be able to understand what are the key value propositions of the technology, how do they change and what trade-offs have to be made regarding them when the different components are modified. To evaluate the different kinds of blockchains that could exist, I propose a simple set of questions. Answering the following can give an idea about what the blockchain will look like and how it will function:

- What data is recorded on the blockchain?
- How "smart" is the blockchain?
- Who can participate?

- How is consensus reached?

The data stored can be simple or complex. In the case of bitcoin and various cryptocurrencies, it is payment transactions, but any asset that is digital could be stored on the blockchain. Examples could include ownership records, land titles, copyright, health records etc. The data could either exist purely in digital form, or be in reference to real world assets that exist separately.

The blockchain can be "smart" to varying degrees. Bitcoin includes a simple scripting language, allowing users to embed various conditions in transactions, but it was Ethereum and the Turing-complete programming language contained therein that enabled a truly smart blockchain. This in turn enables "smart contracts", where the logic is automatically executed on the blockchain (Buterin 2013). A simple example of this could be a bet, where payment is automatically made when the outcome is known, or purchase of an asset where both money and ownership are stored on the blockchain and automatically change hands without the need for an intermediary.

Participation in a blockchain can be either open or permissioned. In an open blockchain like bitcoin or Ethereum, anyone can participate (Buterin 2016). In a permissioned blockchain the participants are known or preapproved, which could be the case for example in a blockchain used by a consortium of banks. This is often a point of contention, many arguing that private blockchains are sufficiently different that they shouldn't be considered blockchains at all (O'Connell 2016).

An important aspect is the consensus mechanism. This is the way the participants decide on valid transactions, or what is the correct state of the distributed ledger (Swanson 2015). Several designs exist, each with its own trade-offs. One example is the proof of work scheme used in bitcoin, where participants expend computing power to solve mathematical problems. The trade-off in this case is between security and cost – proof of works has been shown to be very secure, which is needed in a trustless environment where the other participants are not known, but

it's also very power-consuming. The details of this will be explored in more detail in a later section.

It can be argued that at this early stage, a clear definitions is not yet required. Some of the difficulty with definitions arises from the fact that blockchain is both a technological and economic innovation. It relates to areas like database architectures, but also to more abstract concepts like trust and value. Common terminology can be expected to emerge over time, and it will help with the discourse. For the purpose of this thesis, we will consider blockchains from a broad viewpoint, accepting both bitcoin and non-bitcoin, as well as public and private systems to be blockchains. A sufficiently broad definition is presented by Pascal Bouvier (2015), which allows for many different kinds of implementations to be observed as blockchains:

> *Definition of Blockchain: A type of distributed ledger that comprises two objects; transactions and blocks. Transactions are the data and blocks are the records that order the confirmation of the data.*

1.2 Problem formulation

Digital innovations can be described as going through the following four phases: discovery, development, diffusion and impact (Fichman et al. 2014). Glaser (2017) places blockchain currently in the second phase, development, the first phase having happened around 2015 when the technology gained mainstream awareness. The relevant managerial questions to ask in the development phase are: "What constitutes the digital innovation's core feature set?" and "To what potential organizational uses can it be put?".

The first goal of this thesis is to understand blockchain technology, including the different ways it can be set up, its value propositions and limitations. Because the technology is so new, understanding how it works at least on a basic level is the first step to understanding how it can be used and what implications it can have for existing and upcoming business models. This presents the following research question:

How does blockchain technology work, and what are its value propositions and limitations?

The second goal is to find out what the technology can be used for. The large amount of hype has led to the potential of the technology perhaps being overestimated, and being applied in many areas regardless of whether it actually provides new value in that specific area. It has even been called a solution looking for a problem to solve, especially when it comes to private blockchains. But it can be expected that not all the use cases are fully understood yet, and it will not be known if the technology can solve a particular problem until someone tries it and attempts to build a solution. Understanding the technology can however at least give an idea about what's viable, giving the second research question:

What are suitable uses for blockchain, and what are the things that have to be considered when deciding on a use case?

1.3 Methodology and paradigmatic assumptions

Saunders, Thornhill & Lewis (2009) define an exploratory study as a means to finding out "what is happening; to seek new insights; to ask questions and assess phenomena in a new light". It is useful when the particular nature of a problem is not known, and therefore this methodology is used is used in this thesis, as blockchain is a relatively new phenomenon and the concepts are still emerging.

Exploratory studies are often based on available literature and secondary research, which is also the approach taken in this thesis. As discovered in the literature review, blockchain has been studied to a certain extent in academic literature, but because the technology is new and rapidly evolving, not all of the current issues are covered. Therefore various practitioner and consultancy papers will also be used to build an understanding of the technology and the current issues and opportunities related to it.

As with any research project, the paradigmatic assumptions have to be considered. The positivist philosophy is based on the assumption that social reality can be quantified and be used to produce law-like generalisations (Saunders 2009). Blockchain, however, is inherently a human construct, and is dynamic, I.E. understanding of it can change over time, place and culture. Therefore the interpretivist approach is assumed in this study, as the phenomenon is not even clearly defined yet. According to Orlikowski & Baroudi (1991): "Interpretive studies assume that people create and associate their own subjective and intersubjective meanings as they interact with the world around them". Orlikowski & Baroudi go on state that in this approach generalization to a population is not the goal, but instead it is to understand the deeper structure of a phenomenon. This is in line with the exploratory nature of this study, and definitive answers are not necessarily expected to be found.

The credibility of the study must also be assessed. According to Saunders et al. (2009), this should be done by scrutinizing the reliability and validity of the research. The former can be understood as the degree to which the methods used will produce consistent results, and the latter as the degree to which the findings really answer the questions the research is about. In a study based on secondary research it is important to try gain a holistic understanding of the phenomenon. As we shall see, in blockchain there are various philosophical standpoints on what the technology is and what it should be, some of them taking a very strong approach on the expense of pragmatism. Every effort has been taken to avoid bias and compare these viewpoints where necessary, and not to over represent one over the others.

1.4 Literature review

Blockchain has been researched from an academic perspective since it first entered the scene. Much of the early research was naturally focused on bitcoin, its technical and economic aspects and role as a currency. Yli-Huumo, Ko, Choi, Park & Smolander (2015) have done a systematic review on the current state of research on blockchain technology. Their findings show that as of 2016, a majority of research is

still focused on bitcoin, only 20% dealing with other blockchain applications. Much of the bitcoin-related research is focused on the security, privacy and scalability challenges of the system. The authors note that added focus on blockchain technology is expected to result in increased number of studies in the near future, and suggest various directions for this research.

Because of the time it takes for academic research to go from inception to publication, it is perhaps not surprising that much of the current published academic research is focused on bitcoin. Focus has only in the last few years shifted to more general applications of blockchain. Many of the issues are however covered by various professional literature and consultancy papers, including reports from companies and organizations exploring the use of blockchain technology, like Accenture, Deloitte, IBM, Morgan Stanley and the European Central Bank. While not peer-reviewed, these reports provide valuable insights into some of the questions not yet covered by scientific papers. They also provide a higher level view of the strategic expectations for blockchain, although care has to be taken to assess the reliability of these source materials, as not all of them display a sound understanding of the technology, leading to claims that are unfounded or not based on the real aspects of blockchain.

It is perhaps telling of the blockchain phenomenon that most of the literature used in this thesis was published within the last two years. 2016 saw also the publication of the first issue of Ledger, a peer-reviewed scholarly journal that focuses specifically on blockchain-related issues.

2 Blockchain technology

In this section I will explore how blockchain technology works in order to build an understanding of what's possible with the technology. This will be done by looking at the technical aspects of the bitcoin blockchain, and extrapolating from these the aspects that are common to blockchain architectures. Lastly, the key value propositions and the trade-offs they come with will be identified.

2.1 Underlying technologies

In order to understand how blockchain works, it is necessary to understand some of the underlying technologies, an overview of which will be provided here. A level of technical understanding is assumed from the reader, but these initial explanations should help better understand blockchain. None of these technologies are particularly new –they predate blockchain by decades– and they also have various applications outside blockchain use. To a regular user of many modern IT applications, they are the "hidden" workhorses that make the technology run.

2.1.1 Public key cryptography

Public key cryptography is an encryption scheme that uses two keys: a public key and a private key. Public keys are used to encrypt data, and can be disseminated widely, while private keys are used for decryption, and as the name implies are not shared. The two keys are linked, meaning that messaged encrypted with a particular public key can only be decrypted with the according private key. In order for the scheme to be secure, it should not be possible to calculate a private key based on the related public key (Salomaa 2013).

2.1.2 Digital signatures

A digital signature is based on public key cryptography, and is used to verify the authenticity of a message. Its purpose is to act much like a real signature: when a document or message is signed with a private key, its authenticity can be verified with the corresponding public key. It's also used to ensure that the message received corresponds to the original message that was sent, and hasn't been altered along the way (Salomaa 2013).

2.1.3 Cryptographic hash

A cryptographic hash is a function that takes an input of any particular message or set of data, and maps it to a fixed-length output. The action is performed

one way, meaning that unlike in the aforementioned public key cryptography, the output can't be decrypted and the original message revealed. There are many different hashing algorithms available, and they have various different uses in computing today, for example storing passwords securely. One such algorithm, which is also used by the bitcoin blockchain, is SHA-256. It can be exemplified by producing the hashed output for the inputs "test message1" and "test message2". Although the inputs differ only by one character, the outputs are completely different:

Input	Output
`test message1`	`b7c6fd34d91f01d2a3e9322cc7e9fa72f83254d8fe770693775038fa065e7141`
`test message2`	`545a5a70fc06a7d211f1144f4631160d76104c7116c0f9e0cc8cd917f1b8c29d`

Some requirements for an ideal hash function are (Abidi & Kahri 2014):

- It is deterministic, meaning that the same input will always result in the same hashed output.
- It is one-way, meaning that it's infeasible to calculate the input from the output. The only way is to try all the different possible inputs.
- It should be collision-free, meaning that it's infeasible to find two different inputs resulting in the same output. In theory the possibility exists, but is extremely small.

These properties have several advantages that make hash functions useful for blockchains. This will be explored further in the following section.

