

MSc Program

Environmental Technology & International Affairs



CONTINUING
EDUCATION
CENTER



diplomatische
akademie **wien**
Vienna School of International Studies
École des Hautes Études Internationales de Vienne

Opportunities of Blockchain for Industry 4.0 – Energy Impacts for the Industrial Internet of Things

A Master's Thesis submitted for the degree of
“Master of Science”

supervised by
Dr. Klaus Rapp

Nicholas James Lieb
Student-ID Number 0931250

Vienna, 04/06/2018

Affidavit

I, **NICHOLAS JAMES LIEB**, hereby declare

1. that I am the sole author of the present Master's Thesis, "OPPORTUNITIES OF BLOCKCHAIN FOR INDUSTRY 4.0 – ENERGY IMPACTS FOR THE INDUSTRIAL INTERNET OF THINGS", 88 pages, bound, and that I have not used any source or tool other than those referenced or any other illicit aid or tool, and
2. that I have not prior to this date submitted this Master's Thesis as an examination paper in any form in Austria or abroad.

Vienna, 04/06/2018

Signature

Abstract

Information and communication systems (ICT) are an essential part of today's society, with strong negative impacts on the environment due to its energy consumption. However, its utilization is believed it can have indirect positive impacts on resource consumption reduction through new concepts to generate, allocate, distribute, share and use energy environmentally-friendly. A new major technology, blockchain, receives a great amount of attention regarding energy efficiency, since it allows distributed network applications and therefore is believed to revolutionize computing. Especially applications for industrial production, with utilization within an Industrial Internet of Things (IIoT) have been waiting for a technology leap to make industrial processes more efficient. This Master Thesis investigates the energy efficiency of a blockchain architecture believed to be suited for use in an Industrial Internet of Things based on the ME³SA model for sustainable software compiling. The results show that while the blockchain claiming to be IoT-optimized, IOTA, performs well in terms of energy consumption, the hardware utilization is very high, leaving no room for other applications to run in sideline. Recommendations include artificial enlargement of the network until a sufficient number of participants is reached, reducing CPU utilization, and a reduction of prerequisites to run the blockchain network on smaller 'smart' objects. Further research could include comparative analysis as well as investigation under operating environment.

Acknowledgements

Dear Master Thesis,

It has been a difficult two months between the two of us. You made me research a topic I had little to no knowledge about and made it hard for me to understand the complexity of the subject of higher computing. Furthermore, you tortured me with the nicest weather I have experienced since moving to Vienna, from which I had no benefit. But the more the topic substantiated and the more I started to understand about it, the more I can conclude in peace with you. You made me smarter, more focused and let me think more about the essentials of life I am striving for. Thank you for this. But writing you would never have been able without my brother Alexander, who firmly believed that going for a Master Degree is a good thing for me and supported me by acting as a guarantor for the loan I have obtained to pay my student fees. Thank you, your belief will be never forgotten.

Table of contents

Abstract	iii
Acknowledgements	iv
1. Introduction	1
2. Theoretical Part - ICT and energy consumption	3
3. Industry 4.0	5
3.1. Internet of Things (IoT)	7
3.1.1. The Industrial Internet of Things (IIoT)	12
3.1.2. Environmental considerations for IIoT	14
4. Obstacles of Blockchain Utilization	16
4.1. What is Blockchain?	16
4.1.1. Technical requirements	17
4.2. Application Potentials – Use Cases	18
4.3. Industrial Use Cases for Blockchain	20
4.4. Footprint of Blockchains	25
4.4.1. Energy footprint	25
4.4.2. Mining hardware	28
4.4.3. Ecological footprint for Bitcoin	30
5. Obstacles for IoT-Optimized Blockchains	32
5.1. Scalability and network speed	32
5.2. Consensus Mechanisms	35
5.2.1. Proof of Work	35
5.2.2. Proof of Stake	37
5.2.3. Variants of Byzantine Fault Tolerance	41
5.2.4. Proof of Elapsed Time	44
5.2.5. Proof-of-Authority	45
5.2.6. Proof of Luck	46
5.2.7. Proof-of-Burn	46
5.2.8. Proof-of-Space	47
5.2.9. Network Consensus Mechanism	47
5.2.10. Hashgraph	48
6. IoT-optimized blockchain - The Tangle	52
7. Epistemological interest	56
8. Empirical Part - Research Question	58
9. Methodology	59
10. Case Study: Smart object energy efficiency assessment with an applied blockchain	65
10.1. Case Study: Full node on a Windows personal computer	65
11. Conclusion/Outlook	70
Bibliography/Sources	73

1. Introduction

Blockchains have received a great deal of attention in recent times due to the astronomical rise of Bitcoin, the first-ever blockchain, and the associated earnings of investors. There is much more potential behind the blockchain technology, and in order to utilize this potential, it makes sense for the academic community to further elaborate research on the topic. There have been a variety of fields potentially benefitting from blockchain, one of which is the application in the Industry, particularly under the framework of Industry 4.0. One field which garnered great interest is the Industrial Internet of Things (IIoT), which is further investigated in this Master Thesis.

One of the concerns raised in the debate about blockchain is the high energy impact the technology already has today. It is believed that the energy demand for Bitcoin today is as high as the energy demand of Argentina (Digicomonist 2018), which is an obstacle for the implementation of blockchain. There are actions within the community to reduce the energy impact of blockchains, some very successful. It is of interest to investigate what initiatives are driving the blockchain world towards lower energy impacts, since it removes a bottleneck in the technology that could lead to further implementation in our processes of the future. Therefore, we investigate the energy performance of a blockchain optimized towards utilization in the Industrial Internet of Things.

The goal the written Master Thesis on this specific topic should fulfil is to give an overview of blockchain projects and their specific technological deviations and to give insight why the technology is interesting for a future Industrial Internet of Things. The structure of the thesis will follow the basic principles of scientific research for this type of thesis. After the introduction, the theoretical part in Section 2 will provide us with background knowledge about energy impact of information and communication technology. Section 3 will give definitions on Industry 4.0, the Internet of Things and in particular the Industrial Internet of Things. Here, we will also discuss the challenges to accomplish an Industrial Internet of Things. Section 4 discusses the obstacles of blockchain utilization in general, describing the technology in detail and further spotlighting use cases as well as environmental considerations. Section 5 will further elaborate the obstacles for IoT-optimized blockchains and will further discuss the technical background of blockchain in detail. In Section 6, we will

present the Tangle, an IoT-optimized blockchain that was developed for this specific use case. After we have further elaborated the epistemological interest of the Master Thesis, summarizing what we have learned in the theoretical part and substantiated our interest for the empirical part, we lay down a research question in Section 8, describe the methodology how we intend to answer the research question in Section 9, before we conduct a case study applying the IoT-optimized blockchain and taking measurements for energy impact considerations. The Conclusion/Outlook Section aims to answer the research questions in a satisfying manner and elaborate on challenges during the research. Furthermore, inferences are drawn from the results we obtained from the research and an outlook on possible future research is given.

2. Theoretical Part - ICT and energy consumption

Information and communication technology (ICT) is a major pillar today. It does not only have major impacts on professional and social life, it is also one of the most important drivers of economic growth. However, economic development with steady increases in productivity and consumption has also led to exponential usage of natural resources (Vaidya et al. 2018). The energy consumption caused by ICT is difficult to estimate, but studies for the European Union and the U.S. estimate the electricity consumption from ICT to be between 4.3 and 8% of total electrical demand in 2008 (European Commission 2008, Laitner et al. 2008), which is certainly noteworthy. It deserves attention because the increase is rapid – for the EU, it was believed to be 50% more in 2023 from the 2008 baseline. Efforts have been made to address the issue, like low-energy ICT systems. Not only environmental aspects dominate the discourse – cost of operation, heat dissipation through processors, the operating lifetime of battery-supplied devices are all part of the problem.

On the other hand, the ICT sector is believed to make significant impact to reduce resource consumption to mitigate the impacts of climate change. Actors hope ICT can enable new concepts to generate, allocate, distribute, share and use energy in resource-efficient and environmental-friendly terms. Growing numbers of large infrastructure systems were optimized for lower power consumption, and the general potential for large-scale simulation and control play a critical role. To underpin this statement with numbers, the energy productivity index shows is decreasing for steadily since the 1970s (Blok et al 2015). The trend can be perceived as the decoupling of energy demand and economic growth (Mattern et al. 2010). ICT is a positive driving force for energy productivity, Laitner et al. (2008) estimate one kW used by ICT environments save 10 kW economy-wide through productivity gains. ICT, in this context, should be a tool for indirect energy use optimization and conservation.

ICT is critically important in enabling paradigm shifts in energy sector. The view of energy as a precious resource, the deregulation of the energy market, distributed generation and cooperation efforts have all been highly influenced and triggered through ICT utilization. Recent advancements to find ways for computing efficiency such as networking, embedded systems and automation building have been of relevance in this respect (ibid). This Master

thesis circles around the question how advancements in technical concepts influences the energy demand of information and communication technology.

Important in this context is to mention that Koomey's Law is getting more significance in recent years. Many people are familiar with Moore's Law – a doubling of computational power every 18 months in hardware development. Synonymously, energy consumption for computing improves just as fast as processing power. A study by Dr. Jon Koomey revealed that for “a fixed computing load, the amount of battery you need will fall by a factor of two every year and a half”. With knowledge of this law, it is possible to develop more mobile computing and sensing applications, as it is certain that energy efficiency will continue to steadily improve (Brynjolfsson 2011). Better tubes at first, and secondly smaller transistors have contributed to reduced power usage, and there is every reason to believe that this trend will continue at least for as long as Moore's Law is intact, which is believed to be 5 to 10 years (Koomey et al. 2011).

With smaller power requirements of computer-based devices, the software side is a source of energy efficiency efforts. The equally important aspect is how energy and resources can be conserved through efficient behavior of ICT hardware. Sustainable software development and technology, in this mirror, should take negative and positive impacts because of a software product over its whole life cycle into account (Dick et al. 2011). Software goes through the same life cycle as haptic products, from development and distribution to usage, and from there eventually to deactivation and disposal. As for haptic products, their release has first-order, second-order and even third-order effects on the environment, which should be considered. Overarching indicators used to measure the sustainability of architectural software concepts by researchers so far were energy behavior of software, capacity and resource utilization. Energy behavior measures the degree to which the energy consumed by an application meets requirements. The capacity describes the degree to which the maximum energy consumption limits of an application meet requirements. Resource utilization tells the degree to which the utilization of resources used by an application meets requirements (Bruntink et al. 2014). These indicators will be critical in assessing the blockchain technology in its environmental terms for this Master Thesis.

3. Industry 4.0

Industry 4.0 is a conceptual idea for manufacturing introduced as a strategy by the German government, promoting the computerization of manufacturing in the country (BMBF 2016) around 2010. However, the concept has received widespread attention in all economies of the Northern Hemisphere since then, including the European Union, the US, China, India, and other Asian countries. The term refers to the belief in a “Fourth Industrial Revolution”, with the first three stemming from mechanization, electricity application and the emergence of IT. The fundamental difference between Industry 4.0 and previous revolutions is the opportunity for proactive guidance actors possess. The vision is global interconnectivity, full automation and augmented human-machine interaction, connected cyber-physical systems sharing information to trigger actions. It is assumed that Industry 4.0 will bring about improvements in the industrial process within manufacturing, through engineering material usage, supply chains and product lifecycle management (Gilchrist 2016).

It is difficult to find a clear and precise definition of Industry 4.0. Examples include it can be “best understood as a new level of organization and control over the entire value chain of the lifecycle of products, it is geared towards increasingly individualized customer requirements. This cycle begins at the product idea, covers the order placement and extends through to development and manufacturing, all the way to the product delivery for the end customer, and concludes with recycling, encompassing all resultant services” (Gilchrist 2016). The basis for Industry 4.0, in this understanding, is the availability of all relevant and non-relevant information in real time by connecting all instances involved in the value chain. This connection, of people, systems and things creates dynamic, self-organizing, real-time value added throughout the industrial sector. Optimization can be directed towards different criteria such as costs, availability and resource consumption (Medhi 2016). Another definition singles out dependencies for Industry 4.0 frameworks, like the digitization and integration of horizontal and vertical value chains, the digitization of products and services and the introduction of innovated business models. Another definition defines *Industry 4.0 as a collective term for technologies and concepts of value chain organization, fostering optimization through digital facilitation* (Marr 2016). This last definition is being suitable for the purposes of this paper and will be therefore used in the context of the Master Thesis as definition for Industry 4.0.

It is believed that although advanced digital technology is already utilized, the movements a prospect of complete human-independent production systems yearn are pleasant. It will lead to greater efficiencies and change traditional production relationships among suppliers, producers and customers. According to Boston Consulting Group, there are 9 technology trends forming the building block of Industry 4.0:

- **Big Data and Analytics:** Collection and comprehensive evaluation of data from many different sources, from production equipment and systems to enterprise- and customer-management systems.
- **Autonomous Robotics:** Robots will eventually interact with each other and work safely side by side with humans.
- **Simulation:** Simulations will be extensively used in plant operations to leverage real-time data and mirror the physical world. It will allow operators to test and optimize production processes for product lines.
- **Horizontal and Vertical system integration:** With Industry 4.0, companies, departments, and functions will become more cohesive, as cross-company, universal data integration networks evolve and enable automated value chains.
- **Cybersecurity:** Industry 4.0 creates the need to protect critical industrial systems and manufacturing lines from cyberthreats.
- **Cloud Computing:** Production-related undertakings require increase data sharing across sites and company boundaries. Cloud computing allows the application of such data.
- **Additive Manufacturing:** Additive manufacturing methods, such as 3D printing, gain increasing relevance in the industry, as the demand for individual components and customized products increases.
- **Augmented reality:** Augmented reality systems support a variety of services, such as selecting parts in a warehouse and sending repair instructions. In the future, augmented reality provides workers with real-time information to improve decision making and work procedures.
- **Industrial Internet of Things:** Industry 4.0 also means that products, sometimes even unfinished, are enriched with embedded computing. This will allow field devices to communicate and interact, decentralizing decision-making and analytics (Boston Consulting Group 2018).

For Industry 4.0 to function, implementers must consider design principles to avoid shortcomings in a holistic interconnected and automated system. The first, obvious principle is interoperability between machines, devices, sensors and people to connect and communicate. Information systems, furthermore, must possess the ability to create virtual copies of the real world by enriching all models with sensor data, in short complete transparency of the information they process. The ability of assistance systems to support humans in making informed decisions is of utmost importance. Industry 4.0 systems should enhance production processes and human needs. Finally, cyber-physical systems need to be able to perform tasks autonomously and make their own decisions based on input, by default be decentralized (Hermann et al 2016).

Wolter et al. (2015) outlined the challenges for Industry 4.0, which should not be underestimated: Total interconnectivity can cause IT security issues, greatly aggravated by the inherent need to open closed systems. There is a high reliability and stability need for critical machine-to-machine communication and processing, including latency times. There are needs to maintain integrity of production processes, as well as holistic system design, and the protection of industrial know-how. The lack of adequate skill sets, which is one of the greatest concerns. The threat of redundancy, a general reluctance to change by stakeholders, the loss of jobs, unclear jurisdictional circumstances, unclear economic benefits, these factors are of concern and have to be addressed in some manner before Industry 4.0 can become a reality (Wolter et al. 2015).

3.1. Internet of Things (IoT)

Mattern and Floerkemeier (2010) thought of the Internet of Things as a representation of a vision in which the Internet extends into the real world, embracing everyday objects. The devices can be controlled remotely and act as physical access points, making the Internet an omnipresent force of human life (Mattern et al 2010). The vision of IoT is based on the idea of full integration of devices due to diminishing size, falling prices and lower energy consumption. Due to advanced features, devices can perceive the context in which they are operating. Built-in network capabilities allow them, in further consequence, to communicate with each other and with people, generating substantial added value for their users. Furthermore, everyday objects can inform its users constantly about its state by collecting up-to-date information. The ability to process this information to make objects more efficient for our everyday use in an automatic, rapid and informed manner opens up new opportunities for

utilization and goes together with the opportunity for new business processes, delivering substantial economic and social benefits.

The term “Internet of Things” was first coined by Kevin Ashton of the Auto-ID Centre at the Massachusetts Institute of Technology (MIT), in an understanding of “a standardized way for computers to understand the real world” (Schoenberger 2002). From the technical angle, it is not the result of single novelties or innovations, but instead several complementary developments which provided capabilities to connect the physical and virtual world. The capabilities are foremost based on the ability of objects to communicate and cooperate because they utilize Internet resources with technologies such as GSM, UMTS, Wi-fi, Bluetooth, ZigBee and other wireless standards related to Wireless Personal Area Networks (WPAN). Above that, the capability these devices can be uniquely identified gives them a link to information about the particular object, which observers can learn from. The ability to sense objects’ surroundings, record it and forward it for processing is similarly employed. Embedded information processing can be used to process and interpret sensor information, or give products memory of how they may have been used in the past. Furthermore, smart things can be localized, which gives their users freedom what spatial capacity can be given to objects (Mattern et al. 2010).

The idea of an Internet of Things raises manifold expectations for enterprises as well as users: From a commercial side, increased efficiencies, reduced costs and more targeted service, new business models involving smart things and associated services, the general increase in life quality, smart assistance systems for increased safety, etc. (Mattern et al 2010).

The major trend concerning IoT in recent years is the explosive growth of connected devices. However, it deserves mention that previous predictions for 2020, heralding 50 billion connected IoT devices until 2020 were highly exaggerated. A study conducted in 2016 estimated the number of IoT devices between 7 billion and 17 billion devices, depending on inclusion of computers and smartphones. Expectations now estimate 20 billion connected devices at most (Nordrum 2016). A wide range of applications means that specifics can be different but basic characteristics are shared by IoT devices.

Due to Gautier (2010), the Internet of Things will establish a “semantic web”, which means any innovation in computer sciences can complement the Internet of Things. Consequently, no common standards are needed, things and people will be adaptive to new standards. Gautier refers to this behavior as *event-driven architecture* (Gautier 2010), with no finality.

However, the IPv6 protocol is believed to play a major role connecting devices to IP networks. IoT is a complex system, and will be likely be perceived as a chaotic environment in its totality, since bold predictions claim that 50 to 100 trillion objects one day could make up the human environment, with 1000 to 5000 trackable objects surrounding people closely (ibid). Other trends which were not foreseen for IoT were ambient intelligence and autonomous control, which are also gradually integrated into the concepts of the Internet of Things.

The challenges progressing from the Internet of Computer to the goal of a remote Internet of Things must be done one step at a time. There are several key challenges before IoT can be realized, with some overlaps with Industry 4.0:

- The scalability issue is implicit in every application of computer power. Since IoT potentially has a larger scope than the Internet of computers, the functionalities should be adjusted. On the other hand, things predominantly operate locally, so small- and large-scale applications need to be possible.
- Demand to arrive and operate: Mobile objects which are sporadically used need to establish connections on-demand, organize and configure to suit local environments.
- IoT needs to secure interoperability between physical things, although smart objects are likely to have very diverse information to process. Therefore, standards are needed.
- Suitable services need automatic identification from their dynamic environments, this requires appropriate semantics. Information on the product needs to be available.
- More extensive software complexity is needed to ensure smart objects can be reasonably managed.
- IoT applications need to be robust against both small and infrequent communication packages, as well as huge volumes of data to process.
- Sufficient data interpretation is needed to determine local context of IoT operating devices. The generation of useful and processable information is key for IoT.
- The Internet of Things needs special and customized security and privacy frameworks, depending on its application field. With it, objects can allow selective access, prevent communication with other objects and protect its information from adversaries.
- There needs to be implicit fault tolerance towards malicious behavior.
- The power supply of smart objects is probably the hardest hurdle to overcome. Some sensors, like passive RFIDs, need no external energy source, however this is not true for a number of hardware components. They are problematic due to size, weight and

maintenance requirements. A partial solution to this could be the usage of piezoelectric and pyroelectric material.

