dena
German Energy Agency

Future Energy
Lab

Rethinking Blockchain's Electricity Consumption

A Guide to Electricity-Efficient Design of Decentralized Data Infrastructure

Legal Information

Publisher:
Deutsche Energie-Agentur GmbH (dena)
Chausseestraße 128 a
10115 Berlin

Tel: +49 (0)30 66 777-0
Fax: +49 (0)30 66 777-699

E-Mail: info@dena.de

Internet:
www.dena.de
future-energy-lab.de

Authors:
Philipp Richard, dena
Moritz Schlösser (Project Leader), dena
Hendrik Zimmermann, dena

Vincent Gramlich, Fraunhofer FIT
Felix Paetzold, Fraunhofer FIT
Prof. Dr. Jens Strüker, Universität Bayreuth & Fraunhofer FIT

Conception & design:
die wegmeister gmbh

Last updated:
September 2023

Please cite this publication as follows:
Deutsche Energie-Agentur (Publisher) (dena, 2023) “Rethinking Blockchain’s Electricity Consumption – A Guide to Electricity-Efficient Design of Decentralized Data-Infrastructure”

Federal Ministry
for Economic Affairs
and Climate Action

This publication is issued on behalf of the Federal Ministry for Economic Affairs and Climate Action. The German Energy Agency (dena) assists the Federal Government in various projects to implement the energy and climate targets in the context of the energy transition.

Table of Contents

Preface

One thing is for certain: blockchain technology is one of the most hyped-up technologies of the past decade. People were talking about how it was a digital revolution, a global phenomenon even – one that would change how humans live and interact. Blockchains were going to dismantle the concentration of power in the hands of large corporations with regard to the Internet. They would enable individuals to own their data once again and in turn fix the Internet through the introduction of Web 3.0. And that was not all: cryptocurrencies were going to become an alternative to traditional currencies because they would not be controlled by central institutions like (central) banks or governments. Hopes were raised that millions of people living in developing countries or nations run by autocrats would be provided with access to a fair financial network that was steeled against attempts at external control.

But what is the situation today? Blockchains are no longer being mentioned in the German government's digital strategy, trust in cryptocurrencies has been severely damaged after the collapse of the cryptocurrency exchange FTX in November 2022 and criticisms of the economic, social, but above all, the ecological sustainability of the technology are growing louder and louder.

Proof-of-work mining of new blocks, in particular, consumes vast amounts of electricity and materials. Critics see no real benefit in the technology justifying its expenditure of resources. Cryptocurrencies are considered too volatile to be used as a real means of payment and only serve as speculative objects. Web 3.0 is even said to have 'dystopian potential' with regard to the ownership of personal data and monetization thereof.[1] Allegedly, the blockchain is unsuitable for other applications because, for example, one of its core properties, immutability, conflicts with the "right to be forgotten", a defining aspect of the General Data Protection Regulation (GDPR), and thus with the limitations of the 'real' world.

If we apply the theory of technology development according to the Gartner hype cycle[2], blockchain technology has now passed the Peak of Inflated Expectations and reached the Trough of Disillusionment. But where are we heading? To answer this question and to help further crystallize what blockchain technology can be used for along the Slope of Enlightenment, a sober and scientific view on the benefits and the electricity consumption of blockchains is required. This, we hope, will help to cool down the sometimes rather heated debates between advocates and opponents of the technology and to determine what level the Plateau of Productivity will ultimately take.[3]

The present study focuses on one of these topics and provides a new, valuable basis for the discussions around the electricity consumption of blockchains. This study contains a guide that makes it easier to design blockchains in the most energy-efficient way possible and according to the requirements of specific use cases. Our findings can serve as a basis for comparing the climate and environmental impact of a blockchain as well as its other properties, such as performance and IT security, with those of alternative network solutions, thus enabling an informed decision on what technology to use.

[1] As Dr. Malte Enegele stated during a Bundestag hearing on Web 3.0 and the Metaverse: https://www.bundestag.de/ausschuesse/a23_digitales/Anhoerungen/921548-921548

[2] https://www.gartner.com/en/research/methodologies/gartner-hype-cycle, last accessed on 07.08.2023

[3] https://www.gartner.com/en/research/methodologies/gartner-hype-cycle, last accessed on 07.08.2023

The guide to the electricity-saving design of blockchains is the result of intensive collaboration with the staff of Fraunhofer FIT, for which we would like to express our gratitude. We also want to extend our thanks to the Federal Ministry for Economic Affairs and Climate Action, which has funded this study.

We are convinced that the use of digital technologies is essential for the success of the ongoing energy transition, but also for transformations in other sectors besides the energy industry. At the same time, we advocate minimizing the power consumption of digital technologies. Therefore, this publication is intended as a means to encourage a critical examination of blockchains, as well as alternative network architectures, with regard to their impact on the climate, the environment, and people on the basis of scientific criteria. Only in this way can the benefits of digitalization be weighed against its ecological and social costs.

Moritz

Philipp Richard
Head of Division, Digital Technologies and Start-up Ecosystem -
German Energy Agency

Moritz Schlösser
Expert Digital Technologies -
German Energy Agency

Executive Summary

Digitalization is one of the most significant technological, social, economic, and political transformations of our time. This megatrend has far-reaching implications for all areas of life. Digital technologies are now essential to the vast majority of our industries, as well as to our coexistence in a global society. They are an unprecedented tool for solving both new and long-standing challenges. The downside, however, is not only social, but also environmental, as energy consumption increases due to the growing requirements for the underlying data infrastructure. In addition, the digital economy is characterized by centralization and the accumulation of power, which can pose a threat to open markets and democratic systems.

Promises and Drawbacks of Blockchain Technology

Blockchain technology is often cited as a solution to this ever-increasing consolidation of the Internet in the hands of a few individual players. By design, blockchains provide a decentralized, tamper-proof, and transparent way to store and exchange data. This technology provides an alternative to centralized control of data, whether in the financial sector, the technology industry, or other areas where data is a critical asset. Because it distributes control across the network, blockchain technology avoids data silos and single points of failure, improving data security, availability, and network reliability compared to traditional data infrastructures.

However, blockchain technology has faced significant criticism for its high energy consumption. For example, Bitcoin, the most prominent blockchain network, has become one of the largest energy consumers in the world, consuming approximately 36 percent of Germany's electricity. This has led to questions about the sustainability of blockchain technology. It is important to understand that such high energy consumption is a feature of the Proof of Work (PoW) consensus mechanism used by Bitcoin, which creates trust in decentralized transactions by requiring a significant investment in the form of electricity. However, there are energy-efficient alternatives to the PoW consensus mechanism, such as Proof of Stake (PoS). The smart contract platform Ethereum moved its consensus mechanism from PoW to this non-PoW consensus mechanism in September 2022, an event also known as "The Merge". As a result, the network's electricity consumption was reduced by 99.998 percent, proving that extreme electricity consumption is not an inherent feature of blockchain technology (Crypto Carbon Ratings Institute 2022c). This means that electricity consumption can be reduced through conscious network design.

Addressing Gaps in the Literature

Despite the benefits of these developments, there are gaps when it comes to leveraging new design decisions to reduce electricity consumption. The existing academic and practitioner literature provides some guidance on how to design to reduce a blockchain network's electricity consumption, but it focuses primarily on the decision between PoW and non-PoW blockchains. As such, it does not consider other design options that may impact a network's electricity consumption. In addition, the literature does not provide the tools needed to analyze the use case for which the infrastructure will be used.

Our study aims to fill these gaps in the literature. We begin by identifying the requirements of a use case for its architecture, including but not limited to reduced electricity consumption. Based on our findings, we provide a set of guiding questions to derive these requirements and develop a design guide for an electricity-efficient blockchain network that meets the requirements of a given use case.

Identifying the Use Case's Data Infrastructure Requirements (Chapter 2.1)

The study begins by deriving the data requirements of the use case, which leads to the identification of five essential infrastructure requirements that define the minimum properties the infrastructure must meet to be suitable for a use case. First, it must provide a level of confidentiality that meets the specific needs of the user and ensures protection from unauthorized data access. Second, the infrastructure must maintain a specified level of integrity to prevent unauthorized modification or deletion of data. Third, it must ensure a defined level of data availability to ensure that systems are accessible and reliable. Fourth, the infrastructure should provide the necessary performance to ensure efficient and timely processing and delivery of data to ensure the seamless operation of the use case. Finally, the infrastructure should be designed to minimize environmental impact, a factor we address in this study by focusing on reducing power consumption. While these requirements represent the minimum criteria, exceeding them – for example, by providing higher levels of integrity than required by the use case – can be beneficial, provided all other requirements are met.

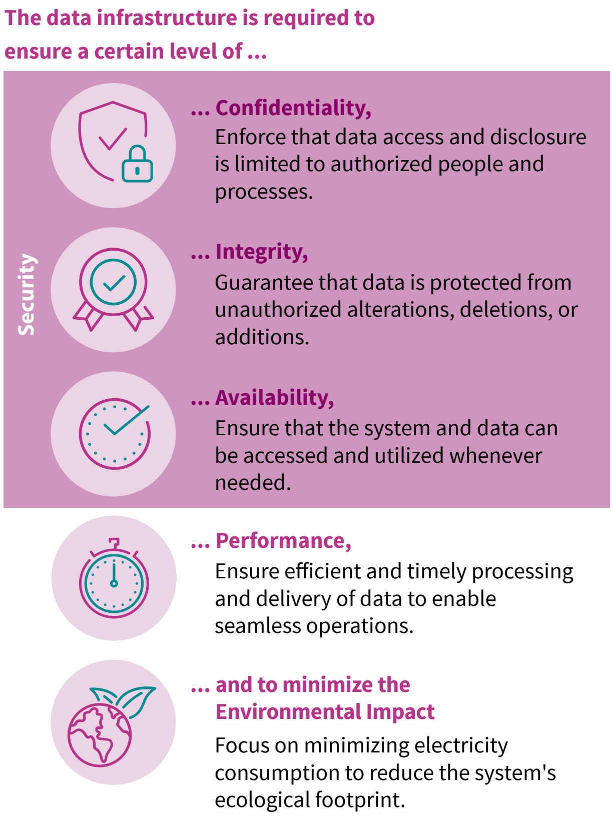

The data infrastructure is required to ensure a certain level of ...

Security

... Confidentiality,
Enforce that data access and disclosure is limited to authorized people and processes.

... Integrity,
Guarantee that data is protected from unauthorized alterations, deletions, or additions.

... Availability,
Ensure that the system and data can be accessed and utilized whenever needed.

... Performance,
Ensure efficient and timely processing and delivery of data to enable seamless operations.

... and to minimize the Environmental Impact
Focus on minimizing electricity consumption to reduce the system's ecological footprint.

Figure 1: The five requirements for the data infrastructure

Guide to Designing an Electricity-Efficient Network (Chapter 4.2)

The guide is divided into two stages: The first stage focuses on thoroughly analyzing the application for which the data infrastructure will be used, and the second delves into the network design:

- Stage 1 focuses on a thorough analysis of the application for which the data infrastructure will be used. We support this process with questions tailor-made to cover the fundamental requirements and boundary conditions of a use case for its data infrastructure in a blockchain-based solution.

- Stage 2 delves into the network design, focusing on a permissioned blockchain network. The multi-step process includes verifying the suitability of a blockchain-based network, selecting the appropriate blockchain type and associated platform, and finally designing the permissioned network. The previously established requirements and constraints support the evaluation of different design options, help to understand their influence on the properties, and ensure that the final design provides an appropriate data infrastructure for the use case.

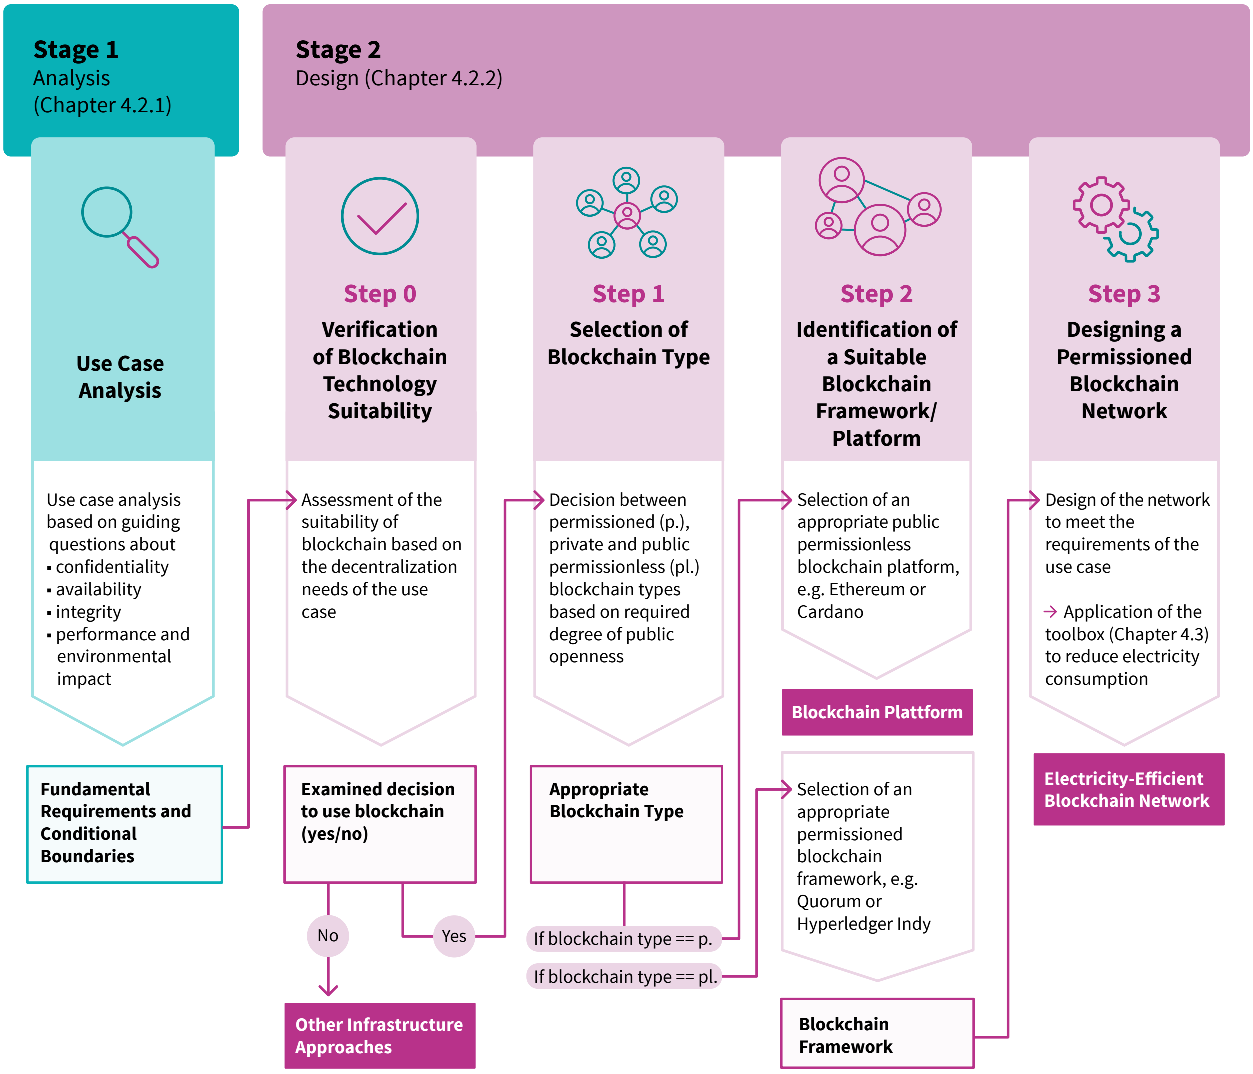
Stage 1
Analysis
(Chapter 4.2.1)
Stage 2
Design (Chapter 4.2.2)
Use Case Analysis
Step 0
Verification of Blockchain Technology Suitability
Step 1
Selection of Blockchain Type
Step 2
Identification of a Suitable Blockchain Framework/ Platform
Step 3
Designing a Permissioned Blockchain Network
Use case analysis based on guiding questions about
▪ confidentiality
▪ availability
▪ integrity
▪ performance and environmental impact
Assessment of the suitability of blockchain based on the decentralization needs of the use case
Decision between permissioned (p.), private and public permissionless (pl.) blockchain types based on required degree of public openness
Selection of an appropriate public permissionless blockchain platform, e.g. Ethereum or Cardano
Design of the network to meet the requirements of the use case
→ Application of the toolbox (Chapter 4.3) to reduce electricity consumption
Blockchain Plattform
Electricity-Efficient Blockchain Network
Fundamental Requirements and Conditional Boundaries
Examined decision to use blockchain (yes/no)
Appropriate Blockchain Type
Selection of an appropriate permissioned blockchain framework, e.g. Quorum or Hyperledger Indy
No
Yes
If blockchain type == p.
If blockchain type == pl.
Other Infrastructure Approaches
Blockchain Framework

Figure 2: The two stages of designing a blockchain network

Stage 1: Analysis (Chapter 4.2.1)

The first stage is to define the use case's data infrastructure requirements. These requirements must be considered from two perspectives. The first perspective examines the fundamental requirements of a use case, such as the expected transaction throughput and the required availability of the system, and aims to determine the characteristics that the final network design must provide. The second perspective establishes constraints that limit the available design options by eliminating impractical or inappropriate design choices. For example, the number of collaborating organizations may limit the maximum number of nodes and thus the size of the network. More information on how to analyze the use case and the guiding questions to help guide the analysis can be found in Chapter 4.3.

Stage 2: Design (Chapter 4.2.2)

Step 0: Verification of Blockchain Technology Suitability

Before proceeding with the design of a blockchain network, it is critical to ensure that blockchain technology is the right fit for the use case by considering both the advantages and disadvantages of decentralization compared to a centralized infrastructure. While decentralization offers benefits such as increased transparency, immutability, and trust, it also introduces complexity and operational challenges, for example, higher electricity consumption, since the operation of the data infrastructure is distributed along multiple nodes.

Step 1: Selection of Blockchain Type

The next stage is determining which blockchain type, permissioned or permissionless, is most appropriate for the use case. Again, the model developed by Hunhevicz and Hall (2020) provides a valuable starting point, focusing primarily on whether all participants are known and the degree of needed audibility, especially by public transparency of all transactions. In addition, we highlight the decision model proposed by Belotti et al. (2019), which also considers the trade-offs between the different properties of a blockchain network.

Step 2: Identification of a Suitable Blockchain Framework

When considering the option of a public permissionless blockchain network, it is important to note the changing landscape regarding Proof-of-Work (PoW) networks. Ethereum's switch from PoW to Proof-of-Stake (PoS) has raised questions about the relevance of PoW networks for inter-organizational data exchange. PoW networks consume significant amounts of electricity for consensus, which conflicts with minimizing environmental impact. PoS-based networks supporting smart contracts offer a more reasonable choice for public networks. It is essential to review relevant literature and recent developments, such as studies by Gräbe et al. (2020), Kubler et al. (2023), and Dena (2019), to make an informed decision on the blockchain type.

If a permissionless blockchain network is selected, the next step is to identify the most suitable one. Use existing networks with necessary security and credibility. Refer to Gräbe et al. (2020), Kubler et al. (2023), and Dena (2019) for overviews of relevant factors and platforms. For permissioned blockchain networks, choose a platform that meets specific use case needs. When considering permissioned networks, note that they offer the flexibility to design the network specifically for the needs of the use case, while maintaining an electricity-efficient design. Several frameworks are available for this type of blockchain, such as Quorum, Hyperledger, and Corda, which offer unique features and functionality, each with its advantages and disadvantages (Capocasale et al., 2023).

Step 3: Designing a permissioned blockchain network
The design of a permissioned blockchain network is a critical step in creating a network that effectively meets the specific requirements of the use case. This phase involves making thoughtful design choices to ensure that the network has the desired properties:

- Ensuring the appropriate level of **data confidentiality** is a vital consideration in the network design process. Blockchains are inherently transparent, making confidentiality a complex challenge. This challenge can be addressed through the employment of design option that limit data access to authorized participants by using a permissioned network, leverage private channel capabilities in permissioned networks, or employ techniques such as data encryption or zero-knowledge proofs to obfuscate data while maintaining its confidentiality.

- Maintaining **data integrity** requires an appropriate consensus mechanism tailored to the network participants and security requirements. In permissioned blockchains, consensus mechanisms based on Proof-of-Authority (PoA) are commonly used, as there is no need for Sybil resistance due to the participants being known. The flexiblity of PoA allows for design choices such as assigning voting rights to highly trsted participants or balancing crash fault tolerance and Byzantine fault tolerance.

- Ensuring the **availability of data and services** is achieved through careful design of the network structure. Decentralization is essential in achieving availability by distributing functionality across multiple nodes, reducing dependency on a single node, and avoiding single points of failure. Specifically, the number of nodes should be set appropriately to achieve the desired reliability. In addition, hosting nodes with different providers in different geographies reduces risk and increases the resilience of the network.

- **Performance** includes network throughput and latency. Adjustments to block sizes, block times, and network latency can directly impact network performance. However, balancing performance and other characteristics is essential, as excessive settings can compromise reliability and other network characteristics. Transaction complexity must also be considered to avoid overloading network nodes with redundant computations.

- **Minimizing environmental impact** requires design choices targeted at reducing the consumption of resources such as electricity and computing hardware. This can be done by avoiding oversizing the network and making conscious decisions in favor of data centers that are committed to electricity efficiency, electricity-efficient hardware, and sustainable network infrastructure. Our toolbox provides help to identify the appropriate design choices.

Toolbox for designing electricity-efficient blockchain networks (Chapter 4.3)

Our toolbox provides a comprehensive set of tools for reducing electricity consumption in non-PoW networks. By using these tools, network designers can adapt their systems to meet the requirements of electricity-efficient use cases. To address the challenge of electricity optimization, we have mapped the tools to the associated main trade-offs, allowing network designers to assess their applicability and ensure comprehensive analysis.

The integrity of a blockchain network, which involves immutable data storage, relies on various cryptographic techniques and the associated consensus mechanism, which requires communication between all participants. To reduce electricity consumption, our tools aim to reduce communication and computation complexity: **While the rate of fault tolerance** and **type of fault tolerance** design choices are related to the consensus mechanism, the introduction of execution sharding involves dividing the network's consensus process into several separate shards.

Data and system availability in a blockchain network is primarily facilitated by decentralization. The following tools reduce the degree of decentralization in various aspects and thus also minimize redundant computation. One approach is to decrease the **number of nodes**, directly reducing electricity consumption. **Serverless blockchains** are another form of centralization, offering high availability but introducing potential outage risks associated with these providers. In addition, **rollups** and data **sharding** introduce a degree of centralization within subsets of the network. Rollups consolidate the processing of specific transactions to a single node operator, while data sharding restricts data storage to a group of nodes.

The performance of a non-PoW network is closely tied to the computational load that each node must handle, which directly affects its electricity consumption. The toolbox provides tools to recalibrate the maximum throughput by adjusting **block size** and **block time**, which directly correlates to a nearly linear decrease in computing and storage utilization. Similarly, minimizing **transaction complexity** directly affects the computation a node must perform.

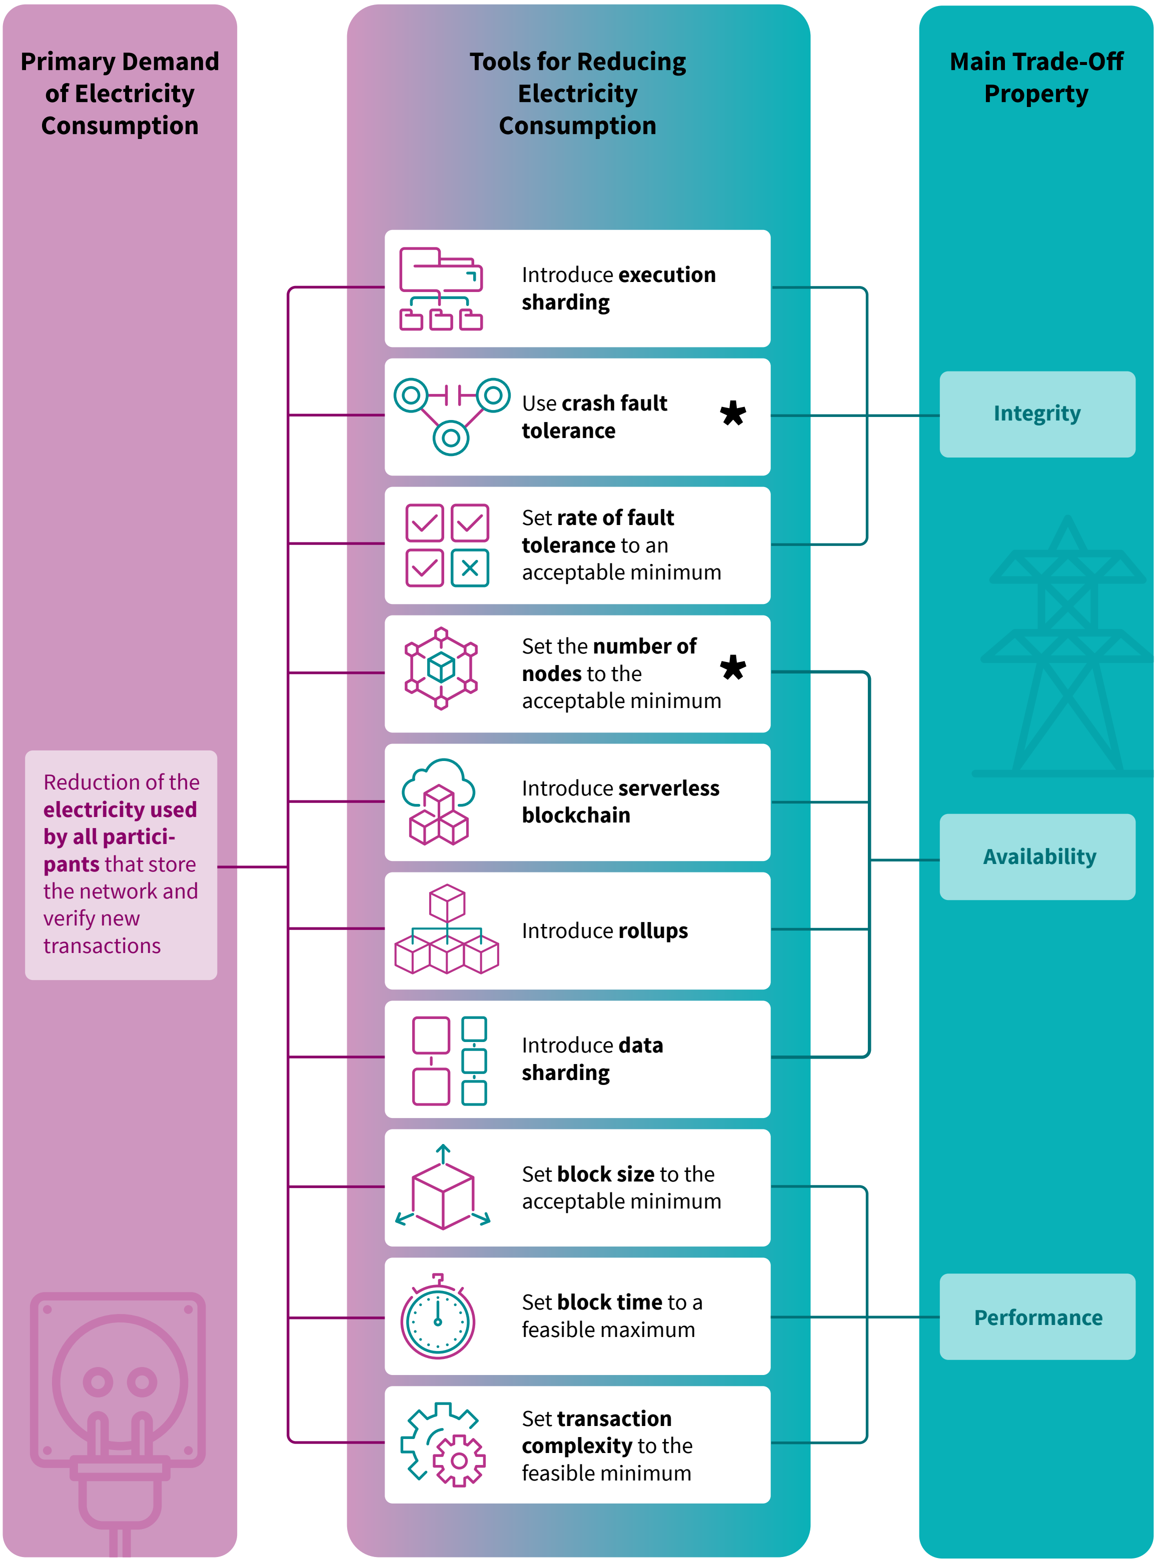
Primary Demand of Electricity Consumption
Tools for Reducing Electricity Consumption
Main Trade-Off Property
Reduction of the electricity used by all participants that store the network and verify new transactions
Introduce execution sharding
Use crash fault tolerance ★
Set rate of fault tolerance to an acceptable minimum
Set the number of nodes to the acceptable minimum ★
Introduce serverless blockchain
Introduce rollups
Introduce data sharding
Set block size to the acceptable minimum
Set block time to a feasible maximum
Set transaction complexity to the feasible minimum
Integrity
Availability
Performance

★ The asterisk marks those design options, which can only be used in a permissioned network

Figure 3: Tools for designing an electricity-efficient non-PoW network as identified by the study

Recommended actions

Based on the results of the study, we present several suggestions for different stakeholders to promote the electricity efficiency and sustainability of blockchain technology:

- As our study does not cover all aspects of blockchain technology electricity consumption, we encourage **researchers** to explore these areas further. Specifically, they could assess the potential electricity savings of the design tools we identified or develop new methodologies to improve the electricity efficiency of a blockchain. We also encourage the development of new frameworks to compare different forms of data infrastructure, allowing for a more comprehensive view of their relative efficiencies. Finally, cross-disciplinary research could help bring different perspectives on the electricity efficiency, potential uses, and benefits of blockchain technologies and determine the circumstances under which additional usage may be justified.

- **Standards organizations and policy makers** could use the results of this research to advance standardization, benchmarking and regulation for blockchain technology. This could include metrics for the electricity consumption or carbon emissions associated with different blockchains, allowing companies or organizations using the technology to calculate their carbon footprint. This work can also be used to evaluate blockchain applications, especially in comparison to alternative data infrastructures.

- **Blockchain framework developers** should also consider the electricity consumption aspect of their software. In doing so, they can incorporate features directly aimed at reducing the amount of electricity consumed. Furthermore, they could contribute to the overall sustainability of blockchain technology by providing practical guidelines for electricity-efficient designs and creating tools for users to monitor the network's power consumption.

- **Both users and operators** of a blockchain-based network should consider various aspects of environmental impact, such as electricity consumption or carbon emissions when choosing a network. Our study allows for such conscious network design. Our study shows that conscious network design can reduce these impacts while ensuring suitability for specific use cases. In this way, users and operators can take advantage of the decentralized infrastructure while enhancing the environmental sustainability of their operations. We also suggest that users demand transparency from network operators about their electricity consumption. This would not only enable an informed choice of networks but also incentivize developers to consider electricity consumption as a priority.

The actions recommended above should be taken collaboratively by the different stakeholders, rather than individually. Further research will certainly fill any remaining knowledge gaps. However, researchers will need to consider the demands of standards organizations and policy makers. Moreover, blockchain framework developers as well as the operators and users of the resulting networks have a unique ability to deliver invaluable insights into the applicability, limitations and remaining shortcomings of tools and regulations for the energy-efficiency of blockchains. We, the German Energy-Agency, hereby encourage all stakeholders who have the power to influence the electricity consumption of blockchains in any way to participate in an 'alliance of the willing' and to join in a coordinated effort to maximize the sustainability of blockchain technology. Such an alliance requires an appropriate ecosystem connecting the different stakeholders, which we would gladly support by acting as an intermediary and organizing the required formats and forums.

1. Introduction

Digitalization – A megatrend

The process of digitization, or digitalization, is one of the most significant technological, social, economic, and political transformations of our time. This megatrend has far-reaching implications for all areas of life. Digital technologies are now essential to the vast majority of industries, as well as to our coexistence in a global society. They are an unprecedented tool for solving both emerging and long-standing challenges. However, these advances come with drawbacks, including psychological health risks associated with high social media use as well as further centralization of the digital economy, posing potential threats to market accessibility and democratic systems. Besides many beneficial implications, digitalization also has less favorable effects. Studies have shown that high social media consumption can impair emotional health and increase the risk of depression, anxiety, loneliness, self-harm, and even suicidal thoughts[4].

Furthermore, the digital economy is characterized by increasing centralization and accumulation of power, with a small number of large companies controlling a significant share of the market. This concentration of power in the hands of a few organizations creates a high level of dependency on these entities, which not only poses a threat to market accessibility and individuals' data sovereignty, but also creates a single point of failure. In such a scenario, the malfunction or collapse of one of these organizations can have far-reaching effects.