2.2 The bitcoin blockchain

On a high level, bitcoin is simply a digital currency that allows transactions to be made without using an intermediary. Like normal fiat currencies, it doesn't have any intrinsic value. It can be used as a medium of exchange for goods and services in

the real world, and it only exists in digital form. The system was first proposed in 2008 whitepaper by Satoshi Nakamoto. The software itself is open source, meaning that anyone can adapt it to their own use, but changes to the bitcoin core protocol can only be proposed by the Bitcoin Foundation – a US based non-profit that supports development of the software. There have been various updates to the bitcoin software along the way. The foundation can propose changes, but for them to be adopted, they must be accepted by a majority of bitcoin users. Therefore, it is possible to change the rules that will be described in this section, according to which the network operates.

2.2.1 Addresses

Similar to an e-mail address, a bitcoin address is used for messages, or in other words to send and receive bitcoin transactions. An address represents the public key, meaning that it must be known to send transactions to it, but doesn't grant access to the bitcoins stored at that address. For this, the private key is required. As the private key provides access to the account, it must be well guarded – revealing it to others would grant them access to the account, and losing it would mean losing access to the account, any bitcoins on it would effectively be lost. (Pilkington 2016)

2.2.2 Transactions

Bitcoin transactions are stored in a distributed ledger on computers in the bitcoin network. In order to send a payment with bitcoins, a user must specify a recipient address and a payment sum. The user then broadcasts the transaction to the network, where it is applied to copies of the ledger on the various nodes in the network. This makes all bitcoin transactions public, as unlike a centralized ledger maintained by a bank or other single institution, the distributed ledger exists all across the bitcoin network. Transactions are encrypted with the user's digital signature, which is meant to ensure that the user making the transaction is the true owner of those bitcoins and has a right to use them. (Nakamoto 2008)

Like when using a traditional ledger, a bitcoin address must have enough money in order to make a transaction. But a bitcoin address doesn't keep an account

balance, and instead when sending money it must reference previous transactions to that wallet. When bitcoins are “spent”, this is then added to the distributed ledger, which in fact contains all bitcoin transactions ever made, all the way to the first one. This prevents double-spending, since once a previous transaction has been used to source a payment, it can’t be used again. (Nakamoto 2008)

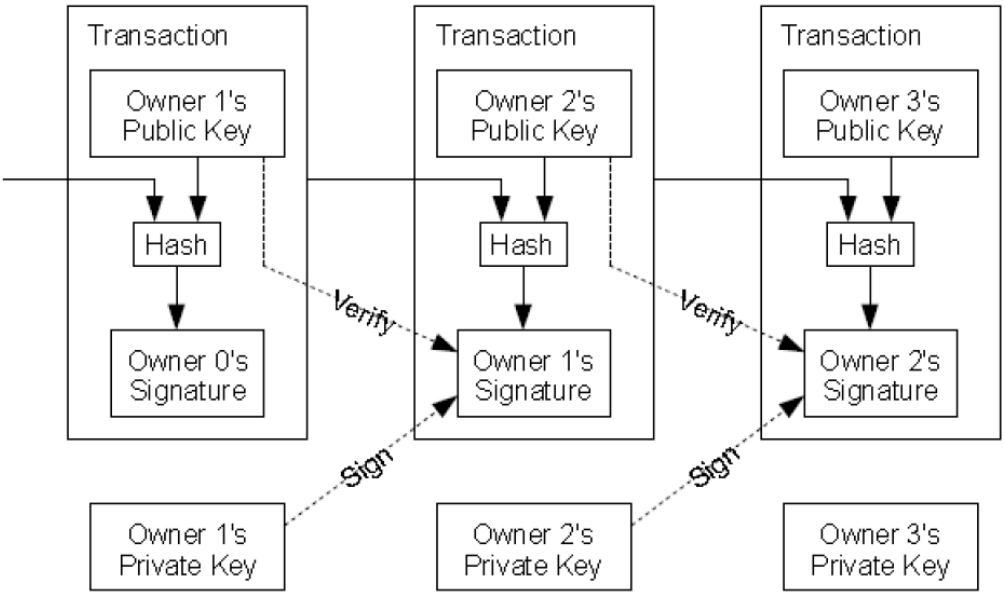

Figure 2: Transactions signed by private keys

2.2.3 Blocks

When bitcoin transactions are made, they are first considered unconfirmed. They are then placed into blocks by nodes in the bitcoin network. A block contains a list of included transactions, and a block header that contains metadata about the block. The transactions in a block are arranged in a data structure called a Merkle tree. This contains the transactions hashed with the aforementioned SHA-256 algorithm, which are then further hashed to provide one root hash for the block. The advantage of this data structure is two-fold:

1. It creates a single identifier (the root hash) that can be checked to make sure that the block is valid. This is useful because the block can contain hundreds of transactions, and verifying each underlying hash separately would be very labor-intensive for the network.
2. It secures all the underlying transactions. As demonstrated before, even a slight change in the input for a hash will dramatically change the output, meaning that altering any single transaction would cause the root hash to change, rendering the block invalid (Nakamoto 2008).

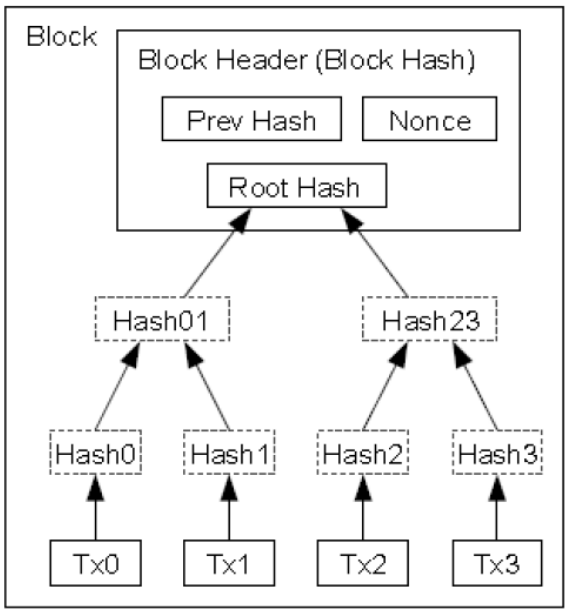

Figure 3: The Merkle tree of a bitcoin block

2.2.4 The blockchain and mining

The blocks are linked together and arranged in order by time, forming the blockchain. In order for a block to be valid, it must reference a previous block, contain only valid transactions have a block header that is below a certain target. The *target* is constantly adjusted by the network to make sure that a new block is created on average every 10 minutes.

To illustrate how the target works, we can consider the following: as described before, output of a hash is a fixed length text string. The target requires the

value of the output to be below a certain number, i.e. have a predetermined number of leading zeroes. Because there is no way to predict how an output will look, the only option is to try many times with different inputs until a correct one is found. This is done by changing the *nonce*, which is a number that is added to the data (see fig. 2). If we take some arbitrary data, for example the text "Hello, world!", and assume the target is to have four leading zeroes, then different nonces can be tried until a solution is found. I have demonstrated this in the following table by running the inputs through the SHA-256 algorithm:

Input	Output
`Hello, world!0`	`1312af178c253f84028d480a6adc1e25e81caa44c749ec81976192e2ec934c64`
`Hello, world!1`	`e9afc424b79e4f6ab42d99c81156d3a17228d6e1eef4139be78e948a9332a7d8`
…	…
`Hello, world!4250`	**`0000`**`c3af42fc31103f1fdc0151fa747ff87349a4714df7cc52ea464e12dcd4e9`

Figure 4: Hash output

In this example different nonces were tried starting from 0, and it took 4,250 tries for a correct nonce to be found. This establishes the mechanism by which bitcoin transactions are confirmed and the network reaches consensus. This consensus mechanism is commonly referred to as *Proof-of-work*, and is the underlying way that enables bitcoin transactions without trusted intermediaries or third parties (Becker et al. 2013). Changing any prior transaction in a previous block would render that block and all following blocks invalid.

The incentive for bitcoin network nodes to confirm transactions in the aforementioned manner and to thereby extend the blockchain comes from a feature in the system called *mining*. When a block is formed, it yields a small amount of bitcoins called the *block reward* for the node that found the correct nonce and formed the block. This mining is the only source of bitcoins being added to circulation, offering a financial incentive for the miners to validate and secure the network.

Increased competition has led to specialized computer hardware being designed to make the mining as efficient as possible (Kroll et al. 2013).

Individual miners can pool their resources and share the rewards accordingly. This decreases the volatility of rewards. If a miner commands a miniscule portion of the total mining power, then they will not find the correct block very often, but when they do they get to keep the whole block reward for themselves. However, if they pool their resources with other miners, they increase their chances of winning, but only get a portion of the block reward according to the resources they contribute. (Walch 2015)

Pooling has led to concerns over increased centralization of the bitcoin network. This could enable the theoretical 51% attack, which assumes that if a malicious party controls over 51% of the network, they could reverse transactions and control which new transactions are added to the blockchain. As of December 2016, just five of the largest mining pools control over 50% of the bitcoin mining power (Blockchain.info 2016). They are however disincentivized from tampering with the network by the fact that they are heavily invested in bitcoin, and any security breach, whether real or perceived, would cause a decrease in the value of bitcoin.

The number of bitcoins awarded from successfully completing a block is halved approximately every four years. Eventually the reward will shrink to zero, and the built-in limit of 21 million bitcoins in existence will be reached. Because the total number of bitcoins is limited, and some are removed from circulation through private keys being lost, the currency may end up deflating. (Pilkington 2016) It is in fact designed to mimic the supply of gold – there is a limited total amount available, and as we get closer to that limit, extraction becomes more labor-intensive (IIF 2015).

When the total limit of 21 million bitcoins is reached, there will still be an incentive for nodes to process transactions through transaction fees. Nodes have a right to prioritize which transactions they process first, and they can do so according

to the transaction fees attached to them. If there is no fee attached, the transaction might never get processed. (Nakamoto 2008)

2.2.5 Forks

Because there are many nodes in the network and they are all competing to form a new block at the same time, two or more of them could potentially simultaneously form a block. Both blocks would be broadcast to the network, suggesting that they be added the chain and the transactions in them be considered confirmed. The blockchain then momentarily branches, until a new block is formed which breaks the tie. The transactions from the discarded block are placed back in the pool of unconfirmed transactions, where they are again available to be used to form new blocks. This means that transactions further back in the blockchain are more 'secure', as it is less likely that their chain will be overtaken by a different branch. (Walch 2015) This mechanism is called forking, and has several implications for blockchain applications, which will be discussed later.