- Short-range communication still needs to be possible even though other smart objects intervene. A typical example to overcome the hurdle is NFC, which uses inductive coupling.
- Establishing wireless technologies for communication with low energy consumption during operation.

The RFID (Radio Frequency Identification) technology works with electromagnetic fields for communication and is expected to be the backbone of the Internet of Things. A stationary reader typically communicating wirelessly with small battery-free transponders attached to objects to identify them. The development of the technology is reflected significantly in the cost reductions (Material Handling Industry, 2018). High cost pressure and battery absence means that RFID communication protocols cannot be based on established Internet protocols due to limited resources. Everyday objects with RFID will therefore not behave in the same way as Internet nodes. It is likely an optimized wireless protocol will be used for the last few meters if adverse conditions are present in the physical world (Mattern et al 2010).

If everyday objects shall be addressed by the Internet of Things, ideally sensors should not resort to special communication protocols like RFID. Mattern et al propose an Internet node behavior of smart objects, making use of the IPv6 protocol with 128-bit addresses. The function would enable developers to incorporate objects into global interoperability, network-wide data packet delivery, data transport across different media and network management. However, one has to keep in mind the enormous resources required to run such a system, both for processor capacity and energy consumption. A solution could be the connection through proxies or gateways, which again has ramifications on architectural complexity, maintenance and operation, initiating costs. Other microchip protocols for tracking, like ZigBee, also require significant amounts of energy, with sensors needed to be equipped at least with AA batteries (Mattern et al 2010).

A logical development for the Internet of Things is leveraging the Web as an infrastructure for smart objects. The formats used can be understood not only by machines, but also people. Interaction of things between a normal web browser and a person allows a variety of explorations in the world of smart things and its relationships, possibly enabling other functionalities additional to display on the web.

The European Commission, regarding the Internet of Things, has released an action plan to lay out its ideas how IoT may change European society and how what framework they see inevitable for the realization of it. First and foremost, European policy-makers and public authorities address the issue in a way to ensure IoT technologies and applications will stimulate economic growth, improve individuals' well-being and address some of today's societal problems (European Commission 2009). They defined 14 lines of action, along which the Commission intends to propose legislation. These are:

- Definition of a set of principles underlying the governance of IoT
- Continuous monitoring of privacy and protection of personal data questions, guidelines to operate IoT devices in compliance with privacy and data protection
- Initiation of the technical and legal discussion on a “right of silence of the chips”, the possibility to disconnect from networked environments at any time.
- Identification of emerging risks, related to trust, acceptance and security
- Monitoring of the development of IoT infrastructures
- Development of standards and Standards Mandate for the European Commission
- Financing research and development in the area of IoT, including microelectronics, non-silicon based components, energy-harvesting technologies, smart networks, semantics, novel applications etc.
- Fostering public-private partnerships for IoT on all scales.
- Support of innovative pilot projects which deliver strong benefits to society, such as e-health, e-accessibility, climate change, or bridging digital divide
- Fostering institutional awareness by informing European Parliament, European Council, Economic and Social Committee, Committee of the Regions, Data Protection Working Party and others about IoT developments.
- International dialogue to agree on relevant joint actions and sharing of best practices
- Usage of RFID technology in recycling lines as part of regular monitoring of waste management industry
- Monitoring the introduction of IoT-related technologies to measure the uptake in the economy
- Establishment of a multi-stakeholder mechanism to work along the lines of action and carry out the tasks laid down (European Commission 2009).

3.1.1. The Industrial Internet of Things (IIoT)

The Industrial Internet of Things (IIoT), a specific form of IoT, goes back to a term coined by General Electric and known as “Industrial Internet”, the “Internet of Everything”, among other designations, and stands for corporate efforts to holistically interconnect industrial hardware, middleware and software to collect data and use the results gained from interrogating large data sets through advanced analytics to achieve operational efficiency and accelerated productivity (Gilchrist 2017). By implementing IIoT, enterprises are provided with a way to get better visibility and insight into operations and assets through machinal network amalgamation and integration. It is important to consider vertical IoT strategies, such as consumer, commercial and industrial forms of the Internet from the broader, horizontal concept of the Internet of Things. They have different target audiences, technical requirements and strategies to achieve their respective ends. The consumer market has the highest market visibility, with smart homes, personal device connectivity, among other examples. The commercial market, alike, benefits from high marketability with provided services such as finances, ecommerce, which focus on consumer history, performance and value. IIoT, on the other hand, is a vertical focused on enterprises and ranging from small- to medium and large-scale businesses (ibid).

The reason why IIoT is pushed across lies in the potential development IIoT possibly could facilitate: The last 15-20 years have seen stellar growth rates of the business-to-consumer sector via Internet, particularly trading in retail, media and financial services. Success stories such as Amazon, Netflix and PayPal have given hope to similar developments and growth to industry, in this context denoting manufacturing, agriculture, energy, aviation, transportation and logistics. Since two thirds of global GDP are generated through industry, the issue is of importance for the sector (ibid).

The innovation history of IIoT started with Ethernet use on the plant floor at sensor level. Additional early use of IIoT were internet protocols as means to merge IT and operational technology (OT). Today, production information can be accessed worldwide. The embracement of open-system architectures will challenge the proprietary networks that were dominating industrial automation. Continual improvements in speed, security and reliability will ultimately lead to an interconnected industry. Systems of the manufacturing environment can get linked to people needing access to this information (Hoske 2016).

However, the IIoT is still in its outset considering the possibilities enterprises already have through advanced technology. Despite the long existence of the Internet, industrial leaders are hesitant to commit to IIoT. It is the result of the uncertainty how IIoT would affect existing industries, value chains, business models, workforces and ultimately productivity and products (ibid). Well-known technologies which are subject to IIoT have been implemented for nearly a decade now, such as machine-to-machine communication and collaboration, as well as advanced sensing. They are not utilized interconnectedly yet since the collection of vast quantities of data for historical, predictive and prescriptive information contributed greatly to increasing revenues in the industry. Possible network effects were not explored further because the earnings indicated no need to do research in the area (Gilchrist 2017).

To illustrate the implications of implementing IIoT systems, a good approximation is the “power of 1%” rule for industry. The term expresses that only 1% in savings of operational costs or reduction of inefficiencies can have significant impacts. As an example, in aviation, fuel savings of 1% relates to \$30 billion Euros. The same holds true for the Oil & Gas Industry, agriculture, transportation and healthcare industries (ibid).

Key IIoT technologies are the already-mentioned advanced sensor technologies and machine-to-machine communication. Sensors produce not just more data for a component, but also different type of data, instead of just being precise. They can be utilized for predictions, self-comparison with similar applications to determine configuration needs and environment adjustments and self-awareness. It poses all prerequisites for self-diagnostics. It enables the instrumentation of machines and processes (ibid). How IIoT is utilized is strongly depending on the use case for implementation. Requirements for manufacturing, for instance, differ greatly from transportation, which also differs from healthcare. However, what IIoT systems can offer are potential solutions for all vertical industries, by utilizing sensor technology, wireless communications, networking, cloud computing, and Big Data analysis (ibid).

Manufacturers consider the changing conditions in production a great challenge, as they are used to process materials into products to sell them. For them, the sticking point in making the digital leap is more cultural than anything (Economist 2016). The two major goals for implementation of IIoT, as laid out by Neuberg (2016), have to be the *upscaling of productivity* and the *downscaling of complexity* in industrial processes. The approach must be “wrap and reuse” rather than “rip and replace”, since it allows greater business control and

will drive the evolution toward smart enterprises, which are more efficient, safer and sustainable.

There are three major areas where IIoT utilization could pay dividends after implementation. In *asset-performance management*, data analytics, wireless sensors and cloud connectivity will improve asset performance. The tools allow easy gathering of data and simultaneous conversion into actionable information. Preventive maintenance, energy management and condition-based can thereby be largely improved (Neuberg 2016). *Augmented operation* with mobile devices will lead to increased productivity, with another impact being the development of user-centric plants rather than machine-centric manufacturing. Finally, *smart-enterprise control* will ascertain tight integration of smart-connected machines and assets, facilitating more flexible and efficient manufacturing. This will also help corporations in reducing complexity for production (ibid).

According to Beyerer and Usländer (2016), the first shift for producers connected to the caesura coming from IIoT realization is from products to services. The data generated delivers the raw material for new services, which will be more profitable than the products they are based on. The Economist also believes that a related change will be the race for “platforming” – software foundations upon which services and applications can be built (Economist 2015). The issue of platforming and opposing concepts, such as process dissemination through blockchain will occupy a major portion of this Master thesis.

3.1.2. Environmental considerations for IIoT

IIoT systems are becoming more complicated with growing scales. This facilitates a number of significant challenges that need to be considered, such as energy consumption. A major argument to adopt IIoT in the early discussions was, in fact, the reduction of resource consumption and carbon emissions of industrial systems. However, IIoT systems themselves consume a considerable amount of energy for their purposes and lead to a larger carbon footprint. On the other hand, the systems typically consist of devices with minor power needs, with batteries being heavily utilized to run these devices, which is an operation limiter for IIoT systems (Wang et al. 2016).

There are various factors which can be tackled to lower energy consumption for IIoT systems. In the IIoT domain, data collection is reliant on massive sensor nodes and devices. Thus, a way to make the system more energy-efficient is the optimization of sensors, processes and communications between IoT devices. Furthermore, so-called wireless sensor networks and

their respective topological structure can be deployed to lower energy demand. The disentanglement of complex mesh or hybrid structures or the abolition of hierarchical networks can either reduce complexity or improve productivity to establish sustainable IIoT systems (ibid).

4. Obstacles of Blockchain Utilization

4.1. What is Blockchain?

The blockchain technology is fundamentally based on the idea of allowing digital information to be distributed. Information held on blockchains exist as shared, and continually reconciled, database. This allows many cases of utilization of data. Since the database is not stored in any single location, all records are transparent and easily verifiable. Hosted by millions of computers simultaneously, the data is accessible to anybody with Internet access (Rosic 2016).

To illustrate the data structure of blockchains, one has to think of it as a chain of data packages, where blocks comprise of multiple information entities, like transactions. The blockchain is extended by appending blocks of information onto the block last computed. By that, it represents a complete ledger of information storing history. Fundamentally important is that blocks are validated by the network through cryptographic means (Nofer et. al 2017).

Each block in the most fundamental blockchain data structure contains a timestamp, a hash value of the previous block, called *parent*, and a *nonce*, a random number needed to be computed for verification of the hash. The nonce value will receive specific attention in the *Blockchain Footprint* chapter. The concept ensures integrity of the blockchain through the first block generated, the *genesis block*, since hash values are unique and prevents malicious acts by design, since changes in the chain would change the hash value (ibid).

Blockchain presents a technology of *built-in robustness*. Storing blocks of information and distributing them identically across the network allows *no single entity to take control* of the blockchain and *no single point of failure* to occur. To achieve identical information across the network, the blockchain lives in a *state of consensus* that by design automatically checks with itself however it is intended by its developers. It can be seen as a *self-auditing ecosystem* of digital value, as the network reconciles every event people want to record on a blockchain. The blockchain means *full transparency of data* within the network (Rosic 2016). Blockchain is fundamentally based on the idea of *decentralization*. Anything recorded on the blockchain is a function of the network. Blockchains have no central authority, the network operates on a

peer-to-peer basis (ibid). The possible network effects possible from these ideas are now in the beginning of being investigated.

The most important concepts being implemented are:

- A fully decentralized distributed database, operating on a peer-to-peer basis,
- A list of fully transparent information,
- A technology that is based on the idea that its information is automatically updated for everyone, creating a state of consensus
- A technology that, however, still protects information through cryptographic means.

4.1.1. Technical requirements

Blockchain networks in their most basic design consist of *nodes*. Nodes are computers connected to the blockchain network using a client. Nodes are the backbone of the blockchain structure. They perform the task of validating and redirecting the information inserted into blocks, additionally they update the network after a state of consensus is reached. For doing so, nodes must carry a copy of the whole blockchain on their memory, but this condition is widely removed from new blockchain designs. Every node administers the blockchain and joins the network voluntarily (Rosic 2016).

To implement a distributed validation and redirection information system like blockchain, there is a need for a sufficient system of proof that information has been validated and redirected and put into a block. In the early stages, the mechanism of choice was the proof-of-work system like Hashcash (Nakamoto 2008). Through proof-of-work, a node can verify that it has contributed with its computing power to the network, display it in the network and the other nodes confirm the efforts, consensus is reached and the blockchain is updated. Once the CPU effort has been expended to satisfy the need of the proof-of-work, the block generated cannot be changed without redoing the work. Because coming blocks are chained after another, the work to change a specific block would include redoing the work for all blocks after it (ibid). Proof-of-work also solves the problem of determining representation for majority decisions. The principle of “One-CPU-One-Vote” applies, the majority decision is represented by the longest chain because the most computing power was invested in it. If the majority of CPU power is controlled by honest nodes, the honest chain will outpace malicious chains in the system, thereby creating a network where no trust is needed (ibid), allowing transfer of data assets all over world on a peer-to-peer basis. The consensus mechanism is the

process of a majority of network validators come to agreement in the state of a *ledger* (Nofer et al. 2017).

The block creation is referred to the term “mining”, the block creators are called *miners*. Miners are rewarded for validating blocks. It should be noted that the process of generating new blocks implies performance problems if blocks are added to the network at a high rate. This obviously reduces the pace of the system, and shows to be a real bottleneck in blockchain application.

The distributed ledger system allows peers to interact with trustworthiness without intermediaries. The absence of intermediaries, known as disintermediation, furthermore fosters data security. The current method of third party systems to establish trustworthy relationships, in an economic sense, implies a risk of security breaches, which becomes obsolete with the usage of blockchain and increasing user security (Nofer et al. 2017). Rückeshäuser refers to distributed ledgers as “democratization of data” (Rückeshäuser 2017). It allows the creation of autonomous business models, based on the execution of Turing-complete codes, like smart contracting. The technology offers a variety of technical design options. The prerequisites depend on factors such as the area of application, network access, and the size of the network. A major differentiation of ledgers is between *permissioned* or *permissionless* ledgers, describing the boundaries of participation blockchain developers may have inducted (Rückeshäuser (2017). Another important factor is the information-ordering precondition of blockchains, we distinguish between inherent chronological order and non-chronological orders.

4.2. Application Potentials – Use Cases

There is a controversial debate whether the application of distributed ledgers may be justified by possible economic benefits (Rückeshäuser 2017). The quantification of potential of blockchain implementation is a topic of relevance, but will not be covered by this Master Thesis. Most notably, the potential of distributed ledger may present itself in different ways than anticipated, depending on the industry and value proposition by corporations. The fields that were identified in the academic community were:

- Digital currency: The most well-known application field for blockchain technology are digital currencies. An electronic coin can be defined as a chain of digital signatures. Each owner of coins can transfer the coin to the next by digitally signing a

hash of the previous transaction and the public key of the next owner, adding it to the end of the coin. A payee verifies the signatures to verify the chain of ownership (Nakamoto 2008).

- Smart Contracting: Distributed ledgers allow coding of simple contracts which are executed when specified conditions are met. At the current level of development, smart contracts can be programmed to perform simple functions, such as kick-in clauses for benchmark achievements.
- Sharing Economy: The sharing economy is already proven success, and distributed ledgers could bring the next evolution to this field by eliminating intermediaries, as it allows direct interaction between peers.
- Crowdfunding, product development: The success of crowdfunding suggests that people want to have direct influence in product development. Peer-to-peer crowdfunding has the potential to create crowd-sourced venture-capital funds.
- Governance: Distributed ledger systems could have great impacts in public participation and organizational decision-making in general. In practice, governance of corporations or municipalities could become fully transparent.
- Supply Chain Auditing: For consumers with ethical minimum standards, the traceability of materials, possibly financing conflicts, is of great importance. Distributed ledgers provide a way to certify the supply chain of products, like sustainably harvested fish.
- File Storage: Decentralizing storage has obvious benefits. It protects sensitive information from getting lost, speeding up file transfer and streaming bandwidth. It is a necessary load relief for current content-delivery systems.
- Prediction markets: Event probabilities are more precise with an increasing underlying body of opinions. Distributed ledger seems to be destined for utilization in this field. The “wisdom of the crowd” applies here.
- Intellectual property protection: The Internet of today is full of free content, leaving copyright holders without deserved royalties. Distributed ledgers present a way of protecting copyright and automation of sale of creative work online, eliminating the risk of file copying and redistribution.
- Energy Management: Blockchain can give way to self-generated energy redistribution and monitoring, as well as certifying renewable energy, which experiences increasing demand.

- Identity Management: There is a need for identity management on the Internet, as the remedies of web commerce are well-known and still a point of failure for judicial systems. The possibility of encryption in distributed ledgers allows online identity solutions to be found.
- Anti-money laundering, know your customer practices: AML and KYC practices currently suffer from labor-intensive multi-step processes which are required. Cross-institutional client verification and transaction monitoring could be enabled by distributed ledgers.
- Land—title registration: Publicly accessible ledgers, in fact, can make all kinds of record-keeping for backtracking purposes. Property titles are important in this context. They are susceptible to fraud, and land registry could help avoiding that.
- Stock trading: Peer-to-peer stock trading could make the already-fast stock market instantaneous, and intermediaries get removed from the process. Some stock exchanges already experiment with blockchain applications for their services, such as ASX and JPX.
- Internet of Things (IoT): Distributed ledgers allow the participation of IoT devices in any sort of network provided for them. They expand the possibilities how we can engage machines for human utilization, with application fields not limited to mass-scale automation systems, machine-to-machine economy, predictive maintenance and data analytics (Rosic 2016). The use cases for blockchain for IoT are further elaborated in this paper.

4.3. Industrial Use Cases for Blockchain

For implementation of blockchain for IoT schemes, it is of high importance for architectural considerations to visualize the locations of blockchain nodes in the system. It has significant impact on bandwidth, computation and space requirements. Business logic and data can also be of influence in this sense: For instance, states of smart contracts can also be part of a blockchain, which would allow tracking of IoT services and states of physical assets. Placing too much logic and data onto a blockchain can lead to poor performance. Finally, the mechanisms of cyber-physical integration need to be considered for architectural purposes. For example, what are objectives of transactions on-chain which constitute a real advantage in comparison to off-chain exchange? Taking these considerations for design into account, Liao et al (2017) came up with four typical architectural styles for IoT blockchain realizations.

They call architectures Fully centralized, Pseudo-Distributed Things, Distributed Things and Fully Distributed. When an architecture is fully centralized, it does not deploy blockchains at all. It is vendor-specific and controlled through a websocket protocol to communicate in a cloud. For Pseudo-Distributed Things, the endpoint nodes of a chain are managed in the cloud by an endpoint container component, and users as well as smart objects can interact with the blockchain through a chain gateway, the integration magnitude is low. Distributed Things describes a high level of blockchain integration in IoT processes, the IoT gateway is not necessary anymore as the smart device is directly controlled by the corresponding smart contract in this architecture. This is of special advantage if tasks of blockchain endpoints need to be performed in a low-capacity environment. Fully Distributed refers to full deployment of blockchain, also in end user devices. The blockchain in this architecture typically has no payment logic implemented, users pay directly to smart contract addresses (Liao et al 2017).