This trend toward centralization underscores the need for a more decentralized data infrastructure that is less dependent on these central actors, but as we transition to these decentralized systems, we must consider their rising environmental impact. In the last ten years the data center capacity in Germany increased by 90 percent[5]. The annual growth rate of mobile data volume transmitted was 23 percent in 2022[6] and the global number of devices connected to the internet (IoT-devices) is forecasted to increase by about 75 percent from 16.4 billion to 29.7 billion between 2023 and 2027[7]. This surge of computational power, data transmission, and interconnected devices is accompanied by an increased consumption of resources such as rare earths, steel, copper, water, and energy for manufacturing as well as electricity for the operation of the digital infrastructure.

Saving energy to contain climate change

On December 12, 2015, 195 nations committed themselves to mitigating climate change and limiting the global temperature rise to preferably 1.5°C as compared to the pre-industrial era. In order to reach this goal, the amount of climate-damaging gases emitted must not be larger than what is absorbed by carbon sinks. One of the major measures to decrease greenhouse gas emissions is the decarbonization of the energy sector or as it is called in Germany: The Energiewende ('energy turnaround' or 'energy transition').

An essential key to a successful energy transition is replacing carbon-intensive fossil energy sources with renewable energies like hydroelectric, solar, wind, and geothermal power. This replacement requires extensive investments and restructuring of the energy system. Moreover, renewable energy potentials are limited. Therefore, the second pillar of the energy transition is reducing net energy consumption, which is either pursued by forgoing usage of energy (sufficiency) or its improved usage (efficiency). Thus, a high share of energy free from greenhouse gas emissions in the energy mix can be achieved and the economic burden reduced, all according to the principle: The cheapest kind of energy is the one we do not have to produce.

The European Union aims to reduce primary energy consumption by more than 40 percent compared to its 2007 projections. Achieving this ambitious goal requires efforts across all industries and areas of life. The options for reducing primary energy consumption range from the modernization of power plants to innovative industrial processes to improved building insulations and more energy-efficient appliances. Likewise, there are many ways to reduce the amount of energy needed to operate digital infrastructures.

One starting point would be to increase the energy efficiency of the digital infrastructure itself. For example, the operation of server infrastructure generates large amounts of waste heat that can be used for district heating. Another approach is to influence the energy consumption of digitalization through the conscious design of digital applications. For instance, personal computers and mobile phones have long provided the option to reduce the screen brightness or to switch to an energy-saving mode, which terminates superfluous processes and throttles the processor power. Additionally, decisions on data administration, data amounts and data redundancy directly influence the extent of the required hardware and its energy consumption.

Blockchain and the debate on its energy consumption

Blockchain technology was conceived to enable decentralized, tamper-proof, and transparent data storage. By distributing data, the technology avoids data silos and single points of failure and, thus, ought to improve data security, data availability and the reliability of the network in comparison to traditional networks. Generally, the operation of databases or the execution of digital financial transactions is the responsibility of central entities such as banks, companies or public institutions, which network users are reliant on. Establishing the integrity of a digital service is key in order to win customers. One possibility to ensure the trustworthiness of a system, is for system owners to install independent control bodies, which, for example, can be formed through democratic processes. However, in the absence of democratic structures or the presence of general mistrust, blockchains offer an alternative to centralized solutions. In the case of blockchains, consensus mechanisms establish the trust in the

[4] https://www.helpguide.org/articles/mental-health/social-media-and-mental-health.htm, accessed 08.08.2023

[5] https://www.bitkom.org/sites/main/files/2023-05/BitkomStudieRechenzentreninDeutschland2023.pdf, accessed 08.08.2023

[6] https://www.bundesnetzagentur.de/SharedDocs/Pressemitteilungen/DE/2023/20230602_JB_TK2022.html, accessed 08.08.2023

[7] https://iot-analytics.com/number-connected-iot-devices, accessed 08.08.2023

correctness of the stored data. Therefore, central entities and even the mutual knowledge and trust of the network participants can be foregone.

The first use case for the blockchain-technology – the cryptocurrency Bitcoin – was launched in 2008. Ever since, it has become a major object of speculation, but also one of the largest energy consumers in the world. Its annual electricity consumption amounts to 205 TWh, equivalent to about 36 percent of Germany's electricity consumption (Cambridge Centre for Alternative Finance 2022). Criticism of the electricity consumption of the Bitcoin blockchain is often formulated as a general criticism of the electricity consumption of all blockchains. High resource costs are an inherent feature of the Proof-of-Work (PoW) consensus mechanism used for the Bitcoin cryptocurrency. Trust into transactions executed decentrally is created by the investment of large amounts of electricity. However, more electricity-efficient alternatives to PoW-consensus mechanisms exist, which utilize another scarce resource instead of electricity in order to ensure consensus. In the case of the Ethereum blockchain, a certain amount of the native cryptocurrency Ether has to be deposited in order to participate in the block creation and is retained in case of failures or malicious actions.

The smart contract platform Ethereum switched its consensus mechanism from PoW to the non-PoW consensus mechanism Proof of Stake (PoS) only in September 2022, an event also known as 'The Merge.' As a result, the blockchain's electricity consumption was reduced by 99.998 percent (Crypto Carbon Ratings Institute 2022c). Apparently, the choice of the consensus mechanism substantially influences the electricity consumption of a blockchain. However, non-PoW blockchains have also been criticized for their increased electricity consumption in comparison to centralized systems due to the redundant data storage and execution of transactions. In the last several years, various methods were developed in order to reduce e.g. redundancy, transaction complexity or the size of the hardware required. The resulting tools were initially conceived in order to increase the blockchain's throughput but can also be used to decrease their electricity consumption. Unfortunately, many of these methods lead to conflicts between reducing the environmental impact of a non-PoW blockchain and other aspects, like the security and performance of the system. Consequently, blockchains need to be designed consciously, taking into consideration the resulting trade-offs in order to make sure that the electricity consumption is minimized while the other requirements of the use case on the network are also satisfied.

Existing guidance for blockchain design

A widely accepted description of the conflict between different aspects of blockchain technology was introduced by Vitalik Buterin[8], the founder of Ethereum, in 2017 in a blog article that covered the concept of sharding. According to Buterin, the idea of the "blockchain trilemma" encompasses the aspects of decentralization, security, and scalability, which are described as binary – an aspect is either fulfilled or not. The central claim is that blockchain designs could only meet two of the three requirements. For example, both Bitcoin and Ethereum were considered to be decentralized and secure, but not scalable (more nodes did not mean higher throughput) and having low throughput, typically only a few transactions per second.

Despite its lasting popularity, the Blockchain Trilemma is not applicable to the use-case-specific design of blockchains. Due to new methods meant to improve the scalability of blockchains, the design process of blockchains is no longer bound to binary decisions between the fulfillment or non-fulfillment of different aspects but takes place on a much more fine-grained spectrum. The Blockchain Trilemma was not conceived to portray this field of tension, in which network solutions can be placed today. Moreover, the classic trilemma only considers the integrity or rather the resistance of a network against malicious takeovers. It is technology-focused and neglects other security aspects important for use cases, like the liveness of a system and the availability of its data. Lastly, no consideration is given to a system's electricity consumption. **In order to allow for a conscious, energy-efficient design of blockchains, the requirements of any given use case on its network solution have to be revised and their interconnections must be described in a way that accounts for the various design parameters of blockchains.**

The academic and practitioner literature offers some guidance for the conscious design of blockchains. Hunhevicz and Hall (2020), Wüst and Gervais (2018) and Beck (2019) provide a suitability check of the blockchain technology for a given use case and offer support in choosing the right blockchain type (permissioned or permissionless, private or public). Examples of publications explicitly concerned with the electricity consumptions of blockchains are Ramesohl et. al. (2021), EU Blockchain Observatory (2021) and Reetz (2019). They mainly focus on the decision between PoW and non-PoW-blockchains and only Rameseohl et. al. (2021) briefly discusses other design options suitable for reducing electricity consumption, such as sharding or side chains. What the publications have in common is that they demand a use-case-specific decision for or against the usage of blockchain technology, as well as the blockchain's ultimate design. However, they do not offer guidance for the analysis of use cases or tangible application examples. **To sum up, the literature lacks an in-depth analysis of the various options for the energy-efficient design of blockchains that accounts for the entirety of its requirements and their interdependencies beyond the choice of the consensus mechanism. Additionally, the existing literature is not sufficiently use-case-oriented. A use-case-centered approach could greatly facilitate the application of the guidance provided.**

[8] https://vitalik.ca/general/2017/12/31/sharding_faq.html, accessed 08.08.2023

1.1 Objectives

The central goal of this publication is to close the existing gap in the literature and thus expand the available guidance for the energy-efficient design of blockchains. Moreover, we aim to enhance the understanding of a blockchain's design parameters, its electricity consumption, its further characteristics, and their interdependencies. As a result, the study will enable more efficient use of blockchain technology and all its potential while minimizing its electricity consumption.

To achieve these goals, the first question this publication ought to answer is:

- What are the possible requirements of a use case on its blockchain, besides minimal electricity consumption?

To extend the existing guidance, the core element of this study is to develop a guideline for the electricity-efficient design of blockchains that takes into account the requirements identified by the first question and also answers the following two questions:

- How can the requirements of a specific use case on its blockchain and the boundary conditions for its design be determined?

- How can a blockchain be designed to satisfy its use case's requirements and minimize its electricity consumption?

Due to the potentially higher electricity consumption of blockchains in comparison to centralized databases, the deployment of this technology has to be use-case-specific and well thought out. The assessment and the comparison of the usefulness and the environmental impact of blockchains and their alternatives is not part of this publication. However, the guidance provided here will enable the electricity-efficient design of blockchains and thus the conscious, well-founded and appropriate deployment of network solutions in a way that accounts for its use case's requirements and its social, ecological, economic and environmental consequences.

1.2 Research Methodology

In order to answer the questions described in the last chapter, we applied a three-part approach:

- a systematic literature review,

- expert interviews, and

- expert workshops.

By combining these methods, we aimed to gain comprehensive insights into the topic, identify knowledge gaps (cf. Webster and Watson 2002), and develop a guidance framework for designing an electricity-efficient data infrastructure based on blockchain technology.

Literature Review

We conducted a systematic literature review (SLR), following the recommendations of Kitchenham et al. (2009), to gather knowledge on the environmental impact of blockchain networks. We searched six well-established databases (AISeL, IEEE Xplore, Nature, ScienceDirect, SpringerLink, and Web of Science) and the ArXiv and SSRN preprint databases using multiple synonyms for the keywords "blockchain", "cryptocurrency", and "sustainability". This search yielded approximately 7,000 articles, which we screened for relevance based on titles, and, finally on a full text basis. We considered publications that addressed the electricity consumption of blockchain technology or provided approaches to reducing such consumption and also included papers more generally aimed at improving blockchain efficiency. After applying our exclusion criteria, 95 relevant publications were identified. We conducted a snowball search (forward and backward searches), including gray literature from practitioners, which resulted in 10 additional sources. In total, we included 105 publications in the final review.

Expert Interviews

Building on the findings of the systematic literature review, we conducted expert interviews with practitioners and academics, all of whom have published in the field of blockchain technology. These interviews allowed us to delve deeper into specific issues, gain valuable insights, and validate our understanding of the topic. In addition, the expert opinions helped fill gaps we had identified in the literature and provided us with practical knowledge to support our findings.

Workshops with Experts

To further refine our understanding and generate actionable solutions, we organized two workshops with experts from diverse backgrounds, including blockchain developers, researchers, and consultants. In the first workshop, we discussed the myths surrounding the electricity consumption of blockchain technology to lay the groundwork for developing our framework. In the second workshop, we collaboratively explored approaches to designing a decentralized data infrastructure and validated the guidelines presented below.

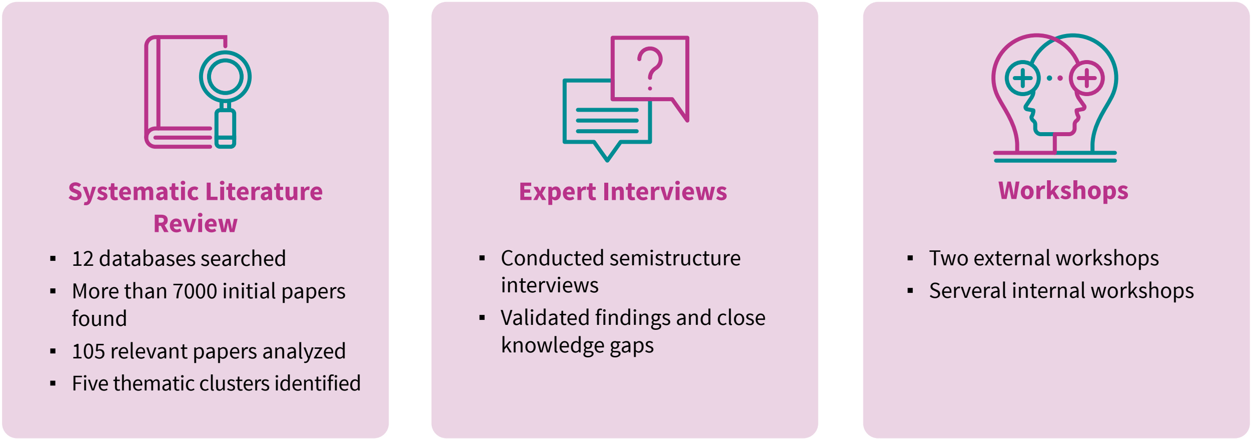

Systematic Literature Review

- 12 databases searched
- More than 7000 initial papers found
- 105 relevant papers analyzed
- Five thematic clusters identified

Expert Interviews

- Conducted semistructure interviews
- Validated findings and close knowledge gaps

Workshops

- Two external workshops
- Serveral internal workshops

Figure 4: The three methodological pillars of the study

1.3 Study Structure

In Chapter 2, we offer some theoretical background on the topics of data governance and blockchain, which provides the theoretical foundation for the rest of the study. Then, in Chapter 3, we present the current state of knowledge on blockchain electricity consumption through our literature review and then provide a visual representation of the parameters that influence the electricity consumption of a blockchain network. In Chapter 4, we present our guide and toolbox for designing electricity-efficient blockchain networks. Chapter 5 presents case studies in which we apply our toolbox, and finally, Chapter 6 concludes the study.

2. Theoretical Background

This chapter provides an in-depth examination of the essential concepts involved in designing a blockchain-based data infrastructure specifically tailored to the needs of a use case. We begin by deriving the data requirements of the use case, providing an understanding of what properties the underlying data infrastructure must provide. We then explore the fundamental concepts of blockchain technology, providing an understanding of how the technology can be applied to meet the use case's requirements.

2.1 Data Infrastructure

The phenomenon of digitalization has increasingly highlighted the importance of data management and the design and implementation of data infrastructures. This focus not only recognizes the importance of data, but also underscores the dynamic evolution of different types of data infrastructures that must respond to the development of new use cases and their ongoing demands on the data infrastructure. In the following, we explore this interrelationship between use cases, data, and ultimately their data infrastructure requirements.

Data as a Crucial Component of Use Case

Understanding the importance of data in various use cases requires recognizing it as a strategic asset. Because data does not have a value of its own, organizations need to focus on the value it adds in achieving the specific goals of the use case. For example, data can drive functional requirements, enable data-driven actions, or support informed decision making within a particular use case. A well-designed data infrastructure is critical for this (Khatri and Brown 2010), as it ensures that the data can be used as intended by the use case and thus helps the use case realize its full potential.

Use Case and Data Infrastructure

The data infrastructure, consisting of hardware, software, and network layers, is the foundation of any digital use case. It acts as the backbone that allows for the processing, storage, and management of data specific to the functional needs of the use case. Each layer impacts the overall functionality, and all elements must be considered as a complete unit. For example, implementing redundant network infrastructure increases network resiliency, while deploying high-performance hardware can improve the amount of data processed.

The relationship between the data infrastructure and the use case is bidirectional, with each enhancing the capabilities of the other. The specific requirements of the use case drive the essential properties and functionalities required of the data infrastructure, enabling new capabilities and opportunities for generating and consuming data within the use case. At the same time, as the use case evolves to include higher volumes of data or the need for advanced collaboration capabilities, it introduces new requirements and considerations that must be addressed in the design and implementation of the data infrastructure (Weill 2004). Therefore, the alignment between data infrastructure and the use case is critical to ensuring successful implementation and achieving the intended results. By understanding and addressing the specific requirements of the use case, organizations can establish an efficient way to handle data and thus meet the data's requirements (Abraham et al. 2019).

Designing Data Infrastructure Based on Use Case Requirement

One approach to deriving data infrastructure requirements is to analyze the requirements of the data itself, as outlined by Panian (2010). These requirements include various aspects such as accessibility, availability, consistency, verifiability, and security. Satisfying these aspects ensures data quality, which makes the data consistent with its intended use and thus ensures the functionality of the data for the use case.

A specific level of **information security** must be met to ensure that an organization's data is stored and processed securely. A common framework for deriving IT security requirements is the CIA Triad, which effectively addresses the primary security aspects related to data handling and the underlying IT systems. The security requirements for the IT infrastructure are therefore structured around three core principles: Confidentiality, Integrity, and Availability of the data (Samonas and Coss 2014). This straightforward approach makes the model universally applicable and suitable for various industries and organizational contexts. Its flexibility allows for easy adaptation to new situations and technologies. As a result, this framework provides a solid foundation for intuitively deriving the properties an infrastructure must possess to meet the specific security requirements dictated by the use case.

Based on the principles, three use case requirements can be defined for the IT infrastructure. Depending on the data, the use case defines a **specific level for each of these as a requirement:**

- **Confidentiality** requires that data is protected from unauthorized access. Depending on the use case, the required level can range from protecting general information to restricting sensitive data.
- **Integrity** refers to the required protection of data against unauthorized modification and deletion, and ensuring that information is complete and accurate. The required level can vary from allowing minor inaccuracies to requiring absolute accuracy.
- **Availability** focuses on accessibility and reliability of the infrastructure so that users can access data and services when needed. Levels of availability can be derived, ranging from non-time-sensitive access to immediate, critical access to data and services.

In addition to meeting IT security requirements, the data infrastructure must provide a certain level of **performance**. It must be able to consistently process and respond to requests regardless of system load, as timely and reliable access to data stored in the infrastructure is critical to the proper execution of processes. Performance can be described by two metrics: throughput, which describes the transactions performed within a given period and may vary based on load peaks (e.g., more transactions during working hours), and latency, which indicates the time between requesting and processing a task (Kannengießer et al. 2021; Sedlmeir et al. 2021a)[9]. Specifically, time-critical use cases, such as real-time data processing, can be affected by latency and may therefore have special requirements in this area. For instance, a healthcare application that involves real-time patient monitoring requires low latency to ensure timely updates and alerts. A high-throughput e-commerce platform, on the other hand, needs to handle a large number of transactions during peak shopping periods without experiencing performance bottlenecks. In addition, the potential growth of requirements over time must be taken into account. Therefore, the performance should meet today's requirements and provide enough flexibility to allow the organization to respond adaptively to future situations (Abraham et al. 2019).

Minimizing the **environmental impact** of IT systems is the final requirement for designing an IT infrastructure. As more data is stored and processed, the need for hardware increases, as does the direct environmental impact. Data centers, which house servers and other integral components of the IT infrastructure, are substantial consumers of electricity. The International Energy Agency, in their 2020 report, disclosed that data centers and data transmission networks collectively devoured approximately 200 TWh of electricity. This figure accounts for around 1% of global electricity use. The environmental impact goes beyond the electricity needed to operate these systems. A comprehensive perspective should also include, among other things, the manufacturing process and recycling of hardware, as IT equipment production requires raw material extraction and processing. In addition, the ongoing technological advances are shortening the lifespan of used hardware, and improper recycling often results in e-waste, exacerbating the environmental impact. Although we recognize these multiple facets of environmental impact, the focus of this study is on the electricity consumption of the infrastructure, due to its immediate and measurable effect on operations.

A use case outlines five key requirements that a data infrastructure must meet (Figure 5). First, it must meet the required levels in **confidentiality, integrity,** and **availability**. Second, it should provide the performance to ensure efficient and timely processing and delivery of data to provide seamless operation of the use case. Third, the infrastructure should be designed to **minimize environmental impact**, a factor we address in this study by focusing on reducing of electricity consumption. While these requirements represent the minimum criteria, exceeding them – for example, by providing higher levels of integrity than required by the use case – can be beneficial, provided all other requirements are met.

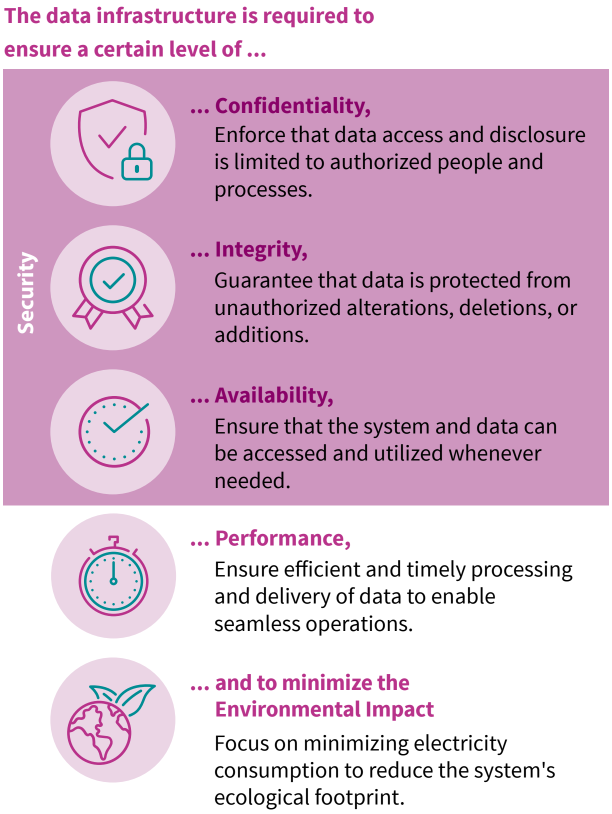
The data infrastructure is required to ensure a certain level of ...

... Confidentiality,
Enforce that data access and disclosure is limited to authorized people and processes.

... Integrity,
Guarantee that data is protected from unauthorized alterations, deletions, or additions.

... Availability,
Ensure that the system and data can be accessed and utilized whenever needed.

... Performance,
Ensure efficient and timely processing and delivery of data to enable seamless operations.

... and to minimize the Environmental Impact
Focus on minimizing electricity consumption to reduce the system's ecological footprint.

Figure 5: The five requirements for a data infrastructure

Considering Centralization and Decentralization in Data Infrastructure

As digitalization progresses, use case requirements evolve and become more complex. Especially in scenarios that involve collaboration across multiple organizations, the structure of the data infrastructure must be carefully considered. The decision to implement a centralized or distributed approach can have a profound impact on the final design of the use cases, as well as the way in which organizations collaborate.

When multiple autonomous organizations work together and share common goals, collaboration can be mutually beneficial: Organizations can learn from each other, pool resources and use them efficiently, and ultimately improve or develop products and services (van den Broek and van Veenstra 2015). Depending on the use case, different forms of decentralization are possible. Especially in a more globalized world and with more interorganizational settings, new needs for interorganizational cooperation have emerged (van den Broek and van Veenstra 2015). This introduces a complex set of requirements, including data sharing and the need for a common underlying infrastructure. Therefore, the choice between a centralized or a distributed platform and

[9] In this study, we focus on latency because it provides an intuitive way to navigate the stages of the design process.

infrastructure then becomes a significant factor (Lee et al. 2018) The degree of decentralization within the data infrastructure can be tailored to specific needs and considerations, serving as the technical foundation for various use cases. In a centralized setting, a single entity or platform owner takes complete control and responsibility for the data infrastructure; in a distributed form, this role is shared among all stakeholders (Lee et al. 2018). Conversely, in the decentralized system, data is stored in a distributed manner, such as in a blockchain network. Some use cases may take a hybrid approach, integrating elements of both centralized and decentralized systems to meet specific data exchange requirements. Innovative technologies are emerging as part of the investigation into new forms of data infrastructure. Blockchain technology in particular stands out for its unique approach to ensuring decentralization while meeting the requirements of many different use cases.

2.2 Blockchain

Satoshi Nakamoto first proposed the technological concept of a blockchain as part of his vision for a decentralized payment system, famously known as Bitcoin (Nakamoto, 2008). Rather than rely on centralized intermediaries to facilitate transactions, the system uses technological and cryptographic measures to establish trust in a thoroughly decentralized network.

Foundations of Blockchain Technology

In a blockchain network, data is exchanged through transactions. The sender signs the transaction with their private key to ensure authenticity and immutability. These transactions are stored in blocks, where the number of transactions depends on the size of the transaction and the specified block size. Nodes within the network store a local copy of this chain and propagate transactions throughout the network. When a new block is created, the system chronologically adds it to the global ledger, which essentially acts as a distributed database across participating nodes. Each block in the ledger is linked to the previous one by a hash pointer (the black arrows in Figure 6). A hash can be thought of as a unique fingerprint of the block's contents. By including the previous hash as a hash pointer within a block, a chain is created that goes back to the first block in the ledger. If someone modifies the contents of a block after it is added to the ledger, the block's hash changes, breaking the chain because the modified block's hash no longer matches the hash stored in the next block. As a result, the older a block is, the more difficult it is to modify a transaction stored in it, since this would require changing the hash of all subsequent blocks in order for the modified transaction to be considered valid by the network. This principle ensures that the information stored in the ledger can be considered immutable and therefore tamper-proof (Tschorsch and Scheuermann 2016). In addition to propagating transactions, nodes also participate in a consensus mechanism. This mechanism ensures the overall integrity and consistency of the blockchain by finding agreement among network participants on new blocks, their transactions, and their order.

Participants in a Blockchain Network

Typically, there are multiple entities in a blockchain network that interact with the network in different ways and thus contribute to the functioning of the network in different ways:

- **Nodes** are vital components that store valid blocks transmitted within the network. It is essential to distinguish between full nodes, which make the entire ledger available, and light nodes, which store only a portion of the ledger. The more nodes participate, the more distributed; hence, the network becomes more decentralized. As the number of full nodes increases, so does data redundancy, ensuring that as long as there is at least one honest node that stores valid data, the entire ledger can be replicated from it.

- A subset of nodes, **block producers**, contributes to the consensus mechanism by acting, for example, as miners in PoW or validators in PoS. Their participation enhances the consensus mechanism's viability and increases the network's security by increasing the effort required to execute an attack successfully. Furthermore, when consensus participants are selected to create a block, they may receive block rewards and transaction fees.

- **Other participants** interact with the data stored on the blockchain via the nodes. The interactions can take the form of transactions, such as sending and receiving coins, executing smart contracts, or saving or receiving data stored on blockchains.

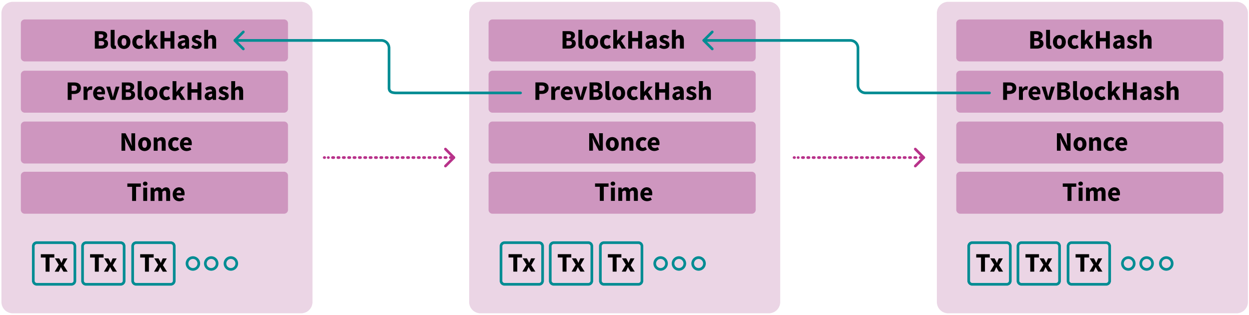

Figure 6: Simplified Blockchain (Tschorsch and Scheuermann 2016)

The Throughput of a Blockchain Network

The block size and the block time primarily determine the transaction throughput of a network. The block size indicates how many transactions fit into a block, and the block time is the average time it takes for a new block to be added to the ledger. The Bitcoin network, for example, has a fixed block size of 1 MB and aims to create a new block every ten minutes. On the other hand, Ethereum proposes a new block every 12 seconds and adjusts the network's throughput by varying the block size depending on the network demand. However, increasing the throughput by increasing the block size and block time has trade-offs. For example, processing more transactions in a shorter time requires more computing power and higher configured hardware for the nodes, leading to higher centralization of the nodes, as potential network participants are excluded due to the additional hardware requirements (Kannengießer et al. 2021).

Transaction Complexity and Smart Contracts

In a blockchain network, the complexity of a transaction – defined primarily by the computational effort required to execute it – can vary depending on several factors. Transactions interact with and modify the data stored on the blockchain, which can be new information or requests for change. The simplest form of transaction is typically the transfer of cryptocurrency between two parties. Smart contracts – self-executing contracts with the terms of the agreement written directly into the code – have made more sophisticated forms of data exchanges possible. These contracts automatically execute and enforce contract rules when predefined conditions are met, enabling transparent, secure, and decentralized program or contract execution without a central authority or intermediary. This increased complexity increases the computational load on nodes, degrades network performance, and – in a public permissionless network – raises transaction costs (Sedlmeir et al. 2022a).

Blockchain Types

Since the advent of Bitcoin, different blockchain types have emerged with varying degrees of decentralization. These can be classified into two dimensions, as proposed by Beck et al. (2018). The first dimension distinguishes whether participation in the consensus mechanism is open to everyone (permissionless) or permissioned, meaning that only certain and pre-selected entities are allowed to participate. The second dimension defines whether the transactions are accessible to the public or only to a specific group (see Figure 7).

Based on this distinction, we can identify three distinct types of blockchain designs. First, in a **permissionless** public blockchain, such as the one proposed by Nakamoto for Bitcoin, anyone can freely join without any barriers, participate in the network in all parts, and leave the network again without any barriers. However, this open design requires additional security measures, which we will discuss in the following.

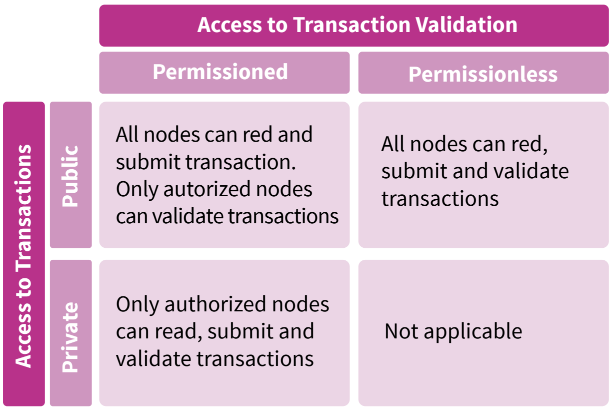

		Access to Transaction Validation	
		Permissioned	Permissionless
Access to Transactions	Public	All nodes can red and submit transaction. Only autorized nodes can validate transactions	All nodes can red, submit and validate transactions
	Private	Only authorized nodes can read, submit and validate transactions	Not applicable

Figure 7: Blockchain Types (Beck et al. 2018)

Depending on the use case, allowing only a select group of authorized entities to participate in authorizing valid nodes can be more appropriate in the form of a **permissioned blockchain**. This approach can lead to a more efficient network design, resulting in higher network performance, among other benefits. The access to transactions can be private, limited to a specific group, or public, depending on the need for audibility or third parties' intended use of the data. For example, Sovrin[10] and Ripple[11] are public **permissioned blockchains** where anyone can view and utilize the information stored on the chain, but only a pre-selected group of nodes has the authority to transmit and validate transactions. On the other hand, several blockchain frame-works, such as Hyperledger Fabric or Quorum, can serve as private **permissioned blockchains**, where both the permission to validate transactions and the ability to view information on the chain are strictly controlled and limited to a pre-selected group.

[10] www.sovrin.org

[11] www.ripple.com

Consensus Algorithm and Consensus Mechanism

The consensus algorithm plays a crucial role in maintaining the integrity of the data and transactions that are stored on the blockchain. It enforces specific criteria that new transactions must meet before they are accepted and incorporated into the ledger, effectively safeguarding the data's validity. For example, based on the algorithm, a sender's availability to complete a transaction can be validated by confirming that they have sufficient coins to complete the intended transfer. One of the defining aspects of a consensus mechanism is its fault tolerance – its capacity to function correctly even when certain parts of the networks malfunction or act maliciously. One key measure of this capacity is the **fault tolerance rate**, which refers to the maximum proportion of the network that can contradict the network without compromising the system's functionality or integrity. When the level of contradiction across the network exceeds this rate, the system's ability to reliably perform operations is compromised. Fault tolerance is categorized into two types: **Crash fault tolerance (CFT**) ensures that the network continues to function correctly even if a certain percentage of consensus participants stop operating due to network splitting or node crashes. The stronger security guarantee is provided by **Byzantine fault tolerance (BFT)** which ensures the system's functionality even if a certain percentage of malicious and malfunctioning nodes are part of the consensus.