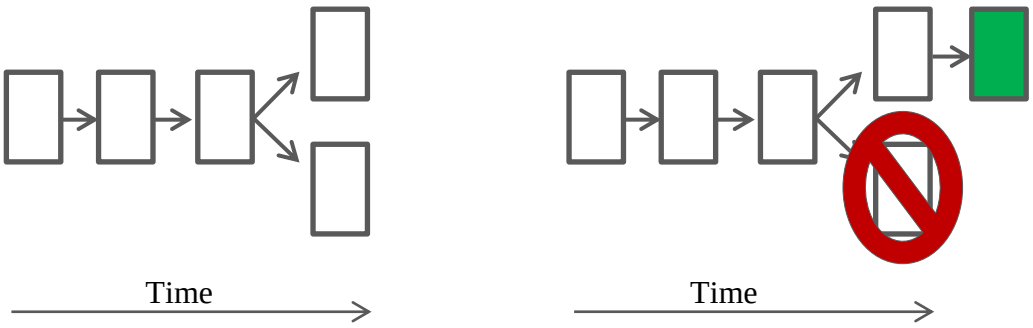

Figure 5: The blockchain forking

The design of the bitcoin system means that it has several features which make it an attractive payment mechanism. First, no trusted intermediaries or third parties are required. A normal user can send money to someone on the other side of the world almost for free and relatively quickly. Second, by following proper procedures, a much higher degree of anonymity can be achieved than with most

traditional payment methods. Coins are associated with bitcoin addresses, not private individuals. (Pilkington 2016) Transactions are however public and it is possible to connect them if some parties in the transaction chain are known. Additional steps can be taken by users to prevent this. Bitcoins also have a higher divisibility than fiat currencies, allowing for transactions of down to 0,000 000 1 bitcoins at a time. This could be useful for various cases where small microtransactions are required (Walch 2015).

2.3 Breaking down the components

Now that the bitcoin blockchain has been studied, in this section we can gain an understanding how the technology can be used in other blockchain applications.

2.3.1 Consensus mechanism

The above described proof-of-work is just one of many ways that consensus on the blockchain can be reached. In a permissionless and trustless network like bitcoin it has the advantage of security, but on the other hand it is very power-consuming to operate (Forte et al. 2016). Since only the winning miner or pool of miners who have found a block get to collect the block reward, the work of other miners towards that block is essentially wasted. As soon as a block is found, everyone then starts working from the beginning towards the next block. This leads to a great waste of resources, and it has been estimated that the entire bitcoin mining network currently uses as much electricity as the country of Ireland (Forte et al. 2016). Besides electricity, there is also waste in the hardware itself. Due to increased competition and decreasing margins, specialized hardware is often used that is designed to solve the bitcoin puzzle and not do anything else. As this hardware is upgraded, the old components have little use outside bitcoin mining pools (Ekblaw et al. 2016).

One proposal to make better use of proof-of-work-based mining is to utilize the computing power for something useful (Swan 2015). Instead of calculating SHA-256 hashes that end up being wasted, the work could be done towards socially beneficial distributed volunteer projects like SETI or Folding@home – the goals of

which respectively are to search for extraterrestrial intelligence, and to power medical research by simulating protein folding. The challenge with this approach is that these "puzzles" don't necessarily fulfill the requirements for a good proof-of-work puzzle. One of these is that the solution must be hard to find, but easy to verify, which is the case with hashing-based puzzles. (Narayanan 2016) Another problem is that basing the puzzle on data from a centralized source such as SETI could potentially compromise security of the network.

Another alternative to proof-of-work is a scheme called *proof-of-stake*. Proof-of-work works essentially like a lottery, where participants' chance of winning is proportional to the mining power they contribute. In the case of bitcoin the hardware has become very uniform, and any differences in efficiency are minimal and usually related to the cost of electricity. The premise of proof-of-stake is that instead of using money to buy hardware that is then used to buy "votes", participants could instead stake their money directly, and consensus could be reached through a kind of virtual mining. The obvious advantage of this approach is that it removes the waste of hardware and electricity from the equation (Narayanan 2016).

In a proof-of-stake system participants would have to be incentivized to behave honestly because they are invested in the system, and doing otherwise would decrease the value of that blockchain, the same way as bitcoin mining pools tampering with the bitcoin blockchain would. Proof-of-stake includes a fair share of drawbacks, which are actively being worked on. It is however not a proven concept yet the same way that proof-of-work is, considering it has been used on the bitcoin blockchain since 2009 without a major security incident. The Ethereum blockchain is planned to be moved to proof-of-stake sometime in 2017 (Buterin 2016a).

In a permissioned blockchain there is no need for some of the incentives that are present in the aforementioned examples. Some of the security mechanisms that are needed in an open trustless environment are not needed, and participants can be given the right to validate transactions based on their known identity or contractual relationship with the network. A blockchain operated by a group of companies could

operate on this principle, known as *proof-of-authority* or *proof-of-identity (*Pass & Shi 2016).

The above described are just some examples of how consensus can be achieved on the blockchain. Ultimately the exact mechanism depends on the use case, and all of them have various trade-offs associated with them. On an open blockchain security is critical because the possibility of malicious participants is higher than in a known network, even if it makes the network slower and more costly to run. The consensus mechanism is one of the key components of the technology, because it dictates the way the participants decide what is the correct version of the shared database.

2.3.2 Transparency

From our example of the bitcoin blockchain we can establish that all transactions on the blockchain are visible to all users. This is an important feature, because establishing the validity of a transaction is based on referencing previous transactions, and these have to be known for verification to be possible. From the perspective of different use cases, this can be either a positive or a negative feature. For bitcoin users, this makes it necessary to take extra steps to protect one's anonymity and payment history. Although addresses are not directly linked to a person's identity, with the whole history of transactions being visible, even just one known address being available can make it possible to connect the dots and reveal a person's transaction history (Reid 2013).

This transparency has implications for various use cases. It can be beneficial in a blockchain used for supply chain or provenance tracking use, where the state and history of various objects should always be known. Also for regulatory and auditing use it can be beneficial for various actors to be able to inspect the blockchain for the true nature of events, as opposed to having to go through books and records in the traditional way (Pilkington 2015). But in some cases where blockchain participants also compete with each other, as could be the case with multiple banks operating a

shared blockchain, sharing all transactional data would not be viable because it would give others access to data that provides competitive advantage

There are two ways in which transparency of data on the blockchain could be reduced. The first is storing it encrypted or only storing reference data while the real data resides outside the blockchain. This however introduces the requirement to store either the encryption key or the actual data in a central location, weakening the case to use a blockchain system in the first place.

The second way to reduce transparency is the cryptographic concept of zero-knowledge proof. This entails masking the transactions in a way that they are provably valid without revealing the contents of the transaction (Garman et al. 2013). Such systems have been proposed for electronic voting, where it should be made certain that a person has a right to vote, and can only vote once, without revealing who they voted for. In the blockchain world the concept is used by the cryptocurrency Zcash launched in 2016, which works much like bitcoin but with the promise of real anonymity. The technology is still in its infancy, but represents an interesting experiment into mitigating one of the perceived drawbacks of blockchain.

In a cryptocurrency use case like bitcoin, transparency has several implications. Bitcoin is sometimes used for illegitimate purposes like ransomware, money laundering and tax evasion. Regulators have woken up to this fact, and in the United States the IRS has filed for permission to identify bitcoin owners at the bitcoin exchange Coinbase in order to investigate potential tax evasion (Phillips 2016). Likewise, the EU parliament has proposed legislation that would require bitcoin users to be identified in order to prevent money laundering and terrorism financing (Coleman 2017). From a societal perspective these are clearly beneficial goals, but they come at the expense of privacy. It can be argued that privacy also has benefits for users: it can provide a monetary safety valve for people living under aggressive capital controls and rapidly inflating currencies, such as is the case in Venezuela, and protect people from oppressive regimes that could illegally seize their assets. A

completely anonymous cryptocurrency would affect both scenarios: on one hand it would protect legitimate users, but one the other hand it would also aid criminals.

2.3.3 Settlement finality

An important aspect of value transactions is the concept of settlement finality. This refers to the idea that once a transaction is completed, it's completed for good and can't be changed. This is important from a business perspective, as decisions have to be made based on the assumption that a transaction is final. Especially in the financial sector it's crucial, but even a small business must be able to decide if they should deliver goods without risking the payment for them being reversed. Various legal frameworks exist for ensuring settlement finality and guaranteeing the smooth operation of payment systems even in cases where one party is insolvent, for example EU Directive 2009/44/EC. Problems in settlement finality could cause a domino effect of problems for other participants in the system. There is debate over whether public or private blockchains can act as an adequate settlement mechanism.

Tim Swanson (2016) argues that public blockchains can't act as an acceptable settlement mechanism. Looking at the previously described example of the bitcoin blockchain, obstacles to settlement finality arise from two things: the consensus mechanism and the possibility of forks. The older a transaction is, the less likely it is to be reversed, and an accepted rule of thumb is usually to wait six confirmations, i.e. one hour for considering a transaction settled. Settlement is therefore probabilistic, and as more blocks are added to the chain, the chance of the transaction being reversed approaches zero. In extraordinary circumstances even relatively old transactions can be reversed, as was the case when a bug in the blockchain was fixed in 2010 causing half a day's worth of transactions being reversed. The acceptable level of certainty depends on the size of the transaction: 99% chance of settlement finality might be acceptable in the case of payment for a cup of coffee, but not for signing over real estate or a car.

According to Vitalik Buterin (2016c), from a philosophical standpoint there is no system that offers true 100% settlement finality. Whether a system is based on blockchain or is fully centralized, extraordinary circumstances could arise in either case which introduce settlement risk to the system. Ultimately it is the role of courts to decipher intent if problems arise and rule on ownership. As discussed earlier, strict legal requirements exist regarding the degree of settlement finality that must be ensured by a system used for payments. Therefore it seems unlikely that a public blockchain could be adopted for uses like interbank payments or financial markets. A private blockchain on the other hand could be built to specifications that meet these criteria, and find use in areas where public ones can't be used.