Christidis (2016) states that while blockchains have advantages, especially when it comes to adjustments in a use scenario, the flaws are still significant: Privacy-wise, he says by identifying patterns and connections between addresses, interested parties can draw informed inferences and utilize it to their advantage. In an IoT scheme, this has to be considered. Christidis, therefore, suggests safeguards such as new keys for every transaction, not use the same blockchain for all transactions in a permissioned blockchain or to attain transactional privacy through homomorphic encryption or zero-knowledge proofs (Christidis et al 2016). Under these circumstances, they also shed light on other constraints, like the miner set, since it could cause censorship through mining prevention, the volatility of tokenized assets, the need for secure communication and file exchange services on blockchains and DNS services to reduce latencies (ibid).

De Castro (2017) openly claims that the IoT and IIoT will experience the biggest gain from blockchain implementation for specific use cases, including finance and identity applications. He praises the opportunities for the pace of chains of command, which would dramatically increase with blockchain in contrast to centralized control. Peer-to-peer messaging bypasses inefficient central databases and allows autonomous communication. Furthermore, the property of eliminating a single point of failure reduces the implications of failure of a single entity to a minimum. He believes peer device communication will become the new normal for large device deployments (De Castro 2017).

For IoT, it is believed that security will play a key role going forward. The data structure of blockchain does not allow many of the known attacking methods for computer systems, as no

single entity can take over the network. According to estimates from Gartner, 20% of businesses will deploy security solutions for the protection of their IoT devices and services, with a need to broaden the scope of security strategies in direction of including also these devices. Costs of regulatory compliance and governance needs create tremendous burdens for enterprises and should also not be underestimated (ibid).

While the vision of De Castro allows a first insight how the industry could possibly utilize blockchain, there are several other fields where it makes sense to deploy distributed networks as a backbone for industrial use, most notably in manufacturing. Sandner et al (2017) provide an exhaustive list of fields for the manufacturing industry where blockchain is of high interest. They see potential in supply chain management and auditing, with the Internet of Things and 3D printing being other specific areas where blockchain could disrupt established manufacturing processes.

According to Sandner, PwC developed a scheme to determine beneficial blockchain use cases, and what prerequisites for deployment may be. Therefore, four of the six requirements for blockchain implementation need to be met. For PwC, multiple parties have to share data in some context, multiple parties have a need to update data for a process, there is a requirement for verification, intermediaries in these processes add resource expense, interactions are time-sensitive and a form of transaction interaction is utilized for blockchains to make sense.

Databases with traditional software overwhelmingly still fulfil most requirements, the conditions need to be met to facilitate processes through blockchain. Decentralized systems only make sense in an environment of multiple parties involved where the sharing and timely data reliably is necessary to advance processes. Blockchain can serve as means for disintermediation. Intermediaries, after all, cause costs, increase complexity and provoke delays. Less intermediaries translates to more efficiency, if implemented correctly (Sandner et al. 2017).

Furthermore, the functionality of blockchain determines its use as a tool to manage and secure digital relationships as part of a system of record. Relationships through economic cooperation should not be underestimated. The digital

Assessment of usefulness framework

	Multiple parties share data	Multiple parties update data	Requirement of verification	Intermediaries add cost and complexity	Interactions are time sensitive	Transaction interaction
3D printing	✓		✓	✓		✓
Asset sharing	✓	✓	✓		✓	✓
Digital product memory	✓	✓	✓	✓		
Distributed manufacturing	✓	✓	✓	✓		✓
Machine communication	✓	✓			✓	✓
Quality documentation	✓	✓	✓	✓		
Supply chain tracking	✓	✓	✓	✓		✓
Verification of spare parts	✓	✓	✓	✓		

world presents opportunities to expand these relationships and have economic cooperation without underlying ties, but cannot replace them. Trust issues between parties are a predominant underlying condition to apply blockchain. Institutions establishing trust between parties can be removed from the process (ibid). The need for irreversibility of data in processes is also a valuable property for blockchain application, like controlling and general documentation.

However, as long as decentral and collaborative processes are not state of the art technically and culturally, blockchain will only serve as a medium, whereas other centralized technical solutions commonly used today have greater advantages. The greatest challenge to blockchain application are unclear legal and regulatory ramifications, lack of confidence from enterprises, lack of standards, technical issues, scalability, and latency. A prerequisite for blockchain use in certain areas, therefore, is to tackle these topical issues. Since holistic solutions addressing all obstacles are rare today, rapid technological progress after the Proof-of-Concept phase is low, preventing faster adoption of the technology (Sandner et al 2017).

To validate potential applications and the assessment of expected developments, the phase of blockchain evolution needs to be determined. Sandner et al utilized the Gartner Hype Cycle to do so. It is a widely-used tool to classify maturity, adoption and social application of technologies. It shows the expectations and attribution towards new technologies since emergence. The cycle consists of five phases: After being triggered, the expectations immediately increase greatly and unrealistic attributions are made. When the expectation cannot be met, the hype decreases, eventually falling into the trough of disillusionment. An increasing number of people understanding the technology lead to a slow increase in attention, and in a slope of enlightenment the technology is finally understood by a sufficient pool of experts. Finally, mainstream adoption and productivity increases are observed. For blockchain, it is believed it has not passed the trough of disillusionment yet. However, it is believed the first major use cases will be running in three to five years (Sandner et al 2017).

The use cases, broadly speaking, can be separated into three types of utilization: Supply Chain Management and Digital Product Memory, Industry 4.0 applications, and 3D printing. For Supply Chains, a PoC showed that blockchain could be deployed to track containers during shipping. The goal was to reduce the substantial amount of paperwork connected to shipment. Vendors, like shippers, freight forwarders, ocean carriers, parts and customs could potentially be connected to the blockchain, significantly dropping shipment costs. Another


implementation scheme showed that certifications and other important product information for the supply chain can be displayed on a blockchain. Every product would receive a “digital passport” proving authenticity. Additionally, the Proof-of-concept showed the waste of valuable materials which can also be mitigated through product life cycle documentation on the blockchain (Sandner et al. 2017).

Similar to supply chain management, the idea of blockchain implementation for an Internet of Things (IoT) was also tested exhaustively and proven as a concept. Issues of identification of IoT devices and the vulnerability during this process were one field of application. The use case arose due to the expensiveness of current authentication from authorities. Specifically, devices were equipped with a digital identity which cannot be manipulated and automatically updated. Another use case proposed a usage of blockchain to timestamp data. Apparent is also a marketed use of data collected from IoT devices, with a potential to utilize that data for other marketable purposes (ibid).

Finally, 3D printing, itself a field attracting a great deal of attention, can make use of blockchains. For example, the establishment of a 3D supply chain, fully automated by blockchain, is thinkable. Smart contracts can automatically negotiate pricing, customer conditions and logistic service. The information can also be utilized for product recycling. Another project showed that 3D printing could utilize blockchain for point-of-use systems to save logistic and inventory costs (ibid).

To evaluate disclosed use cases for the manufacturing industry and their time horizon for implementation, Sandner et al created a Cross-Potential matrix for above use cases identified. The variables for this scheme were the *time to market* and their *potential for the manufacturing industry*. These variables stem from the fact that the majority of use cases are

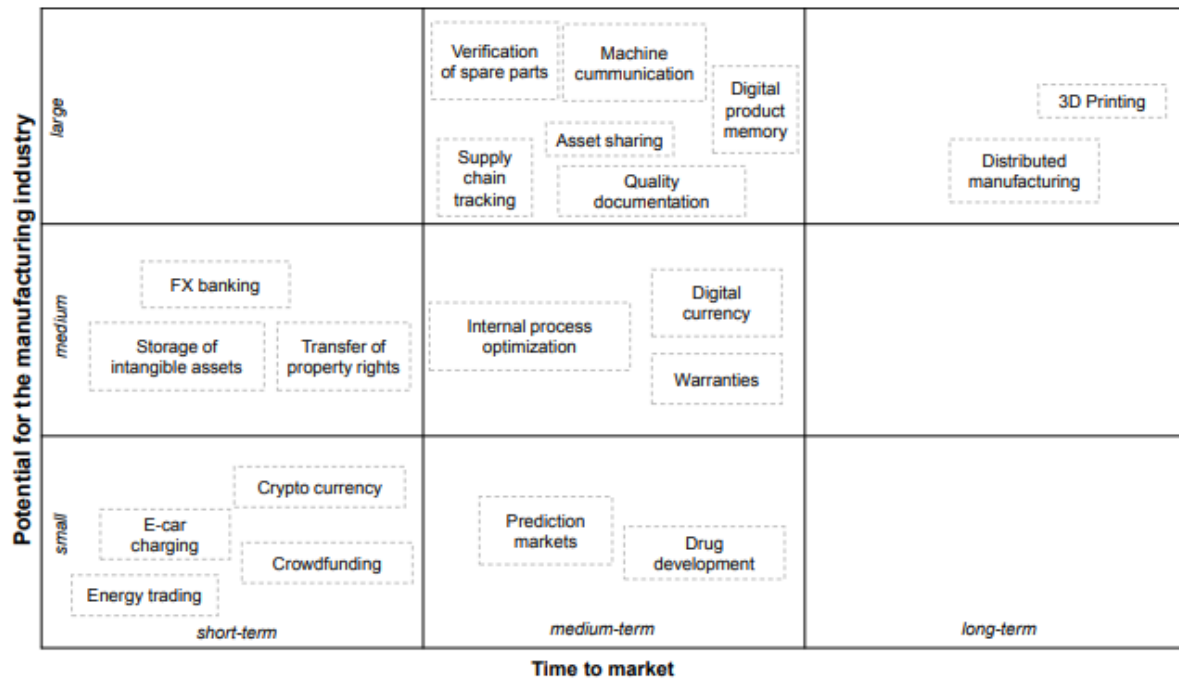
Assessment of blockchain in the Gartner Hype Cycle



only Proof-of-Concept and are yet to enter the market. Use cases are driven by hype without chances of implementation, and its real impact is expected to develop long-term. The potential for the manufacturing industry, therefore, is hard to evaluate. Some concepts have not even been proven to function, some

use cases could have cross-sector potential and be transferred from other application fields.

Distributed manufacturing and 3D printing are believed to have the most impact. But due to technical barriers today, the commercialized use of blockchains in the manufacturing industry will not be reached soon. The greatest challenge, however, still is the process towards a more open culture for the new technology (Sandner et al 2017).



Generally, short-term potential, although with small impact for the manufacturing industry, are energy trading, crowdfunding, cryptocurrencies and E-car charging. Other short-term potentials which could be of more interest in manufacturing are storage of intangible assets, FX banking and property rights storage. Medium-term, drug development and prediction markets play a small role for manufacturers, but warranties, internal process optimization and digital currencies may be use cases which could be more interesting for manufacturers. Applications which are believed to spur the manufacturing industry medium-term are supply chain tracking, asset sharing, quality documentation, machine communication, digital product memory and spare part verification. Long-term, distributed manufacturing and 3D printing will see a blockchain backbone, the researchers believe (Sandner et al 2017).

4.4. Footprint of Blockchains

4.4.1. Energy footprint

The energy footprint of blockchains was first investigated by O'Dwyer and Malone (O'Dwyer & Malone 2014) with a study on Bitcoin, the first known blockchain. This came with the

increasing public interest in the cryptocurrency. The research generally focused on mining, the creational process of Bitcoins. For Bitcoins, valid transactions are collected into blocks and added to a ledger by linking it to previous blocks. The network forms a common view by validating transactions, and the block is finally added to the chain of blocks by finding a signature linking the transactions, which is called the nonce value. The nonce value satisfies an equation of a hash function with SHA256 encryption. There is a good deal of computational power needed to find a solution to the equations. On the other hand, by *mining* a block the Bitcoin network member who eventually completed the task is rewarded with the Bitcoins created from mining, so the process is incentivized. To get a feeling how much computational power is needed to fulfil the criteria for bitcoin mining, let us regard the equation

$$H(B.N) \prec T(1),$$

where B represents the string of recent transactions, N is the nonce value, ‘.’ is the interlinkage operator and H is the Bitcoin hash function, and T is the target value. In the case of Bitcoin, H is represented by the equation

$$H(S) := SHA256(SHA256(S))(2).$$

SHA256 is the hash function. It consists of 64 characters, where each character (in hexadecimal notation) represents 4 bits of information. As a result, the string represents 256 bits of information. For Bitcoin, the information contained in the hash are all open transactions B, with the amount desired to be transferred, and nonce value N. Through computational power, a value for N is randomly or systematically found to satisfy equation 1. When N is found, the block is completed, is sent to the Bitcoin network for verification and upon completion added to the blockchain. Completing the block results in reward for the block creator (in 2018, the reward was 12.5 Bitcoins, with a halving of the reward occurring after creation of 210,000 blocks). Mining refers to finding a value N to satisfy value T for the network.

The rate of Bitcoin creation, by definition, is limited by the network choice of the target value T. This value depends on participants in mining activities as well as their computational power they deploy and is defined by the difficulty D:

$$D = \frac{T_{\max}}{T} (3),$$

where the largest value for T_{\max} is $(2^{16}-1)2^{208} \sim 2^{224}$. The hash function of Bitcoin is chosen in a way that it can attain a value between 0 and $2^{256}-1$. As a result, there is a probability to find a nonce value N satisfying equation 1 of

$$p = \frac{T}{2^{256}} = \frac{T_{\max}}{D2^{256}} \approx \frac{1}{D2^{32}} \quad (4).$$

The probability shows how much effort in terms of calculation is needed to find the nonce value. However, upgrading computational systems can speed up the process. If our system computes hashes at rate R , the expected time to create a block is

$$E[t] = \frac{1}{p} = \frac{D2^{32}}{R} \quad (5).$$

To give an example, if your hardware allows you to calculate 1,000,000 hashes/s, and difficulty D is 4,250,217,920¹, then

$$E[t] \approx 1.8 \times 10^{13} \text{ s} \quad (6).$$

This means that stacking up hardware to allow more computations over time (increasing R) would result in finding the nonce value N faster, which is predominantly done by miners to increase their chances to receive the reward. However, the difficulty D is recalculated every 2016 blocks, with the underlying goal of creating a block every 10 minutes to guarantee gradual issuing of Bitcoins. At what rate the last 2016 blocks were created determines the new difficulty, as it is calculating the estimated hash rate of the Bitcoin network and by that estimating how long it will take to create the next 2016 blocks. For instance, if the resulting estimated difficulty is two times harder than the previous difficulty, to find the nonce value will be two times harder or easier. There is an increasing trend in difficulty in the Bitcoin network due to the growing number of participants in the mining competition.

The limitations of Bitcoin mining, as a result, is the hash rate of the hardware as well as the cost of operation of mining hardware. Hash rate R is measured in millions of hashes per second (Mhash/s). In combination with power usage P (W), the energy efficiency of Bitcoin mining hardware can be determined by

$$\varepsilon = \frac{P}{R} \left(\frac{W}{Ghash} \right).$$

4.4.2. Mining hardware

Mining hardware has seen four generations of hardware, with each generation aiming at increasing \mathcal{E} by increasing R and reducing P . Initially, miners used computers only with Central Processing Units (CPU) for general purposes to perform mining. The SHA256 encryption performs 64 rounds of encryption for each character, involving XOR, multiplex and majority operations. Each round of computation is dependent on the next round, which is why parallelization of computation to determine the nonce value is very difficult and the mining process very wasteful for the performance of multicore CPUs, both in terms of calculating power and energy consumption. (Taylor 2013).


In October 2010, an open-source mining software called OpenCL was released on the internet, and quickly optimized for the purposes of SHA256 computation by miners. The real breakthrough in usage was an application programming interface for OpenCL to use the GPU and control its parameters in response to temperature and user-specified tuning parameters. Through the API, miners were enabled to adjust voltages of the GPU to reduce costs or increase hash rates, and adapt the GPU parameters in a way to maximize throughput. Since Bitcoin mining does not make use of GPU bottlenecks like the memory system or floating-point calculations, a mining system with GPUs can be pushed beyond normal bounds of reliability.

An optimization in the systems for mining occurred only months later, when Field Programmable Gate Arrays (FPGAs) were utilized for nonce value computation. FPGAs are intended to be customized for the specific purpose the user needs it for, which is why there are several techniques how to program the hardware for personal benefit. For the case of Bitcoin mining, loop unrolling, a technique to optimize the execution speed of a program, was applied, at the expense of the binary size of the software. In this circuit, FPGAs are toggled in a way to enable the usage of different circuit applications for each of the 64 rounds of hash computation, each with separate pipeline registers. The search for the nonce value would proceed down the pipeline, computing each round singularly and thereby finding one nonce value per cycle. This allowed very efficient allocation of computation between FPGA registers. Additionally, it propelled innovation from miners, as they first started developing custom motherboards one would find in usual personal computers to minimize unnecessary costs, by removing components like RAM and I/O. Instead, miners focused on sufficient power supply and cooling to run the hardware all-time. Besides not having many advantages

towards GPUs, the reason they were implemented was the reduction of energy consumption to one-fifth, improving cost-benefits for miners.

In early 2013, a further development of hardware, assembled from an initiative from within the Bitcoin community, remarkably funded only through online fora, was introduced: Application Specific Integrated Circuits (ASIC). It contained dedicated circuitry optimized to perform computations very efficiently. The first microprocessors from this generation needed only 1.05 Volt while performing at 4.2W per Ghash/s. To put that in perspective, ASICs are roughly 40 times more efficient in mining than GPUs, while they are 4.4 times cheaper per Ghash/s. It followed a single double pipeline system, which basically means a doubling of computation power from FPGAs. The design of complete mining units with plenty of these circuits focused on a reduction of power consumption, mainly for cooling.

In general, we can assume that advancements for mining hardware will continue as long as the reward is economically beneficial. There is ample room for optimization in Bitcoin mining, like reducing system-level power distribution and cooling overheads. It is believed that existing hardware's efficiency alone can be improved by a factor of 4. As the number of Bitcoin miners increases, so does the difficulty to find nonce values and demand for newer hardware. This hardware will have more power and a lower energy footprint, but both are limiting factors.



Source: Vranken (2014)

4.4.3. Ecological footprint for Bitcoin

With increasing need for energy to compute Bitcoins, an increase of the ecological footprint of Bitcoin is inevitable. To determine the ecological impact of Bitcoin mining on a small scale, we want to find the energy cost for the process. Let U be the unit cost for a Joule of energy, then the energy cost for finding a block is

$$C_e = E[t]PU \approx \frac{D2^{32}PU}{R} = D2^{32}\varepsilon U.$$

For the above described hardware generations for Bitcoin mining, ε differs greatly. The average ε currently (April 2018) ranges between 0.1 W/Gh and 0.3 W/Gh. Through the energy cost of finding a block, we can evaluate the ecological impact of Bitcoin mining. Wood (2014) assesses the impact by comparing mostly-utilized ASIC hardware purchasing statistics and data for carbon emissions footprint of electricity generation. However, he does not cite any sources where the data was retrieved from.