In a decentralized network like Bitcoin or Ethereum, participation is open and therefore unpredictable. As such, consensus cannot rely on a simple vote distributed to each network participant, as the system would be vulnerable to a Sybil attack: In a **Sybil attack**, a single entity can flood the network with fake identities because there is no cost associated with joining the network, thus gaining disproportionate influence in the voting process (Douceur 2002). To mitigate the risk of Sybil attacks in this trustless environment, it is essential to implement a Sybil-resistant consensus mechanism within the network. Such a mechanism ensures that voting power is tied to a scarce, digitally verifiable resource that is difficult to generate. This requirement imposes a direct cost on participants in the consensus process. In permissioned networks, a mechanism to resist Sybil attacks is not required as all participants are known. This allows for a wider choice of consensus mechanisms and a more efficient way to archive consensus. The way consensus is reached is one of the critical design decisions that affects not only the security of the network, but also, among other things, the transaction throughput and electricity consumption of the network. The following section provides an overview of the most common consensus mechanisms, which are summarized in Table 1.

Proof of Work (PoW) is a consensus mechanism first introduced by Bitcoin and used by various public permissionless networks. Initially proposed as part of the Hashcash network designed to limit email spam (Back, 2002), the mechanism requires a node to allocate resources to solve a cryptographic problem posed by the protocol. The challenge of creating a block in PoW is to find a "nonce" – a random number that, when combined with the block's data, produces a hash that is below a certain target set by the network. This process is resource-intensive and requires nodes (called miners) to try numerous nonces until they find one that meets the criteria, which requires a significant investment of computing power. Conversely, validating a new block must be efficient and easy for other nodes to perform (Luu et al. 2015). Participation in the consensus mechanism is called mining, and the participating nodes are miners, who are compensated for participating in the consensus mechanism. Instead of having equal voting rights for each participant, miners must invest a scarce resource – in this case, computational power, i.e. hardware and electricity. The computing power is typically measured in hashes per second (H/s). The more hash power miners contribute, the more likely they find an acceptable solution to the puzzle and finally earn a reward for participating in the consensus mechanism (Tschorsch and Scheuermann 2016). The Bitcoin network allows the use of hardware optimized explicitly for the mining process, known as Application-Specific Integrated Circuits (ASIC). These devices provide higher hardware efficiency than general-purpose computing systems such as Central Processing Units (CPUs) and Graphics Processing Units (GPUs). As a result, a miner can generate more hash power with the same amount of energy. Hardware efficiency is measured in hashes per joule (H/J). Some networks, such as Ethereum before it switched to Proof of Stake, and Monero, have deliberately chosen ASIC-resistant algorithms that are designed to perform best on general-purpose CPUs and GPUs. However, the increase in hash performance provided by the miners does not permanently reduce block time or increase transaction throughput because the puzzle difficulty is adjusted to maintain an average block time of ten minutes. This is because the consensus mechanism periodically adjusts the difficulty of the puzzle based on the hash power of the network, for example, to maintain an average block time of ten minutes in the bitcoin network (de Vries and Stoll 2021).

Proof of Stake (PoS) (King and Nadal 2012) has been proposed as an electricity-saving and more performant alternative to Proof of Work. Instead of trying to solve a hash puzzle by providing computing power and electricity, participants in the consensus mechanism, called validators, provide some form of collateral ("stake"), usually in the form of the native cryptocurrency. A validator's chance of being selected to validate the next block, and thus receive a reward, increases proportionally with the amount of cryptocurrency staked. In addition, some PoS systems include a mechanism whereby validators can lose a portion of their staked cryptocurrency as a penalty for acting dishonestly or failing to validate transactions in a timely manner. In September 2022, Ethereum switched from a PoW to a PoS consensus mechanism.

Proof of Elapsed Time (PoET), a consensus mechanism developed by Intel, was designed as a scalable alternative suitable for enterprise environments (Chen et al. 2017). Like PoW, consensus participants try to solve a cryptographic puzzle, and the winner gets to propose the next block. However, unlike PoW, there is no high computational overhead and thus no excessive electricity consumption. This is achieved by having the nodes compute a random wait time at the beginning of a cycle, and being inactive until the wait time is over, until the node with the shortest wait time can propose the next block. However, to prevent a malicious actor from artificially reducing the wait time and the opportunity to propose a new block, the consensus mechanism must be executed in a specialized processor, called a Trusted Execution Environment (TEE), which ensures the secure and reliable execution of the algorithm and protects it from external interference and manipulation.

Proof of Authority (PoA) designates a select group of participants as authorities who can participate in the consensus mechanism. Network participants trust these authorities to act honestly since their reputation and accountability are at stake. If these authorities were to act dishonestly, they would risk losing their credibility. Who qualifies as an authority depends on the use case. For example, in a regulated environment, a government agency could serve as an authority. Alternatively, in a situation with a limited number of known, equal, and trusted parties, authority could be distributed among all network participants (Jennath and Asharaf 2020).

	Proof of Work (PoW)	Proof of Stake (PoS)	Proof of Elapsed Time (PoET)	Proof of Authority (PoA)
Electricity Consumption	Very high	Medium / Low	Medium / Low	Low
Consensus Resource	Computational Power	Capital as Collateral	1 TEE = 1 Vote	Authority
Openness	High	High / Medium	Medium / Low	Very low
Ability of Consensus Participation	All participants	Coin stakeholder	Participants with Intel CPU	Participants specified by the blockchain's governance
Miner/validator Election	Hash power	Coin stake	First participant whose wait time is up	Vote or random selection
Applications / Examples	Bitcoin, Ethereum (before Merge), Monero	Ethereum, Solana, Polkadot	Hyperledger, Corda	Hyperledger, Quorum, Corda

Table 1: Characteristics of different consensus algorithms (based on Lei et al. (2021))

The Properties of Blockchain-Based Data Infrastructure

Blockchain technology offers a promising solution for coordinating collaboration and establishing trust in distributed environments. It has been widely explored to address the challenges faced in various domains. By consciously designing a blockchain network, it is possible to leverage its inherent properties and meet the requirements outlined in the previous chapter. To further explore how blockchain technology can provide the necessary infrastructure properties, see the table below. It illustrates how a blockchain network can provide the environmental impact, performance, security, and availability properties that meet the specific requirements of the use case.

Property		Short Explanation	Extended Explanation
Environmental Impact	Electricity consumption	Electricity consumption of the blockchain itself	The electricity consumption of the blockchain can be approximated with the combined electricity consumption of the nodes and consensus participants in the network. To achieve a holistic view of the network's electricity consumption, other participants that do not operate a node, like internet service providers or stakeholders only interacting with the blockchain, e.g., transacting or retrieving information, need to be accounted for.
Performance	Latency	Average time between sending and confirmation of a transaction	Latency refers to the time a sender must wait between requesting a transaction and receiving confirmation of the transaction's non-reversible inclusion in the blockchain (also known as "time to finality"). This waiting time can vary based on the network's maximum throughput and the demand for transaction space, which is determined by the average complexity and number of transactions at a given time. The throughput of a blockchain depends on the block size, which defines the capacity of each block, and the block time, which determines the average time until a new block is proposed. The larger the transaction size, the fewer transactions can be executed in a single block.

Property		Short Explanation	Extended Explanation
Security	Confidentiality	Limits data access to authorized entities	In a blockchain, there is complete transparency between all nodes. However, one way to ensure the confidentiality of sensitive data is to select only authorized entities as node operators within the network, resulting in a private permissioned blockchain. On the other hand, if no sensitive data, such as publicly available information, is exchanged, or if the data is secured by encryption or other anonymization techniques, it can be stored on a public blockchain.
	Integrity	Ensures data accuracy and validity	One of the core features of blockchain technology is the immutability of the data stored on-chain, which guarantees the integrity of the stored data. This immutability is achieved through the consensus mechanism, which ensures that only valid and verified transactions are added to the blockchain ledger. The level of data integrity is defined by the type and degree of fault tolerance and the cost of obtaining the appropriate percentage of voting power provided by the consensus mechanism, which also ensures that only valid data transactions are added to the ledger. Another indicator of the system's integrity is the blockchain's decentralization, which primarily depends on the number of full nodes and consensus participants within the system. To accurately assess the network's level of decentralization and the number of nodes, other factors that determine the system's fail-safety, such as diversity of hosting, network providers, or client implementation, should be considered.
	Availability	Robustness of system and data access	Blockchain's decentralized nature ensures availability of systems and data by eliminating single points of failure. As the number of nodes in the network increases, so does the system's robustness. The network can tolerate some failed nodes thanks to the elimination of any single points of failure. The number of nodes required to be functional depends on the fault tolerance level chosen; the higher level, the more failed nodes can be tolerated. Similarly, because only one honest node is needed to retrieve data, data availability increases as the number of nodes increases.

Table 2: Properties of a blockchain

3. Understanding the Electricity Consumption of Blockchain Technology

In this chapter, we present the current state of research on the electricity consumption of blockchain technology. To determine the current state of knowledge, we conducted a systematic literature review of existing analyses. The findings not only form the basis for the subsequent parts of this study but are also meant to serve as an aggregated knowledge resource for both academia and industry.

Sections 3.1 through 3.3 unpack the multiple thematic threads that emerged from our literature review. Section 3.1 delves into the quantification of blockchain electricity consumption, comparing the electricity requirements of proof-of-work (PoW) and non-PoW networks. Section 3.2 presents different technical approaches to reducing the electricity consumption of blockchain technology, and Section 3.3 shifts the focus from technical to economic and policy approaches. Subsequently, Section 3.4, drawing on the insights from the literature review, presents a graphical representation of the main parameters influencing the electricity consumption of a blockchain network, depending on the consensus mechanism. The chapter concludes with an excursus that focuses on one specific technological implementation – rollups – and analyzes how this method improves the network's electricity efficiency.

3.1 Quantifying the Electricity Consumption of Blockchains

Due to a blockchain network's decentralized nature, no single source of electricity consumption can be directly measured. Instead, the network's electricity consumption consists of all participants' combined electricity consumption (O'Dwyer and Malone 2014; Stoll et al. 2019). Moreover, these participants and their electricity consumption are not homogeneous but can vary greatly depending on factors such as their role in the network, their hardware, and location-specific characteristics such as electricity prices, infrastructure, or the data center's climate region (de Vries 2020; Gallersdörfer et al. 2020). This poses particular challenges for determining total electricity consumption. In the course of our research, we identified several papers on the environmental impacts, such as CO_2 emissions and the amount of e-waste generated by decommissioned hardware devices (see e.g. de Vries and Stoll (2021); however, due to the focus of our study, these papers have not been included. Analysis of electricity consumption in the literature mainly focuses on the consensus mechanism used in the blockchain network, as this has the most significant impact on electricity consumption and the composition of consumption factors. Therefore, this next section follows the distinction between different consensus mechanisms and is divided into PoW and non-PoW mechanisms.

3.1.1 Electricity Consumption of PoW Networks

Most quantitative analyses of blockchain electricity consumption focus on the PoW consensus mechanism, especially within the Bitcoin network. While there are different approaches to determining electricity consumption, all articles we identified concur that a PoW blockchain's electricity consumption originates from the electricity used to solve hash puzzles or other computationally intensive tasks to gain voting rights in the consensus mechanism. Hence, the literature emphasizes electricity usage for mining activities, generally neglecting other network-related electricity consumers like node operators that do not participate in the consensus mechanism, as their share of the overall electricity consumption is considered insignificant.

Challenges for determining PoW electricity consumption

Determining the electricity consumption of a PoW network directly is not feasible due to the decentralized nature of a permissionless blockchain. Independent entities can enter the network and participate with their subsystems. Therefore, an unknown number of miners can make individual decisions about their mining operations, including the choice of mining devices and the specifications of the data centers. These factors significantly influence the electricity used for operating the network (Sedlmeir et al. 2020b). Nevertheless, the design choices made for the network can lead to certain constraints that affect the most electricity-consuming activities - mining operations (de Vries 2020). Based on these insights and observable actions in the network, one can estimate the network's electricity consumption. In the following, we present the most commonly used two methods to determine the network's electricity consumption based on a technological and an economic approach.

PoW-EC – Method: Technological Approach

One method for estimating electricity consumption involves analyzing the publicly observable hash rate of the network and estimating the electricity miners consume to generate this hash rate. Then, the electricity consumption can be approximated using the network's hash rate multiplied by the electricity used per hash by the miners:

$$\text{Total electricity consumption} = \text{Total hash rate} \times \text{Electricity used per hash}$$

In this formula, the total electricity consumption is in watts, the hash rate is in hashes/second, and the electricity consumption per hash used by the miners is measured in hashes per joule (H/J). As the model relies on two relatively stable variables, the network's hash rate and the electricity efficiency of the used mining hardware, it is argued to provide a reasonably robust estimation (Crypto Carbon Ratings Institute 2022c; Sedlmeir et al. 2020b; Stoll et al. 2019).

A total of 13 publications in the identified literature used this approach to determine the electricity consumption of the Bitcoin network (Coinshare 2022; de Vries 2018, 2019, 2020, 2021; Gallersdörfer et al. 2020; Krause and Tolaymat 2018; Küfeoğlu and Özkuran 2019; Mora et al. 2018; O'Dwyer and Malone 2014; Sedlmeir et al. 2020a, 2020b; Shi et al. 2021; Song and Aste 2020).

In addition, with Cambridge Bitcoin Electricity Consumption Index (CBECI) (CCBECI 2022) and Digiconomist (Digiconomist 2022), we include two sources from the grey literature that related work has frequently cited. Furthermore, the method was applied to Ethereum when it still used PoW (de Vries 2022; Gallersdörfer et al. 2020; McDonald 2021; Sedlmeir et al. 2020a, 2020b; Qin et al. 2021; Shi et al. 2021; Zade et al. 2019), and additional PoW-based cryptocurrencies among those with the highest market capitalization (Gallersdörfer et al. 2020; Krause and Tolaymat 2018; Sedlmeir et al. 2020a, 2020b). In addition, Song and Aste (2020) modified this approach to determine the electricity costs of Bitcoin mining.

Hardware efficiency

While observable data can determine the hash rate on the network, electricity used per hash by miners, i.e., the average efficiency of the mining hardware, used cannot be measured directly (Sedlmeir et al. 2020b). Different solutions to this challenge have been proposed in the literature. For example, several publications calculated the lowest theoretical electricity consumption by assuming that miners use only the most efficient hardware available, providing a robust estimate of the lowest possible bound of the electricity consumed (Coinshare 2022; de Vries 2020, 2021, 2022; Krause and Tolaymat 2018; Sai and Vranken 2022; Song and Aste 2020; Sedlmeir et al. 2020b; Vranken 2017). However, this approach has been criticized for presumably significantly underestimating actual electricity consumption (de Vries 2020; Koomey 2019). Therefore, some studies have attempted to obtain a more accurate estimation of hardware efficiency based on the distribution of hardware, e.g., using sales figures from hardware manufacturers or the estimated lifetime of certain device types. However, these approximations introduce new uncertainties due to insignificant empirical data (Sai and Vranken 2022). Additionally, some authors included the Power Usage Efficiency Factor as a multiplicator in order to incorporate additional electricity consumption beyond the miner's devices, such as the cooling of the hardware or the network infrastructure of the data centers (Coinshare 2022; Crypto Carbon Ratings Institute 2022c; McDonald 2021; de Vries 2018, 2019, 2020, 2021; Stoll et al. 2019).

PoW-EC – Method: An Economic Approach

The second approach is based on the assumption that all mining participants act economically and rationally. Thus, the expenditures of miners for participating at the consensus mechanism will not be higher than their expected revenue (Lei et al. 2021). Accordingly, the method approximates an upper bound of a PoW cryptocurrency network's mining-related electricity consumption. The approach was used to determine the electricity consumption of Bitcoin (de Vries 2018, 2019, 2021; Gonzalez-Barahona 2021; Küfeoğlu and Özkuran 2019; Sedlmeir et al. 2020a, 2020b; Shi et al. 2021; Stoll et al. 2019; Vranken 2017), Ethereum (before the switch to PoS) (Sedlmeir et al. 2020a, 2020b; Shi et al. 2021) and additional PoW-based cryptocurrencies among those with the highest market capitalization (Sedlmeir et al. 2020b).

The relation between income and expenses is represented as follows:

Total mining income ≥ Total mining costs

Total Mining Income

Total miner income represents the value disseminated by the network to incentivize miners to contribute resources. This participation in the consensus mechanism is crucial to keep the network operational. A miner who successfully solves a hash puzzle for a block receives both transaction fees and the 'block rewards', where the amount of these rewards are dictated by the design of the consensus mechanism. Similar to the hash rate used in the technical approach, the total mining revenue can be unambiguously determined because its distribution in the network is verifiable. (de Vries 2021; de Vries and Stoll 2021; Gonzalez-Barahona 2021; Küfeoğlu and Özkuran 2019; Sedlmeir et al. 2020b; Stoll et al. 2019).

Furthermore, since both transaction fees and block rewards are paid in the network's native currency, the exchange rate directly impacts mining income and thus electricity consumption. For example, Sedlmeir et al. (2020b) observed a sharp drop in electricity consumption after bitcoin's value collapsed in 2020, and a spike in electricity consumption after bitcoin's value recovered. The formula for total mining income is as follows. Please note, that both parameters are already converted into a fiat currency, such as the Euro or US Dollar, for ease of understanding.
Total mining income=Block reward+Transaction fees

Total Miner's Costs

Much like hardware efficiency, the total costs for miners are not directly observable. This is primarily because the expenses associated which each mining operation are unique. Therefore, it is necessary to make assumptions regarding the miner's cost composition to estimate the network's electricity consumption. In doing so, most studies identified electricity consumption as the primary determining factor for how much hash power a miner contributes (McDonald 2021). In other words, miners decide to adjust the hash rate they offer – and hence the electricity consumed by their mining devices – based on the anticipated revenue and prevailing electricity prices. However, due to the lack of empirical data on the miners' actual electricity expenses, different prices are assumed in the studies we analyzed, introducing some degree of uncertainty in the results (Sai and Vranken 2022).

Many of the studies analyzed assume that miners spend their entire revenue on electricity consumption. However, such assumptions may lead to an overestimation of electricity consumption, as they fail to consider additional costs such as variable costs (e.g., infrastructure and personnel) and long-term costs (e.g., investments in new hardware) (Gola and Sedlmeir 2022). To address this, several researchers incorporate these additional overhead expenses as fixed costs, effectively reducing the revenue a miner can spend on electricity. Based on the required hardware investment costs, de Vries (2018) and Qin et al. (2021) estimate that miners spend about 60 percent of their total income on electricity.

Accordingly, the miners' costs can be expressed as follows:

$$\text{Total miner's costs} = \text{Electricty consumption} \times \text{Electricity prices} + \text{Fixed costs}$$

Putting both elements together, we end up the following equation.

$$\text{Block reward} + \text{Transaction fees} \geq \text{Electricty consumption} \times \text{Electricity prices} + \text{Fixed costs}$$

Finally, with some simple transformations, we can derive the upper bound of the overall network's electricity consumption:

$$\text{Electricity consumption} \leq \frac{(\text{Block reward} + \text{Transaction fees}) - \text{Fix costs}}{\text{Electricity price}}$$

3.1.2 Electricity Consumption of Non-PoW Networks

Several publications agree that the electricity consumption of Non-PoW networks is significantly lower than that of PoW networks (Heinonen et al. 2022; Kohli et al. 2022; Platt et al. 2021; Crypto Carbon Ratings Institute 2022c; de Vries 2022). However, there is a relative scarcity of publications quantifying the electricity consumption of non-PoW blockchains. Thus, our systematic literature review yielded only eight results, with only two peer-reviewed publications (de Vries 2022; Platt et al. 2021; Shi et al. 2021), five white reports published by the Crypto Carbon Ratings Institute (Crypto Carbon Ratings Institute 2022a, 2022b, 2022c; Gallersdörfer et al. 2022), and the analysis of CBECI (2023) based on the results of the Crypto Carbon Institute.

Determining Electricity Consumption in PoS Networks

These publications unanimously agree that the high electricity consumption of PoW systems primarily stems from the computational demands of solving hash puzzles to secure voting rights in the consensus mechanism. Therefore, by substituting computational power with other scarce resources, such as staked currencies in Proof-of-Stake (PoS) systems, the primary driver of high electricity consumption is effectively eliminated, resulting in non-PoW blockchains having a significantly lower electricity demand than their PoW counterparts. However, it is important to note that the overall electricity consumption of different non-PoW networks will vary, mainly depending on the computational load each node has to bear as well as the number of active nodes in the network. (CBECI 2022, de Vries 2022; Gallersdörfer et al. 2022; Crypto Carbon Ratings Institute 2022a, 2022b, 2022c; Platt et al. 2021; Shi et al. 2021).

Another point of agreement between all authors is that the total electricity consumption of a non-PoW network can be approximated by multiplying the number of active nodes with the nodes' average electricity consumption:

$$\text{Electricity consumption} = \text{Active nodes} \times \text{Average electricty consumption per node}$$

In contrast to PoW mining, the average electricity consumption of a node is primarily driven by idle load consumption. This idle consumption mainly depends on the hardware utilized, the lower boundary of which is determined by the minimum hardware requirements to operate a node. Network design decisions influence these minimum hardware requirements. This encompasses transaction complexity, optional functions like stake delegation, and other factors (Gallersdörfer et al. 2022). For example, Solana nodes execute multiple separate smart contracts in parallel, which enables the high performance of the network but also requires a significant amount of computational resources from the nodes.

Interplay Between Number of Transactions and Electricity Consumption

Platt et al. (2021) compared the idle consumption of nodes to their extra electricity consumption caused by the computational load. They inferred that a rise in transactions at first only minimally contributes to higher electricity consumption due to the negligible computing load an individual transaction generates. Therefore, an increase in network throughput with constant node hardware decreases the average electricity consumption per transaction since the idle consumption of the network nodes can be allocated across more transactions (Platt et al. 2021). Nevertheless, the literature underscores that a direct comparison of networks based on this metric is inadequate due to discrepancies in the networks' decentralization or transaction complexities (Gallersdörfer et al. 2022). For instance, a payment transaction has lower computational complexity and electricity consumption than executing a complex smart contract (Platt et al. 2021; Gallersdörfer et al. 2022).

Contrary to the number or complexity of transactions determined by the blockchain design, the number of nodes participating in a permissionless network cannot be directly controlled. However, elevating hardware requirements for a specific network can curtail the number of active nodes as higher investment costs deter potential node operators from network participation, resulting in lower decentralization (Gallersdörfer et al. 2022). Thus, higher hardware requirements can have a dual impact on electricity consumption: On the one hand, they lead to increased electricity consumption by individual nodes. On the other hand, they can reduce the total number of nodes in the network by raising the barriers to entry.

Rieger et al. (2022) is the only publication that examines the electricity consumption of permissioned networks. Their work used a benchmarking framework to compare different networks, each with different nodes and consensus mechanisms. Their results suggest that the electricity consumption for each transaction is primarily influenced by the degree of decentralization of the network and the set fault tolerance of the consensus mechanism. Consistent with the existing research, they found that non-PoW consensus mechanisms consume significantly less electricity, in line with the consumption levels of conventional IT systems.

The Crypto Carbon Ratings Institute's Approach to Electricity Analysis

In its reports, the Crypto Carbon Ratings Institute analyzed nine PoS networks, including Algorand, Avalanche, Cardano, Polkadot, Solana and Tezos (Gallersdörfer et al. 2022), Tron (Crypto Carbon Ratings Institute 2022b) and Ethereum after the "Merge", i.e. its transition from PoW to PoS (Crypto Carbon Ratings Institute 2022c). In addition, the Institute studied a network built with Polygon that uses its own PoS infrastructure in the second layer, while Ethereum is used as the main network or first layer. They also compared Polygon's electricity consumption before and after Ethereum's shift to PoS (Crypto Carbon Ratings Institute 2022a). A more detailed discussion of this follows in the next chapter.

While the number of active nodes can be retrieved from already available data sources, no such source exists for the electricity consumption of nodes for the different networks. Thus, the authors analyzed the electricity consumption of the nodes, considering the variation in different node software and hardware configurations. To do this, the Crypto Carbon Ratings Institute used a controlled test bed environment to measure the electricity consumption across different hardware configurations. These ranged from a Raspberry Pi to a high-end computer equipped with an AMD Ryzen Threadripper 3970X CPU. The authors installed the node software of the networks under investigation on these devices. However, some hardware in the pool could not meet the requirements of some networks due to the varying hardware demands of different node software. These devices were, therefore not tested for those respective networks.

The nodes were operated for at least one day to measure electricity consumption under realistic conditions, considering varying levels of transaction throughput. Just like in PoW networks, determining the hardware distribution of network participants is not straightforward, so estimations have to be made. In this context, the Crypto Carbon Ratings Institute assumed the overall hardware distribution for network participants by using all hardware configurations capable of running a network's node software. However, this approach introduces some uncertainty into the projected electricity consumption.

Differing Hardware Requirements Across Networks

Of the networks studied, Avalanche, Algorand, Cardano, and Tezos specify only low requirements for the hardware, which can be met even by entry-level notebooks (Gallersdörfer et al. 2022). However, four of the networks analyzed (Ethereum, Polkadot, Polygon, and Tron) required at least moderately configured consumer hardware to run a node. Finally, Polygon had the highest requirements, which only the most performant devices in the hardware pool can meet. As shown in Figure 8, these hardware requirements directly influence the network's overall electricity consumption. For instance, the Solana network consumes the most electricity with its high node requirements. In contrast, Polkadot and Tezos, with their modest hardware requirements, have lower electricity consumption per node. In addition, fewer nodes are in these networks, resulting in lower overall electricity consumption. (Gallersdörfer et al. 2022).

Electricity Consumption on a Per-Transaction Basis

When considering electricity consumption on a per-transaction basis, the efficiency benefits due to higher-level hardware configurations and a higher transactions throughput become

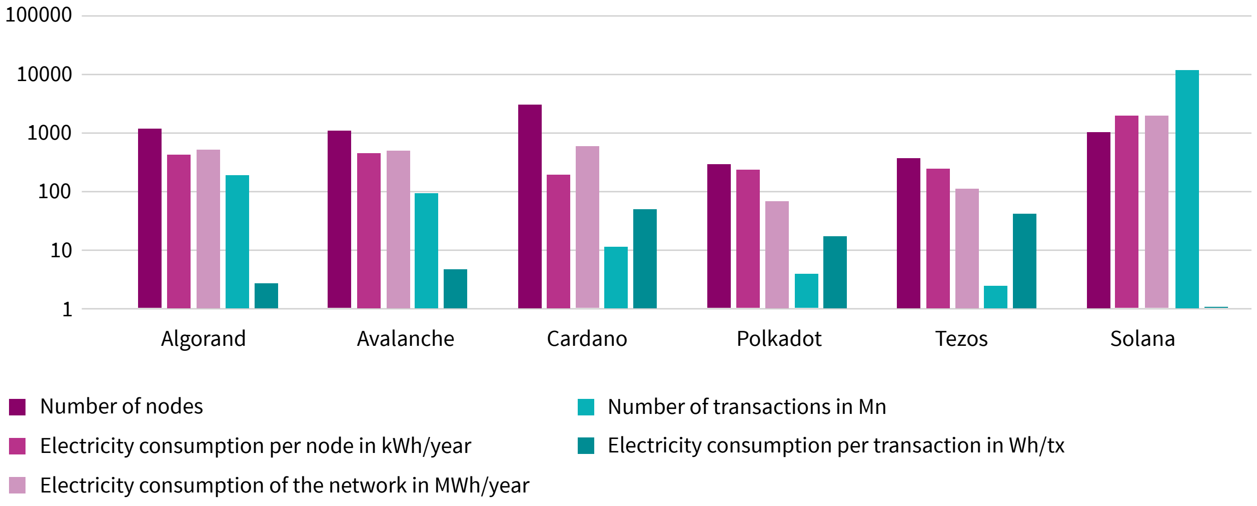

Figure 8: Comparison of the PoS networks studied in Gallersdörfer et al. (2022) by the following parameters: number of nodes, electricity consumption per node in kWh/year, the electricity consumption of the network in MWh/year, and the number of transactions in millions. Note the logarithmic scale of the Y-axis.

apparent: Solana, which processed a total of 11.8 billion transactions at the time, consumes significantly less electricity per transaction than Polkadot, which handled a total of 4 million transactions (Crypto Carbon Ratings Institute 2022c). The negative correlation between the number of transactions in the network and the electricity consumption per transaction, as Platt et al. (2021) demonstrated in their theoretical model, is also evident in the other networks examined. In simpler terms: When a network processes more transactions, the electricity required for each transaction decreases because the electricity cost of maintaining the network infrastructure is spread out over more transactions.

3.2 Technical Approaches to Reducing Blockchain's Electricity Consumption

The literature in this section agrees with the findings from literature in previous chapters that the computations carried out to determine the voting influence in PoW are responsible for most of the electricity consumption in the large networks that use this consensus mechanism. The literature summarized here tries to tackle this issue by proposing the usage of alternative consensus mechanisms or adjusting the PoW consensus mechanism. These concepts can be divided into three groups. The first group contains consensus mechanisms that aim to reduce the consumed electricity by using a different Sybil resistance mechanism, i.e., coupling voting power in the consensus mechanism to a scarce resource other than computational power, as is done in PoS. The second group suggests changes to the PoW consensus mechanism that aim at reducing its computational and, therefore, its electricity costs. Finally, the third group of consensus mechanisms is very similar to PoW, but instead of proving computational power by solving hash puzzles, participants need to solve extensive computational problems that are important to other applications. This approach aims to create additional value with the computational power used for participating in the consensus mechanism.

3.2.1 Leveraging Consensus Mechanisms with Lower Electricity Consumption

The literature referenced in this subsection tries to tackle the electricity consumption caused by the extensive computations associated with PoW in two different ways. The first approach is to replace PoW with another mechanism for Sybil-resistance that utilizes a digitally verifiable scarce resource other than the computational resources used in PoW. All of the literature agrees that the Sybil resistance mechanism in PoW is the main reason for blockchains' electricity consumption, and therefore, blockchains using another mechanism for Sybil-resistance will have lower electricity consumption. Between these other consensus mechanisms, which exclude the electricity-intensive PoW Sybil resistance mechanism, there is no significant difference in the electricity consumption highlighted by the literature. The different mechanisms proposed to replace PoW vary widely but some are more popular than others, and most can be grouped into a few categories if abstracted to the right level.

PoS

The most popular Sybil resistance mechanism besides PoW, in real-life applications as well as in the literature, is PoS (Gundaboina et al. 2022; Heinonen et al. 2022; Kohli et al. 2022; Wen et al. 2020). Additionally, suggestions have been made for some adaptations of the PoS mechanism by adding a second variable that also reflects the stake over time (Król et al. 2019) and honest participation, or considers the waiting time since the last selection for block creation (Marangappanavar and Kiran 2021). While these proposals claim to create a fairer system for reward distribution and entail less electricity usage than PoW, it is neither mentioned in the literature nor is it to be expected that these adaptations to the PoS consensus mechanism will have any influence on its electricity consumption. This means that, in terms of electricity consumption, they can be viewed as being equivalents to regular PoS options.

Transition from PoW to a consensus mechanism with lower electricity consumption.

Ethereum, the second biggest cryptocurrency, switched from PoW to PoS on the 15th of September, 2022, which reduced its electricity consumption by 99.988% (Crypto Carbon Ratings Institute 2022c). However, de Vries (2022) argues that this resulted in a share of the hash power and thus electricity consumption used for participation in Ethereum's PoW migrating to other PoW blockchains, diminishing some of the electricity savings, especially in the short term.

Even before Ethereum's successful transition, some researchers investigated whether Bitcoin could change its consensus mechanism. According to Heinonen et al. (2022) it is plausible for the Bitcoin network to switch to most of the less electricity-intensive mechanisms. They only categorize a few exceptions like proof-of-elapsed time as implausible because they are either too close to a permissioned system or do not bring significant improvements, as the hardware would still need to be hoarded. Kostal et al. (2018) highlight that the Bitcoin network might even be forced to switch to another consensus mechanism in the future. When the mining rewards get too low (because of the halvings) and a rise in transaction fees does not compensate for this effect, this could lead to the economic security of Bitcoin reaching such low levels that the transition to another consensus mechanism might be deemed appropriate by a majority of the community. On the other hand, Gola and Sedlmeir (2022) emphasize that a change from PoW to PoS may be challenging to attain, as it would require support from the majority of the PoW miners, who would effectively end their business model by voting in favor of the change. Additionally, it would need to be ensured that the safety assumptions for PoW also hold after the transition to PoS, i.e., honest participants controlling at least 51% of the hash power and the stake.

Like PoA

Proof of authority builds on the concept of the participants of the consensus mechanism being known to prevent Sybil attacks. While this is often applied in permissioned blockchains, there are also suggestions of integrating registration processes into permissionless blockchains to register consensus participants (Yu et al. 2019; Wen et al. 2020) or a combining registration with the depositing of stake as it is done in PoS (Solat 2017).