2.3.4 Smart contracts

One useful aspect of blockchain technology is the concept of smart contracts, and many of planned use cases rely on this feature. The concept has existed since before blockchain technology, blockchain has made improvements in making them more feasible and enforceable due to the immutable and distributed nature of the technology. Several uses for smart contracts have been proposed. These range from automatic payments of dividends, to assurance contracts and trade finance. Many of the blockchain use cases observed in the next chapter rely on smart contracts to some degree.

A smart contract can be defined as a contract attached to software in a way that the terms of the contract are automatically executed (Kolvart, Poola & Rull 2016). They can range from simple to complex, and have potential especially in the financial sector. The main advantage over a normal contract is the improved efficiency, since contract terms are self-executing. This could allow for leaving out many of the parties traditionally present when making contracts, like banks, lawyers and consultants. There are however many questions to be answered still, regarding both the technical and legal aspects of such contracts.

The bitcoin blockchain includes a simple scripting language that can be used for basic smart contracts. A simple example of this could be a transaction that requires two out of a possible three signatures to be completed, setting up multi-party verification. Bitcoin script has several limiting features however –it does not allow loop functions for example– and was never intended to be a smart contract platform (Kumerasan & Bentov 2014). Other blockchains, like Ethereum on the other hand were built with smart contracts in mind. Ethereum includes a Turing-complete programming language, meaning that it is computationally universal, and therefore generic enough to solve any reasonable programming problem (Morini 2016). This allows for much more sophisticated smart contracts than bitcoin script.

Smart contracts could be used in many cases that require transactions based on some known logic or conditions. A simple example could be facilitating a purchase of a used car – an everyday scenario where a lack of trust can exist between parties. The seller wants to make sure they get paid before signing over ownership of the car, and the buyer wants to make sure they get what they are paying for. If both the money and title to the car were stored on a blockchain, then a simple smart contract approved and digitally signed by both parties could make the two assets change hands without a trusted intermediary. The agreement is then validated, and when it is disseminated to the blockchain it can't be modified or canceled. From here various more complex implementations can be designed.

One of the issues regarding smart contracts is that since they are based on programming, everything will have to be explicitly defined. From a contractual perspective this is a good thing, as the clearer the terms of the contract are, the less room for interpretation there is in the terms. In the real world however, contracts often include legal expressions that carry the weight of history and precedent, but are still up to interpretation to a certain degree (Swan 2015). Therefore the more complex the business and technical case where the contract is being applied, the more complex and harder to interpret the resulting program will become.

Another aspect of smart contracts is the legal interpretation of them. In a legal sense, a contract is an agreement between parties to enter into some relationship or conduct some transaction. Smart contracts are an extension to this, but according to some views they can't be interpreted as legal contracts, and would instead exist in parallel to them (Kolvart et al. 2016). Therefore the enforceability of them in a potential dispute may be on loose footing. Different legal systems also have different definitions for what constitutes a contract, and so parties acting across legal boundaries would have to consider the requirements of both systems when setting up their smart contracts. Existing work towards common standards, like the Principles of European Contract Law applied across EU should however make this easier (Kolvart et al. 2016).

2.3.4.1 The DAO

An example of the real-world issues that can happen in smart contracts is the case of the DAO. The DAO, short for Distributed Autonomous Organization was a project on the Ethereum blockchain with the goal of creating a decentralized crowdfunded organization where the participants could decide what the pooled money would be used for. Various projects were proposed, both commercial at not-for-profit ones. The rules of the DAO were programmed into a smart contract, and the project gathered over 150 million USD worth of Ether, the native Ethereum blockchain currency, in funding (Atzei et al. 2016).

Despite a thorough review process before launch, The DAO was eventually subject to a security exploit allowed an unknown hacker to extract the funds stored in the smart contract. The loophole in the underlying code had not been discovered until the project was live and funded. A discussion followed over what should be done, and the Ethereum community decided to fork the blockchain back to before the DAO was launched, effectively rolling back all changes and reverting to an earlier state (Bradbury 2016). This is significant because the security issue was in the DAO and not in the Ethereum code itself, which was just the platform that the DAO smart contract ran on.

The case exemplifies several concerns in the concept of a blockchain smart contract. First, it is extremely difficult if not impossible to design software without any bugs in it. Despite thorough testing and quality assurance procedures, many of these are not found until the software is run in real-world conditions it is built for. In normal IT systems these can be fixed in updates, and developing such systems is an iterative process. In the DAO on the other hand the governing rules could not be changed because of the immutable nature of the blockchain. In hindsight it is easy to see how this combination of faulty software managing a record amount of crowdfunded money could lead to a disaster

The second issue relates the precedent the case sets for dispute resolution. As mentioned, the vulnerability was in the DAO and not the Ethereum platform, and further, the terms stated that the underlying code of the DAO was the rules that the organization would abide by. Therefore the hacker acted in accordance with the rules, even if not the spirit of them. Arguments were that the platform should not be rolled back even for such a large loss, because it would compromise the decentralized and immutable nature of the blockchain, and weaken trust in it (Atzei et al. 2016). The record amount of money almost certainly had something to do with the decision, and comparisons to bail-outs of banks were made. Despite this, the community eventually opted for the rollback, setting an important precedent for future blockchain projects.

The DAO rollback also highlights an important point for all blockchain implementations: the blockchain is immutable only when it's decentralized. If the participants can come together, there is nothing stopping them from changing the underlying rules or rewriting the history of transactions.

2.3.5 Comparing public and private blockchains

Blockchain started with bitcoin, a fully open and permissionless network where anyone can participate. As discussed in the previous chapter, an open blockchain has several aspects that make it useful for a cryptocurrency payment system: decentralization, pseudonymity of participants and resistance to central

control. Over the last several years more or less private blockchains with controlled access have emerged both as an idea as well as examples being built by various companies in the space. Since the requirements for these private blockchains are different than those for a decentralized cryptocurrency, as are the restrictions placed on the various stakeholders that engage with them, it's only natural that they are configured quite differently from their public counterparts. In fact, many of them have little to do with cryptocurrencies, besides the shared origins of the technology.

Reception of private blockchains has been mixed. On one hand, they have been well received by companies who wish to build use cases on them and can't for some reason use the public blockchains. On the other hand, they have been said to "not solve any major problems and not having a high chance to succeed" (Rizzo 2015), and being to open blockchains what closed Intranets were historically to the open Internet, in that they may have some value, but will not lead to any kind of revolution or disruption (Scott 2016). Many of these views seem to arise as much from philosophical standpoints as from purely practical ones, considering private blockchains a last-ditch effort by dinosaurish middlemen to resist disruption and disintermediation in their industries.

On the other hand, approach to open blockchains like bitcoin has in some cases been equally unenthusiastic from the private sector, companies shying away from the negative reputation of it but investing heavily in private blockchains. Bitcoin expert Andreas Atonopoulos has equated this to "the horse-carriage association of America announcing that they will adopt the core technology of the automobile: the pneumatic tire". However potentially viable use cases exist for both, although the open blockchain in bitcoin is more mature, having been around longer.

From an openness perspective blockchains can broadly be categorized in 3 different groups (Buterin 2016b):

- Public blockchains – ones where anyone in the world can participate, having access to both read the data as well as submit their own

transactions. The economic incentives and cryptography ensure security and immutability of the ledger. These can be considered fully decentralized.

- Consortium blockchains – ones where only pre-approved participants may send transactions and read the transaction data. This could be for example a consortium of 15 banks, where 10 have to agree for a transaction to be valid. These can be considered partially decentralized.
- Private blockchains – ones restricted solely to a single organization, for example a company. These can be considered centralized, although considering an inside view of an organization there may still be various parties present with different interests.

As we can see, read and write access to the blockchain can be decided on separately. In some cases they might be restricted to the same level, while in others read access could be extended to a larger group, like an auditor or regulator or even the general public. This could be the case for example in a land registry blockchain, where write access is restricted, but ownership of land being public knowledge, read access is available to everyone. The line between private and consortium blockchains is often blurred, and in many cases they are both considered to be a part of permissioned blockchains, while completely open blockchains form the other permissionless group, splitting the technology into just two groups.

Permissioned blockchains have the advantage that the organization running them has the ability to modify them. This means that if necessary and agreed on by the parties, they can modify transactions or change records (Parker 2016). In some cases this is an absolutely necessary requirement. For example in many financial system or public sector applications, such as payment systems or land registries, control is needed to be able to correct fraudulent or illegal transactions. Although these should not be allowed in the first place, it's still possible that for example funds

that are first thought to be legally acquired are later found to be profits from criminal activity, and there has to be a way for authorities to seize them.

Since the participants in a permissioned blockchain are known, transactions don't necessarily have to be verified by all nodes, leading to lower transaction costs. There is also less risk of a majority group taking over the network, as could happen in a permissionless blockchain if the network becomes too concentrated (Parker 2016). Both of these advantages result from the fact that in a permissioned blockchain there is at least some trust between the parties, even if they may not fully trust each other. Therefore not all the security mechanisms are needed, which work very well in an open blockchain but make the network heavier to run. Parties would still be bound by contractual obligations as well as a desire to avoid the potential reputation loss that would result from fraudulent activity.

If permissioned blockchains offer a degree of control for the different parties, permissionless ones can be considered truly decentralized in that no group owns the blockchain, or has the ability to alone make changes (Buterin 2016b). This can be seen in public blockchains like bitcoin and Ethereum, where updates and modifications to the technology can be proposed by somewhat central parties (the bitcoin and Ethereum foundations), but adopting them is ultimately up to the community. This means that if a change doesn't get accepted by the majority, it is not adopted. Therefore no central party has the power to force changes onto the platform. While control is an advantage and even a requirement in some cases, in cryptocurrency use this lack of central control only serves to strengthen trust in the system.

Another clear advantage of open blockchains is in network effects. By both first mover advantage and being open to anyone, bitcoin has already reached some major network effects, like being usable in most countries in the world (Buterin 2016b). If we contrast this with a consortium blockchain run by a group of banks, in order to use it a person would still have to be a customer of one the participating banks, limiting the potential for growth and network effects.