The knowledge of the current network hash rate would allow us to calculate it ourselves. However, there are many sophisticated calculators on the Internet working with this exact formulas and data. Of all of them, the Digiconomist Bitcoin consumption index is the most

cited calculator (Digicomonist 2018). At the time of research (April 2018), Bitcoin's current estimated annual electricity consumption was 61.5 TWh/year, thereby contributing to the global energy consumption at a 0.27% percentage. The estimated annual carbon footprint is 30,200 kt CO₂ for Bitcoin mining. There is no specific source declared which underlying data is used to determine the carbon footprint. One can assume that carbon footprint data for electricity generation for China and other states contributing predominantly to collective Bitcoin mining were taken, since the carbon footprint heavily relies on a country's energy mix. The forecast for Bitcoin's energy consumption is to double until 2019, which would match the total energy consumption of Argentina. (Badkar 2018)

There are many other concerns when it comes to the sustainability of Bitcoin. Through the mechanisms of Bitcoin creation and attachment to the Bitcoin blockchain, Proof-of-Work (PoW), real-time transactions seem not to be within reach, and the growing size of the blockchain, leading to reasonable power consumption of the network for consensus purposes (Vranken). Multiple alternatives have been proposed to address the energy consumption of the proof-of-work mechanism. Additionally, whether we're dealing with a public, private or of a consortium type of blockchain also determines the mechanisms how information included in a blockchain is validated, and thereby influences its ecological footprint.

5. Obstacles for IoT-Optimized Blockchains

The application of blockchain for IoT applications in general is not as straightforward as one would think. There are a few obstacles to consider before thinking of use cases of IoT blockchain, the three main challenges are:

- Generally, high resource requirements for consensus mechanisms in use,
- Scalability issues and reduced pace of the network due to consensus need, the bottlenecks of a blockchain network to avoid malicious behavior, such as double spending (Dorri, Kanhere, Jurdak 2017).

5.1. Scalability and network speed

Scalability is one of the most pressing challenges blockchain technology is facing in the light of complete acceptance in applications. Scherer claims this is mostly due to implications of permissionless blockchains because, in contrast to permissioned blockchain, do not have configuration conditions to allow parallelism or partitioning of the network (Scherer 2017). Scalability is, first and foremost, a debate of variable impact in the blockchain world, which is why variable property reshuffling for better scalability is the most-discussed aspect. Proposals have been made to find linear solutions, such as blocksize increase. Blocksize increase, however, can only create breathing room until better scaling solutions are found. Furthermore, in order to increase blocksize, a hardfork would be necessary (Scherer 2017). Hardforks change the protocols of blockchains, and validators or miners have to be updated, which means they can decide whether they want to remain in the network. Since it can cause stability issues, softforks are generally preferred. Softforks are forward- and backward-compatible with the existing blockchain which can provide the same adjustments to the network. Concerns are that a great majority of the network have to agree on softfork to avoid security issues (ibid). An example how such a variable change to the network could be facilitated is Segregated Witness in Bitcoin. Segregated Witness aims to remove signature-related data from transactions to make them smaller in size. It also addresses other issues like malleability of blocks and second layer solutions. Others have pointed to micropayments and if the transactions stemming from should even needed to be recorded on the blockchain. A suggestion was to create micropayment channels in the Bitcoin network. For permissioned

blockchains, scalability is not as dramatic an issue as for permissionless systems, since scalability issues can always be solved within the system without consent.

Others agree that the ability to scale blockchains is highly dependent on a solution for the two bottlenecks transaction processing and state storage. Croman et al (2016) have conducted a study on reparameterization of variables for blockchain networks, particularly blocksize and block computation interval. They argue that scalability issues cannot be tackled without giving up high security in a blockchain network (Croman et al 2016). The study showed that to determine the throughput limit it was observed that blocksize and block interval have to satisfy

$$\frac{blocksize}{X\%throughput} \prec blockspace .$$

The block interval, in general, should not be smaller than 12s for full utilization of network bandwidth, effective in the year 2016 (Croman et al 2016). It should be noted that these guidelines are intuitively chosen, after research was conducted. However, even these limits can only sparsely improve blockchain scalability.

That is why, when considering blockchain design, fundamental redesigns are needed to solve scaling issues, maybe even radical re-architecture (Croman et al 2016). They suggest considering five different layers for blockchain design in terms of scalability:

- Network Planes
- Consensus Planes
- Storage Planes
- View Planes
- Side Planes

Network Planes propagate transaction messages. In this sense, it is a broadcast layer tasked with all communication between nodes. The Network Plane's scope is a little narrower, because nodes only propagate messages representing valid transactions and only these are accepted as inputs. Most blockchain networks, by measurement, do not fully utilize their underlying network bandwidth, making the Network Plane a bottleneck. A natural way of improvement is thus the avoidance a full reception and validation of a transaction before propagating it to the network. Another way to enhance scaling is to avoid the propagation of

all transaction, only to be transmitted again to be included in block mining (Croman et al 2016).

In the Consensus Plane, academics have pointed to the trade-off between security, speed and bandwidth. The fact every node has to process all transactions and store the entire state of ever account balance poses a problem. As a result, a blockchain cannot process more transactions than a single node, slowing itself by default. A possible solution to this bottleneck is the creation of subsets of nodes validating subsets of transactions, known as *sharding*. The idea is that by dividing the global state of accounts into smaller entities, shards. Shards validate transactions in their shard network, and transaction across two shards can be achieved with a “debit” (Scherer 2017). Others acknowledged the aforementioned GHOST protocol to be a significant improvement for PoW consensus.

The Storage Plane provides several ways to be implemented, as it provides obstacles. Storing ledgers is resource-intensive, Bitcoin for example requires nodes to store the entire Bitcoin ledger for all applications, a memory size of 160 Gigabytes (reference needed). This is obviously inefficient, possible solutions include sharding the storage of UTXO data structure or Distributed Hash Tables (Croman et al 2016).

View Planes consider that validator or miners are not obliged to operate on the full ledger storing the blockchain history. They locally compute on a *view* of ledgers called “unspent transaction inputs”, or UTXO. Views of the blockchain via replication, meaning a local update of the blockchain on the operating nodes could improve scaling. Croman et al also suggest that views can be outsourced to a third-party-computation-provider. It would mean nodes do not have to store entire blockchains to operate full ledgers (Croman et al 2016).

Side Planes contemplate the operation of sidechains with specific off-chain functionalities. They could be used as payment networks where transactions go along the node line of pre-coded “collateral” channels. These channels occupy a designated data set that can be exchanged until the data is finalized (Croman et al 2016).

As we have learned, consensus plays a critical role to assess the efficiency of blockchains in use. Since in this plane the most efficiency gains can be expected, different consensus mechanisms for various use cases, in our case the Industrial Internet of things (IIoT) deserve further illustration.

5.2. Consensus Mechanisms

5.2.1. Proof of Work

Nakamoto suggested to implement a proof-of-work system for timestamp servers based on the ideas of Hashcash. Proof-of-work initially begins with the scanning for the nonce value through SHA-256, which we have discussed in our previous chapter. The hash begins with zero bits. “The average work required is exponential in the number of zero bits and can be verified by executing a single hash” (Nakamoto 2008). For the timestamp network, Nakamoto explains, once the CPU work has been done to satisfy the PoW prerequisites, the block cannot be changed without redoing the work. Later blocks added to the existing chain would also have to be redone by PoW in order to change the block in question. As the chain grows, the probability of an attacker to have the CPU needed to launch an attack on the blockchain diminishes exponentially with the number of added blocks. If there are conflicting blocks in the network, the block which can present the most computing effort represented by the longest chain is added as the next block, while the other blocks are represented in the blockchain in a sidechain. The difficulty of finding the nonce value is evaluated every 2016 blocks (Nakamoto 2008).

If a nonce value is found by a contributing node, it has the right to mine the next block. All pending transactions between the validation of the last block and the block in question are affected, they are validated by the node who found the nonce value. The node then assembles a list of transactions from the period in question, the block. As soon as a block of transactions is generated, the node runs the process of encrypting the transactions into one array of digits and characters, the hash. The hash not only contains the nonce value and the, now approved, transactions, but also information about the hash value from the last block created, and a timestamp of its creation point in time. This important piece of information references the blocks to each other, forming the blockchain (Floyd 2018). To complete the process, the mining node has to conduct the coinbase transaction to receive mining reward. It is the first transaction in the new block. The coinbase transaction cannot be spent before being confirmed around 100 times through the creation of new blocks, equaling roughly 17 hours before it can be spent (Morrow 2014). The concept is also known as Nakamoto Consensus (Tobin 2018).

By the time of development, the PoW mechanisms' intent, first and foremost, was security. The defense of malicious attackers to overpower the network was of the highest priority, predominantly to gain trust in the Bitcoin ecosystem. Back in 2009, PoW was in fact without alternatives as a consensus mechanism (Buterin 2013). However, the consensus mechanism, in terms of computing efficiency and, as a result energy efficiency, is not optimal at least. The computation of the nonce value takes up quite some computational work, and is today done by specialized hardware solely designed for this purpose. Furthermore, more computing power results in higher chances of finding nonce values, which has led to concentration of mining efforts to only a few actors. As one can assume, the benefits of the blockchain technology are not utilized under this setting, as has been pointed out in some studies (Dorri et al. 2017). Another obstacle for PoW are the capacities needed to maintain the distributed ledger locally, since memory is used for storage and continuous update of the ledger.

A way to overcome the hurdles of centralization of mining and to redistribute PoW to more actors is the mechanism initially used by Ethereum, named Ethash. The nonce value, in this case, is called a seed. In contrast to Bitcoin, where the value is used only for validation, the seed allow to compute random cache, from which a dataset of 1 Gigabyte can be generated, with the property each item in this dataset depends on small numbers of items from the cache. The mining process here requires taking random items of the dataset and hash them together. The verification is done with low memory using cache to regenerate specific pieces of the dataset needed. The mining process is highly randomized (ibid).

GHOST protocol

The obstacle of blockchain storage and update for real-time verification is a real bottleneck for the Proof-of-Work consensus. There are several initiatives to overcome the hurdles to propagate block data more quickly among miners, like the FIBRE network (Higgins 2016). It was introduced as a way to sidestep the problem of latency – the time it takes data packages to be transferred to one another. But the main challenge of conflicting blocks in case of network latencies remains, and the capacity for additional transaction processing is very much needed.

Sompolinsky et al. (2013) therefore suggested an alternative to the longest-chain rule, called GHOST. It changes the conflict resolution usually obtained in a PoW consensus blockchain. To understand GHOST, we have to visualize a blockchain as an ever-growing tree, with the longest chain being the trunk of a tree and sidechains from previous conflicting transactions

as branches. To implement new modifications and features onto a blockchain, all peers of the network have to agree on them. If a substantial number of peers, but not all, reach a consensus on the implementation, they can initiate *forks*, a secession of the blockchain at a certain point. At each fork, the GHOST protocol selects the root of the blockchain which represents the most computing effort done by the network and forms a subtree of it.

An orphan, or stale block, is created when two nodes try to validate a block at the same time in case the network suffers from overload due to various reasons already outlined. GHOST includes the stale blocks, called *uncles*, in the calculation of which may be the longest chain computing the highest cumulative difficulty, the longest chain. The rules to form a subtree are as follows: A block representing the root of the blockchain then has to specify its number of uncles in the network. Furthermore, an uncle included in a block must be the direct child of the new block, but cannot be the direct ancestor of the block being formed. Also, an uncle must be different from all other uncles in previous blocks (Madeira 2018). Through this process, the single validation of all blocks, be it stale blocks or others, is guaranteed, while improving the throughput of the network by creating many long chains. This increases the network speed and reduces overload delays, by that also contributing to energy efficiency. The GHOST protocol was utilized in the Ethereum up to the point of the implementation of Casper, which will be discussed in detail in this chapter at a later stage.

With that being said, the depiction of the ledger, with real-time update on the state of the chain for all peers and the need for unanimous consensus of this state represent real obstacles for implementation of blockchain, since it is inefficient in its hardware utilization and network properties. Therefore, alternatives have been researched over the years for various use cases. The most well-known alternative is Proof-of-Stake.

5.2.2. Proof of Stake

Proof-of-work helped to give birth to the blockchain technology itself and one its fields of utilization, cryptocurrency. However, the very nature of PoW means that the cryptocurrency is highly dependent on energy consumption. It introduces significant cost overhead in operation. Furthermore, fear arose that more energy consumption could put pressure on transaction fees to sustain the inherit level of security, since the higher cost is not incentivizing. King and Nadal (2012) developed a concept that would still enable decentralized cryptocurrency without depending on high energy requirements, calling it

Proof-of-Stake (PoS). In general, it describes a proof of ownership of the private key, gaining access to the information on the blockchain.

In the initial model, the block generation was based on the idea of having two different types of blocks, proof-of-work blocks and proof-of-stake blocks. The proof-of-stake in the new type of blocks is a transaction named *coinstake* (similar to Bitcoin's mining transaction *coinbase*). The coinstake transaction is a transaction a block miner pays to himself, while gaining the right for *Kernel input*. It is required to meet the requirements of the hash target protocol under PoS, making the verification similar to PoW of the Bitcoin protocol. The difference is that hashing operations need to be computed with limited search space, precisely one hash per unspent cryptocurrency output per second. It makes the process of hashing more efficient and reduces the energy consumption.

Furthermore, the concept of Coin-Age was introduced for the effective functioning of PoS. Coin-Age describes how long coins of a certain cryptocurrency were held times the period. The hash targets the stake kernel must compute is a target set per unit of Coin-Age consumed in the kernel. As a result, the more Coin-Age is consumed in the kernel (the older the coins were brought back into circulation) the easier it is meeting the hash targets. (King 2012) In other words, a user's stake of ownership of cryptocurrency is utilized to verify a block of the blockchain. Coin-Age is also used for determining the next validator in the network. Older and larger sets have greater probabilities of signing the next block, and by doing that, the Coin-Age returns to zero, a process happening with all coins over time, eventually. This type of consensus is called *chain-based Proof-of-Stake* (Buterin, Griffith 2017).

Instead of users spending big amounts of money to equip themselves with specialized hardware for the mining process, the user can simply buy the cryptocurrency and use it to purchase higher probabilities to create blocks in the blockchain system by validating transactions. In principle, PoS pseudo-randomly selects validators, ensuring no validator can predict its turn. The Coin-Age mechanism allows some predictability.

Besides *chain-based Proof-of-Stake*, another way of consensus-finding under Proof-of-Stake is taking advantage of the *Byzantine Fault Tolerance* (BFT). BFT is based on research from Castro and Liskov (1999), and first executed in the scheme of blockchain technology by Tendermint (2014). The principle behind it is rather simple: Byzantine Fault Tolerance is a characteristic of a system which tolerates failures similar to the Byzantine Generals' Problem, and the most class of failure modes in hierarchical systems. The Byzantine Generals' Problem

refers to a theorem formulated by Leslie Lamport, Robert Shostak and Marschall Pease (1982).

The assumptions were based on how to find consensus if an army under a General wanted to attack an enemy, but Lieutenants were unable to communicate with each other and could only succeed if all Lieutenants would attack simultaneously. Furthermore, the theorem specifically addresses the possibility of traitors who would agree to attack at an agreed time but would lie about it. In this scenario, there is still a need to find consensus even if there were traitors. Under Byzantine Fault, the majority vote is taken. To refer that back into computer science, the algorithm to reach consensus is based on the value of majority of the decisions taken.

For the Tendermint consensus, this specifically means that validators (equivalent to Lieutenants) are users with currency bonded by posting a bond transaction. Similar to the chain-based version of PoS, the stake bonded allows for voting on new blocks to complement the blockchain. There is at least a 2/3 majority vote of validators needed for a block to be added to the chain. In contrast, it means the blockchain is resilient up to 1/3 of Byzantine participants (Kwon 2014).

There are several obstacles arising from this consensus-finding: Baliga puts emphasis on a problem of naïve PoS algorithms, the ‘Nothing-at-Stake’ problem. It refers to implementations not providing incentives for nodes voting on the correct block. As a result, nodes possibly could vote on multiple blocks supporting multiple forks. This behavior increases the voters’ chances of winning rewards as there is nothing to “expend” in doing so. This stands in contrast to PoW, where nodes would need to split up resources in order to vote on multiple forks. To efficiently implement PoS, developers need to address or sanction Nothing-at-Stake. The proposed PoS algorithm for Ethereum, called Casper, is a hybrid protocol between PoW and PoS and described in detail later in this chapter.

Another strong argument against Proof-of-Stake is made by Poelstra (2015). He claims correctly that it is true holders of a cryptocurrency under PoS are incentivized to agree on each extension because they:

- Are randomly chosen and therefore unlikely to be in collusion,
- Even if stakeholders would collude, they do not want to undermine the system because they would lose their own value,
- They have limited capacity to cause damage to the system anyway, since the next selection of stakeholders will only choose one reasonable block history for extension.

To define a cost function that measures stake in a cryptocurrency, we need an already existing consensus history to identify malicious behavior. Since there is an incentive to verify many blocks with one's stake (Nothing-at-Stake), the verifying parties may “grind” through many potential blocks, only publishing one that results in themselves being the signer for the next block, and taking control of the blockchain in the process. This behavior is known as the *Stake-grinding attack* (Poelstra 2015).

Another problem Poelstra identified is that of *Costless Simulation*. It refers to the issue that coins bonded against stake only exist within the blockchain to which the stake belongs at the time of the signing. This allows the creation of “cheap histories”, the continuation of the chain in a way that favors the signer because the signer could ultimately choose between alternating histories. Even if stakeholders are bonding large market values of crypto, dishonest or deterrent behavior to the blockchain cannot be determined until coins are signed. Along the lines of verification, there is ample room for malicious behavior.

Casper protocol

The Casper protocol is an overlay atop a proposal mechanism for blocks which want to be added to the blockchain and was proposed by Buterin and Griffith (2017), the engineers behind Ethereum. As there was growing concern about the efficiency of the currency aiming to be the backbone for smart contracting, the Foundation put emphasis on refining the Proof-of-work scheme the Ethereum blockchain was created on. Until this date (April 2018), the protocol has not been implemented. However, in terms of future energy consumption of the blockchain technology it is worth examining the protocol and which weaknesses it is addressing.

According to Buterin and Griffith, the Casper protocol provides accountability, dynamic validation, defenses and modular overlay. The basic idea behind Casper is that stakeholders can stake a portion of their currency holdings for validation. Following this, blocks will be mined as known from the PoW mechanism. The validators place a bet with their stake on the validated block, hoping it will be added to the chain. If the block gets appended, the validators will receive a reward proportionate to the bet. On the other hand, if the validator acts in a malicious manner, Casper can identify the behavior and will penalize validators by retaining the validators' entire stake (Rosic 2017). Because the size of penalty exceeds the reward from validation, this type of PoS provides more security incentives than PoW.

This is possible due to the introduction of checkpoints and heights, where each block is required to deliver information on. If a validator, who is impeccably identifiable, provides ambiguous checkpoint and height information on a block, the protocol can identify it as an effort of Nothing-At-Stake. The two fundamental properties of Casper, in this regard, are *accountable safety* and *plausible liveness*. *Accountable safety* refers to conflicting checkpoints for the blockchain. There is only one checkpoint which can be validated as the extension of the chain unless 1/3 of the validators violate a slashing condition, which would result in the loss of one third of the total deposit. *Plausible liveness* describes that regardless of previous events, if 2/3 of all validators validate an extension of the chain, it is possible to finalize a new checkpoint without miners violating slashing conditions. The evidence of any attempted violation is stored onto the blockchain as a transaction, at which point the malicious validator loses his deposit and a finder's fee is granted to the discoverer of evidence (Buterin, Griffith 2017).