Others

Other consensus mechanisms like proof of retrievability (also called proof of space or proof of space time), where storage is used as a scarce resource (Kohli et al. 2022; Wen et al. 2020), or proof of elapsed time (also called proof of TEE-stake or proof of hardware) where only hardware (i.e., the trusted execution environment in CPUs) is used as a scarce resource (Milutinovic et al. 2016; Wen et al. 2020), also shed off PoW's high electricity consumption but still cause a certain level of electricity consumption in addition to the required hardware and the resulting e-waste. An approach to fully forgo a Sybil resistance mechanism is presented by Jacquet and Mans (2019), where a new block is created if the hash value of the included transactions is below a certain threshold. Whether this could lead to electricity consumption similar to PoW and to additional security issues resulting from spam attacks where transactions are issued for the sole reason of reaching a desired hash value remains unexplored.

There are also other changes proposed in the literature that relate to the mechanism of reaching consensus and not the Sybil resistance. These are, among others, using the gossip protocol on a hashgraph (Kohli et al. 2022), probabilistic consensus mechanisms like Fast Probabilistic Consensus (Kohli et al. 2022; Wen et al. 2020), or mechanisms based on Byzantine agreement (Kohli et al. 2022; Wen et al. 2020). Since these implementations mainly use some form of PoS as a Sybil resistance mechanism, they have significantly less electricity consumption than PoW. If and how these changes in the consensus mechanism itself influence the electricity consumption, i.e., how the electricity consumption of the systems above compares to traditional PoS systems, remains an open question in need of further research. However, it is doubtful that these modifications would be relevant to any scenario where participation in consensus is not computationally heavy for any given node.

3.2.2 Strategies for Reducing the Electricity Consumption of PoW Networks

Another class of proposed consensus mechanisms aims at adjusting PoW to reduce the network's electricity consumption. For this purpose, reputation or credit models are introduced, which can be based on stake and past behavior when participating in consensus (Alofi et al. 2021b; Alofi et al. 2021a; Alofi et al. 2022; Xue et al. 2018; Wen et al. 2020; Wang and Gem Lina 2022) or based on a trust graph between network participants (Bahri and Girdzijauskas 2018). The idea behind this is to either reduce the difficulty of the hash puzzle for more reputable or trusted individuals (Bahri and Girdzijauskas 2018; Xue et al. 2018; Wang and Gem Lina 2022; Ouaili et al. 2022) or select only a subset of miners (Alofi et al. 2021a; Alofi et al. 2021b; Alofi et al. 2022). Other forms of selecting only a subset of miners to participate in the creation of the next block are by demanding a certain stake from miners (Monem et al. 2020), creating two rounds of mining, where only a certain group of winners of the first round is allowed to participate in the second round (Lasla et al. 2020) or randomly determining a set of public keys that are eligible for mining the following block (Lundbæk et al. 2018). Additionally, Castellon et al. (2022) propose changes to the hash calculation algorithm to reduce the electricity consumption for achieving a given difficulty level.

There are two misconceptions shared by most of these approaches. The first is, that a lower difficulty of the hash puzzle or an increase in electricity efficiency that allows achieving the same difficulty with lower electricity consumption would lead to lower electricity consumption of the network. While this is true for the creation of a single block, a lower difficulty would just lead to a faster block generation with constant electricity consumption, or even increase electricity consumption for most blockchains, as the number of issuances for block creation would increase, leading to higher incentives for miners (see, e.g., Gola and Sedlmeir 2022). This line of thinking is also undermined by all the literature items listed in the previous chapter, which determine an upper bound to the electricity consumption that only relies on economic factors and is independent of the difficulty of the hash puzzles (see e.g., Sedlmeir et al. (2020a) and Lei et al. (2021)). Thus, electricity consumption depends on the incentives for participation in the consensus mechanism and not on the number of participants or the difficulty of the hash puzzle. Another misconception is, that mining power can easily be bound to a single account. When selecting only a subset of miners for the creation of the next block, there is no trivial way to stop mining pools from creating multiple addresses, e.g., staking at multiple accounts, and then shifting all computational power to the selected addresses. Furthermore, the introduction of staking or similar concepts raises the question if there are still relevant differences in security assumptions to a pure PoS consensus mechanism and, if not, how the additional computational effort of the hybrid approach can be justified.

3.2.3 Utilizing Useful Computation in Consensus Mechanisms

Besides proposing new consensus mechanisms to decrease the electricity consumption of a blockchain, there is also a significant number of literature items that propose adjustments to the PoW consensus mechanism to use the computations for solving problems other than the hash puzzles in the PoW algorithm proposed by Nakamoto (2008). Therefore, the literature that will be covered in this section does not primarily intend to decrease the electricity usage of the blockchain but instead tries to define computationally expensive problems to be solved so that the computations that determine influence in the PoW consensus mechanism – and thereby the hardware and electricity – are utilized in a “more useful” way.

Motivation

Chatterjee et al. (2019) describe the PoW algorithm as a specific instance of distributed problem solving and conclude that the system and its computational resources could be leveraged to solve extensive computational problems that already exist in other areas. This view is also what motivates many other similar proposals (Chatterjee et al. 2019; Shoker 2018; Chenli et al. 2019; Li et al. 2019; Qu et al. 2021; Chaurasia et al. 2021; Shibata 2019; Talukder and Vaughn 2021; Toulemonde et al. 2022). Another subgroup of the literature proposes that solving a certain class of computational problems, even if the specific instance has no use outside of the blockchain, would create a strong financial incentive to develop more efficient algorithms for this problem class, which would then be beneficial for all other applications facing this problem class. Examples of this would be solving NP-hard problems (Loe and Quaglia 2018; Syafruddin et al. 2019) or training deep learning models (Chenli et al. 2019; Li et al. 2019; Baldominos and Saez 2019). An additional benefit of changing the problem class solved in the PoW algorithm is that an incentive for the creation of more efficient hardware for these problems would be created, and then dependency on the manufacturers of ASICs, hardware specialized in solving hash puzzles, would be reduced (Chatterjee et al. 2019; Loe and Quaglia 2018).

Problem Description

The problems proposed in the literature form a wide range (Heinonen et al. 2022). Many researchers agree on the usage of NP-complete problems to ensure the complexity of the problems and the equivalent computational effort needed to solve different instances of these problems (Chatterjee et al. 2019; Loe and Quaglia 2018; Shibata 2019; Syafruddin et al. 2019). Proposed examples include Protein Folding (Chatterjee et al. 2019), searching prime numbers of the form 2^p-1 (Chatterjee et al. 2019) or the Traveling Salesman (Loe and Quaglia 2018; Syafruddin et al. 2019). Other suggestions are matrix-based computations (Shoker 2018; Wei et al. 2022) which are closely tied to the training of deep learning models (Chenli et al. 2019; Li et al. 2019; Baldominos and Saez 2019), federated learning (Qu et al. 2021), or machine learning (Wei et al. 2022). Chaurasia et al. (2021) propose a very blockchain-specific problem of calculating private keys corresponding to specific and “interesting” public keys and addresses, while Dong et al. (2019) suggest generalizing the work required in PoW to a broader spectrum of IT- and cloud-services like network hosting, or provision of storage or computing resources.

Besides these novel problems, some literature items still incorporate the solving of hash puzzles in their proposed consensus algorithm, either to construct a problem instance from the solution of the hash puzzle (Loe and Quaglia 2018) or to enable miners to still utilize their existing hardware by giving them a choice between solving a hash puzzle or an instance of the new problem class (Chatterjee et al. 2019).

Proposing and Solving Novel Problems

Regarding the proposal of problems to be solved through the consensus mechanism, two different options can be found in the literature. The first is that anyone can propose a problem or use the service (Chaurasia et al. 2021; Chatterjee et al. 2019; Chenli et al. 2019; Dong et al. 2019; Qu et al. 2021; Shibata 2019; Shoker 2018; Toulemonde et al. 2022; Bizzaro et al. 2020). In such systems, the proposer must pay a fee for the proposal itself and offer a reward that is given to the solver. The amount of the reward can either be chosen by the proposer (Chatterjee et al. 2019; Chaurasia et al. 2021; Chenli et al. 2019; Shoker 2018) or the pricing is done algorithmically depending on supply and demand (Dong et al. 2019). In many scenarios, the proposer can also specify a time period for which the problem will be available (Chaurasia et al. 2021; Chatterjee et al. 2019; Shoker 2018). The other option for problem proposition is to generate the problems automatically by already embedding the problem creation into the consensus algorithm (Loe and Quaglia 2018; Syafruddin et al. 2019; Talukder and Vaughn 2021; Ball et al. 2018).

Open Questions

This section of the literature reveals several unresolved questions concerning the security and functionality of the proposed systems within blockchain technology. Some of these systems reintroduce a centralized institution responsible for managing and coordinating problem assignment (Qu et al. 2021; Wei et al. 2022) or ensuring data availability of the solutions found (Li et al. 2019). However, it is unclear how this impacts the system’s security and its level of decentralization. A recurring question is whether verifying proposed problems requires significantly less computational effort than solving them. Although hash puzzles fulfill this requirement, not all problems can assume this. The absence of this asymmetry between solving and verification might require additional complexity to incentivize nodes for verification, leading to potential system vulnerabilities. It is equally important to ensure uniform difficulty for all problems. Variation in the computational effort required to solve problems could affect block time, possibly influencing the security assumptions traditionally associated with Proof of Work. The impact of such block time variations remains unknown.

Another open issue pertains to the scalability and parallelizability of problem-solving. While hash puzzles are fully scalable, this property might not apply to all problems. Scalability in this context means, for example, that one ASIC has 1/10th of the chance of ten ASIC-Devices to find a possible solution, since it is based on trial and error of testing inputs and calculating their respective hash. If this property is not given, i.e., a supercomputer consisting of 1000 GPUs is more than 1000 times more likely to find a solution before a computer working with a single and identical GPU, centralization effects will occur, giving participants with the highest computational power the entitlement to create all blocks instead of just the share of all blocks corresponding to their share of the computational power of the overall network. Lastly, the economic security of a PoW network remains unaddressed, a crucial aspect of the network’s security assumptions. Traditional PoW posits that economic security is about 51% of all mining costs. However, replacing the problem of solving hash puzzles with “more useful” problems could reduce the cost of achieving 51% network control, given these problems’ inherent economic value. This significant issue is yet to be resolved.

3.3 Economic and Policy Approaches to Reducing Electricity Consumption

The literature items in this section analyze economic or political changes and their effect on the electricity consumption of blockchain networks and use these analyses to propose policies to reduce electricity consumption, emissions, and e-waste. The research we found agrees that precise information and analyses on the environmental impact (electricity, carbon, and e-waste) of blockchains is needed for sound policy decisions (Lei et al. 2021), including location-based differentiations of miners. This geospatial information is essential for determining carbon emissions, as these depend on electricity consumption and the carbon intensity of the mining facility's electricity sources, which, in turn, strongly depend on the miner's location (Truby et al. 2022). Equipped with this information, policymakers should be aware of the dangers and ready to intervene in order to cause a shift to a more beneficial use of the technology (Truby 2018; Gola and Sedlmeir 2022).

While the analysis below focuses on PoW networks, the majority of researchers agree that switching to a Sybil resistance mechanism other than PoW would be a significant advance in reducing the impact of the respective network (de Vries and Stoll 2021; Erdogan et al. 2022; Truby 2018; Truby et al. 2022; Gola and Sedlmeir 2022). Besides switching to a different Sybil resistance mechanism, the most straightforward option for reducing the electricity consumption of PoW mining proposed by the literature is an **increase in electricity prices** (de Vries 2021; Gonzalez-Barahona 2021; Qin et al. 2021; Truby 2018; Truby et al. 2022).

Reduce electricity consumption by reducing the ratio of costs for electricity to other costs

Another way to reduce the electricity consumption of PoW mining is by reducing the ratio of electricity costs to other mining costs, in particular hardware costs. Since miners have a fixed potential income, determined by the network's mining incentives and transaction fees, increased spending on hardware costs would result in decreased spending on electricity and, therefore, decreased electricity consumption. This could be achieved by increasing hardware costs (Gonzalez-Barahona 2021), which could be achieved by instating special taxes on mining hardware (de Vries 2021; Truby 2018; Truby et al. 2022). Another possibility to achieve this effect is by making hardware more efficient while maintaining or increasing its acquisition or amortization cost (Gonzalez-Barahona 2021; Qin et al. 2021; Truby et al. 2022; Mohsin et al. 2020; Polemis and Tsionas 2021).

Reduce electricity consumption by decreasing miner income

A third vector to reduce electricity consumption proposed by the literature is decreasing miner revenue, which would reduce their electricity expenses. This could be done by reducing transaction fees (Gonzalez-Barahona 2021) or by reducing mining incentives (e.g., by accelerating the halving in Bitcoin) (Gonzalez-Barahona 2021). Another possibility would be to introduce new taxes and enforce existing ones. Special taxes on mining activities could make the businesses less profitable, and taxes paid for transacting in a PoW system would reduce the transaction fees, i.e., the miner revenue (Badea and Mungiu-Pupazan 2021; Erdogan et al. 2022; Truby 2018; Truby et al. 2022; Gola and Sedlmeir 2022).

An indirect impact on mining revenues could also be achieved by banning certain PoW currencies, restricting their trading, or introducing and enforcing taxes on capital gains from digital assets based on PoW systems. This would likely cause them to depreciate in value, reducing miners' revenues, as they are primarily measured in the native currency of the network (Badea and Mungiu-Pupazan 2021; de Vries 2021; Truby 2018; Truby et al. 2022; Gola and Sedlmeir 2022).

Prohibiting mining operations

The literature also covers existing bans and crackdowns on large-scale mining operations by cutting off electricity or confiscating hardware. Examples listed are in Iran or Quebec (Canada) because of the destabilization of the electricity grid (de Vries 2021; Truby 2018) or in New York State or China (Truby et al. 2022; Gola and Sedlmeir 2022). These bans can reduce electricity consumption since existing mining operations try to seek the cheapest source of electricity. By banning them from their existing choices, mining operations migrate to other locations with more expensive electricity sources, reducing their electricity consumption. Prohibiting mining operations can also affect the carbon emissions of the network, decrease them if the banned location has electricity with an above-average carbon emission intensity, or increase them if the intensity is below average (de Vries et al. 2022; Jiang et al. 2021). De Vries (2021) points out that there are boundaries to policy options since, in theory, only a laptop and an internet connection are needed for mining. However, in practice, large-scale mining operations are responsible for most of the mining power since they benefit from economies of scale. Such large-scale operations, which are sometimes even operated by publicly listed companies, can be targeted by policies quite well as the examples above show (de Vries 2021).

3.4 Modeling Blockchain's Electricity Consumption

In this chapter, we present a comprehensive model to illustrate the multifaceted factors that influence the electricity consumption of blockchain networks. This representation is intended to provide a deep dive into the complexities involved, providing researchers and practitioners with insights to understand the interrelationships of the parameters influencing them, and to identify potential opportunities to reduce the electricity consumption of blockchain networks. We developed two different versions of the model, each with a different level of complexity. The first model captures all the factors that could potentially affect the electricity consumption of blockchain networks. The second model builds on these results and focuses on the most important components to offer the results in a more straightforward manner for the reader. The development of these models was informed by insights from our systematic literature review, which allowed us to identify key relationships and knowledge gaps. We then undertook an iterative refinement process, incorporating feedback from expert interviews and workshops conducted to ensure the comprehensiveness and validity of our findings. In this way, we were able to make sure that the model reflects not only the findings of the literature review, but also the latest insights from industry experts and academic researchers.

In this study, we present the condensed version of our model. By emphasizing the most influential parameters, we aim to provide the reader with a comprehensive yet digestible understanding of the underlying parameters that affect electricity consumption.

Structure of the Model

As shown in Figure 9, our model is divided into two main sides: Proof-of-Work (PoW) on the left and Non-PoW on the right. This visual distinction allows us to directly compare the electricity consumption characteristics of the two types. Each side is further divided into three distinct sections. The first section outlines the key drivers of electricity consumption for the main electricity demand. The next section breaks down these drivers into component equations derived from our systematic literature review. The outer section then links these components to the electricity consumption parameters we identified, which are highlighted in yellow. We also use asterisks to indicate those parameters that are exogenous and cannot be controlled by network design. Exogenous parameters include political factors such as regulations or taxes, market influences such as electricity or coin prices, and technological developments such as more efficient hardware. We have also marked with an asterisk those parameters that are only controllable in permissioned systems. In the following, we will first examine the PoW side in detail, before moving on to the non-PoW side.

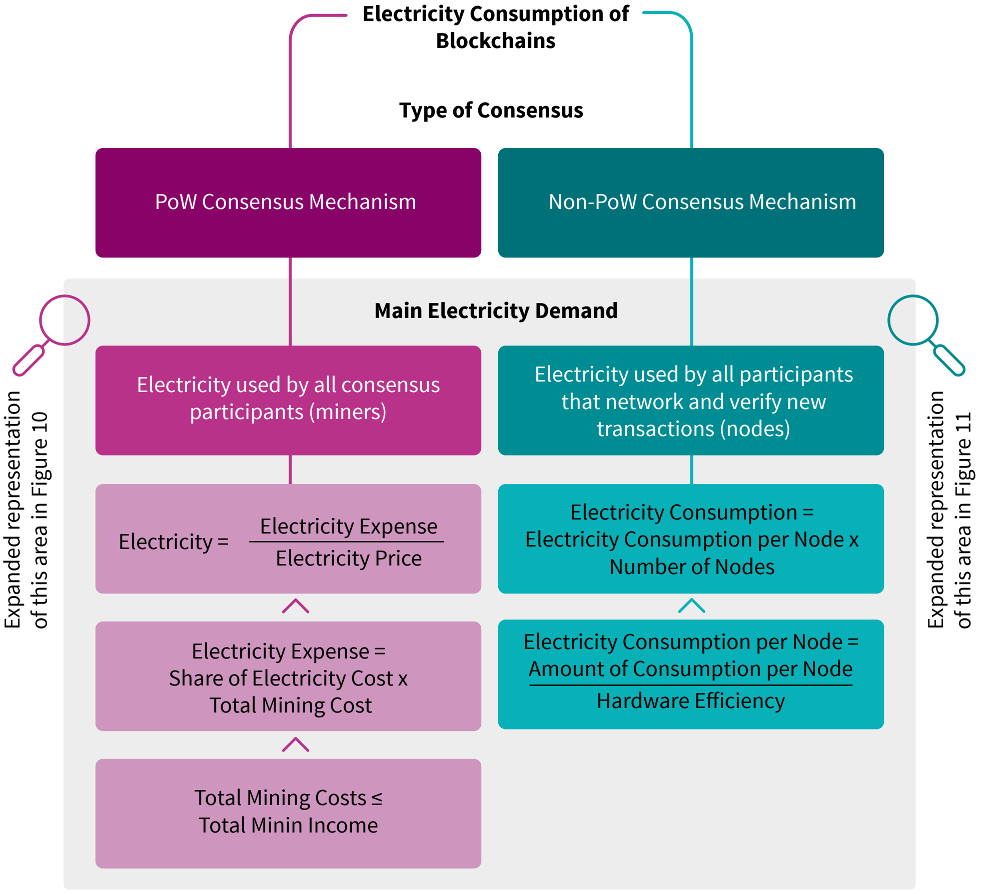
Electricity Consumption of Blockchains
Type of Consensus
PoW Consensus Mechanism
Non-PoW Consensus Mechanism
Main Electricity Demand
Electricity used by all consensus participants (miners)
Electricity used by all participants that network and verify new transactions (nodes)
Electricity = Electricity Expense / Electricity Price
Electricity Consumption = Electricity Consumption per Node x Number of Nodes
Electricity Expense = Share of Electricity Cost x Total Mining Cost
Electricity Consumption per Node = Amount of Consumption per Node / Hardware Efficiency
Total Mining Costs ≤ Total Minin Income
Expanded representation of this area in Figure 10
Expanded representation of this area in Figure 11

Figure 9: The influencing drivers on the electricity consumption of a blockchain network

3.4.1 Electricity Consumption in PoW Networks

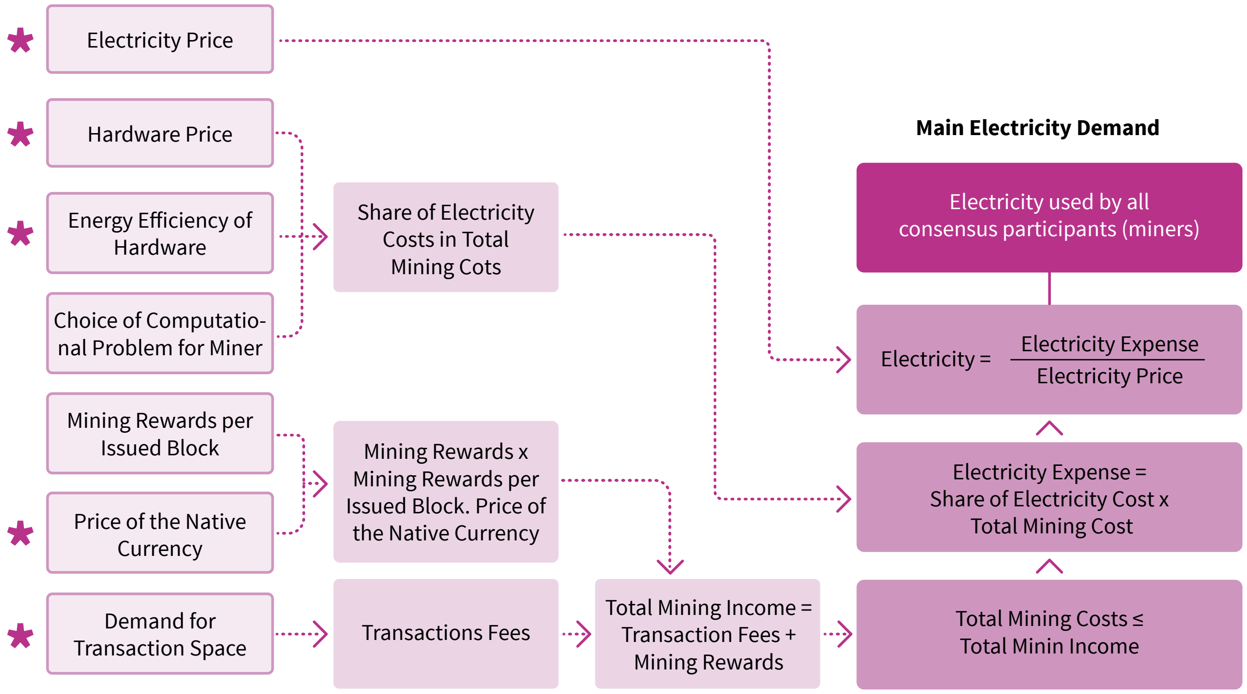
*
Electricity Price
*
Hardware Price
*
Energy Efficiency of Hardware
Choice of Computational Problem for Miner
Mining Rewards per Issued Block
*
Price of the Native Currency
*
Demand for Transaction Space
Share of Electricity Costs in Total Mining Cots
Mining Rewards x Mining Rewards per Issued Block. Price of the Native Currency
Transactions Fees
Total Mining Income = Transaction Fees + Mining Rewards
Main Electricity Demand
Electricity used by all consensus participants (miners)
Electricity = Electricity Expense / Electricity Price
Electricity Expense = Share of Electricity Cost x Total Mining Cost
Total Mining Costs ≤ Total Minin Income

Figure 10: Parameters that determine the electricity consumption of a PoW network

Key Assumptions

In PoW, miners consume electricity to solve complex mathematical problems called hash puzzles. The more electricity they consume, the greater their chances of proposing a block to the network and earning a reward. Compared to the activities of miners, the electricity consumption of other system components, such as nodes or third parties, can be considered negligible. Ethereum's transition from PoW to PoS further illustrates this: electricity consumption was reduced by 99.98 percent due to the absence of hash mining, showing the electricity consumption of other entities, such as full nodes, to be insignificant.

Estimating Network Electricity Consumption

The methodology for estimating electricity consumption follows a top-down economic approach outlined in Chapter 3. This approach allows for a comprehensive perspective by considering both economic and external factors. The process comprises three key elements: first, calculating the miner's total income, which determines the maximum cost of the mining operation; second, determining the portion of the cost allocated to electricity expenses; and finally, calculating the actual electricity consumed.

Total Miner's Income and Total Miner's Costs

Microeconomic theory suggests that a rational miner will never spend more than the revenue generated by the operation, and at the same time, in a competitive market, the margin (i.e., the difference between revenue and expenses) is minimal. Therefore, **the miner's total income** – the sum of mining reward and transaction fees – sets an upper bound on **the miner's total costs** and thus their maximum contribution to the network. For each valid block, a miner receives a reward from the consensus mechanism, thereby incentivizing electricity investment in the consensus process and compensating miners for their investment in hash power. The reward amount depends on the network configuration and can change over time. For example, the Bitcoin network halves the number of coins issued as a reward approximately every four years to prevent currency inflation. Miners anticipate this reduction in income and respond by reducing mining effort and electricity consumption, as evidenced by the most recent halving, which reduced hash power by 20% (Sedlmeir et al. 2020b). However, the long-term effects of this reduction and subsequent halving remain uncertain, as noted in the literature and by several experts.

Mining rewards are paid in the network's native currency, so the value depends on the **current price of the cryptocurrency**. Higher market demand increases the coin's value, potentially mitigating the impact of reduced mining rewards. Conversely, falling coin prices negatively impact miners' income, making mining operations potentially unprofitable, as seen in cases like the 2022 Bitcoin price crash, which led to reduced mining activity. The second source of revenue for miners comes from **transaction fees** paid directly by blockchain users. Higher network activity results in higher transaction fees because the demand for transaction space increases while its supply, i.e., the blockchain's capacity to process transactions, stays constant. Miners can choose which transactions to include in a new block, prioritizing those with the highest fees to maximize revenue. Like cryptocurrency prices, fees paid result from a market mechanism between users and miners, limiting the ability to directly influence transaction fees through network design. Intuitively, reducing average transaction fees could be achieved by offering higher throughput, such as increasing block size or block time. However, the long-term impact on transaction fees is unpredictable, and these changes could negatively impact network functionality and reliability.

Share of Electricity Expenses

The share of electricity expenses depends on several factors. Miners cannot allocate their entire income to electricity as they must cover additional operating expenses. The higher these operating costs, the lower their budget for electricity.

In our model, we have included three factors that affect the share of electricity costs, because they are directly related to the Proof of Work consensus mechanism: **the choice of a computational problem**, hardware costs, and the **electricity efficiency of mining equipment**. It is important to note that additional cost elements, such as personnel costs, data center rent, or operational financing costs, also affect the share of electricity costs but are outside the scope of our model. When designing the consensus mechanism, choosing the type of **hash puzzle or computational problem** to be solved is critical from a network security perspective. For example, networks such as Monero or Ethereum (pre-merge) feature ASIC-resistant hash puzzles, encouraging miners to use mainstream hardware such as consumer GPUs or CPUs. In contrast, the non-ASIC resistance of Bitcoin's hash puzzle favors specialized mining equipment, allowing for greater hardware efficiency. In the existing literature, the choice of the hash puzzle is primarily discussed in terms of network security since it directly influences the consensus mechanism. However, in our interview study, several experts emphasized that the impact of this choice's on the network's electricity consumption should not be underestimated as it directly influences the miner's choice of mining devices. This could affect the proportion of electricity costs, as specialized ASIC hardware has a significantly different hardware-to-electricity cost ratio than generic hardware.

The following two variables – **hardware price and hardware electricity efficiency** – influence each other. Miners can afford fewer machines if the mining hardware gets more expensive. Fewer machines can be run simultaneously if the mining hardware is less efficient. Mining devices with higher electricity efficiency can produce more hash power for the same amount of power. Although the existing literature suggests that simply providing more efficient hardware does not directly reduce a network's electricity consumption, one expert interviewed for this study mentioned that increased hardware efficiency could lead to lower electricity consumption if purchasing more efficient mining equipment requires a higher initial investment.

Electricity Price

The final factor that affects electricity consumption is the **electricity price** for the mining operations. When electricity becomes cheaper, miners use more electricity because it increases the number of hashes they can produce and thereby their chance of receiving a reward. The price miners pay for electricity depends mainly on the location of their mining operation. In addition, such operations are remarkably flexible, not tied to specific locations, and can be established in areas with cheap, reliable electricity. This extreme flexibility is reflected in the ability of miners to respond immediately to changing exogenous circumstances. For example, miners in China have relocated their mining operations several times during the year to follow the fluctuating supply of renewable electricity, such as moving to regions with low-cost hydropower during the rainy season (Gola et al. 2022; Sun et al. 2022). A more recent example is the emergence of the largest mining operations in the United States, driven by low electricity prices in the country, whereby most mining activities are now located in the USA (Coinshare 2022).

The design of a PoW network cannot directly influence how much the miner pays for electricity. One example found in the literature is requiring proof of green electricity as a condition of participation in the mining process. This approach could increase the average cost of electricity for all miners due to the limited supply of green electricity. However, implementing this requirement requires additional governance and a central institution to decide which electricity sources are considered renewable, raising similar issues to those discussed in the context of "proof of useful work". Furthermore, such increased centralized governance could undermine the concept of an open blockchain network, and, with it, the impetus for using electricity as a costly means of reaching consensus.

3.4.2 Electricity Consumption in Non-PoW Networks

Key Assumption

In a non-PoW network, the consensus mechanism does not use electricity as a scarce resource. As a result, the significant component of electricity consumption shifts from the electricity consumed by consensus participants to that consumed by all nodes for storing and processing transactions.

As the background chapter explains, a network has several types of nodes, including full nodes and validators. Both share the same computational cost from processing and verifying incoming transactions and continuously updating the stored ledger. Validators incur additional computational costs that stem from participating in the consensus mechanism by signing blocks and verifying signatures from other validators. In our analysis, we assume an average electricity consumption for a node and include all types of nodes in our subsequent analysis, following the CRCI Institute's methodology.

Estimating Network Electricity Consumption

In a non-proof-of-work context, the network's electricity consumption can be determined in two steps, focusing on the nodes involved: First, you analyze the number of nodes in the network. Since each node is responsible for computing and validating

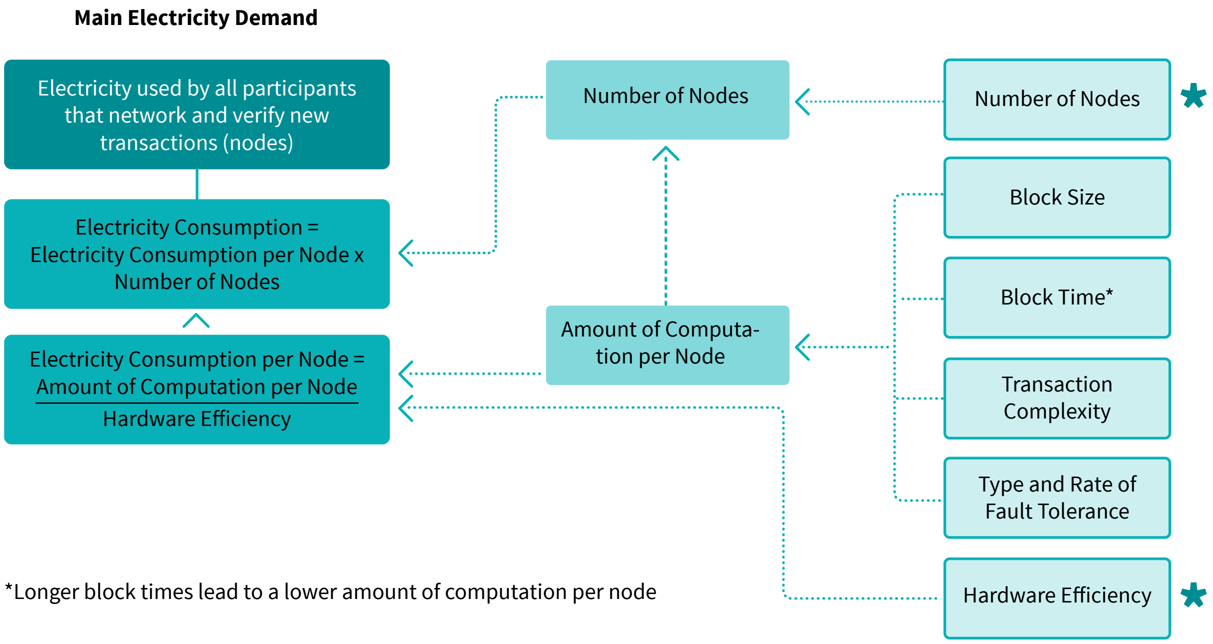
Main Electricity Demand
Electricity used by all participants that network and verify new transactions (nodes)
Electricity Consumption = Electricity Consumption per Node x Number of Nodes
Electricity Consumption per Node = Amount of Computation per Node / Hardware Efficiency
Number of Nodes
Amount of Computation per Node
Number of Nodes
Block Size
Block Time*
Transaction Complexity
Type and Rate of Fault Tolerance
Hardware Efficiency

*Longer block times lead to a lower amount of computation per node

Figure 11: Parameters that determine the electricity consumption of a non-PoW network

identical transactions, each node performs the same computations redundantly.