As we can see, both permissioned and permissionless blockchains have their own set of drawbacks and advantages. At this point it can't be said that one is clearly better than the other, and in some cases the requirements of the use case means the choice is already made.

2.4 Defining the value proposition

Now that the technology is better understood, we can consider the value proposition. This will be done from both a technological perspective and a business model perspective, comparing blockchain with existing database technology and studying how it can improve business processes.

2.4.1 Compared to centralized databases

Blockchains are often compared with centralized databases, because in some cases they can be used to perform a similar role. There are however differences in how the two systems operate, and these should be evaluated when considering a use case. There are also obvious differences in the maturity of the technology: relational databases have a long history, proven track record and large community of people who know their workings (McConaghy et al. 2016). In short, they are the known commodity. Blockchain, on the other hand, is still in an early stage, and many of the capabilities and issues haven't been discovered.

Blockchains can have an advantage over centralized databases when used over organizational boundaries. The advantage stems from the fact that whereas in a central database validity of the data must be enforced by a central authority, on a blockchain it's enforced by the cryptography and logic of the system itself. This protects the data from being erased or corrupted by a malicious actor, while a centralized database has a weak spot when a single human has write access to the data. From an organizational standpoint it can also free up the people and resources that would otherwise be needed to maintain a central database (Greenspan 2016).

A blockchain can also have an advantage in robustness over a centralized database. Because a copy of the database exists across all the nodes in the blockchain network, redundancy is practically built in (Greenspan 2016). Even if nodes are periodically disconnected from the network, once they rejoin they can form an accurate picture of what has happened in their absence. Because of the peer-to-peer nature of the network, multiple nodes can fail before it has an effect on the whole network. In centralized databases redundancy and duplication is also used, but such systems are expensive and difficult to build (Greenspan 2016).

One area where centralized databases would seem to have an advantage is speed. Blockchains have the added burden of running cryptography and processing data on all nodes in the network, whereas writing and reading a centralized database is comparatively fast assuming the data comes from a trusted and authorized source. The issue of data being visible to all participants is also in present in blockchain, which depending on the use case could be considered a problem or a benefit. These characteristics of blockchains and centralized databases can be summarized in the following chart:

	Blockchain	Centralized database
Disintermediation	x	
Robustness	x	
Data confidentiality		x
Speed		x

Figure 6: Comparing blockchain with centralized databases

2.4.2 Business process improvement

Blockchain has been compared to the Internet as an innovation that could improve business processes. Business Process Management can be understood as optimizing a company's business processes in order to save resources (Van der Aalst et al. 2003). In other words it can be described as achieving the same result for fewer

inputs. The Internet made it possible to automate many processes in the early 1990s, but it has been argued that the greatest value of an innovation doesn't come from the ability to automate existing processes, but when it enables fundamentally new ones (Milani, Garcia-Banuelos & Dumas 2016). If this is the case for blockchain, then the question can be asked: how can blockchain improve business processes? This business process improvement can be explored through blockchain in relation to the value-driven Business Process Management framework:

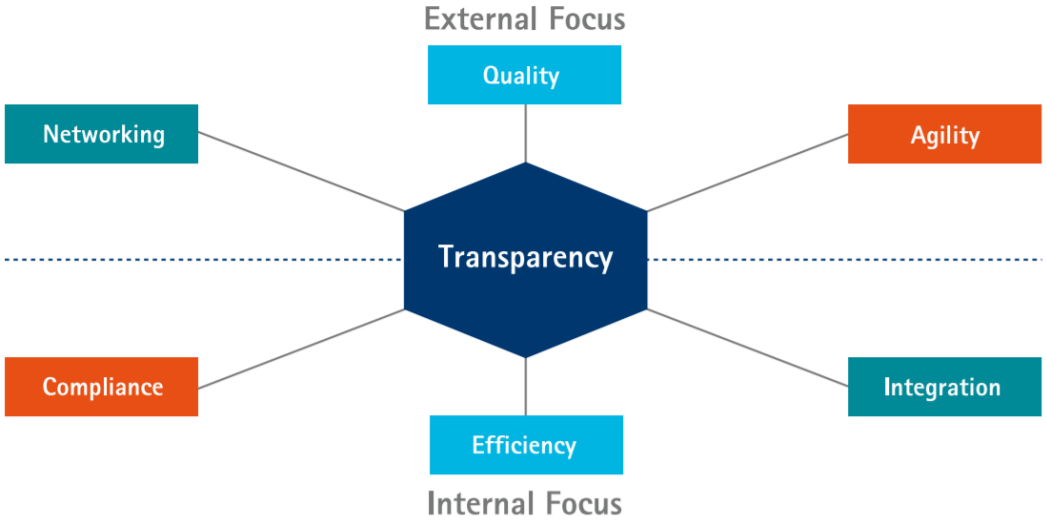

Figure 7: Value-driven BPM framework

The framework consists of three opposing value-pairs which organizations have to balance and through which they can attempt to improve their proceesses (Franz, Kirchmer & Roseman 2012):

- Quality–Efficiency: the choice between focusing on streamlining and efficiency or high quality.
- Integration–Networking: the choice between integrating and developing internal processes against networking and benefiting from external inputs.
- Compliance–Agility: the choice between being highly adaptive against complying with standards and regulations.

Many blockchain use cases are focused on reducing transaction costs, leading to increased efficiency. Especially uses in the financial industry could provide time and cost savings through reducing friction and disintermediating third parties. An example of this is the settlement system, where blockchain could provide shorter settlement times from the current industry standard of T+3 days. On the quality side, improvements from blockchain could come from better data quality, especially in industries like insurance or airlines, where data inputs come from many different sources and discrepancies can lead to bad customer experiences (Milani et al. 2016).

In the Agility–Compliance dimension, blockchain could improve agility especially in the public sector. As an example of this, the UK government is currently experimenting with blockchain-based welfare payments in order to provide more customizable payouts depending on the recipient's situation. Compliance benefits could be realized especially in the financial sector, which is subject to heavy regulatory requirements, which need to be carefully assessed and enforced both internally and externally (Milani et al. 2016).

One proposed use of blockchain is to act as a software connector, which could help organizations be more integrated by reducing data silos and facilitating better communication between systems (Xu et al. 2016). Especially the financial sector is notorious for legacy systems that have been due for an update for years. The same system could also facilitate better networking between organizations in many sectors. In the shipping industry for instance, as goods pass through various parts of the value chain a huge physical paper trail is generated. Shipping giant Maersk is currently experimenting with blockchain to digitize this paper trail, allowing the numerous stakeholders along the way to interact through blockchain (Allison 2016).

3 Blockchain use cases

In this section we will explore some potential use cases of blockchain. These were selected because they are scenarios where multiple parties are present who need

to share information, while being able to verify the validity of that information without necessarily trusting the other parties.

3.1 Use case: Healthcare

At present, various stages of systems exist in different countries for transacting electronic healthcare records (EHR), such as patient information. Some countries have or are developing centralized national EHR systems (notably Estonia, which is also experimenting with using blockchain products for the EHR system), while others, like the United States has moved only fitfully in this direction. Even in the United States various EHR solutions exist though: the Office of the National Coordinator for Health Information Technology lists 175 unique vendors who have supplied certified EHR systems to hospitals in the United States. This is a result of legislation emphasizing the adoption of EHRs over providing infrastructure to support them, which has resulted in "walled gardens" of closed, proprietary EHR systems (Burniske, Vaughn, Shelton & Cahana 2016). This is especially a problem in a large country like the United States with a large and fragmented healthcare system – smaller countries like the Nordics or Estonia have been better at avoiding this situation. Regardless, the presence of private healthcare providers with competing interests means that various levels of barriers to creating a nationwide system exist in many countries.

Blockchain technology could have the potential to overcome this problem because it could handle the presence of multiple writers and absence of trust between decentralized and diverse players in the healthcare industry.

3.1.1 Strength: Secure Transactions

With a traditional database, a person wishing to obtain an unauthorized prescription has numerous points of entry into the transaction chain: bribing a doctor to write the prescription, attacking an unsecured server to insert a transaction record, or bribing a system administrator. A blockchain approach prevents a record from being altered at any point in the chain back to its origin, meaning that a pharmacy

filling the prescription can rely that it's legitimately granted. With a secure transaction log the only way an unlawful prescription could be filled would be on the originating end, i.e. with the doctor authorizing the prescription, which is not a record-keeping problem as such. From a supply chain perspective blockchain could also secure and verify the quality and origins of pharmaceuticals, ensuring safety for the end-user.

3.1.2 Strength: Patient Privacy

While the classic cryptocurrency use case of blockchain has difficulty building in privacy protections because of reoccurring transactions, as discussed in the previous chapter, an EHR system built on a blockchain would not suffer the same problems because the only parties to a "transaction" of patient information (doctor, patient, hospital, pharmacy) would be ones already authorized to know the patient's identity. This means the approach could comply with legal requirements such for patient privacy. Parties would only need to ensure that they are complying with existing regulations before initiating a transaction. From a patient perspective, having one's information readily available when dealing with the various touchpoints of the medical system would be a definite advantage.

3.1.3 Strength: No Trusted Intermediary

A traditional database implementation of EHR built by a private corporation implicitly assumes that the corporation can act as a trusted intermediary, i.e., all hospitals agree to do business with the same EHR provider and allow them to handle patient information in database transactions. In most European countries such a system would most likely be run by the public sector, but in a more competitive environment like the United States this requires significant market coordination, or an EHR provider to achieve so much market share that they achieve an effective monopoly. Therefore in an environment where the political and legal landscape is not fruitful for centralized public sector solutions, a blockchain approach could overcome some of the issues with the option of too much centralization in the private sector.

3.1.4 Limitation: Incentives for Walled Gardens

Transparency is inherent to a public blockchain. A patient retains access to their entire medical record on the public shared ledger. A hospital which is given access to a patient's private key could access another hospital's medical records for the same patient, which is a positive from the patient's perspective. From the hospital's, however, it is essentially "giving away" information gathered by medical professionals. Cynically, hospitals and healthcare providers may view barriers to interoperability as being in their interest.