5.2.3. Variants of Byzantine Fault Tolerance

Hyperledger Fabric - PFBT and SIEVE

Hyperledger Fabric is another major infrastructure framework for blockchain development and the most popular permissioned blockchain (Baliga 2017). The idea behind is that in a blockchain ecosystem every peer is required to execute every transaction, maintain a ledger and run consensus. This means it cannot support private transactions and confidential contracts. In Fabric, peers have three distinct roles: The *endorser*, *committer*, and *consenter*. Assuming a shared blockchain between participants, the idea is to allow confidential dealings between two participants of the blockchain, so the transaction does not appear in the distributed ledger. The transaction will generate a result, and the ledger needs the rendered equal result from both parties of the transaction. For multiparty transactions, a similar procedure applies (Hyperledger, n.d.).

Fabric currently supports two consensus mechanisms: Practical Byzantine Fault Tolerance (PBFT) and a variation named SIEVE. Byzantine Fault Tolerance is a form of Proof-of-Stake, which has already been examined in the section for Proof-of-Stake mechanisms. SIEVE initially executes all operations speculatively, and in another step, compares the outputs across the processes. The disadvantage is that this protocol only executes one operation at a time, however the throughput can be varyingly increased by the standard method of batching

operations together (Cachin, Schubert, Vukolic 2016). SIEVE can handle non-deterministic executions, meaning that given a particular input, can handle diverging outputs of code. If the protocol detects a minor divergence, the diverging values are sieved out (Baliga 2017). PBFT and SIEVE have not been investigated towards their energy consumption, but it is evident that a randomized miner choice will require more computational power than a mechanism of highest efficiency.

Cross-Fault Tolerance

Cross-Fault Tolerance (XFT) refers to a protocol that simplifies some of the assumptions Practical Byzantine Fault Tolerance is based on. Since Byzantine Faults are designed in a way to resist attacks by powerful adversaries, the complexity of BFT protocols is high and therefore they are less efficient. XFT is built on the notion that a single adversary cannot control the majority of nodes and generate network partitions at the same time. It is designed to provide correct service as long as the majority of blockchain replicas on the participating nodes are correct and can communicate with each other synchronously (Baliga 2017). This consensus mechanism is seemingly energy-efficient in execution, however lacks substantial defense power in case of a malicious attack.

Federated Byzantine Agreement

The Federated Byzantine Agreement (FBA) is meant to work for permissionless blockchains, with an open end to node participation. They specifically eye financial use cases. However, in terms of energy efficiency it is still worth to look at its function. The most popular blockchain projects working with Federal Byzantine Agreement consensus are Ripple and Stellar, both have slightly diverging consensus-finding protocols. But what unites them is the goal of fast cross-border transactions and the elimination of gatekeepers for payments. Participants in such systems are end users, financial institutions acting as gateways and market makers either being users or financial institutions. Gateways are similar to banks, holding funds in fiat money, but furthermore issue an equivalent in the Ripple/Stellar networks onto the global blockchain. Transactions can be verified by nodes by referring to balance on the global blockchain. Market makers provide the required liquidity in the networks, maintaining multiple gateways and multiple currency holdings (Baliga 2017).

The Stellar Consensus Protocol (SCP)

SCP consists of two sub-protocols: A nomination protocol and a ballot protocol. The nomination protocol intends to create candidate values for a slot. Slots are the consecutively named positions in a sequentially applied log that covers the chronology of transactions. Over time, it creates sets of candidate values at every participating node in the blockchain. By that, nodes can combine the candidate values to produce the same output, a compromised value for a slot (Mazières 2016). Nodes cooperate in the form of quorums and quorum slices. Quora are a set of nodes sufficient for reaching agreement. A quorum slice is a subset of a quorum with the ability to convince one node for an overarching agreement. Quora and quorum slices represent real-life business relationships, which means a node can take part in multiple quorum slices and quora. In order to reach global consensus in the blockchain, quora have to overlap (Baliga 2017). When a composite value is agreed upon by the nomination protocol, the ballot protocol is executed. The ballot protocol enables federated voting to commit or abort composite values for the block. When there is consensus on a composite value, the value is inherited for the slot in question. When there is agreement for abortion of a composite value, the value is discarded (Mazières 2016).

Ripple Consensus Protocol

The Ripple Consensus Protocol is defined by the requirement of nodes to provide a Unique Node List (UNL). In principle, the UNL consists of other nodes that are trusted by the given node, and consensus achieved by nodes consulting other nodes in its UNL. As a prerequisite, each UNL needs to have a 40% overlap with other nodes in the Ripple network (Baliga 2017). Consensus is proceeded in rounds. In each round, each node takes all valid transactions it has seen and makes them public in a list called the “candidate set”. The nodes then merge these sets onto their UNL and place a vote on the accuracy of these transactions. If transactions pass a threshold of ‘yes’ votes, they proceed to the next round of voting. In the final round, a minimum threshold of 80% is needed of a UNL agreeing on a transaction. All transactions meeting the requirements are added to the ledger (Schwartz, Youngs, Britto 2014).

Delegated Proof-of-Stake

Delegated PoS is a special form of consensus under the Proof-of-Stake regime. It is divided into two parts: Election of block producers and scheduling production. By electing the block producers, the stakeholders make sure they stay in control of the events on the blockchain because they have the most to lose in case of malfunction. For Delegated PoS, the block producers are called witnesses (dantheman, 2017). Witnesses serve the role of validating

signatures and timestamping transactions, thereby generating the block. Stakeholders can elect any number of witnesses. Each node is allowed one vote per share per witness, which is known as approval voting. N witnesses electable by total approval are selected. It is defined in a way that at least half of all nodes can have confidence there is sufficient decentralization. When stakeholders express their number of witnesses, they have to vote for at least that number of witnesses. The block production has a financial incentive for the witness, the pay rate is set by elected delegates. The witnesses are shuffled continuously to ensure parity within the network (ibid).

Delegates are elected similarly to witnesses. The delegate becomes a second signer on an account which can propose changes to the network parameters, known as the genesis account. The parameters include block size, transaction fees, witness pay etc. After a majority of delegates have approved changes to the network, the stakeholders are granted a review period, during which they may vote out delegates and select new delegates, thereby nullifying the network changes. This mechanism was chosen to ensure delegates cannot collude and take over the network, as well as protecting them against regulations which may apply to the technology in the future (dantheman 2017). The procedure allows network fragmentation, which is inefficient in terms of maintenance of the blockchain. The energy need for this consensus form is high.

5.2.4. Proof of Elapsed Time

Proof-of-Elapsed-Time (PoET) describes a consensus mechanism currently utilized under the blockchain project Hyperledger, initiated by the Linux Foundation in 2015. The project's objective is the advancement of cross-industry cooperation with distributed ledgers. The initial focus of the project is on performance and reliability of these systems to transform technology companies, supply chain applications and the financial sector (Linux Foundation, 2015) with support of global players active in these fields. It also intended to integrate open protocols and standards, including blockchains with different consensus mechanisms. To achieve that, Hyperledger provides digital blockchain infrastructure, similar to a compiler.

Since the blockchain sector is relatively new and the opportunities only researched marginally, Hyperledger provides frameworks with distinguishing properties, with one of them being Hyperledger SawtoothLake, a framework developed by Intel. Interestingly, Intel utilizes a consensus mechanism which is fundamentally different from others, called PoET. It is intended to run in Trusted Execution Environments (TEE), such as Intel's Software Guard

Extensions (SGX). Trusted Execution Environments (TEE) are secure areas of a main processor. It guarantees code and data inside to be protected with respect to confidentiality and integrity. TEE offers execution space where higher levels of security are needed (Galindo 2017). Intel SGX is described by Intel as “a set of new CPU instructions that can be used by applications to set aside private regions of code and data” (Hoekstra 2015). Proof-of-Elapsed-Time exclusively relies on Intel’s SGX.

In PoET, a random *leader* is chosen through a lottery model based on SGX to finalize a block. It uses the model to counter untrusted nodes and open-end participation of nodes in the consensus procedure. To work correctly, PoET has to distribute the *leader* selection randomly across all participating nodes and security that the given leader is not acting maliciously. Here, the Trusted Execution Environment comes into play (Baliga 2017). Specialized hardware components create *attestations* that code was correctly run in a protected environment, allowing a network participant to prove other participants that it is running trusted code for the network. Additionally, trusted code runs in an environment that is private to the rest of the application, which ensures a malicious participant cannot cheat by manipulating the PoET (Riley 2018). Each validator requests waiting time from the code running in the TEE, the shortest waiting time wins the lottery. The validator has to prove it has the shortest waiting time and waited for a designated period of time before it can mine the next block. Through the randomization, the leader role is equally distributed among all validating nodes (Baliga 2017).

PoET as a blockchain consensus mechanism holds many opportunities since the CPUs needed for operation can be found in consumer electronics already, which means the infrastructure for a functioning network is already in place. Additionally, it appears this mechanism is energy-efficient, at least compared to proof-of-work. However, so far, no studies have been conducted as to how much energy is required to run the process of PoET (JP Buntinx 2017).

5.2.5. Proof-of-Authority

We have seen many forms of PoS consensus mechanisms which all have one assumption in common: Those holding a stake in a network are incentivized to act in its interests, and the more stake a participant holds, the higher this interest seemingly is. But it is important to note that same sizes of stake are valued differently by its holders. Reasons can be different financial power, specific interest in the blockchain field, among other reasons. This obstacle

led to the creation of a staking concept where circumstance is replaced with an explicable invaluable resource: reputation.

This is taken account of in the Proof-of-Authority (PoA) consensus model. It is a modified form of PoS, where a validator's identity, including officially issued documentation for an individual, is utilized to perform consensus (POA Network 2017). For the concept to function, there is a need for true identities, eligibility for staking has to be difficult to obtain and the procedure to establish an authority needs to be the same for all validators. The identity validation is done by notaries, who already have identity information freely accessible on the public blockchain and put the verification through via formal documentation (photos, scans, etc.). Even in the event of identity stealing, the public staking allows no single actor to overwhelm the network and act maliciously within the system. To make the eligibility for staking identity hard to obtain, validator candidates have to pass notary exams. Fulfilling the notary requirements and going through the documentation process makes the procedure of gaining authority explicit and unified and independent of the network itself. It establishes integrity, transparency and, ultimately, trust that every participant has the same means to earn status in the network (POA Network 2017).

5.2.6. Proof of Luck

Proof-of-Luck is a mechanism proposed by He, Kanwal, Milutinovic and Wu (2017) and based upon Trusted Execution Environment assumptions. At the beginning of each round, the miner calls for mining the next block. The request is queued, and the participant may wait for a new block, but simultaneously they are allowed to switch to a luckier, alternative block while waiting. If a participant receives a luckier block before their own mining completes, they will not need to broadcast their own block. PoL represents a diverging form of Proof-of-Elapsed-Time, combined with a randomization that is executed in the protocol (He, et al. 2017). In terms of energy efficiency, the Proof-of-Luck does not bring upon a substantial advantage in comparison with other consensus mechanisms.

5.2.7. Proof-of-Burn

For the Proof-of-Burn mechanism (PoB), the idea is that blocks can be burned to reduce the need for great computational resources. The role of burning "coins" is to affirm proof when

the PoB mechanism is used for mining. Coins are therefore sent to a burn address, which is predetermined and has no ownership. After burning, the address can be used to generate blocks again. They generate PoB hashes, and blocks can be generated parallel to the normal PoW mining mechanism. A user can choose whether to mine blocks in the traditional way by investing computational power, or by investing into generation of burn hashes, since the burn hashes needs no further computation to verify the block generation. The team behind Slimcoin, which utilizes PoB supplementary to the Proof-of-Work mechanism, referred to burned coins as mining rigs, since the burn replaces the hardware needed for verifying transactions. However, the duality of the system implies that high amounts of energy are still needed, at least for the PoW consensus (P4Titan, 2014).

5.2.8. Proof-of-Space

Proof-of-Space is based on the same idea for consensus as is the Proof-of-Work scheme. However, PoSpace requires to dedicate memory space rather than computational power for mining. Spacemint, a blockchain project exclusively proposing this consensus mechanism, claim it rewards smaller miners fairly, according to their contribution to the network. PoSpace is said to be ecological, since the dedicated space for mining the node offers access to requires minimal computation. Furthermore, it is economic, since nodes almost always have unused disk space available. Finally, the mechanism is not prone to dedicated hardware, which could cause unfair discrepancies for network participants (Park, Kwon, Fuchsbauer, Gazi, Alwen, Pietrzak 2018). However, this mechanism does not appear suitable for industrial utilization, since industrial computers are designed differently than personal computers. Consequently, memory is efficiently distributed and the implementation of such a system in an industrial network doubtful. For personal use, the obstacle of providing access to personal hardware may be too hard to overcome.

5.2.9. Network Consensus Mechanism

Obelisk is the name of a unique consensus algorithm, the Network Consensus, established by the Skycoin project. The uniqueness lies with the mining process, which is eliminated in contrast to other blockchain applications. The underlying idea is a web-of-trust consensus, cycling the mining incentives and exponentially improving transaction speeds. The network consists of nodes, each node subscribes to a list of trusted nodes. The more subscribers a node can

persuade to follow him, the more influence it holds in the network. The nodes binding the network have a personal blockchain assigned to their point of access, similar to a “public broadcasting channel”, where actions are recorded. Consensus decisions and communication occur through the personal blockchain. The public record left by each node allows the network to react to defection by cancelling subscriptions of untrustworthy or malicious nodes. Furthermore, the community can shift the balance of power in the network if it is too concentrated by changing trust relationships (Skycoin 2016).

The thought process behind the mechanism was well laid out by an anonymous user in the Skycoin network. Participants in the network pursue parity for decisions within the network to move the blockchain forward, and not only one decision is correct. That is the reason many, potentially all, nodes can perform block-making. Due to asynchronous nature of block-making, the nodes produce independent, observation-based and cryptographically-signed *opinions*, since the blocks produced from all nodes will have asynchronous information to process. Through this mechanism, a sample of opinions how the blockchain should move forward is generated, and opinions are only approximations, lowering the likelihood of attacks.

The Network Consensus Mechanism also avoids the following: Nodes are averse to having an opinion for a decision (on which is the next block in the chain) differing from that of their nearest neighbors in the network, or a small set of local neighbors. This behavior is susceptible towards manipulation and malicious acts, which is why it can easily be utilized in an arbitrary way. Without a leader election in the network, there are restrictions to the mechanism to find consensus: The node needs to form its own independent *opinion* based on the statistical analysis of opinions it received, needs to perform verification and fraud detection, has to be sovereign and independent in terms of seeking payments in return for supporting a given opinion in the network and needs to be able to receive raw data and process it to form new opinions, like block hashes (Mill/Anonymous 2016). In the context of energy efficiency, this mechanism has advantages towards other consensus mechanisms, since there is no energy wasted in generating the blockchain.

5.2.10. Hashgraph

Hashgraph follows a new distributed consensus algorithm with the introduction of a new data structure. Any member can create a transaction, which will eventually be put in to a block and

spread throughout the chain. In usual blockchains, the community decides on the continuation of a single long chain. If two nodes create blocks simultaneously, one needs to be discarded. In the Hashgraph mechanism, this is not the case. There is also no harm in the growing structure. This is achieved by directed acyclic graphs (Baird 2016). This concept is explained in detail in Section 6 about the Tangle.

The core concepts of Hashgraph are:

- Transactions can be made by any node, all members receive a copy of it and reach Byzantine Agreement on the order.
- A small group of attackers will have difficulties to unfairly influence network, it can be determined as fair.
- The distribution of information within the network is secured by the gossip protocol, the repeated choice of random nodes and giving them all the information they have.
- It utilizes directed acyclic graphs recording the gossip and its order.
- The hashgraph is spread through the network with the gossip protocol. Information being gossiped is the history of the gossip itself.
- Since every node carries a copy of the hashgraph, nodes can calculate how the sending node might have voted had they executed a traditional Byzantine agreement that involved sending votes. This is called virtual voting.
- Some transactions on the blockchain are chosen to be witnesses. A witness is defined famous if the hashgraph shows received it fast after creation. A Byzantine Agreement algorithm is then run to determine the set of famous witnesses (ibid).

The gossip protocol has the purpose to ship information through the network as fast as possible, since all nodes have to know every event occurring in the network fast. To find agreement on the linear order of events in an asynchronous environment, a Byzantine Agreement protocol needs to be in place. Byzantine fault tolerance protocols without leaders depend on members sending each other votes. To save resources, the hashgraph consensus does not require any votes to be sent through the network. All participants receive an identical list of transactions that includes all submitted transactions. This list is called the “total order” (Graczyk 2018). Since this list will be identical for many participants at the same time, no consensus on the sequence of the hashgraph is needed.

However, if the graphs are not identical or the network needs to agree on an event that may be placed in an earlier location, virtual voting comes into play. Virtual voting utilizes the fact

that the hashgraph of participants may be slightly different, but consistent, meaning they are identical to some point back in history. Through the gossip protocol, any information a participant is not aware of will be sent to him. One node calculates a total order of events by calculating a series of elections, election rounds. A given event can participate in every round, but might not participate in every election round. Votes on any event even can be changed, depending on the election round. Now, because the hashgraph of the participants is consistent, the calculation can assume the vote of other nodes in the network, which is exactly what the protocol does. The calculation is performed locally and votes of other nodes are hypothesized (Baird 2016).

The procedure always follows the same rule: A node randomly picks another node, and gossips about the known events on that node. The receiving node then creates a new event to record that gossip. The first event created between the two nodes is called witness. Only witnesses can send and receive virtual votes. The virtual votes of the nodes that were locally computed are now shown publicly and it is decided whether or not the witness is famous. A witness is famous if many witnesses in the election rounds can see it, and not famous if they are not seen. The protocol runs an election for each witness, and once it is determined famous, the election is over and the hashgraph updated to its newest version. Once the consensus is reached if each witness is famous, it is easy to find a consensus timestamp and total order on older events (Hashgraph, 2018).

Efficiency-wise, the hashgraph protocol eliminates many obstacles. The communication for consensus in the network is very energy-efficient, since information exchange is kept at a minimum, thereby needing no computing power. The gossip protocol allows to send a lot of information quickly through the network at low computing power input, making it even more suitable for energy-efficient applications. However, consensus requires every node to receive an updated hashgraph. With an increasing number of nodes, this could lead to inefficiencies (Graczyk 2018).

To conclude, developers and other enablers of the blockchain industry have addressed multiple obstacles of the blockchain technology, stemming predominantly from the reflections of specific use cases. The use cases are widespread and need detailed consideration for the purpose of improving processes, such as the application in the Industrial Internet of Things. However, the limitations of IIoT together with scalability and consensus issues of blockchain require a combination of the schemes discussed in this Thesis. A project recently gaining

great attention is IOTA, which has specifically targeted to solve the obstacles of blockchain usage as the backbone of the Internet of Things, and with it IIoT. It is discussed here in further detail.

6. IoT-optimized blockchain - The Tangle

A blockchain project named IOTA presented a completely different angle of access to consensus mechanisms in general. In fact, IOTA even claims that their protocol is the next evolutionary step in blockchain development. The tangle is worth investigating in this context because it claims to be the cryptocurrency of the Internet of Things Industry and is specifically designed for this use case. The idea behind it is a rather complex one, with directed acyclic graphs (DAG) used for storage of transactions. Popov notes correctly that the Proof-of-Work mechanism creates two distinct types of participants in the system: Those who issue transactions, and those who approve them. This creates unavoidable discrimination of some participants and justifies that new assumptions needed to be made to utilize blockchain (Popov 2018).