However, their electricity consumption can deviate depending on the hardware used. Thus you need to calculate the average electricity consumption per node by measuring the electricity consumption for different hardware configurations and estimating the hardware distribution in the network. Following the second step, you calculate the average electricity consumption of a single node. This can be obtained by dividing each node's average computational load by the nodes' average hardware efficiency. Each parameter is explained in more detail in the following sections.

Hardware Efficiency

Hardware efficiency determines the electricity consumption of each node for a given workload. In a non-PoW-based network, node operators only need to provide hardware capable of verifying and storing the transaction in a network. The more efficiently the hardware performs this task, the less electricity a node ultimately requires. Thus, more efficient hardware directly leads to lower electricity consumption in the network.

Hardware efficiency depends on the hardware configuration. For example, a highly configured server with a GPU is less efficient than a lightly equipped Raspberry Pi. Which hardware a node uses depends on the network's hardware requirements, which are determined by the peak computational load a node must handle to remain active on the network. As a node's workload increases, so do its hardware requirements and, ultimately, the hardware selected by the node operator, as observed in our systematic literature review in chapter 3.1.2.

Number of Nodes

The number of nodes contributes to the network's decentralization and augments the calculations' redundancy, increasing electricity consumption. Decentralization is one of the key features of blockchain networks. This refers to the distribution of computational resources and participation in the network across nodes in the network. As the number of nodes increases, the security of the network increases, but the electricity consumed by the network also increases (at least) linearly due to redundant computation of transactions and increased communication overhead. The ability to influence the number of nodes in the network depends on the type of blockchain. For example, in a permissioned network, only a selected group of entities can join, limiting the number of nodes in the network. In contrast, this limitation does not apply to a permissionless network, as anyone can freely join and participate.

However, higher hardware requirements create more significant barriers to entry, as already existing hardware may not be suitable and additional hardware investment is required. As a result, increasing hardware requirements generally reduces the number of active nodes, as higher investment costs discourage potential node operators from participating. Therefore, higher hardware requirements increase electricity consumption per node, but reduce the number of active nodes in the network.

Amount of Computation per Node

The amount of computation a node must handle depends on the blockchain's architecture and actual transaction throughput, as these elements dictate the processing requirements, data storage, and network bandwidth needed to validate, store, and distribute transactions within the system. The load on a node can be divided into two segments: The first is the base load required to participate in the network, which includes node-to-node communication even when no transactions are being processed and the network is idle. The electricity consumption during this state depends on the node's hardware, with more powerful nodes having higher idle electricity consumption. The second segment is the additional effort required to validate new transactions; increased network throughput results in higher node electricity consumption from the computational load.

Parameters Influencing the Average Amount of Computation per Node

The amount of computation a node is required to process is partly influenced by the design of the blockchain. When creating a blockchain network, several parameters can be adjusted to minimize a node's computational load and, consequently, the network's electricity consumption. We have identified four parameters based on a literature review and expert interviews.

In addition to idle load, a node's workload is primarily dictated by maximum throughput capacity, which is determined by factors such as **block size**, **block time**, and **transaction complexity**. Both block time and block size determine the maximum number of transactions that can be executed within a given period. These variables also affect the communication required and can increase traffic between nodes. Larger block sizes require nodes to exchange and validate more data within a block interval. A shorter block time reduces each block's interval, meaning the network must reach consensus and exchange messages between nodes more frequently. Moreover, if the block time is too short, it can increase the risk of forking or chain splitting, leading to redundant computation of the same transactions due to orphaned blocks. Therefore, smaller block size and longer block time can reduce network traffic between nodes and the required bandwidth and, thus, the electricity consumption of the individual nodes as well as the whole system.

Transaction complexity significantly affects the computational load on each node in a blockchain network. The more complex a transaction is, the more computation and electricity it consumes. Transaction complexity depends on factors such as the number of inputs and outputs and, crucially, whether the transaction involves a smart contract operation. Reducing potential complexity, for example, by limiting the maximum execution time of a smart contract, can directly reduce the electricity consumption of each node. This approach reduces the maximum electricity consumed by individual transactions and enables node operators to minimize their hardware, as more complex smart contracts induce more complex computation. For example, Solana enables parallel execution of smart contracts, allowing multiple complex transactions to be executed simultaneously. However, this design feature can only be leveraged by devices that enable multi-threading, leading to higher hardware requirements.

The choice of consensus mechanism, along with its corresponding **fault tolerance rate and type**, can affect the computational overhead. For example, a Byzantine fault tolerance mechanism requires additional communication between nodes to protect the network from malicious actors. This additional communication overhead increases as the number of nodes increases due to the growing need for coordination between them. Conversely, crash fault tolerance requires less communication overhead, resulting in lower electricity consumption per node. Similarly, reducing the rate of fault tolerance results in fewer nodes required to validate a block, thus also reducing the average computational cost for each node.

Additional Concepts for Reducing a Network's Computational Load

Furthermore, we identified three additional technical concepts, primarily for enhancing the network's efficiency by reducing the amount of computation a node has to carry out, which are incorporated in our extended model.

Serverless Blockchain

Firstly, the **serverless blockchain** concept proposed by Sedlmeir et al. (2022b) compromises some decentralization features of conventional blockchain networks to take advantage of the superior performance and adaptability of fully managed cloud services. Blockchain nodes operating within a cloud service can dynamically adjust to the current transaction throughput, decreasing idle consumption as the computational resources in the cloud can be utilized by other services during periods of low transaction activity. Moreover, the electricity efficiency of cloud computing surpasses that of consumer hardware, which diminishes electricity use stemming from computational workloads. Yet, serverless blockchains face significant constraints in their choice of possible cloud service providers, creating a heavy reliance on a single or a few entities.

Sharding

In addition to reducing the number of nodes, as discussed above, system redundancy can be minimized by lowering redundancy within nodes. Instead of having all nodes store the entire blockchain and process every transaction, nodes could hold only a certain amount of data or process only a certain portion of transactions. This concept, known as **sharding**, originally comes from traditional databases. In the blockchain context, sharding metaphorically divides a single, homogeneous blockchain into multiple shards managed by different subgroups of nodes. One implementation of this concept, called "Danksharding", in which a particular type of blockchain data storage is sharded among all nodes, is currently being developed for Ethereum (Ethereum Foundation 2021). Danksharding introduces a new data type distributed across only a part of the network using erasure coding, allowing data to be reconstructed when a particular subset of nodes provides their codes. This type of sharding presents an immediate trade-off between the computational load of individual nodes and data availability since data is stored less redundantly. The concept of sharding can also be applied to a blockchain's transaction processing (i.e., the execution layer), presenting a trade-off between individual nodes' computational load and system integrity. However, the sharding of transaction power is still in its infancy due to various technical and security challenges, such as cross-shard communication and open questions regarding the consensus mechanism.

Rollups

The third design option for improving network efficiency and reducing electricity consumption focuses on reducing redundancy by processing transactions in a more centralized subsystem. **rollups** are designed to process transactions on a separate, faster system (called Layer 2) and then transfer the transaction data and a proof of correct transaction execution to the parent blockchain (Layer 1 or the main network) with significantly reduced electricity consumption. This approach allows the transaction to be processed by the rollup's more efficient system while still benefiting from the security of the main blockchain.

Rollups typically do not have a decentralized consensus mechanism and are managed by a single operator. This introduces more centralization, which reduces redundancy and, therefore, electricity consumption. Nevertheless, rollups store references on the main blockchain in the form of fraud proofs (Optimistic Rollups) or validity proofs (ZK rollups) that ensure the correctness of all rollup transactions, preventing malicious behavior from the centralized party. Rollups maintain a strong link to the main blockchain for security purposes. However, they sacrifice some security regarding system and data availability by introducing a central point of failure. Still, the integrity of the data remains intact. Overall, these subsystems can handle transactions that do not require high availability guarantees, allowing the system to handle more transactions while maintaining or reducing the computational load and electricity consumption of the whole network. While transactions are executed off-chain, the Rollup Operator periodically writes proof of the recent transactions to the blockchain. The type of proof depends on the type of Rollup:

- **Optimistic rollups** use a "fraud-proof system": The initial proof consists of the transaction data that is stored on the main chain. Thus, no cryptographic proof of the validity of the individual transactions is presented upfront. The basic assumption is that the executed transactions are valid until someone challenges their validity. Depending on the design, a specifically designed committee or any participant can challenge transaction validity by highlighting the fraudulent transaction to the smart contract that handles the fraud proofs. If this challenge is successful, the Rollup Operator is financially punished, and the challenger receives a reward. Conversely, if the challenge is unsuccessful, the challenger bears the cost of verification. Due to the lack of cryptographic protection, the processing of the transactions and the submission to the main chain do not cause any computational overhead. However, a significant drawback is that transactions are finalized only after a waiting period to allow time for potential challenges.

- **Zero-knowledge (ZK) rollups** write a validity proof to the main blockchain, a succinct proof of the correctness of all processed transactions. Its succinctness allows proof verification, i.e., verification of the correctness of the transactions processed by the rollup, with far less computational effort than verifying all transactions individually. This cryptographic proof, which is quick to verify, allows transactions to be completed immediately without a waiting period. However, the computational resources required to compute these proofs are significant and cause additional computational overhead and, thus, electricity consumption for the rollup operator.

Excursus: The Electricity Consumption of Rollups

In the previous chapter we identified rollups as a potential design choice to reduce network electricity consumption. The key feature of this concept is that transactions are executed more centrally, so the computational effort is performed only inside a more centralized system and not by all nodes in the blockchain network. While the benefits of this solution in terms of increasing throughput efficiency are widely discussed, the literature has yet to explore its potential implications for reducing electricity consumption. The following section is aimed at addressing this issue.

Since most recent projects use ZK rollups, we will focus on the potential electricity savings that can be achieved by using this approach. To do so, we developed an empirical model that shows how rollups interact with the blockchain network, allowing us to estimate potential electricity savings under different assumptions. We relied on data collected by the CRCI Institute to model the electricity consumption of main blockchains depending on their transaction throughput and the number of nodes. For the electricity consumption of the Zero-Knowledge Prover (the rollups operator), we used the Rollup architecture and prover time provided by Polygon Zero (Polygon Labs 2022), a ZK rollup on the Ethereum network. By analyzing different scenarios, such as the number of transactions or the network's decentralization, we were able to derive four insights for implementing electricity-efficient Rollups:

- **Finding 1 – Validity proof generation is the primary electricity consumer:** In a ZK rollup, most electricity consumption is caused by the generation of the validity proof due to its computational complexity. Therefore, depending on the use case, the frequency of such validity proof generation and, thus, of written transactions on the main chain can be reduced, resulting in a higher potential for electricity savings.

- **Finding 2 – Decentralization within the rollup itself has a minimal impact:** A rollup can also be a network in which several nodes process and store the transactions independently. The additional nodes have a negligible impact on electricity consumption, but only if the rollup network has fewer nodes than the blockchain network. However, since proof generation is the most significant driver of electricity consumption, only one of the nodes should perform this task.

- **Finding 3 – ZK rollups become more electricity-efficient as transaction throughput increases:** The additional computational overhead caused by proof generation depends on the number of transactions, but this dependency is sublinear. Thus, the factor of computational overhead from proof generation decreases with an increase in transactions. Therefore, the electricity savings achieved by implementing a rollup become more significant as the total number of transactions processed in the system goes up or the proportion of transactions processed by the rollup increases.

- **Finding 4 – Rollups only offer potential electricity savings if the blockchain is decentralized:** Only if the network has a certain number of nodes can it prevent redundant computations from outweighing the necessary electricity cost for creating the validity proof. Otherwise, the cost of computing the proof exceeds the electricity savings in the network. Thus, introducing a rollup only provides electricity savings if the number of nodes processing the rollup is significantly (at least two or three orders of magnitude) lower than the number of nodes on the main blockchain. Therefore, even for a centralized rollup, the main blockchain needs to have at least around a hundred or more nodes (also depending on the factors introduced above, e.g., the number of transactions).

4. Designing Electricity-Efficient Blockchain Networks

In this chapter, we present a comprehensive guide to the design of electricity-efficient blockchain-based data infrastructures, taking into account the specific requirements of use cases. This guide serves to provide scientifically informed support for the design process, based on the findings of the previous chapter.

First, Section 4.1 provides an overview of existing blockchain guides and frameworks, providing context and foundational knowledge for our approach. Next, Section 4.2 presents our guide, which consists of two key phases: use case analysis (4.2.1) and network design (4.2.2). Section 4.3 introduces our toolbox for designing electricity-efficient blockchain networks, and finally, in an excursus, we revisit the blockchain trilemma.

4.1 Review of Existing Blockchain Guides and Frameworks

The decentralized nature of blockchain technology introduces new trade-offs and challenges that include technological factors as well as legal and economic aspects. In addition, existing assumptions about the technology must be constantly reevaluated as this rapidly evolving field introduces new concepts such as second-layer solutions and sharding. In addition, misconceptions about the technology, such as exaggerated expectations during the technology's hype phase, have led to it being used for inappropriate use cases, which has resulted in unmet expectations or even complete project failures (cf. Labazova 2019).

In the following, we discuss the various frameworks and decision models that are available in the scientific literature to address these challenges and support the development of blockchain-based networks.

Decision Flowcharts and Models

Most of the academic work focuses on the selection of appropriate blockchain types. Numerous network-based frameworks and decision models can help navigate these challenges and guide the design process. Decision flowcharts, for instance, often use graphs where nodes represent closed-ended questions, with the corresponding answers forming the edges, guiding a path to the final decision. Essentially, these flowcharts are divided into two steps: first, assessing the suitability of blockchain in general for a given use case by asking whether a centralized, and therefore potentially more efficient, infrastructure might be better suited to the use case. If blockchain technology is deemed appropriate, the next step is to select the type of blockchain by determining the degree of decentralization required. Wüst and Gervais (2018) were the first to provide a decision framework based on this structure. Subsequently, several decision models were explored in specific use cases. For example, Pedersen et al. (2019) continued the work of Wüst and Gervais (2018) and proposed a ten-step decision framework, with each step derived based on their findings using a case study from the logistics sector. Building on the analysis of other models, Hunhevicz and Hall (2020) developed another use-case-driven model specifically for the construction industry. Despite their focus on specific industries, the questions posed in both publications are general and universal, making them applicable to various use cases and recommended for further consideration.

However, these models provide only an initial classification because the questions are highly abstract. Belotti et al. (2019) address this limitation in their model by asking more detailed questions about the technology's suitability and presenting possible trade-offs, enabling users to find a more tailored blockchain type for the use case. Several non-academic studies provide a more detailed set of questions in a white paper published by the World Economic Forum (Mulligan et al. 2018), but some of the questions are no longer relevant due to the document's age.

Metrics-Based Frameworks for Evaluating Blockchain Networks

Several authors have proposed metric-based frameworks for evaluating and selecting suitable blockchain networks to provide more refined recommendations. Scriber (2018) proposes a conceptual framework that includes open-ended questions, encouraging readers to self-assess the relative importance of different factors. Some studies have developed different methods to operationalize the properties of individual blockchain networks. Gräbe et al. (2020), Gourisetti et al. (2020) and Kubler et al. (2023) demonstrate that different network characteristics, such as transaction throughput and decentralization, can be operationalized and used to rank the network based on these values. In addition, Kubler et al. (2023) developed an interactive tool to simplify usability of the framework, but unfortunately, this tool is no longer operational.

Guides Addressing the Sustainability of a Platform

Only a few studies have included sustainability considerations in their frameworks, and most of them are limited to giving recommendations for non-PoW over PoW networks. For instance, Ramesohl et al. (2021) extended the ten-step model developed by Pedersen et al. (2019) by an additional step to selecting a consensus mechanism. However, their recommendation only discourages using PoW-based networks due to high electricity consumption. Bada et al. (2021) take a similar approach, extending the model of Wüst and Gervais (2018) to propose a sustainable consensus mechanism. However, their decision model concludes with the question of whether renewable electricity is used by the network participants, disregarding earlier decisions affecting the choice of blockchain type. Moreover, their definition of a green blockchain is based solely on its CO_2 emissions, neglecting the differences between the electricity consumed by PoW and non-PoW networks. Existing frameworks often oversimplify the issue by distinguishing only between Proof of Work and non-PoW consensus mechanisms, neglecting the significant differences in electricity usage within the same consensus mechanism, as discussed earlier in the study.

Existing Guidelines for the Design of Blockchain Networks

While some works cover overall guidelines for the design of a network, none cover the aspect of environmental impact. Schellinger et al. (2022) provide a comprehensive toolbox for building GDPR-compliant blockchain infrastructure by implementing data encryption and rollups. Similarly, Xu et al. (2021) present different design patterns for exchanging sensitive data and deciding whether to store the data encrypted on-chain or off-chain. Finally, Six et al. (2022) provide a literature review of various design patterns that provide general, reusable solutions to common challenges, such as integrating off-chain data storage or tokenizing assets.

4.2 Guide to Designing an Electricity-Efficient Network

Our guide aims to fill the identified gap in guidance for designing an electricity-efficient blockchain network, addressing the need for organizations to reduce their environmental footprint and ensure a reliable and efficient data infrastructure. In the following, we present a systematic approach that combines comprehensive use case analysis with thoughtful network design:

- **Stage 1** focuses on a thorough analysis of the use case. We support this process with questions tailored to covering the fundamental requirements and boundary conditions of the use case in terms of a blockchain-based solution for its data infrastructure.

- **Stage 2** delves into the design of the network, a multi-step process that includes verifying the suitability of a blockchain-based network, selecting the appropriate blockchain type and respective platform, and, finally, designing the network if a permissioned type is selected. The previously established requirements and boundary conditions support the evaluation of different design options, help to understand their influence on the properties of the network, and ensure that the final design provides an appropriate data infrastructure for the use case that minimizes electricity consumption.

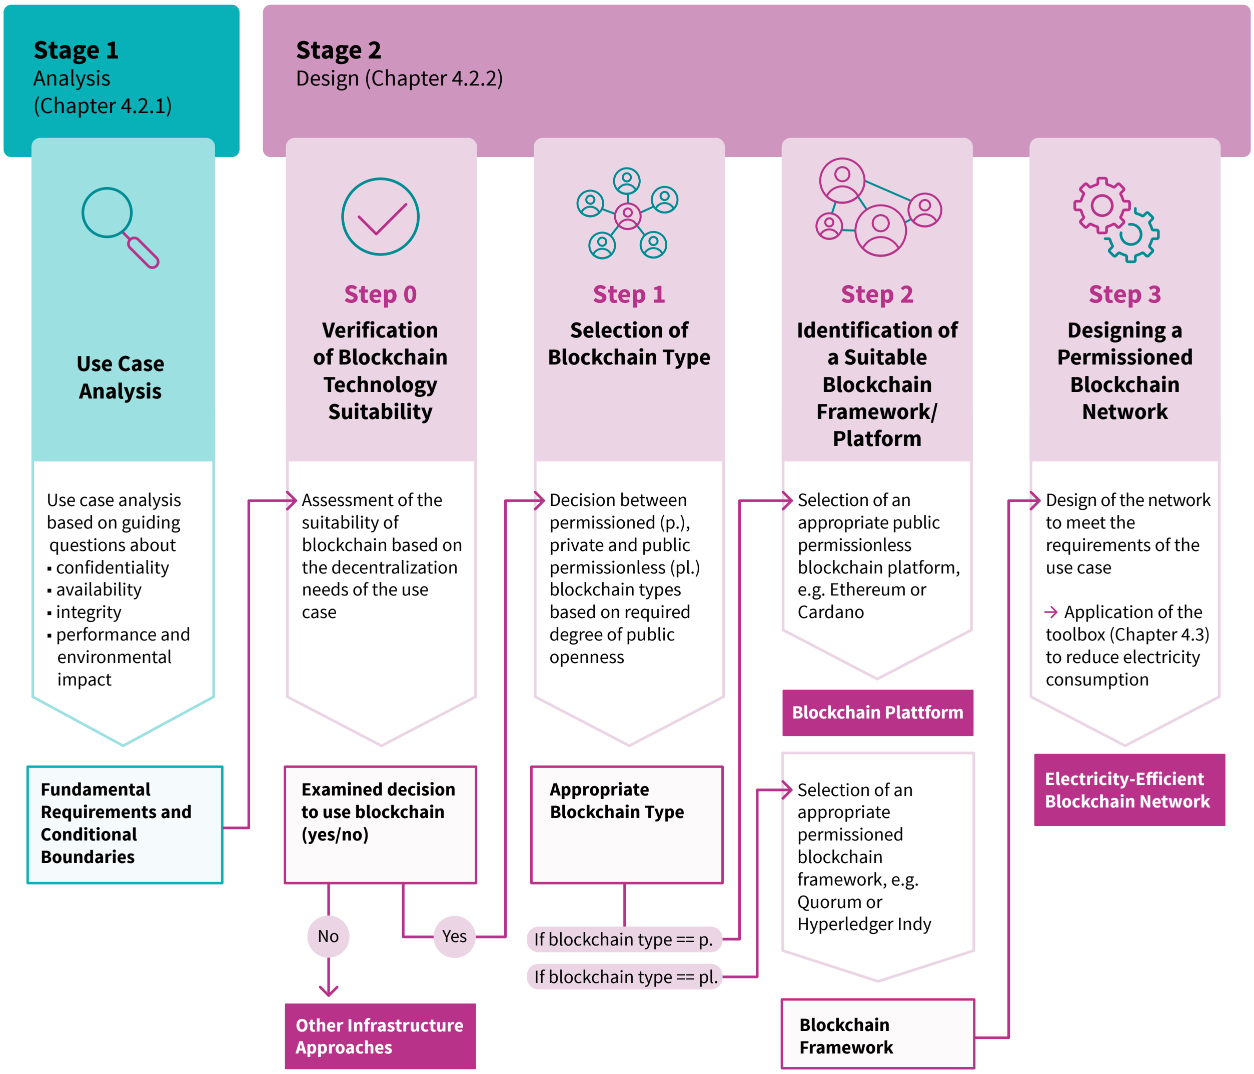
Stage 1
Analysis
(Chapter 4.2.1)
Stage 2
Design (Chapter 4.2.2)
Use Case Analysis
Use case analysis based on guiding questions about
▪ confidentiality
▪ availability
▪ integrity
▪ performance and environmental impact
Fundamental Requirements and Conditional Boundaries
Step 0
Verification of Blockchain Technology Suitability
Assessment of the suitability of blockchain based on the decentralization needs of the use case
Examined decision to use blockchain (yes/no)
No
Yes
Other Infrastructure Approaches
Step 1
Selection of Blockchain Type
Decision between permissioned (p.), private and public permissionless (pl.) blockchain types based on required degree of public openness
Appropriate Blockchain Type
If blockchain type == p.
If blockchain type == pl.
Step 2
Identification of a Suitable Blockchain Framework/ Platform
Selection of an appropriate public permissionless blockchain platform, e.g. Ethereum or Cardano
Blockchain Plattform
Selection of an appropriate permissioned blockchain framework, e.g. Quorum or Hyperledger Indy
Blockchain Framework
Step 3
Designing a Permissioned Blockchain Network
Design of the network to meet the requirements of the use case
→ Application of the toolbox (Chapter 4.3) to reduce electricity consumption
Electricity-Efficient Blockchain Network

Figure 12: Two phases of designing a blockchain network

4.2.1 Stage 1: Use Case Analysis

The design process for a decentralized network presents different challenges than a centralized system, requiring a comprehensive use case analysis to develop an appropriate network design with the required characteristics. To assist practitioners in this task, we will now provide a series of guiding questions that specifically address each of the use case requirements, focusing primarily on aspects relevant to a blockchain-based data infrastructure.

The questions are divided into two sets with distinct perspectives: The first one explores the **fundamental requirements** of the use case, such as the expected transaction throughput and the necessary availability of the system, and aims to determine the properties the final network design must provide. The second set establishes **boundary conditions** that limit the available design options by eliminating impractical or inappropriate choices. For example, the number of participants may impose an upper limit on the network size. Note that the two types of questions are not mutually exclusive. Defining requirements can lead to certain boundary conditions and vice versa.

We have structured the questions along the three parameters of **data security** (comprised of **confidentiality, integrity** and **availability)**, **performance**, and **reduced environmental impact**. We have also numbered these questions for ease of reference. Before delving into these dimensions, we ask some general questions about the data exchanged and the stakeholders involved to set the stage for the analysis that follows.

Stakeholders:

- G-1: Who are the key stakeholders involved in the network, and what are their roles and responsibilities?
- G-2: What level of access and control should each stakeholder have?
- G-3: What resources (financial and technical) can each stakeholder provide?

- G-4: Are there any external stakeholders to consider, such as regulators, auditors, or third-party service providers?

Data and Transactions:

- G-5: What types of data and transactions will be processed?
- G-6: Are there specific privacy or security requirements for the data?
- G-7: How many transactions are to be processed?

Guiding Questions for Security

First, we will present questions that address the security requirements of the data infrastructure. We have labeled each of these questions either "C", "I", or "A", indicating whether they address confidentiality, integrity, or availability, respectively. This help to determine the minimal security requirements for the network.

Fundamental Requirements:

- S-1: What level of reliability and availability does the data infrastructure require? (A)
  - S-1.1 What is the maximum tolerable downtime for the network, and how quickly should the system recover from failures or attacks? (A)
- S-2: Does every transaction have the same availability requirements? (A)
- S-3: Given the potential presence of malicious or faulty participating nodes, how much fault tolerance should there be? (A + I)
  - S-3.1: What share of your stakeholders are honest?
  - S-3.2: What share of your stakeholders are malicious?
  - S-3.3: To what extent can elements of the network be concentrated around individual stakeholders?
- S-4: What are the potential threats to data integrity in this use case?
  - S-4.1 Are there any regulatory or compliance requirements that dictate specific levels of network fault tolerance? (A + I)
  - S-4.2: Are there any regulatory or compliance requirements for data integrity?
  - S-4.3: What number of stakeholders are required to approve a transaction?
- S-5: Do some or all of the transactions have confidentiality requirements? (C)
  - S-5.1: Should there be complete transparency in the network, or should certain transactions only be accessible to certain parties? (C)
  - S-5.2: Will sensitive data that needs to be protected according to legal requirements be written to the blockchain? (C)
  - S-5.3: Are there any organizational or legal restrictions on who can store or access the data? (C)

Boundary Conditions:

- S-6: What kinds of stakeholders will be interested in running a node? (A)
  - S-6.1: Are there any regulatory or compliance requirements that dictate a minimum or maximum number of nodes? (A)
  - S-6.2: Who is eligible to operate a node, and are there any specific criteria or restrictions on node operators?
- S-7: What are the technological capabilities and resources of the actors that will operate the nodes?

Guiding Questions for Performance

The following questions are designed to determine the expected throughput of the network, the maximum latency, and the number of transactions as well as their associated complexity.

Fundamental Requirements

- P-1: What is the estimated number of transactions per second, considering both average and peak transaction volumes?
  - P-1.1: How many parties are expected to use the network, and what will be their typical transaction frequency?
  - P-1.2: Are significant variations in transaction throughput expected?
  - P-1.3: Is there a change in transaction volume to be expected?
- P-2: How computationally expensive is an average transaction?
  - P-2.1: What are the types of transactions performed, and how complex are they?
  - P-2.2: How many smart contract calls are expected and how complex are they?
  - P-2.3: Is there a change in transaction complexity to be expected?
- P-3: What is the desired confirmation time for transactions?
  - P-3.1: Does a transaction always have to be completed within a specific time, or can it also take longer?

Boundary Conditions:

- P-4: Are there any limitations on the part of stakeholders regarding their technological capabilities and resources, such as computing power?
- P-5: Are there limitations to the geographic distribution of network participants that could affect transaction latency or data synchronization?

Guiding Questions for Environmental Impact

Finally, we present questions related to environmental impacts. Due to the study's primary focus on minimizing electricity consumption, other environmental impacts are not directly considered in the following questions. However, these measures in question will have some positive spillovers, such as reducing carbon emissions.

Fundamental Requirements

- E-1: Do specific electricity efficiency targets or regulations need to be considered?

Boundary conditions:

- E-2: Can the use of electricity-efficient hardware and infrastructure be enforced on all participants?
- E-3: Are there restrictions on using specific hardware or equipment that may have a significant environmental impact?
- E-4: Is there a requirement to use specific IT vendors, such as data center providers?

4.2.2 Stage 2: Network Design

The design process we propose is divided into three design steps: selecting the blockchain type, identifying a suitable blockchain platform, and finally, designing the network – if a permissioned network is selected. The toolboxes and frameworks discussed earlier can be used in these steps, as each supports addressing specific aspects of the process. Our toolbox for designing an electricity-efficient network, presented in the following chapter, is to be used in the final step, when the final configuration of the network is being determined. Before starting the actual design, it is necessary to verify the suitability of blockchain technology for the use case.

Step 0: Verification of Blockchain Suitability

An essential part of the design process is verifying that a blockchain-based infrastructure is beneficial for the use case. Although a decentralized network offers certain advantages, it also introduces significant drawbacks compared to a centralized infrastructure, as the distribution and replication of tasks across multiple entities adds complexity and operational challenges (Jagals et al. 2021). For example, in a distributed environment, additional coordination between nodes is required each time new information is written onto the ledger, resulting in increased communication costs compared to a centralized network where a single system performs the task. In addition, the distributed nature of the data requires users to perform separate computations, resulting in higher electricity consumption as multiple actors perform the same computations instead of a single centralized entity. Therefore, evaluating for each use case whether a blockchain-based solution is truly advantageous as opposed to a centralized infrastructure is important. A blockchain-based

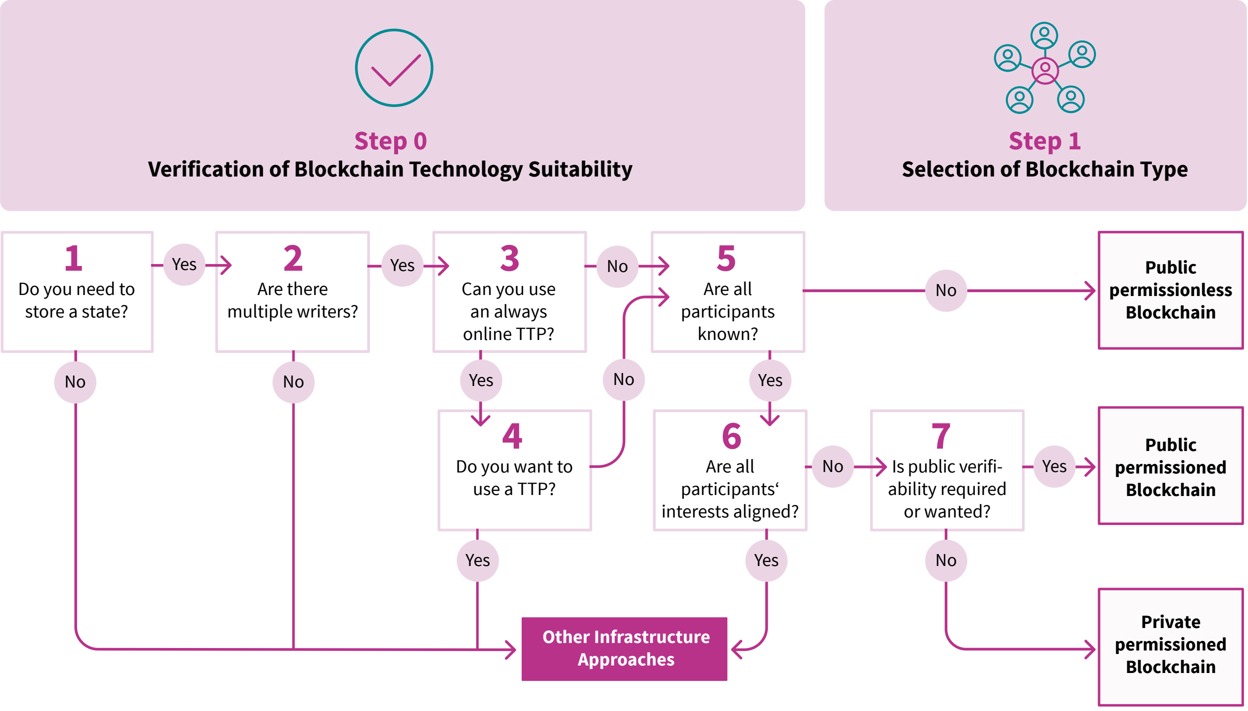
Step 0
Verification of Blockchain Technology Suitability
Step 1
Selection of Blockchain Type
1 Do you need to store a state?
Yes
2 Are there multiple writers?
Yes
3 Can you use an always online TTP?
No
5 Are all participants known?
No
Public permissionless Blockchain
No
No
Yes
No
Yes
4 Do you want to use a TTP?
6 Are all participants' interests aligned?
No
7 Is public verifiability required or wanted?
Yes
Public permissioned Blockchain
Yes
Yes
No
Other Infrastructure Approaches
Private permissioned Blockchain

Figure 13: The decision model adapted from Hunhevicz and Hall (2020)

solution should only be considered if the benefits of the blockchain and its decentralized network outweigh its costs. In the following section, we briefly present Hunhevicz and Hall's (2020) decision model, which serves as a useful initial guide in assessing the suitability of a blockchain-based solution as well as a basis for determining the blockchain type.