This is an economic and political barrier to implementation which may need a political response. However, it is a problem for any unified EHR system, whether based on a blockchain or more traditional databases. The potential economic gains and wins for patient care due to a unified medical record are significant. If the gains of establishing a unified EHR system are worth the effort, then blockchain may actually be easier to implement than a traditional system because of the possiblity for incremental deployment.

3.1.5 Limitation: Private Key Dependency

The primary limitation of a fully public (albeit encrypted) EHR blockchain is that access to patient information is entirely governed by access to a patient's private key. Losing their private key means losing access to their healthcare information; while having it stolen means a third party has access which cannot be revoked short of deleting all patient blocks from the shared ledger, which should not be possible in a true blockchain implementation.

From a user perspective this is a dramatic shortcoming, given the difficulty of remembering and securing passwords for the average user. Moving to a more private blockchain, such as one where only references to centrally stored records are store on the blockchain, or a traditional database would diminish the difficulty by outsourcing it to a trusted intermediary but correspondingly diminish the benefits listed. Yuan, Lin & McDonnell (2015) discuss the possibility of distributing partial keys to a

moderately trusted intermediary network for retrieval, but this still does not deal the problem of theft. On the other hand, in countries with existing digital signing solutions like Finland or Denmark, the blockchain private key could also be tied to this system. Perhaps the best example of a unified nationwide EHR system, Estonia's, also has a universal system for digital identity-verification, to which a private key can be attached with no risk of loss and much less risk of theft. However, this also assumes the existence of a centralized framework and a trusted intermediary (the government) which again obviates the benefits of the blockchain implementation.

Is it possible to accept this risk as a cost of doing business for a well-functioning, unified EHR system based on blockchain? The benefits would have to be weighed against the risks, and as we have seen with the Ethereum DAO incident, accidents can happen. Individuals losing or permanently giving away access to their health records would be politically disastrous for such a system and indeed might entail legal complications due to privacy regulations. No hospital is likely to adopt, much less pay for access, to a record-keeping system in which patients might permanently lose access to their medical records. An implementation which does not in some way mitigate the problem of private key access is almost certainly non-viable.

3.1.6 Conclusion

The application of blockchain technology to healthcare records is technically promising, and could perhaps proceed using small investments to measure incremental returns. However, given the early stage of the technology and the critical nature of any nationwide EHR system, it might not be the best use case for blockchain at this time. However, given that the incremental cost of investment is low, it could still be worth studying and pursuing on a small scale. As an investment, it should be considered high-risk, but with a substantial potential reward. As the technology matures, a system could eventually grow to be a part of the national EHR infrastructure.

3.2 Use case: Supply chain management

When goods are transferred from seller to buyer via a carrier, the core of each step of the transaction from a legal standpoint is the “bill of lading”: a document certifying the goods being transferred from one party to the next and which is agreed to by both parties at the time of transaction. While vertically integrated companies may have centralized databases and highly complex logistical systems for tracking the movement of goods between subsidiaries and affiliates, in a market with many independent buyers and sellers the movement of goods is tracked on the basis of certified bills of lading. Bills of lading are natural tokenized assets: physical transfers can be represented as blocks with the bill of lading being simply a part of the block data on the ledger.

There is no need to maintain a comprehensive tracking system for many commodity goods, but for high-value luxury goods such as pharmaceuticals or electronics, a secure and distributed transaction record can help to prevent fraud and theft. IBM, for example, is promoting a platform for development and testing of blockchain implementations for supply chain tracking. This is also the platform being used to host provenance tracking blockchain Everledger, which will be discussed in the next section.

3.2.1 Strength: Immutable, distributed ledger

The immutable and secure ledger of a blockchain approach is essential for ensuring that the transaction record can be trusted by widely dispersed parties. In the case of high-value goods the incentives for fraud and theft are correspondingly high, and can happen at various points in the value chain. In a traditional centralized database administered by a manufacturer or retailer, individual employees with the ability to do so may be strongly incentivized financially to cheat the system. Even when administered by a third party, the employees at the third party could be bribed. Also paper documents can be easier to forge than properly cryptographically secured digital transactions. An immutable shared ledger solves many of these problems at a

stroke: transactions between parties leave a digital “paper trail” of transaction history that cannot be altered.

3.2.2 Strength: Dispute resolution

Of particular interest to supply chain management is dispute resolution. Disputes about transactions as simple as disagreeing whether a particular shipment had been delivered can delay payments and are often both protracted and time-consuming. The mutually agreed-upon nature of the distributed ledger allows businesses to return to their “last agreed-upon” transaction and investigate where points of disagreement between the parties may have arisen by comparing to their individual records of the series of transactions.

IBM has developed a blockchain implementation of supply chain tracking geared towards dispute resolution, the results of which are startling: with an average of 100 million USD in payments delayed for up to 40 days, they were able to reduce the delay time to less than 10 days, which represents a significant amount of capital freed for other, more productive uses. This is in addition to the person-hours saved by spending less time investigating disputes and ensuring misplaced goods are located. One additional note on this application is that the blockchain implementation here serves as a supplement to a traditional logistics database maintained by individual companies rather than a replacement of it, which can be consulted as a baseline of agreement between parties. This suggests the possibility that small-scale applications could be adopted piecemeal rather than needing to replace existing systems wholesale.

3.2.3 Conclusion

Supply chain management is a very promising application for blockchain techniques with robust interest from companies with significant investment potential. It not only has the potential to secure against unlawful behavior such as fraud and theft but to resolve everyday disputes between parties to a transaction because it is a secure way of sharing a mutually agreed-upon version of events. While business

interest seems to suggest that a distributed ledger of all transactions is not economical, at least for high-value goods the investment is worthwhile. Furthermore the use as a supplement to existing logistical systems means that the scale of the investment needed is not too large for startups to succeed in this area.

3.3 Use case: Provenance tracking

Related to the discussion of supply chain management is the idea of provenance tracking, which is the verification of goods likely to either be unlawful in origin (theft, smuggling, fraud) or unethically sourced (e.g. conflict minerals). Recent efforts have been made in this area including Everledger, a company engaged in provenance tracking of diamonds to prevent criminal fraud, but also to screen for potential conflict diamonds by looking for stones “in regions where forced labor is common or where proceeds from previous sales were used to fund violence” (Nash 2016) i.e. blood diamonds. In this case the need for a secure, distributed and decentralized ledger which allows the chain of control of the shipment to be tracked and stored in an unalterable way is particularly clear.

The fundamental strengths of this approach are the same as for the supply chain management discussion above: a secure distributed ledger ensures that no central authority has employees who can be bribed to alter the transaction record to benefit a criminal. Markets in which provenance is relevant, such as the market for art or diamonds, are based on widely distributed networks of buyers and sellers with oversight structures from many countries representing diverse private interests. A traditional centralized database has difficulty with this problem because employees at the trusted intermediary could always be influenced or bribed to falsify the provenance of the object in question. Blockchain is the perfect approach for such a low-trust, decentralized system. However, this application has a unique drawback.

3.3.1 Limitation: Unique identifiers for tracked objects

The ultimate concern of a provenance-tracking blockchain application is the physical object which the data on the shared ledger represents. Unlike in the case of

supply-chain management, however, the objects are not fungible: two diamonds of equal quality and cut are not equivalent if one is a conflict mineral or stolen and the other is not. Because the key is tracking a particular object, the possibility for fraud remains as long as objects do not have some form of unique identifier intrinsically associated with each one. Many diamonds, for example, have serial numbers engraved matching one on their physical certificate, but these can be removed via polishing.

There are ways of dealing with the problem of spoofing in this context. Everledger, for example, also includes information about the grade, cut, and size of a diamond in the ledger itself, giving them a kind of fingerprint so that diamonds can be associated with a particular record to a greater degree of specificity. It is still possible to re-cut a diamond to fit the details of a particular chain, but as Lomas (2015) points out, potential fraudsters have a natural disincentive to doing so because the loss of size means a financial hit. However, this is a sui generis strategy for diamonds or precious stones in particular that may not be portable to other applications.

3.3.2 Conclusion

The technical aspect of blockchain is an excellent fit for provenance-tracking, and the success of Everledger suggests that further development is warranted. However, new applications in this space will need to deal with the need to uniquely identify objects being tracked so that spoofing in order to re-enter the chain of lawful transactions is quashed. No general solution for this problem is likely to be possible, but if it can be solved for a particular application, the blockchain implementation could easily be competitive with current best practices.

3.4 Use case: Internet of Things

The Internet of Things (IoT) is an emerging concept of massively inter-networked everyday devices which can communicate via the Internet to share information, download updates, and engage in real-time coordination between smart

devices as a loosely unified system. The potential gains are substantial, from greater automation, autonomy and reliability of device operations for networked devices. The eventual market for IoT devices is likely to be very large. Applications such as "Smart" homes, cars, and cities – with networked devices sharing and tracking data have been proposed. However, the Internet of Things is as yet less a technology than a "paradigm" with a number of competing efforts to develop implementations.

There is still a need for robust systems on which peer-to-peer transactions between nodes in an IoT network could be tracked, shared, and authenticated. This is critical to promote system functionality and stability and prevent malicious behavior such as spoofing or theft. Absent some secure peer-to-peer (P2P) method for verifying network interconnections and data transactions between devices it would be possible to trick network nodes into engaging in unwanted behavior or recording garbage data. For IoT to develop into a wide-scale business, industrial and consumer technology, this must be addressed.

3.4.1 Strength: Mutual distrust between nodes

Blockchain may be appropriate for IoT applications precisely because devices on the network cannot be mutually trusted. Although the possibility of spoofing cannot be completely eliminated, securing interactions between devices on the blockchain at least allows for the development of a "canonical" version of events which can be used to troubleshoot or prevent malicious activity from non-trusted nodes.

3.4.2 Limitation: Scalability

There are also reasons to be skeptical of blockchain techniques in IoT applications because of scalability. If the number of devices is small, spoofing becomes problematic because fewer nodes need to be corrupted in order to affect the outcome of whatever consensus algorithm is implemented. But as the number of nodes grows, the devices in question will often be embedded systems much less powerful than the average server cluster or even desktop computer, so they may lack

the ability to host much of the distributed ledger and thereby provide the robustness that a network like Bitcoin can. The solution to this might be to use a secure cloud to host the ledger (this is the approach that IBM's Watson IoT implementation takes) but at that point there may be few advantages to a traditional database.