After all, blockchains are cryptic, verifiable list of documentation of things that have happened in the past. A list is a data structure, and whenever there is an entry written in, it references the previous block and one can verify it all the way back to the beginning, which establishes the sequence, or blockchain (Palmer 2016). The fact the entries have to be written in in sequence is a bottleneck for a number of blockchain use cases, because they require computing resources to be available. Specifically, it limits mobile and IoT devices from being integrated in distributed networks (Milutinovic et al 2016). It causes many scaling issues, and there is a lot of debate within the community which variable in the computing process should be changed to tackle them. Directed acyclic graphs, on the other hand, are unidirectional, but not need to be sequential. IOTA is the most renowned project utilizing this data structure.

A directed acyclic graph can be envisioned as a finite arrow in one direction, but consisting of dots. Usually, the flow of information would follow a direct pattern, from first to last dot. In a directed acyclic graph, the flow of information still has one direction, but information can also leap over dots and hand additional information over to dots which are placed closer to the end of the arrow. In terminology of computer scientists, such a graph has no directed cycles, but a number of vertices connected in a closed chain. In other words, the dots do not follow the pattern of equaling numbers of input and output, but interwoven directions for information flow. There is no consistently-directed sequence such that one can start at a certain vertex and

will be looped back again. However, it is important to note a DAG has topological ordering, a linear ordering of vertices and edges constraining if tasks need to be performed after one another (Thulasiraman, Swamy 1992).

Consensus

IOTA completely removes the mining, and ultimately, the consensus process through a unique mechanism called *tangle graph*. Transactions issued by nodes constitute the site set of this tangle graphs, the set is created by the sequence of approving two transactions randomly in the network when the node wants to issue a transaction itself. The approvals can be thought of as directed edges and are named *tips*. If there is no directed edge between two transactions, but a directed path of length with modulus two between two transactions, we say one transaction indirectly approves the other. The genesis transaction is approved either directly or indirectly by all other transactions. The genesis is an address in the beginning of the tangle with a balance containing all IOTA tokens, and the spread came by sending tokens to several “founder” addresses. The transactions represented on the tangle graph are called *sites*, the network consists of nodes who issue and validate transactions. In short: To issue a transaction, nodes must work to approve other transactions, with no approval if the transactions are conflicting with the tangle history (Popov 2018).

By implication, this also means transactions can be verified several times during their existence. The network makes use of this by attributing a higher level of *confidence* to the transaction, making double-spending attacks almost impossible. For the choice of transactions there are no rules imposed for the nodes. The developers argued in the context of IOTA nodes will follow a reference rule that will be agreed upon by the nodes themselves, since nodes are specialized chips with pre-installed firmware utilized for IoT applications. To issue a transaction, the node chooses two other transactions for approval, the *tips*. The node then solves a puzzle similar to the Proof-of-Work mechanism, with a nonce value found from the concatenated data of these two transactions. The nonce value has a particular form, stemming from the approved transactions. This method for choosing two tips is called the *tip selection algorithm*. The tip selection is done by a weighted random walk from the genesis towards the tips, and ends when a tip is reached. The walk is performed two times, and two tips are chosen. The walk tends to go towards transactions with more cumulative weight, which leads to a higher probability to approve new transactions than old transactions (IOTA Foundation 2018). The walk computation is widely known as Monte Carlo method.

There is no need to achieve consensus *per se* on which valid transactions end up being added to the ledger, meaning it is possible they all can be found in the tangle. In case of conflicting tips, the nodes need to decide on which tip will become orphaned. The tip selection algorithm will be run by the node several times, and will determine which of the tips will be indirectly approved by the issued transaction. Here, confidence again plays a role: If a tip was selected 97 times out of 100 runs of the tip selection algorithm, it was confirmed with 97% confidence (Popov 2018) and likely will be chosen over a transaction with less confidence.

To secure the participation of all nodes in the network, every node has to calculate some statistics, one of which tells how many transactions are received by a neighboring node. If they do not receive any transactions or performs unexpectedly small amounts of transactions, they will be dropped by its neighbors. So even if a node does not issue transactions actively and has no direct incentive to share new transactions approving its own transaction, they have an incentive to participate in the network (ibid).

Scalability

With all that being said, one point seems contradictory to the intended solution of infinite scalability and efficiency the tangle graph is presenting: The tip selection algorithm, as we said, runs its walk from the genesis towards all tips in the network to eventually reach a tip suitable for approval. How is that efficient? That is the reason IOTA, on top of the advanced algorithms that are executed in the ecosystem already, appended additional features to their system. The key for efficiency of directed acyclic graphs is that they are, in contrast to ‘normal’ blockchains, finite. The algorithm achieves this through *snapshots*. Snapshots are a separate technical feature of IOTA essentially capturing all balances at a specific time, pruning the history of transactions leading up to that moment. Snapshotting removes the history of the tangle and information of all data to start fresh. The snapshot generates a new address, which acts like a new genesis address. The history of the previous tangle is discarded (steemhoops99, 2017).

Another obstacle resulting from scalability issues, as we have discussed in detail, is pace of the network, in particular the number of transactions per second possible. The developers of the IOTA ecosystem expect two different regimes of traffic on the tangle graph: Low load and high load. For low load, the typical number of tips is small, and on occasion even becomes 1. The flow of transactions is small, which lowers the latency of the network. Under high load, the number of tips is large, transactions may be queued, which also leads to a low latency of

the network (Popov 2018). In principle, the low load regime network pace is determined by the tip getting approved for the first time in $\Theta(\lambda^{-1})$ time units, where λ is the rate of the incoming flow of transactions. In this regime, the network needs external stabilization in order to function. For the high load regime, the typical timespan for a tip to be approved is $\Theta(h)$, where h is the computation/propagation time for a node. A way to pace the network in a high load regime is the issuance of empty transactions, since they constitute a new tip which can approve transactions. As one can see, the network benefits from higher traffic, and needs to be stabilized during low load.

During low load, the pace of the network is secured by the constitution of a *coordinator* and *milestones*. The reason why it is needed is clear: With a low load of transactions in the network, malicious attackers can create many nodes and malicious transactions can be infiltrated into the system and approve the honesty of one another. In fact, if more than one third of all transactions is malicious (since always two transactions validate a third one), the ecosystem can be hijacked by malicious attackers. Here, the coordinator comes into play: The coordinator is a special node run by the IOTA Foundation and this node is used to directly or indirectly validate transactions. The coordinator checkpoints validate transactions, which are in turn validated by the network. The coordinator issues periodic milestones every minute, which reference to valid transactions (Schiener 2018). The reason for the milestones to exist is that the node could pick any transaction, with the possibility of being malicious and hoping for an honest node to verify the tip. The milestone is signed by IOTA, indicating its honesty. However, the coordinator will be shut down as soon as the network reaches a certain number of transactions per second, since the high load regime is capable of carrying the network (Schiener 2017).

Critics have pointed at some design decisions made by IOTA. For one, IOTA uses a post-quantum-cryptography algorithm, which inflates transaction sizes at a factor of 10. Furthermore, IOTA uses a ternary instead of a binary system. Since all established computer hardware uses a binary system, this decision appears questionable. Also, the small Proof-of-Work needed for consensus has alternatives, as we have discussed exhaustively (Bergmann 2017). However, it is nevertheless of interest how the IOTA system performs in an IoT environment. The empirical part of the Thesis will predominantly focus on this epistemological interest.

7. Epistemological interest

We have learned from the literature review that ICT is not only responsible for energy consumption, but also contributes significantly to the reduction of resource consumption to mitigate impacts for climate change. This is due to optimized large infrastructure systems with ICT systems, potential for large-scale simulation and control, which is shown by the energy productivity index, where a decoupling of energy demand and economic growth can be witnessed since the 70s. ICT can be seen as a tool for indirect energy use optimization and conservation. Contributions to reduction of resource consumption include measures for ICT hardware as well as software and their interaction. The sustainability factor of architectural software concepts like blockchain can be investigated by analyzing energy behavior of software, capacity and resource utilization, with a specific model for assessment provided by Bruntink et al. (2014). These indicators will be adduced to assess the sustainability of blockchain technology architecture.

A special field where ICT could be utilized and could contribute greatly to energy consumption reduction is industrial production. The biggest motion towards more efficiency is Industry 4.0, with its vision of global interconnectivity, full automation, and connected cyber-physical systems sharing information. It will bring about improvements in industrial processes through engineering material usage, supply chains and product lifecycle management. The field of application this Master thesis is concentrating on is the Internet of Things, and more precisely the Industrial Internet of Things, where products are enriched with embedded computing to allow field devices to interact and communicate. In order to establish an environment for the Internet of Things, principles needed to be taken into account are interoperability between devices, sensors and people, must possess the ability to create virtual copies of the real world and perform its tasks autonomously. But there are also challenges such as security issues through opening systems, the high reliability for stability, protection of industrial know-how, lack of adequate skill sets, unclear jurisdictional circumstances, unclear economic benefits and more. Furthermore, the Internet of Things has an implicit scalability issue in every application due to limited resource attachment possibilities, has to operate under the demand of suiting local environments, need more software complexity to manage smart objects, need to have built-in robustness against small and infrequent communication, customized security and privacy networks and have sufficient power supply to provide data. Regarding the Industrial Internet of Things, a good for measure is the "power of 1%" rule,

expressing that 1% in savings in operational costs or reduction of inefficiencies has significant impacts and is an indicator for investment.

The distributed ledger systems of the blockchain technology allows peers to interact with trustworthiness without intermediaries, known as disintermediation. There is a controversial debate whether application of distributed ledgers may be justified by its benefits. The Internet of Things is a possible area of utilization, because it can expand the possibilities how we engage machines, since they could be able to make full transactions with each other and not limit itself to typical use cases such as predictive maintenance and data analytics. The use cases for the Industrial Internet of Things rely on the architectural principles and degree of interconnectivity of the devices engaging in an IoT. However, a great obstacle to overcome is the energy footprint of the blockchain system in its initial form. Due to its high resource requirements stemming from the need for network consensus in a distributed ledger system and scalability issues, resulting in a reduced pace of the network due to the bottlenecks of the technology, which are foremost transaction processing and state storage. There are a number of involved actors in the field who believe that blockchains need to go undergo major redesign efforts for its respective use case. For the Internet of Things, where smart objects of small size are required to communicate with each other, this becomes especially challenging. A project named IOTA developed a blockchain they call Tangle which is believed to suit the requirements of IoT blockchains. It removes important bottlenecks such as sequential entries with the directed acyclic graph, it completely removes mining, finds consensus through transactions issued by nodes which directly constitute the site set and creates transaction confidence, a level of trust for an honest transaction. The algorithm on which the Tangle operates is called tip selection algorithm, choosing at least two transactions and computing a nonce value from the concatenated data of these transactions. All transactions can be found in the Tangle, but their rate of confidence determines how much they contribute to the stability of the tip selection algorithm. If a node does not perform any transaction validations, it will be dropped by its neighbors, creating an incentive to actively participate in the network. The Tangle also features snapshotting, which captures all balances on all nodes in the network at a specific time with the aim of removing the history of the network and create a new blockchain from the snapshot. Since two transactions are needed for approval, the higher load can be witnessed in the network the faster the network becomes. However, the design architecture is still believed to have some flaws to use IOTA as the blockchain for the Internet of Things.

With that being said, this Master Thesis aims to show based on the assessment model by Brutink et al. (2014) how energy-efficient the software architecture of IOTA operates on

various devices within an ecosystem of the Internet of Things, how it interacts with hardware and what obstacles are identified in order to utilize IOTA for smart objects. In this regard, the Master Thesis focuses on answering the following research question:

8. Empirical Part - Research Question

What are the energy impacts of the IoT-optimized blockchain IOTA for smart objects which can be considered parts of the Industrial Internet of Things?

9. Methodology

ME³SA model

To find out how the IOTA ecosystem and specifically the Tangle perform from an energy-efficient standpoint, we have to apply a model to assess the energy efficiency and sustainability of this particular software architecture. To assess software in this context, Bruntink, Kalaitzoglou and Visser (2014) developed a scheme named the ME³SA model for sustainable computing. The model is a progression of the GREENSOFT model by Lohmann et al. (2013), which was proposed to deliver a model to determine green and sustainable software. The GREENSOFT model does not base its assertions on metrics, which makes it somewhat impractical for holistic assessments. In the model, existing software quality standards, like ISO 25010, were taken as reference.

The design of ME³SA follows the Goal-Question metric approach by Basili et al. (1994). The approach defines three levels of measurement: First, a measurement goal is stated. In a second step, questions are raised, are operationalized and have to be answered to reach the goal. In a last step, the metric how the questions can be answered are formulated. By that, ME³SA provides a practical model for evaluation and discovery of improvement opportunities for the energy efficiency of software applications from the viewpoint of an application owner or investigator (Bruntink et al. 2014).

ME³SA is constituted by the indicators energy behavior, capacity and resource utilization, which are operationalized. For energy behavior, the goal is to find out the degree to which the energy consumed by an application meets the requirements (Bruntink et al. 2014). The questions raised to do so are:

- What is the energy consumption of each application component (Q1)?
- How much energy is wasted by an application in the idle state (Q2)?
- How much energy is consumed per unit of work (Q3)?

The energy consumption of a component (Q1) can be calculated as the sum of samples of hardware utilization attributed to specific components, weighted by the power drawn at each sample (ibid). The ME³SA model defines the Annual Component Consumption *ACC* as

$$ACC(c) := \sum_{s \in S} U_c(s) \times P(U_c(s)),$$

with S being the set of samples for measurement, $U_c(s)$ is the CPU utilization attributed to c at sample s , and $P(U)$ is the power consumption in watts at $U\%$ of utilization (ibid). Since measurement only takes seconds to minutes, annual consumption is upscaled by multiplying it with the factor 525,000 (for minute) or by a fraction of 31,536,000 (for seconds).

To evaluate the components' energy usage in idle state, we have to consider that application components have two major states: Running and idle. When a component actively uses hardware components, then it is said to be running, everything else is idle, and energy in idle state is wasted (Bruntink et al. 2014). Since the idle state is relative to running, a metric indicating the relative energy consumption in idle state is needed to measure its energy behavior. The Relative Idle Consumption (RIC) is defined as

$$RIC(c) := \frac{AIC(c)}{ACC(c)},$$

with $ACC(c)$ referring to the annual consumption for component c , and AIC standing for Annual Idle Consumption for component c . RIC values close to 0% indicate less waste of energy. $AIC(c)$ is defined as

$$AIC(c) := \sum \{P(U_c(s)) \mid s \in S, U_c(s) = 0\% \}.$$

To answer questions about consumption per unit of work, we first have to define what the unit of work we want to measure constitutes. Since transactions determine the existence of the IOTA network, we define the unit of work as the work needed to process one transaction over a node (ibid). The Component Consumption per Unit of Work (CCUW) can be further calculated through

$$CCUW(c) := \frac{ACC(c)}{AUW(c)},$$

where AUW describes the annual number of units of work the component processes. It is expressed in kWh per transaction. This makes sense since the relation between work and

energy consumption is well-expressed with this metric. From the numbers computed, costs can be derived and economically expressed (Bruntink et al. 2014).

When we evaluate the capacity of software, we want to know about the degree to which the maximum energy consumption limits of an application meet requirements of the hardware. It is determined by the questions:

- How much does the application require during peak workload (Q4)?
- How much of the theoretical maximum energy budget does the application use (Q5) (Bruntink et al. 2014)?

The quantification of requirement during peak workload is defined as Peak Growth (PGR) and calculated through

$$PGR(c) := \frac{|P(U_{c\Delta}) - P_{100\%}|}{P_{100\%}} \times 100\%,$$

where $U_{c\Delta}$ refers to the maximum (peak) utilization that component c reaches, with $P_{100\%}$ being the theoretical maximum of the hardware.

The answer to determine the theoretical energy budget, we have to consider that it is normal practice to use provisioning hardware to reach peak load. Simultaneously, the power consumption is responsible for determining larger portions of overall energy consumption (ibid). We define the provisioning (PRO)

$$PRO(c) := \frac{ACC(c) - AIC(c)}{|S| \times P_{100\%}},$$

where S again stands for the number of samples.

For resource utilization, the goal is to measure the degree to which the utilization of resources used by an application meets the requirements. In order to find that out, the following questions have to be answered:

- How does energy consumption scale with an increasing workload (Q6)?
- How power-efficient are the host resources with respect to the average workload of the application (Q7)?

- How much of the total energy consumption of the infrastructure is attributed to the application (Q8)?

One of the biggest causes for energy waste is that application do not reduce energy consumption just because workload and utilization are low. This is known as energy proportionality (Barroso et al. 2007) and describes software systems' ability to scale energy consumption. The model calculates the Consumption Near Sweet spot (CNS) as

$$CNS(c) := \frac{CCUW^*(c)}{CCUW(c)} .$$

Sweet spot describes the position in a running system where the hardware is utilized in an ideal manner. $CCUW^*(c)$ is obtained by identifying the minimum energy consumption per work unit throughout the measurement.

By determining the power efficiency of host resources with respect to the average workload, we distinguish between power usage of host resources and ideal power usage. By doing this, we are able to determine the energy budget for current hardware choices (Bruntink et al. 2014). Wong et al. (2012) proposed to measure the difference between ideal power consumption and actual power consumption at distinct hardware utilization levels. Slightly adjusted, the Proportionality Gap ($PG(c)$) is defined as

$$PG(c) := \frac{P(\overline{U_c}) - P^*(\overline{U_c})}{P_{100\%}} ,$$

where U_c denotes the deviation from the ideal power consumption case. P is the power drawn from the component, P^* is the power the hardware should provide.

Finally, to answer the question of total energy consumption of the infrastructure attributed to the application, we have to determine what amount of energy is wasted by host resources to keep up operations of the application. This is defined as Operational Overhead (OPO) and is expressed by the equation

$$OPO(c) := \frac{ACC(c)}{ASC},$$

where ASC is the total annual energy consumption by host resources on which an application is deployed (Bruntink et al. 2014).

System requirements

The ME³SA model by Bruntink et al. is a good way to investigate software architecture in the light of its energy efficiency. To find the hardware to evaluate is a more complex question, since it should fulfill the requirements of the IOTA system, but also of Internet of Things hardware in general. For the IOTA system, a virtual private server (VPS) with a minimum of two cores and 4GB RAM is considered sufficiently equipped for running IOTA on a device (IOTA Partners 2017). The virtual private server is necessary to run a static IP address, one to which neighboring nodes in the IOTA system have access to. Java SE installation is a requirement, too. The use of Java memory should be limited, and memory for IRI file should also be considered. Java is a platform for computing and programming, the IRI repository is the IOTA Reference Implementation and embodiment of the IOTA network specification (IOTA Foundation 2017a). The IRI repository establishes a full functioning node client on a device with a JSON-REST HTTP interface. JSON is a light-weight data serialization format based on JavaScript, a REST interface is one conform with the Representational State Transfer software architecture style (Mellon 2011). With a Static IP address, Java and an IRI repository we fulfill the requirements to run a full IOTA node on a device.

After meeting the software requirements to run a full node, we have to find neighbors to get woven into the network. There are two ways to accomplish this, either by simply asking other node operators to provide their static IP address or by running the Nelson protocol. For either case, port forwarding, ideally UDP 14600 or UDP 14265, is necessary. Operating with Nelson requires the installation of node and the node package manager for IOTA (IOTA Partners 2017). For the sake of simplicity, we asked other node operators to provide their IP addresses, which they thankfully did. Depending on the operating system, we have to take into account some requirements as well. For Windows, we should use a Secure Shell client, a software protocol to operate network services securely over an unsecured network. For Linux/Ubuntu, we can use the console to compile a secure connection through a list of commands (ibid). The IOTA Wallet then has to be downloaded and run by the device in question. With the addition

of the static IP address of our neighbors, ideally four to seven (ibid), we can run a full node on a device.