As shown in Figure 13, the model consists of a number of questions, which are briefly outlined below.

1. **Do you need to store a state?**: Blockchain essentially functions as a distributed database, and the first question to consider is whether the use case needs to store information persistently. If this is not the case, other data infrastructures, such as those that do not store data immutably, are preferable.

2. **Are there multiple writers?**: Blockchain technology enables multiple entities to interact with the data stored on the ledger independently. A centralized database might be more efficient if only one organization writes data, even when it should be accessible to the public.

3. **Can you use an always-online Trusted Third Party (TTP)?** A TTP should be considered if multiple entities wish to collaborate by exchanging data regularly. A TTP serves as an intermediary responsible for executing, logging, and securing transactions, thus simplifying coordination and enhancing regulatory and privacy compliance.

4. **Do you want to use a TTP?**: Even if implementing a TTP is technically feasible, it results in a central entity having control over all transactions, which may not be in the best interest of all participants. Therefore, the question should be asked whether this is a practical solution.

5. **Are all participants known?**: The next question is about the identity of the participants. If they are unknown, a public permissionless blockchain design such as Ethereum or Solana may be the right choice. Otherwise, other data infrastructure solutions are favorable.

6. **Are all participants' interests aligned?**: Finally, a centralized data infrastructure solution may be more appropriate when all network participants trust each other and have aligned interests. Conversely, a permissioned blockchain can provide a shared infrastructure for a trustless environment.

Step 1: Selection of Blockchain Type

The first design step is determining which blockchain type, permissioned or permissionless, is most appropriate for the use case. Again, the model developed by Hunhevicz and Hall (2020) provides a valuable starting point, focusing primarily on whether all participants are known and on the degree of audibility required, especially by public transparency of all transactions. In addition, we would like to highlight the decision model proposed by Belotti et al. (2019), which also considers the trade-offs of the different properties of a blockchain network.

Step 2: Identification of a Suitable Blockchain Platform / Framework

When considering a permissionless blockchain network, the next step is to identify the most appropriate one for the specific use case. While it is theoretically possible to design a new permissionless network from scratch, this approach often requires significant investment. The high level of security required for a public network – where anyone can join, participate, and manipulate the network – makes this a prohibitively expensive option. Such openness requires incentivizing a critical mass of participants to commit their own resources – such as capital in a PoS network or computing power in a PoW network – to achieve a high enough level of security and maintain a degree of decentralization in voting power. Thus, using existing networks is usually more practical, as these provide the necessary security and credibility to foster acceptance among network participants.

Example for a use case using a permissionless network
Siemens issued €60 million worth of bonds in 2023 on **Polygon**, a public permissionless blockchain. The use of a blockchain-based infrastructure enables significant efficiencies by making processes leaner, faster and less costly than traditional methods.

The use of a permissionless network over a permissioned network was due to the requirement that it be open to everyone, and the transparency and immutability features inherent in public blockchains added a layer of trust and security to the process. The transmission of transactions via a public network provides a decentralization of consensus that gives investors additional confidence.

Choosing an appropriate public permissionless network from the myriad available options is challenging. Unfortunately, no thorough comparison of public permissionless networks is widely available – neither in the academic literature nor in the blockchain community. Since this guide focuses on the design of a permissioned network, readers can refer to the works of Gräbe et al. (2020) and Kubler et al. (2023) for more information on permissionless networks. Both publications provide an overview of relevant factors, such as performance and decentralization, that are necessary to compare and evaluate different blockchain platforms. In addition, the study by Dena (2019) provides a comprehensive overview of permissioned and permissionless blockchain platforms, providing a solid starting point, although it does not cover more recent developments. Since 2019, the relevance of PoW networks as the foundation for individualized use cases has been increasingly questioned. Following Ethereum's switch from PoW to PoS, the top 10 funded blockchains do no longer include a single PoW network that provides full smart contract functionality. Furthermore, as discussed in Chapter 3, a PoW network consumes significant amounts of electricity based on its consensus mechanism, which is in direct conflict to minimizing the environmental impact of the use case. **Given these two arguments, at the time of writing, we consider a**

PoS-based network that supports smart contracts a reasonable choice when selecting a permissionless network.

When choosing a permissionless network, the options for reducing environmental impact are essentially limited to selecting a network with low electricity consumption. The Crypto Carbon Ratings Institute [12] (CCRI) provides real-time measurements on its website to compare the power consumption of the major non-Proof-of-Work networks. Once a network has been selected, it is advisable to optimize the transactional complexity of the use case running on the network. This will reduce the overall computational load on the network, further minimizing the environmental impact.

Identification of a Suitable Permissioned Blockchain Platform for a Permissioned Network

Having covered the considerations for public permissionless networks, we will now consider permissioned networks only. Unlike permissionless networks, which must operate in a trustless environment capable of serving a wide range of use cases, permissioned networks can be designed to be more efficient. This is possible because the participants are known, and therefore consensus must only be found among that group, rather than an unknown number of untrusted participants.

Several blockchain frameworks are available for building permissioned blockchain networks with different design concepts and functionalities. The work of Belotti et al. (2019) offers a holistic comparison across different dimensions, such as software governance, support, latency, privacy, interoperability, and security.

Permissioned Network Use Case Example

The European Blockchain Services Infrastructure (EBSI) is a joint initiative of the European Commission and the European Blockchain Partnership (EBP) to deliver cross-border digital public services on a blockchain infrastructure. EBSI is a permissioned network, meaning that only pre-approved entities can participate in the network. The decision to design it as a permissioned network is primarily based on security, privacy and control considerations. In an authorized network, the identities of the participants are known, which increases trust and facilitates regulatory compliance. In addition, there is a greater degree of control over who can validate transactions, providing robust security. As a result, the permissioned nature of EBSI ensures the secure and regulated environment necessary for government and public services.

Below, we present the three most popular blockchain platforms for permissioned blockchain networks:

- **Quorum** is a permissioned blockchain platform based on Ethereum and maintained as an open-source project by ConsenSys. Its strong connection to Ethereum means it can incorporate fundamental principles and aspects of Ethereum, allowing it to use the same software libraries and design principles. Initially designed for financial applications, Quorum is adaptable to a broader range of use cases that require high-throughput transaction processing within a known, permissioned group of participants. In addition, Quorum supports transaction-level privacy and network-wide visibility based on the needs of the type of data.

- **Hyperledger** is a multi-project open-source collaboration hosted by The Linux Foundation. Its goal is to foster the development of blockchain technologies across multiple industries. Using a greenhouse approach, the Hyperledger Project provides a collaborative space for developing different blockchain frameworks, such as Hyperledger Fabric and Hyperledger Indy, each with unique features and potential use cases. For example, while Hyperledger Fabric provides a generalized blockchain platform, Hyperledger Indy focuses on the need for self-sovereign identity, limiting its functionality to supporting only operations related to the management and verification of decentralized identities. Due to their different design principles, each framework has distinct advantages and disadvantages.

- **Corda**, developed by R3, addresses the specific needs of financial institutions and their challenges in capturing, managing, and automating legal agreements between identifiable and verifiable parties, even across a distributed network. Unlike other blockchain platforms, the framework does not use blocks. Instead, it uses a structure called a “notary” to validate transactions. Both features allow Corda to maintain privacy by sharing transaction data only with the parties involved, a setting that can be adjusted on a per-transaction basis, allowing consensus at the transaction level rather than requiring the entire network.

[12] www.carbon-rating.com

Step 3: Designing a Permissioned Blockchain Network

The final step in the design process is to design a network that provides the desired properties and effectively meets the requirements. This involves making deliberate design choices to achieve the required properties, but it is critical to recognize that these choices often involve unintended tradeoffs with other properties (Kannengießer et al. 2021). For example, it is possible to increase the number of nodes to achieve higher availability. However, increased decentralization results in higher electricity consumption due to more redundant computation and communication overhead required to reach consensus among nodes. Therefore, it is critical to evaluate the impact of each design choice against the previously established requirements. This evaluation helps determine how well the selected features align with the use case objectives. If there is a mismatch between the network properties and the use case requirements, the design choices can be iteratively adjusted to find a design that ultimately provides the required properties.

Security: Confidentiality

Due to the inherent transparency of blockchains, maintaining data confidentiality is a complex challenge, especially when it comes to data exchange (cf. Sedlmeir et al. 2022a). Organizational and legal requirements, such as those outlined in the GDPR, can restrict data access and storage (see questions S-5).

Primary Goal: Define how sensitive data can be shared

Exemplary design questions:

- Can data be stored on the blockchain without encryption or data obfuscation techniques (e.g., zero-knowledge proofs, homomorphic encryption)?
- Is it possible to handle sensitive transactions separately from the main chain in private channels?
- Should the blockchain be accessible to parties that do not host a node?
- Can sensitive data be stored off-chain and only referenced on-chain?

A simple approach is to ensure that only those authorized to participate in the network are allowed to access the data. In addition, some permissioned blockchain frameworks allow restricted access to certain transactions to only a subset of the network. Thus, decisions can be made on a case-by-case basis and participants can access specific transactions (Guggenberger et al. 2022; Capocasale et al. 2023). For example, a consortium of banks can securely exchange transaction data among themselves while

	Corda	Hyperledger Project	Quorum
Major Use Cases	Specialized in the financial industry, digital asset transactions	Modular platform suited for supply chains, finance, healthcare	General application platform, also suitable for the financial industry
Governance	R3	Linux Foundation	ConsenSys
Smart Contract Language	Kotlin, Java	Various	Solidity, Vyper, and Serpent
Consensus Protocols	PoA, PoET	Pluggable consensus protocols – supports PBFT, Raft and others	QBFT, Raft, Istanbul BFT
Throughput	120 - 1000 of transactions per second	Up to dozens of thousands of transactions per second	Dozens to hundreds of transactions per second
Website	www.corda.net	www.hyperledger.org	www.consensys.net/quorum/

Table 3: Comparison of key characteristics and performance metrics[13] of Corda, Hyperledger Project, and Quorum based on Capocasale et al. (2023)

[13] Please note that the consensus protocols and throughput rates can vary depending on the specific configurations and use cases of each blockchain platform, see for example Guggenberger et al. (2022).

keeping it confidential from other network participants. Other approaches provide confidentiality by storing data off-chain, thereby balancing confidentiality and integrity, and various techniques, such as data encryption or zero-knowledge proof, can obfuscate data, so that it can be stored directly on the blockchain without risking its confidentiality (Schellinger et al. 2022; Xu et al. 2021).

Security: Integrity

Choosing the appropriate consensus mechanism for the use case is a critical decision that affects not only the integrity of the data, but also the functionality of the whole. In permissioned blockchains, consensus mechanisms based on Proof-of-Authority (PoA) are commonly used, as there is no need for Sybil resistance since the participants of the networks are known. PoA allows for flexibility in design choices, allowing for customization and balance between integrity and other properties.

An example of this flexibility is the ability to assign voting rights only to a subset of highly trusted participants, such as government agencies, or trusted and regulated entities, such as banks. For example, in a network with a large number of nodes, consensus efficiency can be improved by concentrating voting power within a highly trusted subset (S-3 + S-4). In this way, consensus can be reached more efficiently while maintaining the integrity of the network. Another critical consideration in network design is the choice between fault tolerance and Byzantine crash tolerance.

Primary Goal: Select and configure a secure consensus mechanism

Exemplary Design Questions:

- How should consensus be found?
- How decentralized should the consensus mechanism be?
- Who should have the authority to validate transactions within the network?
- What incentives can be provided to encourage non-malicious participation in the network?
- What type of fault tolerance mechanism is appropriate for the network?

Security: Availability

If a use case requires high data availability, the network can be designed to be highly decentralized, making the network more resilient against individual node failures. The network structure significantly impacts the system's availability, mainly depending on the degree of decentralization. If a use case requires high availability, the network can be designed to be highly decentralized, making the network more resilient to individual node failures. (S-1+, S-8). As the network becomes more distributed, the dependency on individual nodes decreases, avoiding a single or few points of failure (Gojka et al. 2021). Therefore, as a first step, the number of nodes in the network must be set accordingly to achieve the desired reliability.

However, the resilience of the network does not only depend on the number of nodes; the level of availability of each node is also essential. For example, institutional actors such as corporations or government agencies may provide more reliable nodes due to their greater resources and infrastructure redundancy. In addition, nodes should be hosted by different providers in different regions to ensure diversity and mitigate risk (Keller and König 2014).

Primary Goal: Determine the necessary degree of decentralization

Exemplary Design Questions:

- Who are the eligible node operators?
- What criteria should determine the choice of node operators?
- Is there any data suitable for off-chain transmission?
- Can the nodes be hosted with different data centers and providers?
- What is the minimum number of nodes required?

Performance:

By making conscious design choices, the performance characteristics of the network can be closely aligned to the requirements of the use case. For example, network throughput can be controlled almost directly by adjusting block sizes and block times (Sedlmeir et al. 2021a). However, it is crucial to recognize that excessive adjustment of these values can compromise the network's reliability, as the research of Guggenberger et al. (2022) has demonstrated.

In addition to setting block sizes and block times, other design choices and factors play a critical role in determining the maximum achievable throughput of the network (P-1). The number of nodes, the selected consensus mechanism, and the type of fault tolerance also affect the maximum achievable throughput and should be considered when designing the network. It is essential to carefully consider these factors and balance performance with the other properties of the network. In addition, network latency, or, in other words, the delay in communication between nodes, also affects throughput by increasing the time it takes to reach a consensus.

Primary Goal: Achieve necessary network performance

Exemplary Design Questions:

- Can certain computations be offloaded off-chain to enhance performance?
- Can the network be divided into shards to improve scalability and performance?
- How can network latency be minimized to ensure optimal communication between nodes?
- What measures can be put in place to prevent bottlenecks due to high transaction complexity?
- How can the network adapt to fluctuating workloads and dynamically allocate resources to meet performance demands?

Transaction complexity is another critical factor to consider (P-2). Because all nodes compute each transaction redundantly, the network's performance is limited by these nodes' computing resources. Exceeding their capacity can lead to transaction congestion or even a complete network crash. Therefore, when designing a network, it is important to determine what computations must be performed by all network nodes and identify opportunities to offload computations to reduce the overall computational load on the network.

Environmental impact

Minimizing environmental impact requires thoughtful network design. One approach is to avoid oversizing the network, which would lead to unnecessary consumption of resources such as electricity and computing hardware. To this end, each design decision should be targeted and focused, aiming to meet, but not exceed, the specific requirements of each use case for the network. In an info box in Section 4.2, we illustrate how to determine which design options are suitable for an energy-efficient design that still meets availability requirements.

For this purpose, the following chapter presents our toolbox, which proposes several tools and examines their influence on the network characteristics according to the requirements of the use cases. Further considerations such as choosing a data center committed to electricity efficiency, the underlying network infrastructure, as well as the conscious selection of energy-efficient hardware, are essential factors in reducing the overall environmental footprint of the blockchain network.

4.3 Toolbox for Designing Electricity-Efficient Blockchain Networks

Based on our results in Chapter 3.4.2, we identified 11 Tools that could potentially reduce the electricity consumption of a blockchain network.

For PoW networks, we found two design choices that directly affect individual mining operations. By reducing the amount of electricity consumed, the miner provides less hash power, which reduces electricity consumption at the cost of also reducing network security.

For non-PoW blockchains, we observe a more complex and nuanced set of trade-offs, including environmental impact, security, and performance. To derive these considerations, we analyze the design parameters identified in our toolbox and their impact on the properties of the data infrastructure. This allows us to illustrate their influence on the overall network requirements, ultimately assisting us in identifying the most critical trade-offs of each design option.

The following table summarizes the identified tradeoffs. It shows how each design choice affects the goal of a blockchain network. We used symbols to represent each kind of impact: a '+' symbol indicates a positive impact, '+/-' indicates no or minimal impact, and '-' indicates a negative impact. We also marked those relationships where the impact could be neutral or negative (-/o) and neutral or positive (o/+). Note that the table does not capture all side effects under certain conditions and network configurations. For example, if a design parameter such as block size is set too high or too low, it can affect the overall functionality of the network (Kannengießer et al. 2021).

		Security			
Design Choice	Performance	Confidentiality	Integrity	Availability	Electricity Consumption
Increasing Transactional Complexity	+	o	-/o	-	-
Decreasing Block Time	+	o	-/o	-	-
Increasing Block Size	+	o	-/o	-	-
Increasing Fault Tolerance [14]	-	o	+	+	-
Increasing Number of Nodes	-/o	o	+	+	-
Introducing Rollups	+	o/+	o	-	o/+
Introducing Sharding	+	o	-/o	-	+

Table 4: Overview of various blockchain design choices and their impacts on security, performance, confidentiality, integrity, availability, and environmental impact

4.3.1 Tools for Electricity-Efficient Non-Proof of Work Networks

We found ten different design options for influencing the electricity consumption of non-PoW networks, which we assigned to the integrity, availability, and performance trade-offs. It is noticeable that none of them affect confidentiality. This is because they focus primarily on the consensus mechanism, the transaction validation process, and finally, the network structure and do not concern aspects related to data confidentiality.

Figure 14 provides an overview of the identified tools and their associated primary trade-offs in a permissioned non-PoW network. Below, we list each design option along with relevant information in tables organized by the most significant trade-offs. In the first columns, we present introductory information about the design choice and refer to the chapter where we have presented it in more detail. Next, we offer design questions that provide a basis for evaluating the suitability of the design for the use case. Finally, we describe the tool's effect on the network properties to help the reader better assess the impact.

Non-PoW – Security: Integrity

The integrity of the blockchain network, and, therefore, the integrity of the data, is achieved through various cryptographic techniques and the consensus mechanism. To maintain integrity, nodes perform computational tasks and communicate with each other by validating and verifying transactions in the blockchain network. However, as the number of nodes increases, so does the communication overhead (see Section 4.2.2). Each of the tools proposed below aims to reduce the computational burden associated with ensuring integrity. Design choices concerning the level and type of fault tolerance are related to the consensus mechanism, which can reduce the overhead of finding consensus across the network. The introduction of execution sharding further reduces the number of nodes required to find consensus by dividing the consensus process into separate shards. These design choices reduce the workload for each node, resulting in a reduction in electricity consumption. This effect is significant if the workload can be reduced to the point where less powerful devices are needed, thereby significantly reducing the average energy consumption of a node.

[14] In the table we have summarized "Change Rate of Fault Tolerance" and "Change Type of Fault Tolerance" into one item.

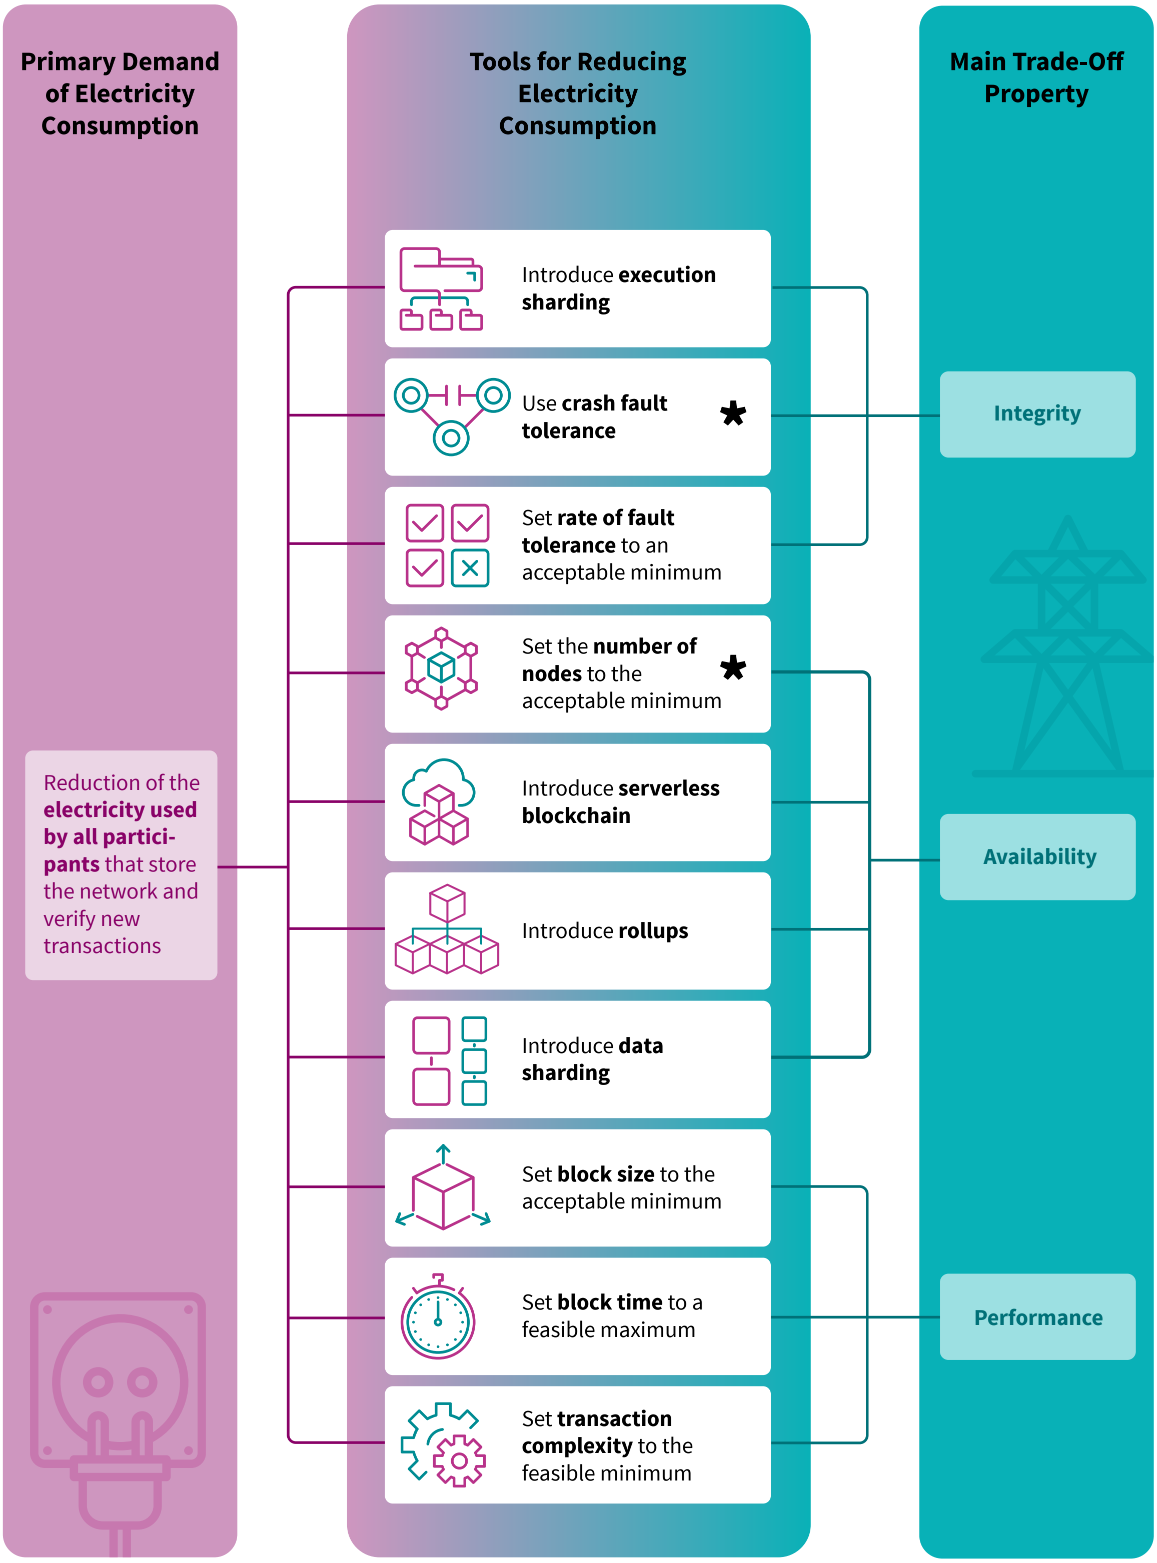

Primary Demand of Electricity Consumption

Tools for Reducing Electricity Consumption

Main Trade-Off Property

Reduction of the **electricity used by all participants** that store the network and verify new transactions

Introduce **execution sharding**

Use **crash fault tolerance** *

Set **rate of fault tolerance** to an acceptable minimum

Set the **number of nodes** to the acceptable minimum *

Introduce **serverless blockchain**

Introduce **rollups**

Introduce **data sharding**

Set **block size** to the acceptable minimum

Set **block time** to a feasible maximum

Set **transaction complexity** to the feasible minimum

Integrity

Availability

Performance

* The asterisk marks those design options, which can only be used in a permissioned network

Figure 14: Tools for designing an electricity-efficient non-PoW network

Tools	Description	Guiding Questions	Trade-Off
Introduce execution sharding	Sharding is a scalability approach that divides a blockchain network into smaller, more manageable units called “shards” with their own processing power. Each of these shards is responsible for processing a specific set of transactions and executing smart contracts, rather than all nodes handling all transactions. As a result, this approach significantly reduces the electricity consumption of individual nodes. ▶ Chapter 4.2.2	Are the integrity guarantees still sufficient? Can I distinctly divide the networks into different shards which can act independently? Would I have too many interactions between the different shards?	Execution sharding weakens integrity because it no longer relies on the majority in the entire network to reach consensus, but only on the majority in each shard. This makes it easier for a malicious party to gain control and provides an opportunity to tamper with or manipulate the data handled by the shard, which is then distributed throughout the network. In addition, the network design becomes more complex due to the added complexity of handling different shards, which can introduce new failure points or vulnerabilities.
Use crash fault tolerance	In a Byzantine fault-tolerant network, nodes collaborate to validate data despite faulty or malicious nodes, leading to increased communication. A crash fault-tolerant mechanism can be employed if all nodes are trustworthy, reducing the communication needed for consensus. ▶ Chapter 4.2.2	Is it certain that no participant will act maliciously? Is there a high number of nodes and a short block time to significantly reduce electricity consumption?	Increases the risk of a successful attack on data integrity, as only one malicious actor is required.
Set the rate of fault tolerance to a sufficient level	The fault tolerance rate specifies how many faulty or malicious nodes the network can tolerate before it faces data inconsistencies or failures. Lowering the rate can decrease the communication between nodes, thereby reducing the computational load in the network. ▶ Chapter 4.2.2	What is the lowest threshold for the share of trusted nodes? Does the network have sufficient transactions and validators to justify electricity savings at the cost of reduced attack resistance?	Reducing fault tolerance in a blockchain network can increase the risk of successful attacks on data integrity since fewer malicious or faulty nodes are needed to compromise the network.

Table 5: Tools that involve a trade-off concerning the integrity property of a non-PoW network

Non-PoW – Security: Availability

Data and system availability in a blockchain network is primarily facilitated by decentralization. The following tools reduce the degree of decentralization in various ways, targeting redundant computation.

A lower number of nodes directly reduces redundant computations within the network, leading to a disproportionate reduction in electricity consumption and less communication overhead between nodes. Serverless blockchains are another form of centralization. They rely on limited cloud service providers for their basic infrastructure. While this approach provides high availability, it introduces potential outage risks associated with these providers. In addition, rollups and data sharding introduce a degree of centralization within subsets of the network. Rollups consolidate transaction processing onto a single node operator, while data sharding distributes data storage across a subset of nodes in the network. These strategies effectively reduce the electricity consumption for individual nodes by streamlining transaction processing and data storage. However, they also present challenges. For example, Rollups can only handle certain types of transactions, and only a limited number of market-ready implementations are available.

Tools	Description	Design Question	Main Trade-Off
Set the number of nodes to the acceptable minimum	The decentralization of a blockchain network is characterized by the number of nodes, each of which must compute the same number of operations. Therefore, reducing the number of nodes sufficient to achieve the required uptime reduces the number of redundant computations and thus the network's electricity consumption. ▶ Chapter 3 .2.2	Can the number of nodes be reduced without compromising the required uptime while maintaining the required availability? Does the node consume enough electricity to merit the reduction?	Reducing the number of nodes results in a more centralized network, which affects network availability by increasing the risk of system failure or data loss. In addition, centralizing the network can lead to a higher risk of transaction censorship because communication is concentrated in fewer nodes, compressing the availability of the services.
Introduce serverless blockchain	In serverless blockchains, nodes are hosted by cloud service providers, which allows computing resources to be elastically adjusted based on current transaction throughput rather than continuously tuned for peak capacity. In addition, the high reliability and availability of cloud services can potentially reduce the number of nodes required. ▶ Chapter 3.5	Is it possible to enforce the use of a serverless infrastructure? Are there significant and predictable variations in transaction throughput?	The effects on the network's availability by limiting the infrastructure to only a few cloud service providers are unclear (Keller and König 2014).

Tools	Description	Design Question	Main Trade-Off
Introduce rollups	Rollups aggregate transactions through a single or few rollup operators, which store proof of their correctness on the main blockchain. Verifying these aggregated proofs is less computationally intensive than verifying individual transactions, reducing electricity consumption for other blockchain nodes. ▸ Chapters 4.2.2 and 4.3	Is there a subset of transactions that do not have very high requirements for availability guarantees? Is the number of network nodes and potential for transaction aggregation in rollups sufficient to justify the computational overhead for the rollup operator?	Rollups centralize transaction processing on a single node or a few nodes, increasing the risk of censorship and introducing potential single points of failure. However, the integrity and long-term availability of transactions remain secure because proofs are stored on the main blockchain, preventing malicious behavior and allowing manual transaction processing.
Introduce data sharding	Sharding is a scalability approach that divides a blockchain network into smaller, more manageable units called “shards”. With data sharding, each shard has its own data storage. The transferred data is only validated and stored within a shard, reducing total electricity consumption. ▸ Chapter 4.2.2	Are the availability guarantees still sufficient when only a part of the network stores the data? Would I have too many interactions between the different shards?	Data Sharding impacts data availability because the data is only held by one part of the network instead of being distributed throughout the network.

Table 6: Tools that involve a trade-off concerning the availability property of a non-PoW network

Non-PoW – Performance

In a non-PoW network, performance is closely related to the computational load that each node handles. Adjusting block size and block time to reduce the network's potential throughput results in an almost proportional decrease in computation and memory utilization. In addition, minimizing transaction complexity directly affects the computation a node must perform. The effect on electricity consumption is marginal as long as only computation and memory usage are reduced. However, it becomes much more significant when these changes result in reduced hardware requirements, allowing the use of more electricity-efficient devices.

Tools	Description	Design Question	Main Trade-Off
Set block size to the acceptable minimum	Network throughput requirements are often based on peak situations, resulting in hardware oversized for average throughput. Therefore, reducing the block size reduces the maximum throughput while reducing the hardware requirements for the nodes, allowing for a more electricity-efficient hardware configuration. ▸ Chapter 4.2.2	Can the network's maximum throughput be reduced? How low can the maximum throughput be?	Reducing the block size reduces the maximum throughput of the network and may increase latency when the network is operating at maximum capacity.
Set block time to a feasible maximum	Block time determines the average transaction processing time for a blockchain. However, not all data needs to be processed immediately, so increasing block time could reduce electricity consumption and coordination overhead, directly impacting computational load. ▸ Chapter 4.2.2	Can the network's maximum throughput be reduced? Can transactions be processed with higher latency? How high can the maximum latency be?	Increasing block time directly decreases the network's maximum throughput and increases the average latency of transaction processing.
Set transaction complexity to the feasible minimum	Reducing the maximum transaction complexity directly reduces the maximum computational load caused by peaks and thus can reduce the hardware requirements of participating nodes. ▸ Chapter 4.2.2	Do all computations need to be performed on the blockchain? Are there calculations executed on the blockchain, or only the information stored?	The same transaction may need to be executed with multiple smart contract executions due to limited transaction complexity. However, these executions are performed sequentially, one after the other, in separate blocks, increasing transaction duration and directly impacting network performance.

Table 7: Tools that involve a trade-off concerning the performance property of a non-PoW network

4.3.2 Tools for Electricity-Efficient Proof of Work Networks

In the case of PoW blockchains, the main influencing factor is a reduction in hash performance, which directly leads to a reduction in electricity consumption. This focus arises because miners dominate the majority of electricity consumption in a PoW network, as other entities within the network contribute only a small fraction of consumption in comparison. Therefore, the proposed measures target the economic decisions of miners by changing their revenue and cost structures to incentivize lower electricity usage (Figure 15.). However, it should be noted that even though we have identified these tools to reduce the electricity consumption of a PoW network, the consumption would still be extremely high compared to a non-PoW network. Therefore, the use of a PoW network should be carefully considered.