By virtue of the breadth of the application, though, there really is no single "use case" which defines the use of blockchain on IoT, so it makes more sense to focus on individual case studies or applications to understand some of the ways that blockchain could be used. As discussed in an earlier chapter, blockchains have the advantage over centralized databases when it comes to decentralization, while centralized databases have the advantage in speed. Since both of these can be required to an extent in an IoT solution, a hybrid approach of both technologies could also be viable.

3.4.3 Example: SolarCoin

Chain of Things, a blockchain IoT consortium, did a case study analysis with using a blockchain to track solar cell power generation and reimburse owners using a digital currency called SolarCoin. Advantages of this approach include that it is highly modular and can be set up and installed even in isolated or poor areas, then used to provide power and generate digital currency for the owners, and the data tracking/currency aspects have all the security benefits associated with the blockchain approach. But while the transaction history is secure once it is written to the (Ethereum-based) blockchain, there remains the possibility that spoofers could the input data in some way: by creating a spoofed data stream or replicating sensors to claim credit multiple times. While Chain of Things proposes to address this problem by centrally registering sensors with unique IDs, this does tend to undermine the virtue of not relying on a trusted intermediary; there is no reason why this approach could not be taken with a relational database and proprietary hardware.

3.4.4 Conclusion

The fertility of blockchain techniques for IoT applications is difficult to evaluate because of the inherent breadth of the concept of the "Internet of Things." While it is easy to see in theory how a distributed network of non-trusted devices could benefit from secure, distributed ledgers with multiple write access, there are also reasons to think that a blockchain approach is sub-optimal for IoT because of problems with scalability.

However, it is clear that substantial investment in both IoT and blockchain applications will continue to occur. Valuations vary dramatically, but the market could be in the billions of dollars in the next decade. While the development of particular IoT related use cases requires both more breadth and more depth than is possible in this study, further study is warranted.

3.5 Identifying areas where blockchain is useful

Based on the above discussed use cases, we can attempt to establish some common criteria for situations where blockchains genuinely add value. According to Glaser (2017), this kind of understanding is desperately needed in the blockchain space: "Despite many discussions, press releases and talks about blockchain technology, few truly and fully understand or can actually describe with certainty the basic or innovative features introduced by blockchain technology". This lack of understanding can lead to situations where blockchain is attempted to be applied to cases that would be better served by existing technologies.

As assessment of criteria that should be considered when starting a new blockchain project is presented by Gideon Greenspan (2015), who proposes a set of 8 criteria which are summarized in the following table:

The database	Databases are used to store information, but blockchain is specifically a *shared database*, that multiple parties need to have access to.

Multiple writers	Building on the previous point, blockchains can be used when multiple parties need specifically to have write access; shared read access with centralized write access can easily be done with existing database technologies.
Absence of trust	The third rule is the absence of trust between parties to a varying degree – if complete trust exists, even shared write access can be done with existing technologies.
Disintermediation	The solution to a database required to have shared write access between mutually distrusting parties has traditionally been the same: using an intermediary. Blockchain can be used when an intermediary is not feasible.
Transaction interaction	Blockchain transactions build on one another, for example in a payment system, and the integrity of these relationships is important to the use case.
Set the rules	Not specifically a condition, but an outcome of the previous points: some rules must be in place to define what transactions are allowed, setting the consensus mechanism.
Pick your validators	To gain the advantages of decentralization, there need to be multiple validators. Not everyone necessarily needs to validate all transactions, but it should be decided which participants should do this and how.
Back your assets	The nature of the assets stored on the blockchain should be known, and specifically how they are backed in the real world.

Figure 8: criteria for blockchain adoption (Greenspan 2015)

Many of the proposed blockchain use cases are built around disintermediation, in fields where there are already intermediaries present, or the lack of one prevents doing business. It is however important to ask, as Greenspan points out, if there is anything wrong with having an intermediary – in some cases it might still be cheaper, faster or safer to use a trusted intermediary. This can also be looked

at by considering what exactly the role of the intermediary in a given field is. Uber and Airbnb are often mentioned as examples of disintermediation, but according to Tapscott & Tapscott (2016), they represent aggregation, not disintermediation: they aggregate drivers and rooms to the point that their platform is valuable to users. But at the same time they act as a middle-man, taking a part of the payment, because facilitation of payments between the different parties is one of the core value propositions.

Greenspan's last point relates closely to the nature of the asset being used on the blockchain. There is an ongoing debate whether blockchains without a native asset like bitcoin can really be considered blockchains at all. From a user perspective this raises the question: what guarantee is there that the asset actually has value? With bitcoin it becomes as much as trust issue as the case of fiat currencies: it only has value as long as everybody believes it does. However, if blockchain is used for real-world assets, then there must be a way to ensure that if the blockchain says someone owns certain assets, they also have a way to claim them in the real world. In the case of private blockchains this will likely come from various contractual obligations.

A higher level analysis of digital market models where blockchain may be applicable is presented by Glaser (2017), who proposes three such market models:

- Multi-sided markets – these are markets where multiple actors are present on the same market mechanism who have different interests. They are characterized by the presence of intermediaries who provide product and information brokerage, such as stock exchanges.
- Collaborative markets – these also have multiple parties present, who collaborate on platforms that facilitate basic exchange of information, for example multiple companies in a value chain sharing a supply chain management system.
- P2P markets – these are natural candidates for blockchain adoption as they include multiple parties and the presence of intermediaries,

whose value proposition is often built around information and facilitating payments.

4 Discussion

4.1 Inflated expectations on blockchain?

Opinions regarding blockchain in the mainstream have gradually moved from being dismissive towards bitcoin to having positive and high expectations regarding its use for businesses and society. The potential to revolutionize many fields is widely acknowledged. However many of the use cases are at an early or even a purely conceptual stage, and even bitcoin, the most mature example, has various scalability questions that will have to be answered before the next level of usage can be reached.

Broadly speaking there are two kinds of thinking abound regarding blockchain: first is the pragmatic approach, often exercised by people who work on the technological side of blockchain, and express a measured optimism while being aware of the limitations. The second kind is the inflated expectations that "Perhaps all modes of human activity could be coordinated with blockchain technology to some degree, or at a minimum reinvented with blockchain concepts" (Swan 2015) that seem overly optimistic at this point. Blockchain now is said to be at the equivalent of the Internet in the early 1990s, but this assumption is easy to say now in hindsight that we know how much the Internet has changed our daily lives. For blockchain however we can't yet know that it will have the same kind of transformative and disruptive effect.

Gartner (2016) placed blockchain at the "peak of inflated expectations" in the 2016 hype cycle, which can be considered a fairly accurate assessment of the current state of blockchain. It is interesting to note however, that Gartner defines the peak as the period where "Some companies take action; many do not", although in some industries, especially the financial sector, one would be hard-pressed to find a company that isn't currently exploring blockchain in one way or another.

Figure 9: Blockchain in the Gartner hype cycle

4.2 Challenges in adoption

There are a number of challenges that currently hinder adoption of blockchain technology, including the early development stage, regulatory limitations and the lack of interoperability and network effects in many areas. These challenges will have to be addressed before we can expect to see blockchains in wider use.

As we have seen with the cases of bitcoin and Ethereum, there are many unexpected technical challenges that come up when building blockchain applications. Although the space is growing, another challenge related to the relative lack of people who are experienced in the field, compared to established areas like centralized databases. This can also create a gap in companies between people who work closely with the technology and the ones making strategic decisions. Building large IT

solutions already has its own set of challenges, but doing it using a developing technology can make it even more difficult.

Many of the fields where blockchain is proposed to be used are heavily regulated, like finance and healthcare. Because of their importance for society, solutions in these areas require a high degree of resiliency, and building them on an unproven technology may be difficult. Questions like anonymity will have to be covered – if all transactions on the blockchain are visible to everyone, how does that comply with privacy regulations? Also if a blockchain is to exist across countries and market areas, how can it be made to comply with regulations of all the different countries involved? And if transactions and assets are permanent once committed to the blockchain, how can regulators deal with issues like seizing illegally acquired funds? Comparisons are often made to the Internet, saying that despite calls for tighter regulation in the early days, it was the fact that the internet wasn't strictly regulated from the beginning that allowed for growth and innovation. In the case of bitcoin, regulators now seem to have taken a different approach. However, once mature blockchain could also help regulators in many areas.

Common standards and protocols will also have to be established to ensure interoperability. A blockchain system might live or die by its network effects, so achieving critical mass will be crucial. Many players are currently exploring the options and applications for various levels of the blockchain technology stack, and if they all branch out in different directions then the lack of interoperability could become a real issue. As discussed in the healthcare use case, the threat of "walled gardens" where companies build their own proprietary blockchains that don't talk to each other exists.

Although these challenges exist, many of them are actively being worked on by people in the field. Even if they are difficult to overcome, the sooner they are addressed the sooner they can be solved and we can better understand what's possible and start seeing blockchain being used.

5 Conclusion

Blockchain began with the bitcoin in 2008. In less than ten years we seem to have reached a point where the potential of blockchain technology to greatly transform business models in various sectors is widely acknowledged. There are however significant differences in views as to how this will happen: according to some, it will be through public and open blockchains like bitcoin and Ethereum, while others think that private blockchains are the way to go. Regardless, blockchain use cases are being explored on both sides of the table. This is in contrast to many industries where traditional market incumbents have attempted to resist disruption by dismissing new ideas or failing to innovate. Instead, in the blockchain space those that stand to be most disrupted, like financial institutions, are the ones investing heavily in the new technology.

One of the goals of this study was to find out how exactly blockchain works, how can it create value, and what sets it apart from existing technologies. The main points can be summarized as potential improvements in efficiency and reductions in costs through disintermediation and enabling value transfer without trusted third parties. How exactly this can happen depends greatly on the way blockchain is applied. The core components of consensus mechanism, the use of permissionless or permissioned blockchain, and figuring out the various potential issues like transparency and settlement finality have to be addressed.

The second goal was to explore potential blockchain use cases and determine the criteria for identifying these. Various use cases exist in healthcare, supply chain, provenance tracking and Internet of Things, and many others besides these. There are a lot of companies, both start-ups and big ones actively building these use cases. It's important however to stop and think what exactly are the benefits of using blockchain in that particular case, and if it could be done better with existing technologies. Otherwise companies run the risk of painting themselves into a corner and attempting to build impossible solutions, or ones that end up not contributing any value.