Measurements

With the node running and a model for assessment of the energy efficiency of the IOTA Tangle blockchain laid out, we have to consider how we take the measurements to evaluate the energy efficiency of this IoT-optimized blockchain. For Windows, a tool named Joulemeter can be utilized. It is able to measure the energy consumption of software applications running on Windows, using a power model and performance counters (Kansal et al. 2009).

10. Case Study: Smart object energy efficiency assessment with an applied blockchain

For the case study, we have utilized two devices which seem suitable for measuring the energy consumption of the IOTA Tangle and its energy efficiency in further consequence. The first, more stable system was the Lenovo Yoga 2, a personal computer with Windows 8.1 64-bit operating system working on an Intel Premium N3520 quad-core processor with 2.17 GHz base latency and 2.42 GHz burst latency. The graphics board consists of HD Graphics (Bay Trail) based on Intel Gen7 architecture and provides four execution units with 854 MHz latency rate, which guarantees only low processing (Lenovo 2015). However, 4GB of random access memory is sufficient to run an IOTA full node.

10.1. Case Study: Full node on a Windows personal computer

Even if it appears unsuitable to run a full node through a personal computer, it makes sense in the light of equipment for the point of use. Around 75% of all computers use a Windows system (Statista 2018), and a significant number of them is used for industrial production. Furthermore, it is likely that the application, the initiation of a transaction process of the IOTA network, although further executed with other systems, will be done from a Windows PC. Furthermore, a personal computer provides results from energy efficiency measurements in great detail.

To find the Annual Component Consumption for our personal computer running an IOTA node, we look at the list of values the Joulemeter software has recorded in the sheets 'CPU Utilization' and 'Power Battery RoD'. A total of 920 values is recorded in these lists, representing a measurement at every second of the virtual machine power measurement. We choose a representative figure and scale up because the CPU utilization seems to be unstable. A representative figure to calculate ACC is figure 875, showing a utilization of 100% of the total CPU. The corresponding measurement of battery consumption is 5.586 Watt. Applying the formula for ACC correctly, we can determine an annual power consumption of 48.93336 kWh, relative to the CPU utilization the Annual Component Consumption is

48.93336 kWh.

For the Relative Idle Consumption of the IOTA Tangle, we first have to determine a value for the Annual Idle Consumption. A representative figure where the CPU utilization is 0% is figure 277. The Idle Consumption in this state is 3.993 Watts per second, corresponding to 34.97868 kWh Annual Idle Consumption. With this value, we determine the Relative Idle Consumption as

71.48%.

The Component Consumption per Unit of Work (CCUW) can be calculated by knowing the number of transactions the node can process. Assuming a transmission rate of 10 Mbit/s and the size of one IOTA transaction from one node to the other being 1650 Bytes (IOTA stackexchange 2017), the maximum number of transactions is $10,000,000 / (1650 \times 8) = 757$ transactions per second. This means 0.00738 Watt per transaction is needed in theory. In reality, we have to assume a lower transmission rate of approximately 1 Mbit/s for our node, which results in 75 transactions per second. We calculate Annual Units of Work delivered being 2,365,200,000. With that in mind, CCUW results in

20.689×10^{-9} kWh/unit of work.

The Peak Growth can now be calculated by knowing the power consumption at peak CPU utilization. Since 5.586 Watt per second corresponds to 100% utilization, it provides us with the value for $P(U_{c\Delta})$. Following our measurement, the maximum power the hosting hardware can provide ranges from 10 – 11.5 Watt. For the sake of simplicity, we use the value of 11.5 Watt for our calculations. This gives us a result for Peak Growth of

51.43%.

For the metric Provisioning, we subtract the Annual Component Consumption from Annual Idle Consumption for the numerator value, which results in 13.95492 kWh. For the denominator, we set the number of samples S as the number of samples where $P_{100\%}$ was measured, which is 15. Through that, Provisioning accumulates to

8.09%

The Consumption near the Sweet Spot is calculated by assuming that a very low number of transactions, in our case 1, is put through in the IOTA network to identify the minimum observed energy consumption per work unit. This results in 31,536,000 Annual Units of Work. Since it is possible to measure the consumption of one transaction, we were able to determine a value of 4.098 Watt, slightly higher than in Idle state. Since CPU utilization is

100%, we can calculate a value of 35.89848 kWh for one transaction running the network. The adjusted Component Consumption per Unit of Work now corresponds to $11.383 \cdot 10^{-9}$ kWh/unit of work, leading to a Consumption near Sweet Spot percentage of

55.01%.

To find the Proportionality Gap PG, we have to calculate the power usage during average load and compare it with the power usage if the software would be fully proportional. Fully proportional in our case means the energy consumption of one transaction, which was measured to be 4.098 Watt. The average load for all our measurements is 4.422401523 Watt, which is why the numerator for the proportionality gap of our full node is 0.3244 Watt. Divided by $P_{100\%}$, we find the proportionality gap percentage to be

2.82%.

The Operational Overhead (OPO) is calculated by upscaling $P_{100\%}$ to the point of what it would consume to run the hosting hardware on 100% for a whole year. The Annual Host Consumption ASC is computed to be 100.74 kWh, in further consequence the Operational Overhead percentage is

48.57%.

Now that we have calculated the indicators from the ME³SA model, how can we put them into perspective? The Annual Component Consumption shows us how the IOTA full node utilizes the hardware provided would it be constantly needed. The high CPU utilization, at its peak utilizing the CPU continuously 100%, tells us how demanding the algorithm provided by the IOTA network is to process for personal computers. In further consequence this means that running the network in sideline of other applications is not possible for the type of hardware we chose, and we can assume this finding applies also for smart objects. It is at least demanding for a smart object to apply IOTA, since its purpose is to support other applications of a smart object with the ability to process transactions of the IOTA network, thereby running other software in sideline by definition.

The Relative Idle Consumption indicates that the consumption of the IOTA full node is very efficient related to the idle state of the device. The higher the percentage for RIC, the less resources it takes up in comparison to normal running. However, it is difficult how the rate of 71.48% compares with other software applications with state-of-the-art efficiency applied. Comparing the result with the case study of Bruntink et al. (2014), we can assume the consumption efficiency is in mid-range.

For the Component Consumption of Work, it is apparent how relative the results are to what we observe. The assumption of 75 transactions per second processing through one full node is realistic, however not in the current state of the network. This is due to the limited number of participating nodes in the network and the vast differences in transmission rates throughout the network, with 1Mbit/s being the throughput rate during our measurement. For other nodes, the assumptions can differentiate vastly, which is why any hypotheses about the economic ramifications of applying an IOTA full node can only hold ground from the perspective of an analyzer and cannot be interpreted as comprehensive.

The Peak Growth metric gives us enlightening conclusions about the IOTA network: While the energy consumption to run the system is low in comparison, the hardware utilization is high. This obviously changes with the activity within the network and depends on what specifications the hosting hardware provide, but can be derived from running the network over popular hardware.

The Provisioning can tell us two things: Either that the hardware is not utilized perfectly by the application or even idle at times. Suggestions for improvement could include a recommendation to utilize hardware resources better. High percentages, on the other hand, could overconsume energy for running (Bruntink et al 2014). 8% seem to be a reasonable percentage for Provisioning, although it needs to be mentioned the figure for our measurement is probably inconclusive to make any statements about the energy efficiency of the IOTA network.

The Consumption near the Sweet Spot determines whether an application consumption and workload during running are proportional, since energy-efficient software would run up energy consumption linearly in relation to the workload. A percentage close to 0% indicates a higher energy consumption than expected compared to the workload, while 100% indicate full energy proportionality (Barroso et al. 2007). 55.01% show that the proportionality between consumption and workload is significantly better compared to other investigated software (Bruntink et al. 2014), but we can still disproportionate utilization.

The Proportionality Gap represents the gap between actual and lowest proportionality in the energy regime of the hardware for average utilization (ibid). Considering that the gap is this low, we can assume that the energy consumption is close to an optimum in the IOTA network. However, there is room for improvement for the proportionality of the network.

The Operational Overhead shows us how efficient the software runs when and how much efficiency is wasted in the process. Derived from our calculations, it can be witnessed the

component is not fully utilized to its potential, however as we have also seen the energy consumption is comparably low for the hardware, so in terms of wasting energy this figure tells us the network provides a lot of overhead due to its small power consumption.

Initially, a second case study more suitable for IoT environments due to size, specifications and costs was planned to be conducted. To conduct research on such an object, we wanted to run the IOTA full node on a Raspberry Pi 3 Model B. It is a single-board computer (SBC), where the Linux Ubuntu Mate 64-bit operating system was installed. It comes with a Broadcom 1.2 GHz processor in the ARM Cortex-A53 microarchitecture. The operating system has to be downloaded onto a microSD card, which then can be inserted into the single board computer (Raspberry Pi Foundation 2016). This means the memory depends on external hardware the user inserts, in our case we inserted a 32GB card. Since the random access memory installed is only 1 Gigabyte, we had to reserve virtual memory for the node to function. We did this by RAM swapping through ReadyBoost onto our SD card (Microsoft 2018), we reserved 4096 MB for the node to function. However, the virtual memory requirement made it impossible to run the network on this device since the processes of the IOTA network are highly optimized towards dedicated random access memory. This became evident during the case study and is important information for our conclusion.

11. Conclusion/Outlook

The first observation from the preparation of the breadboard construction concerns the prerequisites for IoT devices to run IOTA on an object. While 4 GB of random access memory are certainly available and affordable, it raises the question whether the same applies for small smart objects. The smaller the memory, the more costs the implementation of higher memory sizes onto smart objects produce. On the other hand, if memory is desired to be cheap, then the dimensioning of the memory will constrain its application. Assuming the area of use and quantity requires small-dimension memory, the cost of the benefit of implementation may be higher than the actual savings, even by applying the “power of 1%” rule. While an economic estimation of how much the integration of IOTA into IoT environments will cost is difficult at this stage of development, it can certainly be said that the requirements to run the system increase the costs significantly. The system requirements are high, which is why the implementation cost for supply chain applications, machine-to-machine communication and other application fields will be high. However, as was laid out in the theoretical part of the Master thesis, it highly depends on the use case of a blockchain whether the predicted outcome will be realized. In other words: There will be many cases of use for IOTA where the implementation costs will be low in relation to the benefit it provides.

It should be noted that while running the measurements brought us results from which we can derive statements that are very likely to be true, it is inconclusive whether our case study can be seen representatively for the IOTA network energy impacts. This is due to lacking investigation of an object that could be interpreted as ‘smart’ in a narrower sense, meaning something similar to a single-board computer in terms of size, specifications and cost due to the constraints mentioned above. Furthermore, the measurements of our software show fluctuations for both CPU utilization and energy consumption, however in a stable range, which is why the values taken from the measurements can be interpreted as representative for the entire case study. In this respect, it should be said this is also the reason why representative values have been chosen to identify the metrics supporting us to answer the research question. The devices for measurement, however, are representative for the Industrial Internet of Things.

The results for the energy efficiency of the system show that while the CPU utilization of the IOTA network is very high, the energy consumption does not correspond to the use of the hardware. IOTA operates conciliatory towards energy exploitation, and the indicators validate this statement. However, when we are speaking about efficiency, it needs to be said that

efficiency for software also requires proportionate resource utilization of both CPU and energy use. To answer the research question is difficult since the indicators point into different directions. If we interpret 'energy impacts' solely electricity-related, we can say the energy impacts of an IoT-optimized blockchain for smart objects for the Industrial Internet of Things are low. If we take utilization of hardware into account, we cannot give an indeterminate answer, since the IOTA network utilizes all of the CPU it is provided with 4GB random access memory. In terms of capacity, we can say that the resources utilized by this software are also low.

The research showed significant room for improvement for the proportionality between energy consumption and hardware utilization. Recommendations that can be made from our perspective include the artificial enlargement of the network, the reduction of CPU utilization and the reduction of prerequisites to run IOTA on smart objects. Network enlargement is the key towards preventing fluctuations in the IOTA network that could lead to load regimes where throughput becomes an issue, eventually taking the network down. The more devices are connected to the network, the more stable it becomes, which is why an artificial enlargement would lead to better performance of the network on single full nodes. The IOTA Foundation, so far, has not proposed any measure like this, relying on the organic growth of the network, which is currently exponential (Semko 2017).

The reduction of CPU utilization corresponds with the enlargement of the network, since less full nodes have to perform proportionately more transactions. However, since the network is constantly adjusted to the requirements the Internet of Things and particularly IIoT provide, developers should take the reduction of CPU utilization from the perspective of software architecture into consideration. Because this recommendation is at the core of the network, simple measures like a limitation of the number of transactions a full node can process in a defined period or the implementation of a protocol limiting block intervals and block sizes could be measures to implement such changes, since other changes would include changing fundamental properties of the IOTA network.

At last, if the technology will see global implementation is depending highly on the prerequisites of the network to run. The requirement of 4GB random access memory is probably too high that we will see major implementations in the next 2-3 years. But since the breakthrough of the technology is believed to be in roughly 5 years from now, we could see major use cases for the IoT-optimized blockchain IOTA. It should also be noted that the

requirements could become a constraint to develop use cases for the technology, although it cannot be determined whether this is likely or not.

Further research in this area could go into comparative analysis of various blockchains for application in the Industrial Internet of Things and a holistic research under an operating environment of such a blockchain. The ME³SA model would allow to conduct a comparative analysis of various blockchains for a specific use case like the Internet of Things, as it seems to be the most sufficient model to analyze sustainability of software. A suggestion for improvement for the model or a comparative analysis include to assess the space memory blockchains occupy.

An investigation of performance under an operating environment was the initial idea for research for this Master Thesis. It was discarded when we realized the technology and its application are still in a state of experimenting, meaning real-life use cases are very rare and experimental. However, operating environments of IoT could provide insight into the energy potentials of the IIoT and not only estimates for the assessment of the energy impacts of IoT-optimized blockchain. With improved research models on the matter, we will be able to evaluate the software in much greater detail.

Finally, the blockchain technology has only now experienced great attention due to the rise of Bitcoin and the associated earnings of investors. There is much more potential behind the blockchain technology, as was laid out in this Master Thesis. This potential needs to be utilized, since we can think about a variety of fields where the technology could be beneficial. This Master Thesis was a contribution towards spreading the idea of blockchain and why it makes sense for the academic community to further elaborate research on the topic.

Bibliography/Sources

- Ago, Steemhoops99. 2017. "What Is 'Snapshotting' in IOTA? - 快照." Steemit. August 3, 2017
<https://steemit.com/technology/@steemhoops99/iota-snapshot-what-is-it>.
- Badkar, Mamta. 2018. "Bitcoin Energy Demand in 2018 Could Match Argentina – Morgan Stanley." Financial Times. January 10, 2018.
<https://www.ft.com/content/93b22cb1-0346-38be-bebf-d2e676e19621>.
- Baird, Leemon. 2016. "The Swirlds Hashgraph Consensus Algorithm: Fair, Fast, Byzantine Fault Tolerance," May 2016.
<https://www.swirlds.com/downloads/SWIRLDS-TR-2016-01.pdf>
- Baird, Leemon, Mance Harmon, and Paul Madsen. 2018. "Hedera: A Governing Council & Public Hashgraph Network - The Trust Layer of the Internet."
<https://s3.amazonaws.com/hedera-hashgraph/hh-whitepaper-v1.0-180313-2.pdf>.
- Baliga, Arati. 2017. "Understanding Blockchain Consensus Models." Persistent Systems.
<https://www.persistent.com/wp-content/uploads/2017/04/WP-Understanding-Blockchain-Consensus-Models.pdf>.
- Barroso, Luiz André, and Urs Hölzle. 2007. "The Case for Energy-Proportional Computing." *Computer* 40 (12): 33–37.
<https://doi.org/10.1109/MC.2007.443>.
- Bedford Taylor, Michael. 2013. "Bitcoin and the Age of Bespoke Silicon." 2013 International Conference on Compilers, Architecture and Synthesis for Embedded Systems (CASES). IEEE.
<https://doi.org/10.1109/CASES.2013.6662520>.
- Bergmann, Christoph. 2017. "IOTA: Der Missverstandene Coin?" *BitcoinBlog.de - Das Blog Für Bitcoin Und Andere Virtuelle Währungen*. December 20, 2017.
<https://bitcoinblog.de/2017/12/20/iota-der-missverstandene-coin/>.
- Beyerer, Jürgen, and Thomas Usländer. 2016. "Industrial Internet of Things Supporting Factory Automation." *At - Automatisierungstechnik* 64 (9).
<https://doi.org/10.1515/auto-2016-0104>.

- Blockgeeks. 2017. "Blockchain Scalability: When, Where, How?" Blockgeeks.com, 2017. <https://blockgeeks.com/guides/blockchain-scalability/>.
- Blok, Kornelis, Paul Hofheinz, and John Kerkhoven. 2015. "The 2015 Energy Productivity and Economic Prosperity Index. How Efficiency Will Drive Growth, Create Jobs and Spread Wellbeing Throughout Society" Ecofys. <https://www.ecofys.com/files/files/the-2015-energy-productivity-and-economic-prosperity-index.pdf>
- Bundesministerium für Bildung und Forschung – BMBF. 2018. "Industrie 4.0 – BMBF". Accessed May 16, 2018. <https://www.bmbf.de/de/zukunftsprojekt-industrie-4-0-848.html>.
- Boston Consulting Group. 2018. "Industry 4.0 - the Nine Technologies Transforming Industrial Production." <https://www.bcg.com/capabilities/operations/embracing-industry-4.0-rediscovering-growth.aspx>.
- Brynjolfsson, Erik. 2011. "Is Koomey's Law Eclipsing Moore's Law?" Accessed May 15, 2018. <http://www.economicsofinformation.com/2011/09/is-koomeys-law-eclipsing-moores-law.html>.
- Buterin, Vitalik. 2013. "What Proof of Stake Is and Why It Matters." Bitcoin Magazine. Accessed April 20, 2018. <https://bitcoinmagazine.com/articles/what-proof-of-stake-is-and-why-it-matters-1377531463/>.
- Cachin, Christian, Simon Schubert, and Marko Vukolić. 2016. "Non-Determinism in Byzantine Fault-Tolerant Replication." Cornell University. March 2016. <http://arxiv.org/abs/1603.07351>.
- Cachin, Christian, Marko Vukolic, and Marc Herbstritt. 2017. "Blockchain Consensus Protocols in the Wild (Keynote Talk)." Leibniz-Zentrum für Informatik, Wadern/Saarbrücken, Germany. <https://doi.org/10.4230/LIPIcs.DISC.2017.1>.
- Castro, Miguel, and Barbara Liskov. 1999. "Practical Byzantine Fault Tolerance," In: Proceedings of the Third Symposium on Operating Systems Design and Implementation. New Orleans. February 1999. <http://pmg.csail.mit.edu/papers/osdi99.pdf>

- Christidis, Konstantinos, and Michael Devetsikiotis. 2016. "Blockchains and Smart Contracts for the Internet of Things." *IEEE Access* 4: 2292–2303.
<https://doi.org/10.1109/ACCESS.2016.2566339>.
- dantheman. 2017. "DPOS Consensus Algorithm - The Missing White Paper" Steemit.
<https://steemit.com/dpos/@dantheman/dpos-consensus-algorithm-this-missing-white-paper>.
- De Castro, Andre. 2016. "Blockchain And IoT: A Perfect Match?" Blockgeeks.com. December 20, 2016.
<https://blockgeeks.com/blockchain-and-iot-a-perfect-match/>.
- Digiconomist. 2018. "Bitcoin Energy Consumption Index." Digiconomist. Accessed April 15, 2018.
<https://digiconomist.net/bitcoin-energy-consumption>.
- Dorri, Ali, Salil S. Kanhere, and Raja Jurdak. 2017. "Towards an Optimized BlockChain for IoT." In *IoTDI '17: Proceedings of the Second International Conference on Internet-of-Things Design and Implementation*. 173-178. ACM Press.
<https://doi.org/10.1145/3054977.3055003>.
- Economist. 2015. "Machine Learning." The Economist.
<https://www.economist.com/news/leaders/21678786-manufacturers-must-learn-behave-more-tech-firms-machine-learning>.
- Economist. 2016. "The Great Convergence." The Economist, July 21, 2016.
<https://www.economist.com/news/business/21702487-china-aims-lead-world-connecting-factory-great-convergence>.
- European Commission. 2009. "Communication from the European Commission to the European Parliament, the Council, the European Economic and Social Committee and the Committee of the Regions: Internet of Things - An Action Plan for Europe." <https://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=COM:2009:0278:FIN:EN:PDF>.
- Floyd, David. 2018. "The 17 Millionth Bitcoin Is About to Be Mined: What It Means and Why It Matters." Coindesk. April 26, 2018. <https://www.coindesk.com/17-millionth-bitcoin-mined-means-matters/>.