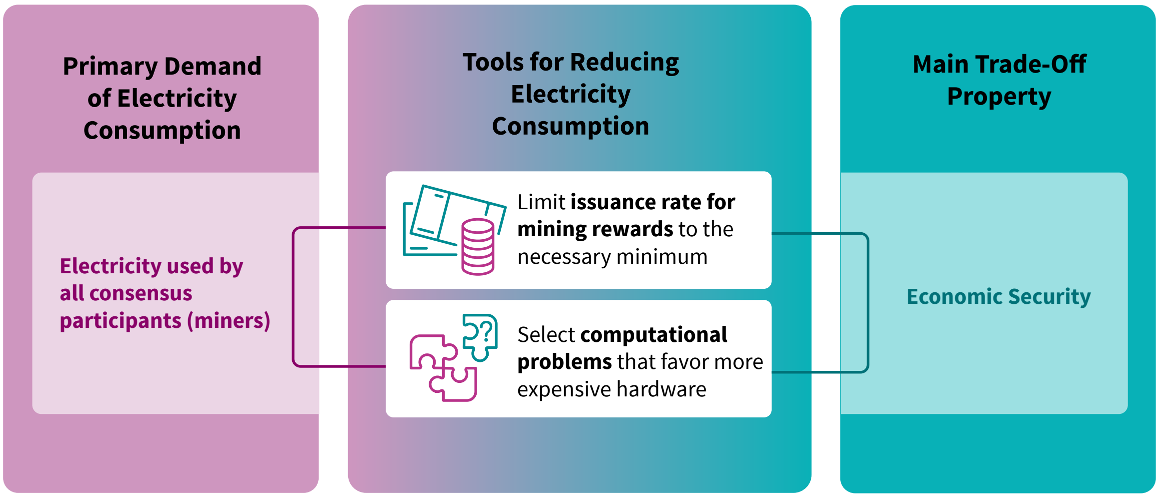
Primary Demand of Electricity Consumption
Electricity used by all consensus participants (miners)
Tools for Reducing Electricity Consumption
Limit **issuance rate for mining rewards** to the necessary minimum
Select **computational problems** that favor more expensive hardware
Main Trade-Off Property
Economic Security

Figure 15: Tools for designing an electricity-efficient PoW network

We have identified two design choices that directly affect the economic decisions of miners and, in turn, influence the network's overall security . Firstly, the rate of issuance for mining rewards can be reduced. This measure impacts the miner's income, likely prompting them to lower their electricity consumption. The second design decision involves selecting a computational problem that increases the hardware-to-electricity cost ratio. This approach encourages miners to invest in more expensive and specialized ASIC hardware, thereby reducing the overall electricity consumption of the network.

Tools	Description	Design Questions	Main Trade-Off
Limit issuance rate for mining rewards to the necessary minimum	The revenue distributed by the network gives miners incentives to invest more resources in the consensus mechanism. Thus, reducing mining rewards can encourage fewer miners to participate in the consensus mechanism with less effort, leading to less electricity used as a scarce resource. ▸ Chapter 4.2.1	Can the guarantee for the system's integrity, i.e., the cost to obtain 51% of the hash rate, be reduced? Are there other significant income streams for miners in addition to mining rewards?	Reducing the miner's income immediately reduces the amount of hash power a miner provides, thus reducing the network's security and making a successful attack on the network more feasible.
Select computational problems that favor more expensive hardware	The computational problem defines the most appropriate hardware. Accordingly, the choice of the computational problem can influence the miner's equipment by favoring specialized ASIC hardware and thus the share of income spent on electricity. ▸ Chapter 4.2.1	Does the higher initial cost of a 51% attack due to the cost of specialized hardware, coupled with the lower operating cost (due to lower recurring electricity costs), change the integrity guarantees of the system? Does this introduce a centralization in hardware production which could lead to a centralization in mining operations?	A non-ASIC-resistant consensus mechanism can reduce security by increasing the concentration of mining power and favoring the vendors of this specialized hardware. It also leads to higher entry barriers since this specialized hardware is expensive and resource-intensive. Such a centralization of consensus power can increase the risk of a successful 51% attack.

Table 8: Tools that involve a trade-off concerning the security property of a PoW network

Excursus: Revisiting the Blockchain Trilemma

The widely accepted Blockchain Trilemma was introduced by Vitalik Buterin, the founder of Ethereum, in 2017 in a blog article[15] that covered the concept of sharding. The trilemma involves the three aspects of decentralization, scalability, and security, which were initially defined as follows:

- Decentralization is achieved when the blockchain can run without depending on a small group of trusted entities. Thus, a sufficient number of nodes is required, which can only be achieved if the hardware requirements for nodes are limited to a certain degree.
- Scalability of a system is achieved when the system itself can process more transactions than an individual node can process. This is desirable because end user's hardware is typically constrained when decentralization is one of the design objectives.
- Security is achieved if the system is resistant to an attacker with resources of the magnitude of the system itself, e.g., in terms of computational (PoW) or economic (PoS) power.

The Interrelationship of the Established Blockchain Trilemma

In the early days of blockchain development and research, these properties were seen as binary categories, i.e., a property was either fulfilled or not fulfilled, and the main blockchain design concepts could only achieve two of the three aspects. For instance, both Bitcoin and Ethereum were deemed as decentralized and secure while not being scalable (more nodes did not mean higher throughput) and having a low throughput of typically only a few transactions per second.

Other blockchains that focused on scalability were also considered secure but achieved high transactional throughput by relying on more performant and, thus, more centrally operated infrastructure. In this sense, these solutions, including private permissioned blockchains, sacrifice decentralization by limiting the participation to a selected set of participants. Multi-chain ecosystems based on multiple application-specific and somewhat connected blockchains could achieve scalability and decentralization, but the individual chains would not meet the formulated security requirements.[16]

Some Aspects of the Trilemma Are Outdated

With respect to scalability, the three aspects of the trilemma are no longer viewed as binary options, but rather as more nuanced spectra with different degrees of satisfaction, allowing for a wider variety of trade-offs between particular levels of these aspects. In addition, the understanding of some aspects of decentralization, and probably even more so of scalability and security, has evolved and become more complex. For example, networks such as Solana or Polygon, which are more scalable and decentralized than the permissionless networks that existed when the original trilemma was formulated, while maintaining a high level of security, have emerged. This raises the question of whether the Trilemma, which is still popular and widely accepted but was conceived as a decision between binary options, is still applicable today as the view on these aspects is becoming more fine-grained to include a panoply of novel findings from new design choices and the current state of research. These new complexities are also reflected in the results of this study.

More specifically, the new complexity is reflected in the analysis of the design parameters, as shown in Table 4 in Section 4.3, which leads to four insights based on the impact of design options on different objectives in a non-POW network:

- **Increasing performance compromises availability:** Design parameters that increase performance and thus lead to higher hardware requirements can decrease the network's availability by reducing the number of participating nodes due to the higher entry barriers, resulting in a more centralized network (Gallersdörfer et al. 2022; Platt et al. 2021). In addition, setting the block size too large and the block time too short can result in not all nodes receiving new proposed blocks, increasing the risk of network fragmentation and compromising integrity (Kannengießer et al. 2021). Similarly, increasing the number of nodes can lead to a similar trade-off between performance and availability (Rieger et al. 2022).
- **Increasing integrity reduces performance:** Consensus mechanisms with higher fault tolerance can reduce the network's performance since additional communication overhead is introduced (Sedlmeir et al. 2021a).

[15] https://vitalik.ca/general/2017/12/31/sharding_faq.html, accessed 08.08.2023

[16] https://vitalik.ca/general/2021/04/07/sharding.html, accessed 08.08.2023

- **Reducing the environmental impact affects performance and security:** Lowering environmental impact conflicts with performance and security. As our analysis of the electricity consumption of a blockchain network has shown, all design parameters increase the overall computational overhead within the network, which directly impacts its electricity consumption.

- **Sharding and rollups can mitigate the trade-offs:** Both techniques aim to improve network efficiency, resulting in better performance with fewer trade-offs. In addition, ZK rollups can enhance confidentiality. However, despite their advantages, these solutions cannot fully address all the trade-offs the revised blockchain trilemma presents.

After reflecting on the four key insights from exploring the interplay between different design options in a non-PoW network, it is clear that a new approach to the Blockchain Trilemma is needed. The newly identified complexities require a more nuanced model that considers the significant trade-offs between performance, security, and environmental impact. In Figure 16 below we present our Revised Trilemma. It is reformulated to incorporate the newly identified trade-offs and provides a more accessible way to designing blockchain networks with electricity efficiency in mind.

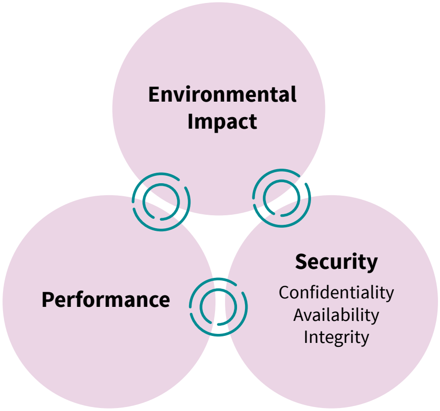

Figure 16: Our new revised blockchain trilemma, applicable to non-PoW blockchains. It displays the area of design choices for blockchains and their effects as a field of tension between environmental impact, performance, and security.

5. Case Studies

Having established the theoretical underpinnings of our guide and detailed its development process, we now move to a more practical segment of this study: the application of our guide to various use cases. The goal of this chapter is to put our guide into practice by demonstrating its functionality in several case studies, each of which addresses a different aspect of digitization and sustainability. Our case studies start with a decentralized electronic prescription system, where we apply our complete guide in detail to a healthcare context. This is followed by a Green Labeling case, which focuses on certifying the origin of electricity, and a Self-Sovereign Identity (SSI) case, which addresses digital identity management. The latter two serve to further illustrate the underlying concepts of the first case.

5.1 Case Study: Electronic Prescription

In the healthcare industry, paper-based prescription management efficiently reduces the risk of patient harm from incorrectly issued prescriptions. In order to digitalize this process and also take a patient-centric approach, Schlatt et al. (2023) present a solution for electronic prescriptions based on Self-Sovereign Identity (SSI) principles and blockchain technology.

The doctor writes the prescription and sends the e-prescription with all the necessary details as a Verifiable Credential (VC) via a secure, end-to-end encrypted channel to the patient's Digital Wallet[17] . The patient can store this credential on their smartphone and present it to the pharmacy of their choice to verify the authenticity of the e-prescription. By capturing the prescription in a verifiable credential, there is no need to store sensitive information in a centralized system or on the blockchain.

Although this approach ensures the integrity of prescription information, it does not provide a means to track the number of times a prescription is filled. To prevent abuse through multiple refills of the same prescription, a token associated with the prescription is stored on a blockchain that serves as the data infrastructure for this use case. The token contains only the prescription ID and a value about the validity of the token, preventing the recording of personal information. Thus, the underlying architecture of the use case consists of two blockchains, one to provide the data infrastructure for the SSI component and the other to manage the tokens. In the following use case and testing of the toolbox, we will focus on the data infrastructure responsible for the tokens, as we have already discussed how to design an SSI network to be electricity efficient. Figure 17 illustrates the infrastructure responsible for token management.

Stage 1: Use Case Analysis

Due to the critical nature and sensitivity of healthcare data, the healthcare sector places high demands on IT services and their underlying infrastructure. Defining these requirements allows us to determine the necessary characteristics of the network to suit the use case. To simplify the complex healthcare landscape, we focus on three stakeholders: pharmacists, physicians, and patients. While each stakeholder needs access to the information about the tokens, only doctors and pharmacists can write and modify them. Integrity is essential for pharmacies, as they are legally responsible for verifying the authenticity of a prescription before dispensing the medication. We will analyze each aspect of the use case's data infrastructure requirements in detail below.

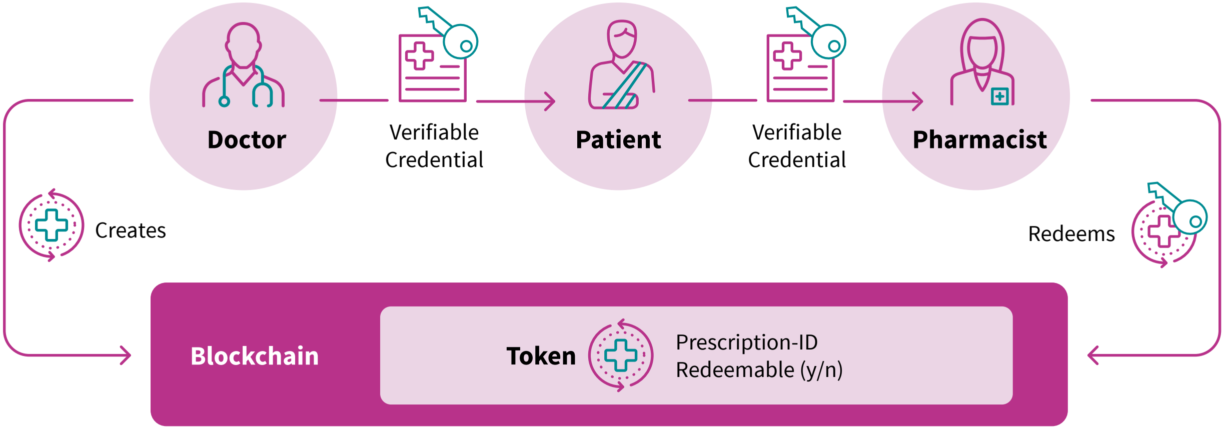

Figure 17: Prescription Token Process - Steps for verifying and invalidating tokens on the blockchain

[17] Self-Sovereign Identity (SSI) digital wallets are secure and private digital locations where individuals can manage and control their own digital identities, storing and selectively sharing personal data and credentials. We introduce this concept in more detail in an additional use case.

Confidentiality:
The strict decoupling of sensitive data from the prescription as a credential and from the tokens means that no sensitive data is processed on the blockchain, reducing the confidentiality requirements to a minimum (S-5).

Integrity:
Prescriptions are sensitive documents that require high integrity to ensure validity and authenticity. Their integrity depends on the pharmacist always being sure

of their authenticity and validity, meaning that the prescription has not yet been filled. To maintain the required level of integrity, **selecting a robust consensus mechanism that can effectively detect and mitigate any malicious activity is critical**, ensuring the reliability of the token and the whole system (S-3).

Availability:
When it comes to availability, high reliability and accessibility are crucial for healthcare applications. As a critical infrastructure, our aim is for the system to achieve 99.99% uptime in a year (S-1). In order to achieve this, we need to determine the required decentralization of the network and possible node operators. Scientific literature suggests **that each pharmacist should operate their own node**, due to their economic interest and regulatory mandate to validate prescriptions (see for example Tayler et al. (2022)). This will ensure that trusted stakeholders in the healthcare ecosystem manage the network infrastructure. Doctors and patients are not considered potential node operators in this use case. The exclusion of physicians is in line with the specific requirements of the use case, and patients, as end users of the system, should not face any barriers or obstacles to participating in this system. Assuming that each pharmacy in Germany operates one node, we estimate a maximum of approximately 18,0000 nodes[18] . Regarding the lower bound of nodes, **the use case does not specify any conditional constraints other than the requirement of decentralization with multiple independent entities.**

Performance:
First, we must assess the expected transaction volume to understand the network's anticipated workload and define the required performance specifications. For context, around 700 million prescriptions are written and filled in Germany each year. Analysis of prescription patterns shows that most transactions occur during regular business hours on weekdays when doctors' offices and pharmacies are open. Conversely, the volume of prescriptions is significantly lower on weekends. For example, about 5 million prescriptions are issued on a Monday, but only about 200,000 on Saturday and Sunday combined[19] . Using these numbers and a simple rule of thumb, **we can roughly estimate that the peak transaction rate during business hours is about 350 transactions per second, and during off-peak hours the transaction rate drops to about five transactions per second** (P-1).

For user acceptance, token generation and validation must be performed quickly. Studies show that longer processing times lead to dissatisfaction, **so the system must achieve a maximum processing latency of less than two seconds (P-9 and P-8)**. Next, let us consider the complexity of the transactions. Since the network only has to manage the tokens, the computational complexity is minimal. **The computation carried out in the network only includes the creation of the token when the prescription is issued and its invalidation by the pharmacy after the prescription has been filled (P-3)**. Since the infrastructure only covers e-prescriptions, no other applications need to be considered in this case study, and for simplicity, we assume no changes in the number of transactions, transaction complexity, and the number of participants in the future (P-3).

Phase 2: Network Design
After defining the requirements for our use case, we can now proceed with the design process. Since we are dealing with a decentralized network, we can assume that a blockchain-based infrastructure is suitable for this particular use case. Next, we follow Hunhevicz's decision model to determine the blockchain type, which asks two questions: The first has to do with the identity of the participants. Since we are in a regulated environment and only known and certified participants host the nodes, we can opt for a permissioned network. This allows us to use a more efficient consensus mechanism without Sybil resistance. The second question relates to the accessibility of transactions and the need for transparency and public accountability. Patients need real-time visibility into the status of their prescriptions to ensure trust and control over their prescriptions. Therefore, we chose a public permissionless network for efficiency and to allow us to customize the structure of the network for our use case. As a blockchain framework, we propose to use Quorum, which is flexible and performant enough for the use case and has been successfully deployed in various public sector applications

[18] https://www.abda.de/aktuelles-und-presse/zdf/, accessed 08.08.2023

[19] https://fachportal.gematik.de/fachportal-import/files/gemSysL_eRp_V1.1.0.pdf, accessed 08.08.2023

Evaluation of Design Choices:

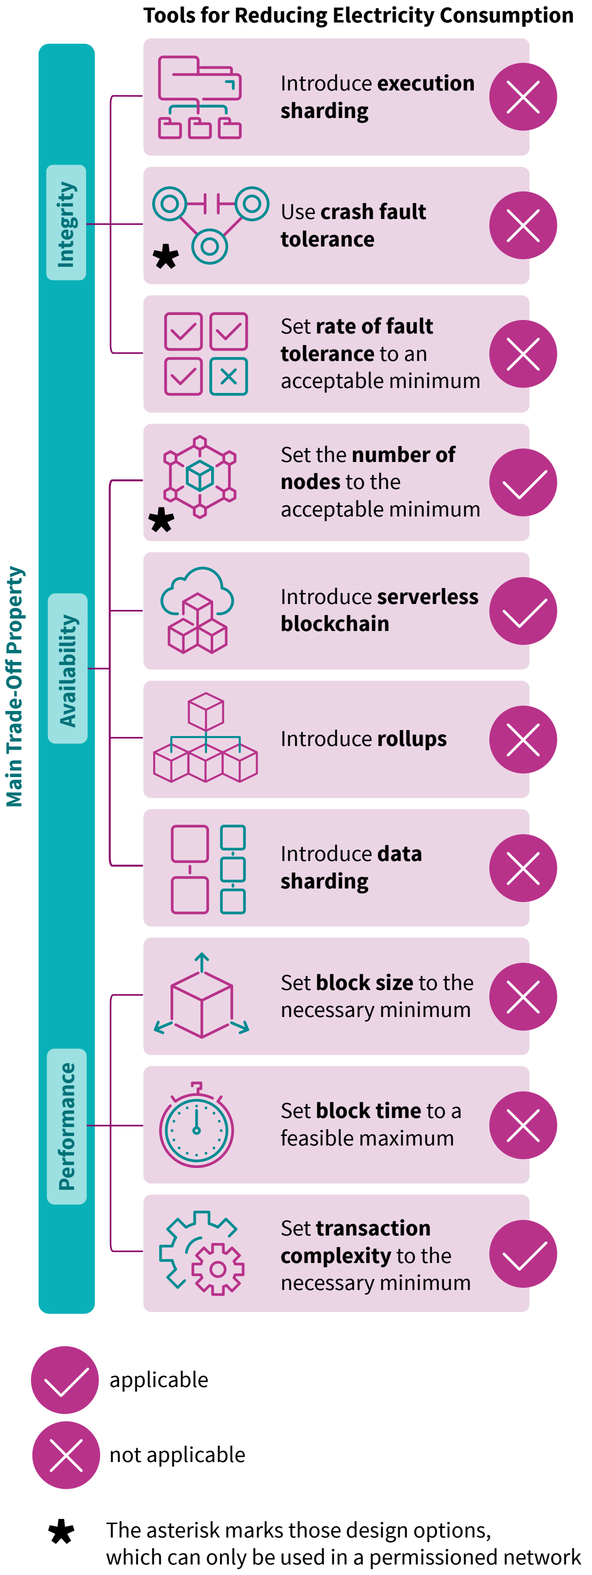
Tools for Reducing Electricity Consumption
Main Trade-Off Property
Integrity
Availability
Performance
Introduce execution sharding ✗
* Use crash fault tolerance ✗
Set rate of fault tolerance to an acceptable minimum ✗
* Set the number of nodes to the acceptable minimum ✓
Introduce serverless blockchain ✓
Introduce rollups ✗
Introduce data sharding ✗
Set block size to the necessary minimum ✗
Set block time to a feasible maximum ✗
Set transaction complexity to the necessary minimum ✓
✓ applicable
✗ not applicable
* The asterisk marks those design options, which can only be used in a permissioned network

Figure 18: Evaluation of design choices for an electricity-efficient design of an electronic prescription system.

Set the number of nodes to the necessary minimum: In this case study, we initially assumed that individual pharmacists would host the nodes, which is also supported by the literature. Taking this idea further and assuming that each pharmacy should host a node, this would result in approximately 18,000 full nodes in Germany. Since each node must be able to handle the high transaction volume, high performance nodes are required, resulting in a high node count network with high electricity consumption. However, we can significantly reduce the number of nodes while still providing the required data availability. In a highly regulated environment such as healthcare, trusted institutions could collectively operate a sufficiently decentralized network. These entities can ensure that their nodes are always highly available, resulting in high levels of availability even with a small number of nodes, as long as they are hosted by different data providers. In our use case, this could be done by the sixteen regional associations of pharmacists, which reduces the network size to this same number.

Introduce serverless blockchain: Transaction throughput varies dramatically over the course of a week, but each node must always be powerful enough to handle transaction peaks. However, this high-performance hardware sits idle during off-peak hours, consuming more electricity than the hardware needed to handle the reduced workload. A serverless blockchain effectively solves this problem. By dynamically scaling based on actual load, the system avoids the need to maintain high-performance servers at all times. This approach allows hardware to be scaled up during periods of high demand, while only allocating the necessary hardware during idle periods, significantly reducing average electricity consumption.

Introduce rollups: In this use case, the introduction of rollups would not lead to a reduction in the network's electricity consumption. This is because after reducing the number of nodes, the network does not have enough redundant computations to compensate for the high computational cost of a ZK rollup, even with the high number of transactions performed within the network (see our excursus in Chapter 3). Furthermore, implementing a rollup may conflict with the requirements of our use case. Prescription tokens would be handled by a central rollup, which contradicts the underlying concept of computing the prescription by a distributed network.

Introduce execution or data sharding: Although theoretically possible, implementing a sharding solution can be challenging. Depending on the implementation, the network would be divided into regional shards that would need to communicate with each other to ensure a patient's choice of pharmacy. Furthermore, the benefits of dividing the network also decrease as the number of nodes is significantly reduced, making sharding unsuitable for this use case.

Use crash fault tolerance: Given the high data integrity requirements of our use case, crash fault tolerance alone is not sufficient. Even with trusted entities responsible for proposing and validating transactions, we must ensure that the network can withstand faulty transactions. For this reason, we require Byzantine fault tolerance in the network.

Set rate of fault tolerance to an acceptable minimum: Due to the small number of nodes, even a considerable reduction in fault tolerance will not provide significant savings and could potentially make the network more susceptible to failures because fewer nodes would be involved. Therefore, we should aim for high fault tolerance rates.

Set block size and block time to a feasible level: The infrastructure must always ensure that it processes transactions within two seconds. Therefore, an adjustment of these parameters in order to reduce electricity consumption would run the risk of not providing sufficient buffers for unexpected peaks in demand to ensure that transactions are processed within this time frame.

Set transaction complexity to the necessary minimum: The transactional complexity handled on the blockchain is minimal, as the infrastructure is only used to create and handle prescription tokens, and the network operates solely for that purpose.

Evaluation

This e-prescription case highlights the importance of considering electricity efficiency in network design and demonstrates how a thoughtful approach can lead to functional and sustainable blockchain solutions. Understanding the unique specifications of this use case is essential, as this can both influence design options and directly mitigate the computational effort of each node.

The case study also shows that the toolbox can address existing misconceptions in the literature, such as the assumption that every doctor or pharmacist needs to run a node. Especially in regulated sectors such as healthcare, there is potential for collaboration between trusted institutions to create a decentralized, yet electricity-efficient data infrastructure that meets the requirements of the use case.

5.2 Case Study: Green Labeling

Green labeling in the electricity sector aims to certify the origin of the electricity consumed by users, ensuring that it is generated from renewable sources. This system combats the deceptive practice of “greenwashing”, where non-renewable electricity is mislabeled as renewable by purchasing certificates. To prevent greenwashing and promote transparency, the use case leverages blockchain technology to establish a tamper-proof, verifiable record of a customer’s green electricity supply and consumption

Set the number of nodes to the necessary minimum: Balancing Availability and Electricity Consumption

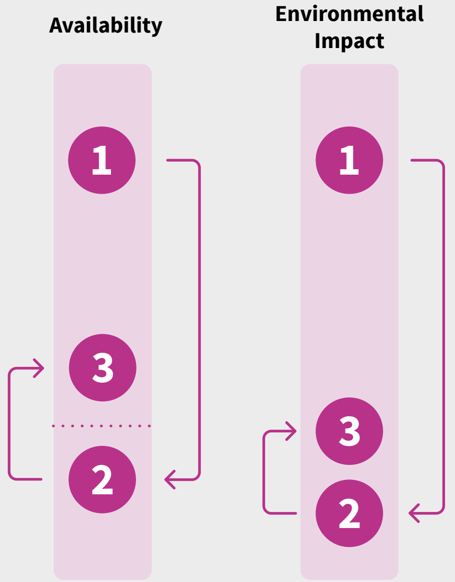

The info box illustrates the decision-making process for determining the required number of nodes, simplifying the trade-off between availability and reduced environmental impact. In the figure above, we have two bars representing the network’s availability (left) and environmental impact (right) properties. Each dot symbolizes a specific **network configuration**, and the higher up the dot, the higher the availability and electricity consumption. Finally, the line represents the required availability level for the use case. Design choices that do not meet the requirements are not viable.

If each pharmacy had its own node, as suggested in the literature, we would achieve an excess in the availability property and this would also result in high electricity consumption due to the extreme redundancy. However, reducing the number of nodes too drastically (2) will jeopardise the required availability. For example, if participation in the network is voluntary, this could result in too few nodes and hence undermine the required level of availability.

Our proposed design choice (3) involves the collaboration of 16 regional associations that can provide high availability nodes to provide the network infrastructure. This collaborative approach provides the required integrity while significantly reducing electricity consumption, thereby balancing availability and electricity efficiency.

(Roth et al. 2022). This will guarantee customers the origin of the electricity they consume, contributing to greater transparency and acceptance of carbon accounting. By decentralizing the data, this use case creates a cross-organizational platform that is accessible to all stakeholders. In this way, the system promotes data-driven value creation that is not limited to individual companies that dominate the market, thus supporting the Energy Transition by encouraging the involvement of all stakeholders. Figure 19 provides a simplified overview of the process.

Stage 1: Use Case Analysis

The basic approach of storing high-resolution consumption and production data on a blockchain is not feasible due to scalability and privacy concerns. In particular, the processing of smart meter data is highly regulated because it contains sensitive personal information, as this data, even when anonymized, can reveal details of an individual's social and economic status (Hinterstocker et al. 2017). **Therefore, this use case demands a high level of data confidentiality** (S-5). To address these issues, only essential data should be written on the blockchain. This can be achieved with a ZK rollup, which proves that a utility does not sell more green electricity than was available during a certain period. This proof is created by the customer's utility company and written to the blockchain. The utility provider processes sensitive data separately from the blockchain, particularly the customer's electricity consumption in the given period. The utility aggregates the electricity consumed by its customers into a single proof, **which is written to the blockchain every five minutes in a single transaction** (P-1). Customers can then access this proof to confirm that they have received the amount of green electricity the utility provider states.

In this use case, the infrastructure serves primarily as a storage repository for later verification by the customer. As such, the remaining infrastructure requirements are relatively modest. The use of zero-knowledge proofs ensures the correctness of the data, as the proof itself verifies the information, **reducing the integrity requirements on the infrastructure.** Similarly, since the original data is stored by the utility company, **availability requirements are also kept low.**

In addition, by aggregating all customer data into a single proof, each utility only performs one transaction on the blockchain every five minutes, regardless of the size of its customer base. This allows the required performance of the network to be kept at a very low level, even if the infrastructure is deployed nationwide (Sedlmeir et al. 2021c). For example, there are about 1500 utility providers in Germany[20]. If all of them participated, **this would result in about 300 transactions per minute, or only 5 transactions per second. Furthermore, the latency requirements are also low**, as the data stored on the blockchain is only used for subsequent verification and is not time-sensitive. **Finally, the complexity of the transaction is low as well, because only the proof, a simple data object, is stored.**

Stage 2: Network Design

Due to its low performance and confidentiality requirements, the use case could be deployed on a public permissionless network such as Ethereum. However, to illustrate the toolbox, we design a public permissionless blockchain using the Quorum framework.

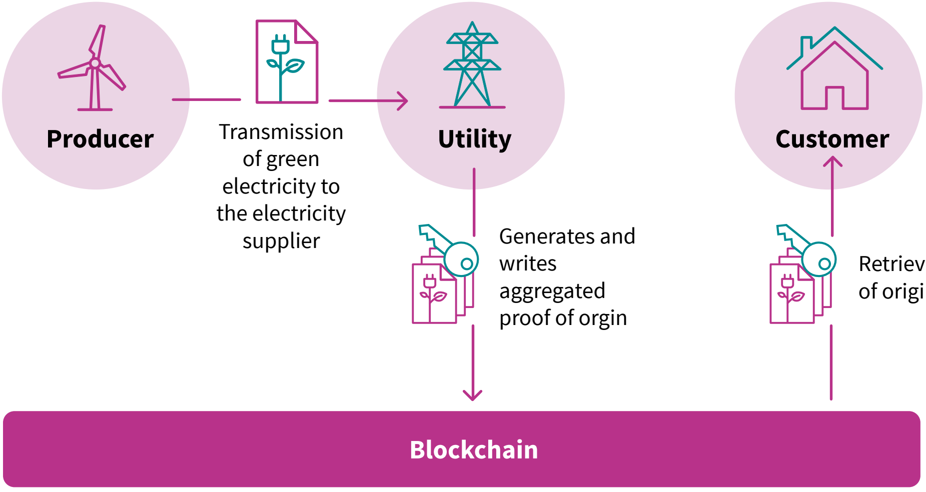
Producer
Transmission of green electricity to the electricity supplier
Utility
Customer
Generates and writes aggregated proof of orgin
Retrieves Proof of origin
Blockchain

Figure 19: A high-level overview of the use case's data infrastructure

[20] https://www.bundesnetzagentur.de/DE/Fachthemen/ElektrizitaetundGas/Monitoringberichte/start.html, accessed 08.08.2023

Evaluation of Design Choices:

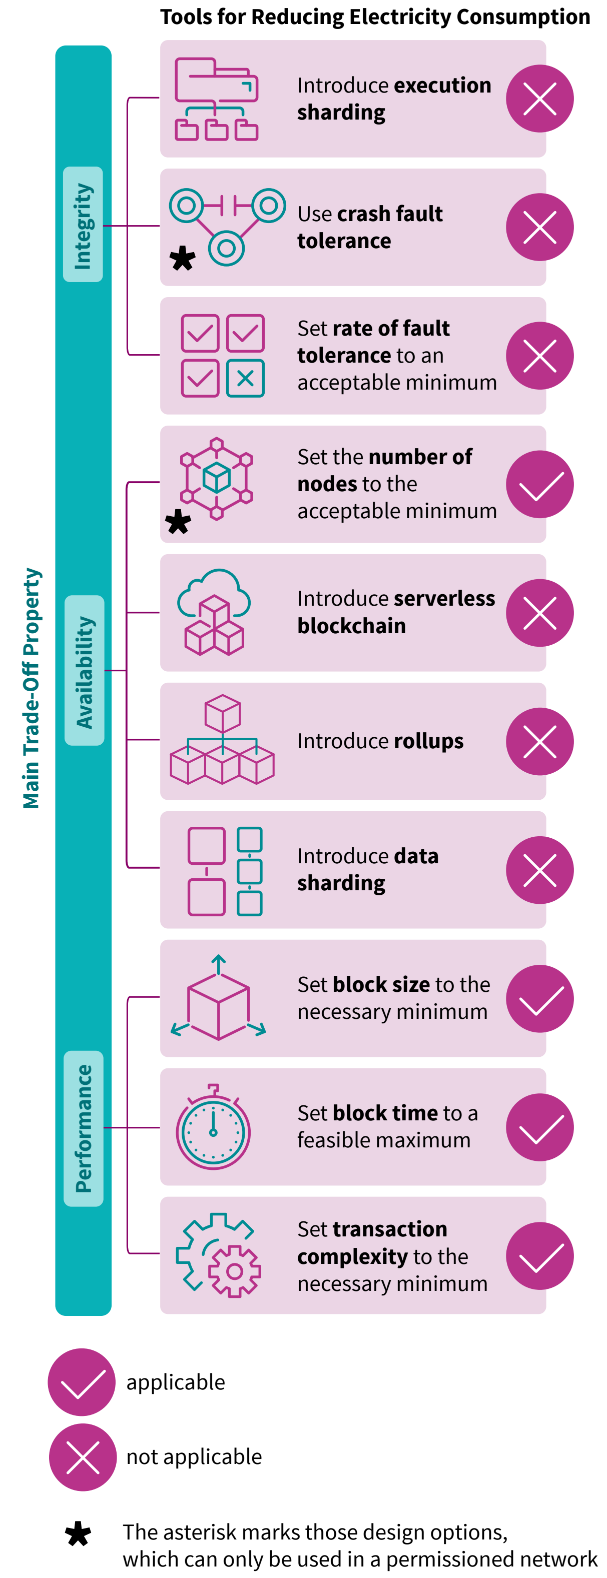
Tools for Reducing Electricity Consumption
Main Trade-Off Property
Integrity
Introduce execution sharding ✗
* Use crash fault tolerance ✗
Set rate of fault tolerance to an acceptable minimum ✗
Availability
* Set the number of nodes to the acceptable minimum ✓
Introduce serverless blockchain ✗
Introduce rollups ✗
Introduce data sharding ✗
Performance
Set block size to the necessary minimum ✓
Set block time to a feasible maximum ✓
Set transaction complexity to the necessary minimum ✓
✓ applicable
✗ not applicable
* The asterisk marks those design options, which can only be used in a permissioned network

Figure 20: Evaluation of design choices for an electricity-efficient design for a Green Label case.