6 References

Abidi, A., Bouallegue, B., & Kahri, F. (2014, June). Implementation of elliptic curve digital signature algorithm (ECDSA). In Computer & Information Technology (GSCIT), 2014 Global Summit on (pp. 1-6). IEEE.

Allison, I. (2016, March 3). Guardtime secures over a million Estonian healthcare records on the blockchain. International Business Times UK. http://www.ibtimes.co.uk/guardtime-secures-over-million-estonian-healthcare-records-blockchain-1547367

Allison, Ian (2016, October) Shipping giant Maersk tests blockchain-powered bills of lading http://www.ibtimes.co.uk/shipping-giant-maersk-tests-blockchain-powered-bills-lading-1585929

Atzei, N., Bartoletti, M., & Cimoli, T. (2016). A survey of attacks on Ethereum smart contracts. Cryptology ePrint Archive: Report 2016/1007, https://eprint. iacr. org/2016/1007.

Bradbury, D. (2016). Blockchain's big deal [financial IT]. Engineering & Technology, 11(10), 44-44.

Becker, J., Breuker, D., Heide, T., Holler, J., Rauer, H. P., & Böhme, R. (2013). Can we afford integrity by proof-of-work? Scenarios inspired by the Bitcoin currency. In The Economics of Information Security and Privacy (pp. 135-156). Springer Berlin Heidelberg.

Blochain.info (2016 Bitcoin mining pools https://blockchain.info/pools

Bouvier, Pascal (2015) Taxonomy is important, consensus computer is the endgame http://finiculture.com/taxonomy-is-important-consensus-computer-is-the-end-game/

Burniske, C., Vaughn, E., Shelton, J., Cahana, A. (2016). How Blockchain Technology Can Enhance EHR Operability. New York: ARK Invest.

Available at http://research.ark-invest.com/hubfs/1_Download_Files_ARK-Invest/White_Papers/ARKInvest_and_GEM_Blockchain_EHR.pdf

Buterin, V. (2013). Ethereum white paper.

Buterin, V (2016c) On settlement finality https://blog.ethereum.org/2016/05/09/on-settlement-finality/

Buterin, V. (2016a) Ethereum Research update. https://blog.ethereum.org/2016/12/04/ethereum-research-update/

Buterin, V. (2016b) On private and public blockchains, https://blog.ethereum.org/2015/08/07/on-public-and-private-blockchains/

Coleman, L (2017) European Union wants to Identify Bitcoin Users, Cryptocoinsnews Jan 01 2017 https://www.cryptocoinsnews.com/the-european-union-wants-to-identify-bitcoin-users/

Deloitte (2016) Bitcoin, Blockchain & distributed ledgers: Caught between promise and reality https://www2.deloitte.com/content/dam/Deloitte/au/Images/infographics/au-deloitte-technology-bitcoin-blockchain-distributed-ledgers-180416.pdf

Dickson, B. (2016). Decentralizing IoT networks through blockchain. TechCrunch. https://techcrunch.com/2016/06/28/decentralizing-iot-networks-through-blockchain/

Ekblaw, A., Barabas, C., Harvey-Buschel, J., & Lippman, A. (2016, September). Bitcoin and the Myth of Decentralization: Socio-technical Proposals for Restoring Network Integrity. In Foundations and Applications of Self* Systems, IEEE International Workshops on (pp. 18-23). IEEE

Fichman, R. G., Dos Santos, B. L., & Zhiqiang (Eric) Zheng. (2014). Digital Innovation as a Fundamental and Powerful Concept in the Information Systems Curriculum. Mis Quarterly, 38(2), 329-343.

Forte, P., Romano, D., & Schmid, G. (2016). Beyond Bitcoin--Part II: Blockchain-based systems without mining.

Franz, P., Kirchmer, M., & Rosemann, M. (2012). Value-driven business process management—impact and benefits.

Garman, C., Green, M., & Miers, I. (2013). Decentralized Anonymous Credentials. IACR Cryptology ePrint Archive, 2013, 622.

Gartner Group (2016) Gartner hype cycle for emerging technologies 2016 http://www.gartner.com/newsroom/id/3412017

Glaser, F. (2017) Pervasive Decentralisation of Digital Infrastructures: A Framework for Blockchain enabled System and Use Case Analysis. Proceedings of the 50th Hawaii International Conference on System Sciences, 2017

Greenspan, G. (2015) Bitcoin vs blockchain debate http://www.multichain.com/blog/2015/07/bitcoin-vs-blockchain-debate/

Greenspan, G. (2016). Four genuine blockchain use cases. http://www.multichain.com/blog/2016/05/four-genuine-blockchain-use-cases/

Greenspan, G. (2016a) Blockchains vs. Centralized databases http://www.multichain.com/blog/2016/03/blockchains-vs-centralized-databases/

Groenfeldt, T. (2016). IBM Trials Blockchain for Supply Chain Dispute Resolution. Forbes.com. http://www.forbes.com/sites/tomgroenfeldt/2016/11/03/ibm-trials-blockchain-for-supply-chain-dispute-resolution/#290aaeb42e98

Kõlvart, M., Poola, M., & Rull, A. (2016). Smart Contracts. In The Future of Law and eTechnologies (pp. 133-147). Springer International Publishing.

Kroll, J. A., Davey, I. C., & Felten, E. W. (2013, June). The economics of Bitcoin mining, or Bitcoin in the presence of adversaries. In Proceedings of WEIS (Vol. 2013).

Kumaresan, R., & Bentov, I. (2014, November). How to use bitcoin to incentivize correct computations. In Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security (pp. 30-41). ACM.

Lomas, N. (2015). Everledger is using Blockchain to combat fraud, starting with diamonds. TechCrunch.com https://techcrunch.com/2015/06/29/everledger/

MacGregor, A. (2016) Can blockchain save IoT from itself? TheStack.com. https://thestack.com/iot/2016/06/02/can-blockchain-save-iot-from-itself/

McConaghy, T., Marques, R., Müller, A., De Jonghe, D., McConaghy, T., McMullen, G., ... & Granzotto, A. (2016). BigchainDB: A Scalable Blockchain Database.

Milani, F., García-Bañuelos, L., & Dumas, M. (2016). Blockchain and business process improvement. BPTrends newsletter (October 2016).

Morini, M. (2016). From'Blockchain Hype'to a Real Business Case for Financial Markets. Available at SSRN 2760184.

Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system.

Narayanan, A., Bonneau, J., Felten, E., Miller, A., & Goldfeder, S. (2016). Bitcoin and cryptocurrency technologies. Princeton University Pres.

Nash, K. (2016). IBM Pushes Blockchain Into the Supply Chain. Wall Street Journal online. http://www.wsj.com/articles/ibm-pushes-blockchain-into-the-supply-chain-1468528824

O'Connell, J (2016) What are the use cases for private blockchains? https://bitcoinmagazine.com/articles/what-are-the-use-cases-for-private-blockchains-the-experts-weigh-in-1466440884/

Parker, L (2016) Private versus Public Blockchains: Is there room for both to prevail?, https://magnr.com/blog/technology/private-vs-public-blockchains-bitcoin/

Pass, R., & Shi, E. (2016). Hybrid consensus: Efficient consensus in the permissionless model.

Phillips, K (2016) IRS Wants Court Authority To Identify Bitcoin Users Transactions at Coinbase, Forbes Nov 21 2016, http://www.forbes.com/sites/kellyphillipserb/2016/11/21/irs-wants-court-authority-to-identify-bitcoin-users-transactions-at-coinbase/#2576464440d1

Pilkington, M. (2016). Blockchain Technology: Principles and Applications. Research Handbook on Digital Transformations, edited by F. Xavier Olleros and Majlinda Zhegu. Edward Elgar.

Rizzo, P (2015) Barry Silbert: Private Blockchains Will 'Capitulate' to Bitcoin, Coindesk.com http://www.coindesk.com/barry-silbert-private-blockchains-will-capitulate-to-bitcoin/

Salomaa, A. (2013). Public-key cryptography. Springer Science & Business Media.

Saunders, Mark N.K. & Thornhill, Adrian & Lewis, Philip (2009) Research Methods for Business Students (5th Edition), Prentice Hall

Scott, A (2016) Andreas Antonopoulos: 'The Open Blockchain Will Change This World', news.bitcoin.com https://news.bitcoin.com/antonopoulos-open-blockchain/

Swan, M. (2015). Blockchain: Blueprint for a new economy. " O'Reilly Media, Inc.".

Swanson, T. (2015). Consensus-as-a-service: a brief report on the emergence of permissioned, distributed ledger systems.

Swanson, T (2016) Settlement risks involving public blockchains http://tabbforum.com/opinions/settlement-risks-involving-public-blockchains

Tapscott, D., & Tapscott, A. (2016). Blockchain Revolution: How the Technology Behind Bitcoin is Changing Money, Business, and the World. Penguin.

Van Der Aalst, W. M., Ter Hofstede, A. H., & Weske, M. (2003, June). Business process management: A survey. In International conference on business process management (pp. 1-12). Springer Berlin Heidelberg.

Walch, A. (2015). Bitcoin Blockchain as Financial Market Infrastructure: A Consideration of Operational Risk, The. NYUJ Legis. & Pub. Pol'y, 18, 837.

Weber, I., Xu, X., Riveret, R., Governatori, G., Ponomarev, A., & Mendling, J. (2016, September). Untrusted business process monitoring and execution using blockchain. In International Conference on Business Process Management (pp. 329-347). Springer International Publishing.

Weske, M. (2012). Business process management architectures. In Business Process Management (pp. 333-371). Springer Berlin Heidelberg.

Xu, X., Pautasso, C., Zhu, L., Gramoli, V., Ponomarev, A., Tran, A. B., & Chen, S. (2016, April). The blockchain as a software connector. In Proceedings of the 13th Working IEEE/IFIP Conference on Software Architecture (WICSA).

Yuan, B., Lin, W., McDonnell, C. (2015), Blockchains and electronic health records. Available at http://mcdonnell.mit.edu/blockchain_ehr.pdf