- Gautier, Philippe. 2010. "3 Questions to Philippe Gautier, by David Fayon". I-O-T : Internet of Things. April 17, 2010.
<http://www.i-o-t.org/post/3questionstoPhilippeGAUTIERbyDavidFayon>.
- Gilchrist, Alasdair. 2016. "Industry 4.0." University of California, Berkeley. Apress.
[#https://doi.org/10.1007/978-1-4842-2047-4](https://doi.org/10.1007/978-1-4842-2047-4).
- Goodman, L M. 2014. "Tezos - a Self-Amending Crypto-Ledger."
https://tezos.com/static/papers/white_paper.pdf.
- Graczyk, Michael. 2018. "Hashgraph: A Whitepaper Review." OpenToken-Medium.com. February 1, 2018.
<https://medium.com/opentoken/hashgraph-a-whitepaper-review-f7dfe2b24647>.
- Hayes, Adam. 2015. "A Cost of Production Model for Bitcoin." *SSRN Electronic Journal*. University of Wisconsin – Madison.
<https://doi.org/10.2139/ssrn.2580904>.
- Hermann, Mario, Tobias Pentek, and Boris Otto. 2016. "Design Principles for Industrie 4.0 Scenarios." In: 2016 49th Hawaii International Conference on System Sciences (HICSS), 3928–37. Koloa: IEEE.
<https://doi.org/10.1109/HICSS.2016.488>.
- Hern, Alex. 2018. "Bitcoin's Energy Usage Is Huge – We Can't Afford to Ignore It." The Guardian. January 17, 2018.
<http://www.theguardian.com/technology/2018/jan/17/bitcoin-electricity-usage-huge-climate-cryptocurrency>.
- Higgins, Stan. 2016. "Bitcoin's 'Nervous System' Gets an Upgrade with FIBRE Network." Coindesk. July 21, 2016.
<https://www.coindesk.com/bitcoins-nervous-system-getting-upgrade/>.
- Hoekstra, Matthew. 2013. "Intel SGX for Dummies (Intel SGX Design Objectives) | Intel Software." Accessed April 26, 2018.
<https://software.intel.com/en-us/blogs/2013/09/26/protecting-application-secrets-with-intel-sgx>.
- Hoske, M.T. 2016. "Finding IIoT Benefits". Control Engineering. 8.

- Hyperledger. n.d. "Hyperledger Fabric." Hyperledger. Accessed April 26, 2018. <https://www.hyperledger.org/projects/fabric>.
- Intel Corporation. n.d. "Introduction — Sawtooth v1.0.2 Documentation." Accessed April 16, 2018. <https://sawtooth.hyperledger.org/docs/core/releases/1.0/introduction.html>.
- IOTA Foundation. 2017a. "Iri: IOTA Reference Implementation. Java. IOTA." Github. <https://github.com/iotaledger/iri>.
- IOTA Foundation. 2017. "Tangle - What Does the IOTA Coordinator Actually Do?" IOTA stack exchange. 2017. <https://iota.stackexchange.com/questions/4/what-does-the-iota-coordinator-actually-do>.
- IOTA Foundation. 2018. "Tip Selection - IOTA Docs." 2018. <https://dev.iota.org/introduction/tangle/tip-selection>.
- IOTA Partners. 2017. "IOTA Full Node Copy-Paste Installation Guide." 2017. <http://iota.partners/>.
- IOTA stack exchange. 2017. "What Is the Max Possible Transactions/Second Rate in IOTA?" IOTA stack exchange. 2017. <https://iota.stackexchange.com/questions/1049/what-is-the-max-possible-transactions-second-rate-in-iota>.
- Jackson Palmer. 2018. "What Is a Directed Acyclic Graph (DAG)? IOTA, Byteball, SPECTRE Reviewed." Accessed May 11, 2018. <https://www.youtube.com/watch?v=hTfSHJGyIG8>.
- JP Buntinx. 2017. "What Is Proof of Elapsed Time?" The Merkle. Accessed April 26, 2018. <https://themerple.com/what-is-proof-of-elapsed-time/>.
- Kansal, Aman, Feng Zhao, Nupur Kothari, and Arka Bhattacharya. 2009. "JouleMeter: Virtual Machine Power Measurement and Management," Microsoft. <https://www.microsoft.com/en-us/research/wp-content/uploads/2009/08/JouleMeterTR.pdf>.
- Koomey, Jonathan, Stephen Berard, Marla Sanchez, and Henry Wong. 2011. "Implications of Historical Trends in the Electrical Efficiency of Computing." *IEEE Annals of the History of Computing* 33 (3): 46–54. <https://doi.org/10.1109/MAHC.2010.28>.

- Koomey, Jonathan G. 2010. "Outperforming Moore's Law." IEEE Spectrum: Technology, Engineering, and Science News. February 26, 2010.
<https://spectrum.ieee.org/computing/hardware/outperforming-moores-law>.
- Kwon, Jae. 2014. "Tendermint: Consensus without Mining," Cornell University.
<https://tendermint.com/static/docs/tendermint.pdf>
- Laitner, John A. and Karen Ehrhardt-Martinez. 2008. "Information and Communication Technologies: The Power of Productivity (Part I)." Environmental Quality Management 18 (2), 47-66. Wiley Online Library
<https://doi.org/10.1002/tqem.20205>.
- Lamport, Leslie, Robert Shostak, and Marshall Pease. 1982. "The Byzantine Generals Problem." ACM Transactions on Programming Languages and Systems 4: 382–401.
- Liao, Chun-Feng, Sheng-Wen Bao, Ching-Ju Cheng, and Kung Chen. 2017. "On Design Issues and Architectural Styles for Blockchain-Driven IoT Services." 2017 IEEE International Conference on Consumer Electronics - Taiwan (ICCE-TW), 351–52. IEEE.
<https://doi.org/10.1109/ICCE-China.2017.7991140>.
- Lenovo. 2015. "Lenovo Yoga | 2-in-1-Geräte, Tablets, Notebooks und Desktops." Lenovo.
<https://www3.lenovo.com/at/de/yoga/>.
- Linux Foundation. 2016. "Linux Foundation's Hyperledger Project Announces 30 Founding Members and Code Proposals To Advance Blockchain Technology | Hyper Ledger Foundation." Hyperledger Project.
<https://web.archive.org/web/20160225023123/https://www.hyperledger.org/news/announcement/2016/02/hyperledger-project-announces-30-founding-members>.
- Lohmann, Wolfgang, Lorenz M. Hilty, Bernard Aebischer, and Göran Andersson. 2013. "ICT4S 2013: Proceedings of the First International Conference on Information and Communication Technologies for Sustainability." ETH Zürich.
<https://doi.org/10.3929/ethz-a-007337628>.
- Madeira, Antonio. 2018. "What Is the GHOST Protocol for Ethereum?" Cryptocompare. Accessed May 20, 2018.
<https://www.cryptocompare.com/coins/guides/what-is-the-ghost-protocol-for-ethereum/>.

- Madsen, Paul. 2017. "I Want Your Vote! (Oh Wait I Already Know It)." Hedera Hashgraph Blog. Medium.com.
<https://medium.com/hashgraph/i-want-your-vote-oh-wait-i-already-know-it-e1faa50b31ad>.
- Malone, D., and K.J. O'Dwyer. 2014. "Bitcoin Mining and Its Energy Footprint." 25th IET Irish Signals & Systems Conference 2014 and 2014 China-Ireland International Conference on Information and Communities Technologies (ISSC 2014/CIICT 2014), 280-285. Institution of Engineering and Technology.
<https://doi.org/10.1049/cp.2014.0699>.
- Marr, Bernard. 2016. "Why Everyone Must Get Ready For The 4th Industrial Revolution." Forbes. Accessed May 16, 2018.
<https://www.forbes.com/sites/bernardmarr/2016/04/05/why-everyone-must-get-ready-for-4th-industrial-revolution/>.
- Mattern, Friedemann, and Christian Floerkemeier. 2010. "From the Internet of Computers to the Internet of Things." From Active Data Management to Event-Based Systems and More, edited by Kai Sachs, Ilia Petrov, and Pablo Guerrero, 6462:242–259. Berlin, Heidelberg: Springer Berlin Heidelberg.
https://doi.org/10.1007/978-3-642-17226-7_15.
- Mazieres, David. 2014. "The Stellar Consensus Protocol: A Federated Model for Internet-Level Consensus." Stellar Foundation.
<https://www.stellar.org/papers/stellar-consensus-protocol.pdf>.
- Medhi, Githarti. 2016. "Creating Value with Industry 4.0 - Master Thesis." Massachusetts Institute of Technology.
<http://hdl.handle.net/1721.1/107606>.
- Mellon. 2011. "Java - What Is JSON REST Interface." Stack Overflow.
<https://stackoverflow.com/questions/5469835/what-is-json-rest-interface>.
- Microsoft Technet. 2018. "Understand ReadyBoost and Whether It Will Speed Up Your System." Microsoft.
<https://technet.microsoft.com/en-us/library/ff356869.aspx>.
- Mill, John Stuart, and Anonymous. 2016. "A Distributed Consensus Algorithm for Cryptocurrency Networks." Skycoin.

<https://downloads.skycoin.net/whitepapers/a-distributed-consensus-algorithm-for-cryptocurrency-networks.pdf>.

Milutinovic, Mitar, Warren He, Howard Wu, and Maxinder Kanwal. 2016. "Proof of Luck: An Efficient Blockchain Consensus Protocol." In: Proceedings of the 1st Workshop on System Software for Trusted Execution, 1-6. ACM Press.

<https://doi.org/10.1145/3007788.3007790>.

Mohammadi, Mehdi, Ala Al-Fuqaha, Sameh Sorour, and Mohsen Guizani. 2017. "Deep Learning for IoT Big Data and Streaming Analytics: A Survey." Cornell University.

<http://arxiv.org/abs/1712.04301>.

Morrow, Jerome. 2014. "What Is a Coinbase Transaction?" CEX.IO Official Blog..

<https://blog.cex.io/bitcoin-dictionary/coinbase-transaction-12088>.

Nofer, Michael, Peter Gomber, Oliver Hinz, and Dirk Schiereck. 2017. "Blockchain." Business & Information Systems Engineering 59 (3): 183–87.

<https://doi.org/10.1007/s12599-017-0467-3>.

Nordrum, Amy. 2016. "Popular Internet of Things Forecast of 50 Billion Devices by 2020 Is Outdated." IEEE Spectrum: Technology, Engineering, and Science News.

<https://spectrum.ieee.org/tech-talk/telecom/internet/popular-internet-of-things-forecast-of-50-billion-devices-by-2020-is-outdated>.

P4Titan. 2014. "Slimcoin: A Peer-to-Peer Crypto-Currency with Proof-of-Burn - 'Mining without Powerful Hardware.'" Department of Computing, Imperial College London.

http://www.doc.ic.ac.uk/~ids/realdotdot/crypto_papers_etc_worth_reading/proof_of_burn/slimcoin_whitepaper.pdf.

Poelstra, Andrew. 2015. "On Stake and Consensus." WP Software.

<https://download.wpsoftware.net/bitcoin/pos.pdf>.

Popov, Sergei. 2018. "The Tangle - White Paper" IOTA Foundation. Berlin.

https://assets.ctfassets.net/r1dr6vzfxhev/4i3OM9JTleiE8M6Y04Ii28/d58bc5bb71cebe4adc18fadea1a79037/Tangle_White_Paper_v1.4.2.pdf.

poulpita. 2014. "Trusted Execution Environment, Millions of Users Have One, Do You Have Yours?" Poulpita.

<https://poulpita.com/2014/02/18/trusted-execution-environment-do-you-have-yours/>.

- Prena, N, Jain Lovee, and N Vedavathi. 2015. "Issues in Internet of Things: A Survey." International Journal for Research in Applied Science & Engineering Technology (IJRASET). <https://www.ijraset.com/files/serve.php?FID=3169>.
- Raspberry Pi Foundation. n.d. "Raspberry Pi 3 Model B." Raspberry Pi. Accessed June 2, 2018. <https://www.raspberrypi.org/products/raspberry-pi-3-model-b/>.
- Ray, James. 2018. "Wiki: The Ethereum Wiki -. ethereum." Github. <https://github.com/ethereum/wiki>.
- Rilee, Kynan. 2018. "Understanding Hyperledger Sawtooth — Proof of Elapsed Time." Medium.com. <https://medium.com/kokster/understanding-hyperledger-sawtooth-proof-of-elapsed-time-e0c303577ec1>.
- Rosic, Ameer. 2017. "What Is Ethereum Casper Protocol? Crash Course." Blockgeeks. <https://blockgeeks.com/guides/ethereum-casper/>.
- Rückeshäuser, Nadine. 2017. "Typology of Distributed Ledger Business Models." Proceedings of the 25th European Conference on Information Systems (ECIS), June 2017, 2202-2217. http://aisel.aisnet.org/ecis2017_rp/140/.
- Sandner, Philipp. 2017. "Application of Blockchain Technology in the Manufacturing Industry." Medium.com. <https://medium.com/@philippsandner/application-of-blockchain-technology-in-the-manufacturing-industry-d03a8ed3ba5e>.
- Scherer, Mattias. 2017. "Performance and Scalability of Blockchain Networks and Smart Contracts," Master Thesis. Umea University 46. <https://umu.diva-portal.org/smash/get/diva2:1111497/FULLTEXT01.pdf>.
- Schiener, Dominik. 2017. "Current Role of the Coordinator. IOTA Guide." Gitbooks. <https://domschiener.gitbooks.io/iota-guide/content/chapter1/current-role-of-the-coordinator.html>.
- Schiener Dominik. 2018. "The Anatomy of a Transaction. IOTA Guide." Gitbooks. <https://domschiener.gitbooks.io/iota-guide/content/chapter1/transactions-and-bundles.html>.

- Schoenberger, Chana R. 2002. "The Internet of Things." Forbes. Accessed May 20, 2018.
<https://www.forbes.com/global/2002/0318/092>.
- Schwartz, David, Noah Youngs, and Arthur Britto. 2014. "The Ripple Protocol Consensus Algorithm," Ripple Labs, Inc.
https://ripple.com/files/ripple_consensus_whitepaper.pdf
- Semko, Roman. 2017. "IOTA: Tangle Growth Update October 2017." Deviota, Medium.com.
<https://medium.com/deviota/iota-tangle-growth-update-october-2017-9fb61a187388>.
- Skycoin Network. 2018. "Skycoin - Business Whitepaper." Skycoin.
<https://downloads.skycoin.net/whitepapers/Skycoin-Whitepaper-v1.0.pdf>.
- Basili, Vic, Gianluigi Caldiera, H. Dieter Rombach, Rini van Solingen. 2002. "Goal Question Metric (GQM) Approach." In: Encyclopedia of Software Engineering, edited by John J. Marciniak. Hoboken, NJ, USA: John Wiley & Sons, Inc.
<https://doi.org/10.1002/0471028959.sof142>.
- Sompolinsky, Yonatan, and Aviv Zohar. 2013. "Accelerating Bitcoin's Transaction Processing. Fast Money Grows on Trees, Not Chains." School of Engineering and Computer Science, The Hebrew University of Jerusalem, Israel. 881.
<http://eprint.iacr.org/2013/881>.
- Statista. 2018. "Betriebssysteme - Marktanteile Weltweit Bis März 2018 Statistik." 2018.
<https://de.statista.com/statistik/daten/studie/157902/umfrage/marktanteil-der-genutzten-betriebssysteme-weltweit-seit-2009/>.
- Swan, Melanie. 2015. "Blockchain: Blueprint for a New Economy." O'Reilly Media.
- Thulasiraman, K., and M. N. S. Swamy. 1992. "Graphs: Theory and Algorithms". Wiley. New York.
<https://onlinelibrary.wiley.com/doi/book/10.1002/9781118033104>.
- Tobin, Jared. 2018. "Byzantine Generals and Nakamoto Consensus." Jtobin.io. Accessed April 16, 2018.
<https://jtobin.io/byzantine-generals-nakamoto-consensus>.

- Ustundag, Alp, and Emre Cevikcan. 2018. Industry 4.0: Managing The Digital Transformation. Springer Series in Advanced Manufacturing. Cham: Springer International Publishing. <https://doi.org/10.1007/978-3-319-57870-5>.
- Vaidya, Saurabh, Prashant Ambad, and Santosh Bhosle. 2018. "Industry 4.0 – A Glimpse." *Procedia Manufacturing* 20: 233–38. <https://doi.org/10.1016/j.promfg.2018.02.034>.
- Vermesan, Ovidiu, and Peter Friess. 2013. "Internet of Things: Converging Technologies for Smart Environments and Integrated Ecosystems." <http://site.ebrary.com/id/10852704>.
- Wang, Kun, Yihui Wang, Yanfei Sun, Song Guo, and Jinsong Wu. 2016. "Green Industrial Internet of Things Architecture: An Energy-Efficient Perspective." *IEEE Communications Magazine* 54 (12): 48–54. <https://doi.org/10.1109/MCOM.2016.1600399CM>.
- Wolter, Marc Ingo, Anke Mönnig, Markus Hummel, Christian Schneemann, Enzo Weber, Gerd Zika, Robert Helmrich, Tobias Maier, and Caroline Neuber-Pohl. 2015. "Industrie 4.0 und die Folgen für Arbeitsmarkt und Wirtschaft: Szenariorechnungen im Rahmen der BIBB-IAB-Qualifikations- und Berufsfeldprojektionen," Institut für Arbeitsmarkt- und Berufsforschung. <http://doku.iab.de/forschungsbericht/2015/fb0815.pdf>
- Wong, Daniel, and Murali Annavaram. 2012. "KnightShift: Scaling the Energy Proportionality Wall through Server-Level Heterogeneity." In: 2012 45th Annual IEEE/ACM International Symposium on Microarchitecture, 119–30. Vancouver. IEEE. <https://doi.org/10.1109/MICRO.2012.20>.