Reduce the number of nodes: Although every utility must run a rollup operator, they do not all need to participate in the network, especially since the availability requirements are relatively modest. Therefore, a relatively small number of nodes operated by reputable entities, such as local authorities, non-governmental organizations, and transmission system operators, should provide a trusted and sufficiently decentralized environment among all stakeholders.

Implement a serverless blockchain: The number of transactions remains fixed, with utilities continuously sending a transaction every five minutes; the benefits of scaling computing power are limited.

Implement rollups: ZK rollups are already implemented in the system to ensure data confidentiality.

Implement sharding: Computational and data sharding may be possible by partitioning the system at the regional level. However, due to the limited throughput and number of nodes, the reduction in computational load on each node is negligible, resulting in only a small potential reduction in electricity consumption.

Switch to crash fault tolerance: When only reliable nodes are participating in the network, crash fault tolerance may be sufficient because the zero-knowledge proof ensures the integrity of the green labels. However, as transaction volume and decentralization decrease, the impact on electricity consumption is limited.

Minimize fault tolerance rate: Similar to crash fault tolerance, a reduction may be possible depending on the node operators involved, but the reduction in electricity savings is negligible.

Set block size to the feasible minimum and block time to a feasible maximum: Due to the very high aggregation of individual transactions, each utility only writes one transaction to the blockchain every five minutes, ensuring high predictability and matching the throughput of the network to the required level by adjusting block time and block size.

Set transaction complexity to the feasible minimum: Minimal transaction complexity is achieved as the blockchain only stores the ZKP.

Evaluation

Indeed, this case study demonstrates how analysis of the use case helps to derive the requirements, and based on that, develop a decentralized data infrastructure that meets those requirements. By carefully considering the needs of the green labeling use case, we were able to derive what data needs to be stored on the blockchain and what data needs to be processed by the utility. Since electricity consumption is highly sensitive data and requires a high level of confidentiality, it can be stored on a blockchain by obfuscating it through Zero Knowledge. In addition, this solution allows a large amount of data to be aggregated into a single transaction, reducing the overall volume of data stored in the blockchain, thereby easing the integrity and performance requirements on the infrastructure. This is where our toolbox can help find the right network configuration. Our toolbox enables iterative refinement of the network design, finding a balance between performance, security, and reduced environmental impact.

5.3 Case Study: Self-Sovereign Identity

As digital services continue to proliferate with digitalization progresses, a secure digital identity becomes critical. One promising solution is Self-Sovereign Identity (SSI), a concept in which users control their decentralized digital identities. Information is exchanged via verifiable, cryptographically signed credentials issued by trusted entities (Sedlmeir et al. 2021b). The Blockchain Machine Identity Ledger (BMIL) project[21], which focuses on the needs of the energy sector, exemplifies the utility of SSI in the context of the machine economy. It proposes a decentralized, blockchain-based system for managing machine identities, eliminating the need for centralized records and benefiting industries with numerous interconnected devices.

SSI allows identity holders to authenticate themselves and provide credentials in a bilateral communication channel without exposing excessive personal information. It uses Verifiable Credentials, where trusted entities, issuers, provide attestations of identity attributes. A verifier can securely validate these credentials while preserving the privacy of the holder, ensuring a secure identity authentication system. Only public information about the issuer is openly stored in the Verifiable Data Registry, which is used to verify the authenticity of the credentials and can be a blockchain-based system (see Figure 21).

Stage 1: Use Case Analyses

The data infrastructure in the SSI use case serves a clear function: to act as a verifiable data registry for storing public information related to the issued credentials, such as the structure of a verifiable credential. As such, the infrastructure is the foundation for the integrity and authenticity of verifiable credentials, even though it does not store personal data. By leveraging the properties of blockchain technology, such as immutability, transparency, and decentralization, the data infrastructure ensures the secure and tamper-proof storage of public information, such as the issuer's public key, which is critical to verifying the authenticity of credentials (S-2 + S-5). This information is accessed each time a credential is presented to verify its integrity, requiring robust IT security for the data infrastructure. A blockchain-based data infrastructure is well suited to meet these stringent IT security requirements due to its inherent properties of immutability, transparency, and decentralization. The data stored in the infrastructure must always be available to ensure credential holders can use the issued credential (S-1).

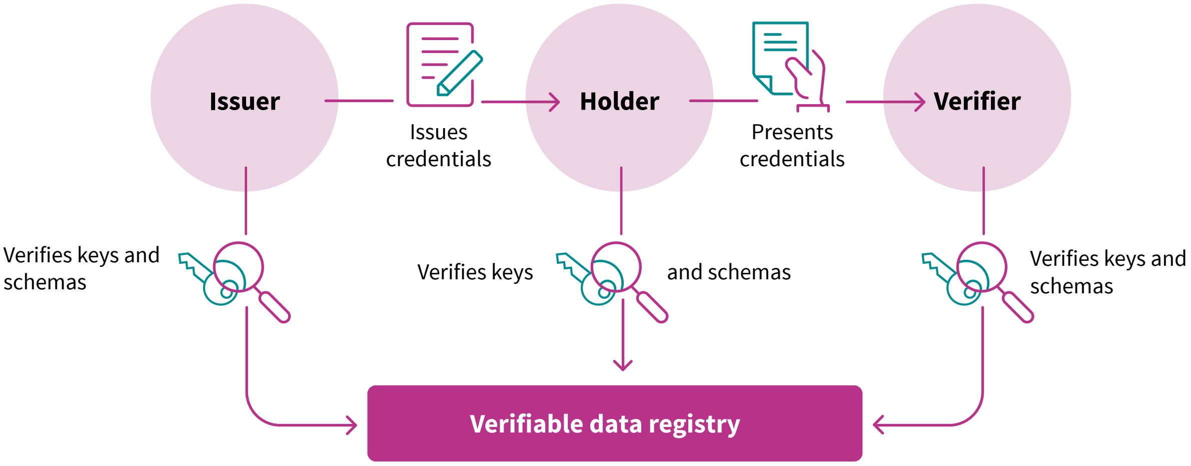

Figure 21. The fundamental architecture of an SSI-System

[21] https://future-energy-lab.de/projects/blockchain-machine-identity-ledger/, accessed 08.08.2023

By retaining personal data in the credential, **confidentiality** of the data exchanged in the credentials is maintained as only public data is stored on the blockchain. Furthermore, by keeping the data stored in the credential off-chain, the amount of data and thus the required transactional workload on the blockchain is reduced. The immutability of the ledger and the consensus mechanism ensures the **integrity** of the credential issuer's information, thereby preserving the validity of a credential. Thus, the SSI concept places high demands on the integrity of the infrastructure, as any successful attack on this information would call into question the credibility of all issued credentials. Consequently, the credibility of all issued credentials depends on the integrity of the data stored in the ledger, such as the signed key used. In addition, the infrastructure should provide transparent and immutable logs of changes to the data. This allows for effective auditing and can help detect and investigate any malicious activity.

The **performance** requirements for this use case are minimal, mainly because only infrequent updates, such as those related to an issuer's public information, are written to the blockchain. On the other hand, creating new issues does not cause any writing to the blockchain when the credentials are exchanged as verifiable credentials. This is demonstrated by Sovrin, one of the most actively used identity ecosystems, which records less than 100 write transactions per day (P1)[22] . However, it is essential to note that the number of read accesses significantly outweighs the number of write transactions, as each credential presentation results in at least one read access. Therefore, the system design must account for a high frequency of read activity.

Stage 2: Network Design

A public, permissionless blockchain based on Hyperledger Indy is the most appropriate option for this use case. Hyperledger Indy, with its strong emphasis on SSI, provides the necessary tools and functionalities to build the underlying infrastructure and meet the specific requirements of the use case. It is widely used in practice within various decentralized digital identity ecosystems, such as Sovrin or IDUnion.

[22] https://sovrin.org/ssi-metrics-dashboards/, accessed 08.08.2023

Evaluation of Design Choices:

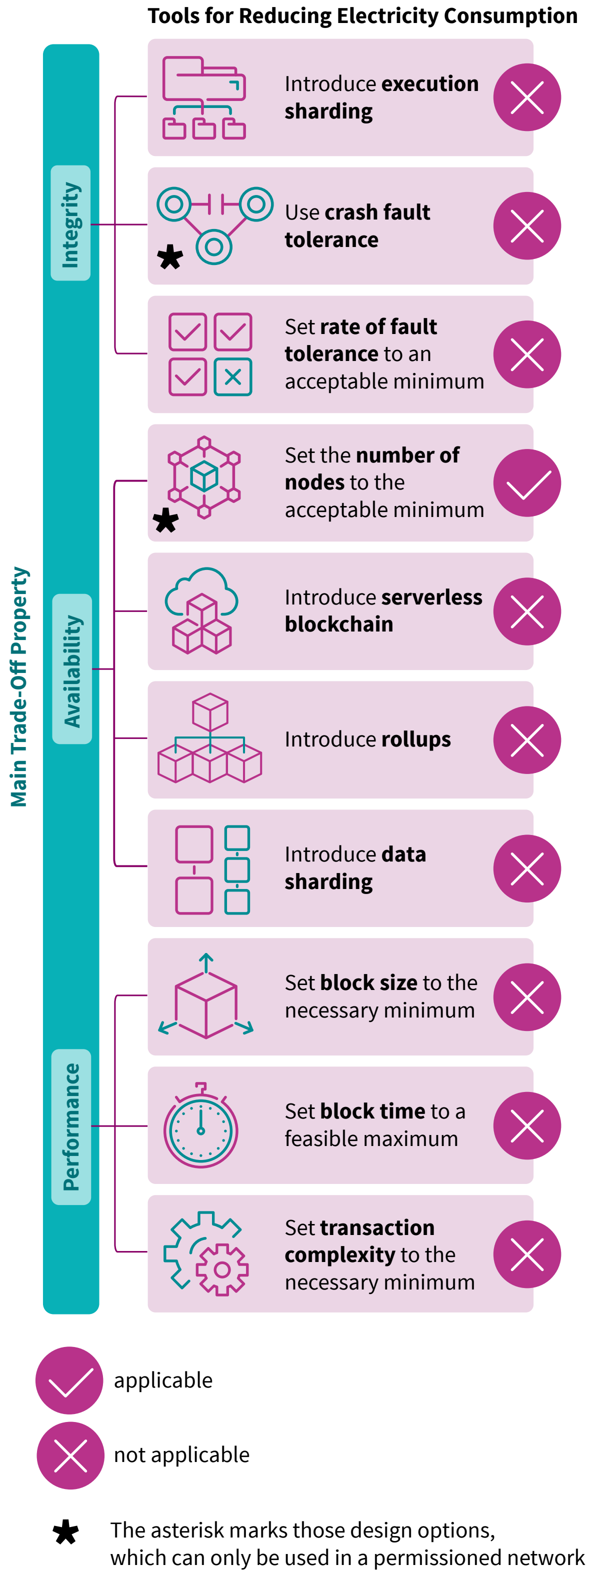
Tools for Reducing Electricity Consumption
Main Trade-Off Property
Integrity
Availability
Performance
Introduce execution sharding ☒
★ Use crash fault tolerance ☒
Set rate of fault tolerance to an acceptable minimum ☒
★ Set the number of nodes to the acceptable minimum ✓
Introduce serverless blockchain ☒
Introduce rollups ☒
Introduce data sharding ☒
Set block size to the necessary minimum ☒
Set block time to a feasible maximum ☒
Set transaction complexity to the necessary minimum ☒
✓ applicable
☒ not applicable
★ The asterisk marks those design options, which can only be used in a permissioned network

Figure 22: Evaluation of design choices for an electricity-efficient network design for the SSI case.

Introduce data or execution sharding: Sharding, which divides the network into small shards, is not recommended for the SSI use case. The small amount of data on the ledger, combined with the need for a high level of availability and integrity, make sharding less suitable as it can negatively impact the overall availability of the system.

Minimize the number of nodes: In an SSI network, hosts are typically highly trusted organizations within the ecosystem, such as banks, educational institutions, NGOs, or corporations. These entities have the necessary technical capabilities and a vested interest in maintaining a self-sovereign identity system. Each of them can provide a high guarantee of the availability of their node, eliminating the need for a decentralized network with a large number of nodes. This is comparable to current implementations of the SSI network, such as Sovrin, which limits itself to 25 nodes.

Introduce serverless blockchain: Due to the limited number of transactions written to the blockchain, a serverless blockchain does not offer direct benefits in reducing electricity. However, adopting a serverless infrastructure can increase the availability of the network as long as the nodes are hosted on different cloud providers.

Implement rollups: Hyperledger Indy does not support rollups.

Use crash fault tolerance and reduce rate of fault tolerance to an acceptable minimum: Given the high integrity requirements and low transaction volume in the SSI use case, the modest savings in computational workload obtained by switching to crash fault tolerance and setting a modest rate of fault tolerance does not justify compromising data integrity.

Set block size to the feasible minimum and block time to the feasible maximum: The selected blockchain framework, Hyperledger Indy, does not allow the block size and block time to be changed, so both values cannot be adapted to the requirements of the use case.

Set transaction complexity to the feasible minimum: There is no option for reducing transaction complexity.

Evaluation

The results of this use case highlight the importance of accurately defining network design requirements in the context of self-sovereign identity (SSI). Hyperledger Indy is tailored to SSI use cases, enabling a data infrastructure that meets high integrity and availability requirements. However, its specialization limits its potential tools for reducing electricity consumption to those that affect the number of nodes and thus network decentralization. However, there are alternative ways to decrease electricity consumption, such as using highly efficient servers, since the nodes have low performance requirements. By carefully considering the specific requirements of the use case and using appropriate tools and configurations, we can achieve an electricity-efficient network design for SSI applications. These findings highlight the importance of aligning the network infrastructure with the unique requirements of the use case to optimize energy efficiency while ensuring the desired functionality and security.

6. Conclusion

As digitalization advances, it places increasing demands on the underlying data infrastructure, including the need for decentralized systems like blockchain technology. Such requirements, however, do not inevitably lead to increasing electricity consumption. Contrarily, our study challenges the common misconception that high electricity consumption is an inherent feature of the blockchain technology. Instead, we have argued that conscious network design, as demonstrated by Ethereum's transition from PoW to PoS, can significantly reduce the network's electricity consumption. Our study explores technological advances that offer promising solutions for significantly reducing the electricity consumption of blockchain networks.

Innovation in Electricity Efficiency for Blockchain Networks

Technological advances continue to provide new design options to reduce electricity consumption, and these are not limited to changes in the consensus mechanism. We have identified several approaches to reducing the electricity consumption of a blockchain network, especially when using a non-PoW consensus mechanism. Our findings have shown that critiques of high electricity consumption in PoW networks cannot be directly applied to non-PoW networks, such as Ethereum or permissioned blockchains. Such knowledge is critical to the proper application of blockchain technology, allowing for informed decisions in designing and selecting an appropriate blockchain network.

While it is undisputed that a decentralized infrastructure may consume more electricity than a centralized server architecture; if a use case requires the unique features of a blockchain network, such as data redundancy or the decentralized consensus mechanism, this additional electricity consumption may be considered worthwhile and with careful design, consumption can remain modestly low. Our study provides guidance on how to provide an electricity-efficient and appropriate infrastructure for a use case, while taking advantage of these unique features of the technology.

Key Contributions and Insights from the Study

Our study makes a meaningful contribution to the current discourse on blockchain technology and its environmental impact, primarily through two novel outputs: a schematic illustration of the main parameters of electricity consumption in blockchain networks, and a guide to designing an electricity-efficient blockchain. By presenting the complex interrelationships of the various parameters in a concise manner, our study provides a comprehensible overview of the various factors influencing the electricity consumption of blockchain networks. In doing so, we are fulfilling one of the recommendations of the expert panel "Fachdialog Blockchain" (Culotta et al. 2022).

Our guide and toolbox also provide valuable assistance in selecting appropriate design options for developing a blockchain-based data infrastructure and our toolbox helps ensure that the infrastructure meets the specific requirements of a use case and promotes its electricity-efficiency. To our knowledge this is the first study to provide guidance on reducing the electricity consumption of a blockchain network, thereby addressing another recommendation identified by the expert panel.

Limitations and Future Scope of Research

A more detailed examination of the aspects that influence the electricity efficiency of the network requires an accurate quantification of its electricity consumption. By measuring the real-world electricity consumption of differently designed networks, future studies could inform a more fine-grained version of the models presented in this study, thus filling this existing gap in the understanding of electricity consumption in blockchain. In addition, our toolbox includes emerging technologies and solutions whose impact on electricity consumption has not yet been thoroughly investigated. This highlights the critical need for future research efforts to quantify the potential electricity consumption reductions achievable through these new techniques. Finally, the use of our toolbox requires a comprehensive understanding of blockchain technology, given the complex interactions between different design choices. Another limitation of our study stems from our intentional focus on the technical aspects of the data infrastructure, particularly electricity consumption. This focus has led us to exclude economic considerations, including potential cost implications of certain design choices, from our analysis. While our study provides a detailed examination of the technical side, a comprehensive analysis requires considering both technical and economic factors. Future studies should address these dimensions to provide a comprehensive view of the issue, covering the economic and environmental suitability of the technology.

Implications for Blockchain Stakeholders

Based on the results of the study, we present several suggestions for different stakeholders to promote the electricity efficiency and sustainability of blockchain technology:

- As our study does not cover all aspects of blockchain technology electricity consumption, we encourage **researchers** to explore these areas further. Specifically, they could assess the potential electricity savings of the design tools we identified or develop new methodologies to improve the electricity efficiency of a blockchain. We also encourage the development of new frameworks to compare different forms of data infrastructure, allowing for a more comprehensive view of their relative efficiencies. Finally, cross-disciplinary research could help bring different perspectives on the electricity efficiency, potential uses, and benefits of blockchain technologies and determine the circumstances under which additional usage may be justified.

- **Standards organizations and policy makers** could use the results of this research to advance standardization, benchmarking and regulation for blockchain technology. This could include metrics for the electricity consumption or carbon emissions associated with different blockchains, allowing companies or organizations using the technology to calculate their carbon footprint. This work can also be used to evaluate blockchain applications, especially in comparison to alternative data infrastructures.

- Blockchain **framework developers** should also consider the electricity consumption aspect of their software. In doing so, they can incorporate features directly aimed at reducing the amount of electricity consumed. Furthermore, they could contribute to the overall sustainability of blockchain technology by providing practical guidelines for electricity-efficient designs and creating tools for users to monitor the network's power consumption.

- **Both users and operators** of a blockchain-based network should consider various aspects of environmental impact, such as electricity consumption or carbon emissions when choosing a network. Our study allows for such conscious network design. Our study shows that conscious network design can reduce these impacts while ensuring suitability for specific use cases. In this way, users and operators can take advantage of the decentralized infrastructure while enhancing the environmental sustainability of their operations. We also suggest that users demand transparency from network operators about their electricity consumption. This would not only enable an informed choice of networks but also incentivize developers to consider electricity consumption as a priority.

The actions recommended above should be taken collaboratively by the different stakeholders, rather than individually. Further research will certainly fill any remaining knowledge gaps. However, researchers will need to consider the demands of standards organizations and policy makers. Moreover, blockchain framework developers as well as the operators and users of the resulting networks have a unique ability to deliver invaluable insights into the applicability, limitations and remaining shortcomings of tools and regulations for the energy-efficiency of blockchains. We, the German Energy Agency, hereby encourage all stakeholders who have the power to influence the electricity consumption of blockchains in any way to participate in an 'alliance of the willing' and to join in a coordinated effort to maximize the sustainability of blockchain technology. Such an alliance requires an appropriate ecosystem connecting the different stakeholders, which we would gladly support by acting as an intermediary and organizing the required formats and forums.

Glossary

Concept	Definition
Application-specific integrated circuit (ASIC)	A purpose-built chip (or device) that is highly optimized for a specific use case, rather than a general-purpose application. Often used for PoW mining, such as in the Bitcoin network.
Availability	The assurance that data and services are accessible when needed.
Block Size	The maximum amount of data that can be contained in a single block on a blockchain. It directly determines the throughput of a network, along with block time.
Block Time	The defined amount of time before a new block, and thus data, is written to the blockchain. The block time, together with the block size, determines the throughput of the blockchain.
Blockchain	A distributed and decentralized digital ledger that records transactions across multiple computers or nodes. Each transaction is stored in a block, which is linked to previous blocks, creating a chain of blocks.
Blockchain Frameworks	Software stacks that allow you to create your own permissioned networks. They allow customization to meet specific needs. Examples include Corda, Quorum, and the Hyperledger project, which consolidates several projects such as Hyperledger Indy, Fabric, and Sawtooth.
Blockchain Platforms	Existing permissionless networks that can be utilized as the underlying data infrastructure for a new use case, such as Ethereum and Polkadot.
Blockchain Type	A classification of blockchain types by decentralization, consensus mechanism (permissioned or permissionless), and data access (private or public).
Byzantine Fault Tolerance (BFT)	A property of a blockchain network that allows it to operate correctly and reach consensus even if some participating nodes are dishonest or exhibit malicious behavior, preventing them from compromising the network’s integrity and functionality.
Confidentiality	The assurance that data access and disclosure is limited to authorized users and processes.
Consensus Algorithm	The specific process used by a network’s consensus mechanism to achieve agreement and determine the next valid block in the chain.
Consensus Mechanism	The algorithm or protocol used by a blockchain network to achieve agreement among participants on the state of the blockchain and validate transactions. It ensures the network’s integrity and security.
Crash Fault Tolerance (CFT)	A property of a blockchain network that enables it to function correctly and reach consensus even if some nodes stop operating due to failures like network splitting or node crashes.
Data Infrastructure	It consists of hardware, software, and network layers specifically designed to manage, store, and process data. Depending on the requirements of the use case, the infrastructure must be designed to provide the necessary characteristics to ensure the functionality of the use case and can be designed in a centralized or a decentralized approach, such as blockchain.

Concept	Definition
Environmental Impact	The impact of digital infrastructure on the environment, most notably in terms of electricity consumption, CO_2 emissions and e-waste.
Hash Puzzle	A computational problem that must be solved in the context of PoW mining in order for the miner to be allowed to propose the next block.
Integrity	The protection of data from unauthorized alteration, deletion, or addition to ensure its accuracy and consistency.
Issuance Rate	The rate at which new digital coins are minted in a network. It is often defined by the network design and plays a critical role in determining the reward for participating in the consensus mechanism.
Layer 1 Blockchain	A standalone blockchain network, complete with core logic and functionality. This includes the consensus protocol and the immutable ledger of transactions that serves as the fundamental building block for all operations on the network. Examples include Bitcoin and Ethereum.
Layer 2 Blockchain	A network built on top of a Layer 1 blockchain to increase scalability and throughput while reducing costs. Layer 2 solutions offload transactions from the main chain using mechanisms such as rollups. A prominent example is Polygon, which uses Ethereum as its main network.
Node	A participant in a blockchain network that maintains a copy of the entire blockchain and participates in the validation and propagation of transactions. A light node downloads only part of the blockchain, while a full node downloads the entire blockchain.
Performance	A property of the data infrastructure that ensures efficient and timely processing and delivery of data to enable seamless operations. This may include aspects like the network's latency, i.e. the time it takes for a transaction to be processed by the network.
Proof of Authority (PoA)	A consensus mechanism used in permissioned blockchain networks where a pre-selected group of nodes with known identities and authority validate transactions and create new blocks based on their reputation or permissions.
Proof of Elapsed Time (PoET)	A consensus mechanism used in some permissioned blockchain networks where participants compete to win the right to create new blocks by waiting for a randomly assigned waiting period, simulating a fair lottery system.
Proof of Stake (PoS)	A consensus mechanism where participants (stakers) validate transactions and create new blocks based on the number of tokens they hold or "stake" in the network. It aims to achieve consensus in a permissionless setting more energy-efficiently than PoW. Prominent examples include Ethereum and Cardano.
Proof of Work (PoW)	A consensus mechanism used in many public blockchains, such as Bitcoin or Dogecoin, where participants (miners) solve computationally intensive puzzles to be authorized to validate transactions and create new blocks, requiring significant computing power and therefore high electricity consumption.
(Mining) Reward	Participation in the consensus mechanism can be rewarded by the network. Usually the reward consists of transaction fees for the executed transactions and newly generated coins, the number of which can be set in the design of the network.

Concept	Definition
Rollups	Rollups aggregate transactions through a single or few rollup operators, which store proof of their correctness on the main blockchain. Verifying these aggregated proofs is less computationally intensive than verifying individual transactions.
Serverless Blockchain	In serverless blockchains, nodes are hosted by cloud service providers, which allows computing resources to be elastically adjusted based on current transaction throughput rather than continuously being tuned for peak capacity. In addition, the high reliability and availability of cloud services can potentially reduce the number of nodes required.
Sharding	A technique that divides a blockchain network into smaller partitions, or shards, in order to increase its efficiency. Sharding can be categorized into data sharding, where data is partitioned, and execution sharding, where transaction processing is divided. This can significantly increase the network's throughput and efficiency.
Smart Contract	Self-executing contracts with predefined rules and conditions that are written as code and deployed on a blockchain. They automatically enforce and facilitate the performance of contractual agreements without the need for intermediaries.
Sybil Attack	An attack on a permissionless network in which a single entity creates numerous false pseudonymous identities to maliciously affect the functionality and integrity of the network. A Sybil-resistant consensus algorithm can mitigate this threat.
Token	A unit of value or digital representation that is used within a specific blockchain system. Tokens can have various functions, including serving as a medium of exchange, representing ownership, or providing access to a particular application or service.
Transaction	A data unit that represents an action or exchange of value on a blockchain. It can involve the transfer of cryptocurrency, execution of a smart contract, or recording of any other relevant information.
Transaction Complexity	The level of computational resources required to process and verify a transaction on a blockchain. It depends on the type and composition of the transaction, including any associated smart contracts or data operations.
Trusted Third Party (TTP)	An intermediary entity relied upon by multiple parties to facilitate and ensure trust in transactions. This reliance on a central entity may introduce vulnerabilities or central points of failure within a system.
(Digital) Use Case	A digital application or use of a digital system to achieve a specific goal or complete a specific task.
Zero-Knowledge Proof (ZKP)	A cryptographic method allowing one party to prove the validity of a statement to another party without revealing additional information.

List of Figures

List of Tables

References

Abraham R, Schneider J, vom Brocke J (2019) Data governance: A conceptual framework, structured review, and research agenda. International Journal of Information Management 49:424–438.

Alofi A, Bahsoon R, Hendley R (2021a) MinerRepu: A Reputation Model for Miners in Blockchain Networks. In: 2021 IEEE International Conference on Web Services (ICWS). IEEE, pp 724–733.

Alofi A, Bokhari M, Bahsoon R, Hendley B (2022) Optimizing the Energy Consumption of Blockchain-based Systems Using Evolutionary Algorithms: A New Problem Formulation. IEEE Transactions on Sustainable Computing:1. doi:10.1109/TSUSC.2022.3160491.

Alofi A, Bokhari MA, Hendley R, Bahsoon R (2021b) Selecting miners within blockchain-based systems using evolutionary algorithms for energy optimisation. In: Chicano F, Krawiec K (eds) Proceedings of the Genetic and Evolutionary Computation Conference Companion. ACM, New York, NY, USA, pp 291–292.

Bada AO, Damianou A, Angelopoulos CM, Katos V (2021) Towards a Green Blockchain: A Review of Consensus Mechanisms and their Energy Consumption. In: 2021 17th International Conference on Distributed Computing in Sensor Systems (DCOSS). IEEE, pp 503–511.

Badea L, Mungiu-Pupazan MC (2021) The Economic and Environmental Impact of Bitcoin. IEEE Access 9:48091–48104. doi:10.1109/ACCESS.2021.3068636.

Bahri L, Girdzijauskas S (2018) When Trust Saves Energy: A Reference Framework for Proof of Trust (PoT) Blockchains. In: Champin P-A, Gandon F, Lalmas M, Ipeirotis PG (eds) Companion of the The Web Conference 2018 on The Web Conference 2018 - WWW '18. ACM Press, New York, New York, USA, pp 1165–1169.

Baldominos A, Saez Y (2019) Coin.AI: A Proof-of-Useful-Work Scheme for Blockchain-Based Distributed Deep Learning. Entropy (Basel, Switzerland) 21(8). doi:10.3390/e21080723.

Ball M, Rosen A, Sabin M, Vasudevan PN (2018) Proofs of Work from Worst-Case Assumptions, Cryptology ePrint Archive, Paper 2018/559.

Beck R, Müller-Bloch C, King JL (2018) Governance in the Blockchain Economy: A Framework and Research Agenda. Journal of the Association for Information Systems:1020–1034. doi:10.17705/1jais.00518.

Belotti M, Bozic N, Pujolle G, Secci S (2019) A Vademecum on Blockchain Technologies: When, Which, and How. IEEE Communications Surveys & Tutorials 21(4):3796–3838. doi:10.1109/COMST.2019.2928178.

Bizzaro F, Conti M, Pini MS (2020) Proof of Evolution: leveraging blockchain mining for a cooperative execution of Genetic Algorithms. In: 2020 IEEE International Conference on Blockchain (Blockchain). IEEE, pp 450–455.

Capocasale V, Gotta D, Perboli G (2023) Comparative analysis of permissioned blockchain frameworks for industrial applications. Blockchain: Research and Applications 4(1):100113. doi:10.1016/j.bcra.2022.100113.

Castellon CE, Roy S, Kreidl OP, Dutta A, Boloni L (2022) Towards a Green Blockchain: Engineering Merkle Tree and Proof of Work for Energy Optimization. IEEE Transactions on Network and Service Management:1. doi:10.1109/TNSM.2022.3219494.

CCBECI (2023). https://ccaf.io/cbeci/index. Accessed 2023-08-08.

Chatterjee K, Goharshady A, Pourdamghani A (2019) Hybrid Mining: Exploiting Blockchain's Computational Power for Distributed Problem Solving. Proceedings of the 34th ACM/SIGAPP symposium on applied computing:374–381.

Chaurasia Y, Subramanian V, Gujar S (2021) PUPoW: A Framework for Designing Blockchains with Practically-Useful-Proof-of-Work & VanityCoin. In: 2021 IEEE International Conference on Blockchain (Blockchain). IEEE, pp 122–129.

Chen L, Xu L, Shah N, Gao Z, Lu Y, Shi W (2017) On Security Analysis of Proof-of-Elapsed-Time (PoET). In: Spirakis P, Tsigas P (eds) Stabilization, Safety, and Security of Distributed Systems. Springer International Publishing, Cham, pp 282–297.

Chenli C, Li B, Shi Y, Jung T (2019) Energy-recycling Blockchain with Proof-of-Deep-Learning. In: 2019 IEEE International Conference on Blockchain and Cryptocurrency (ICBC). IEEE, pp 19–23.

Coinshare (2022) The Bitcoin Mining Network: Energy and Carbon Impact. https://a.storyblok.com/f/155294/x/0c3f3837c8/coinshares_bitcoin_mining_report_jan_2022.pdf. Accessed 2023-08-08.

Crypto Carbon Ratings Institute (2022a) Energy Efficiency and Carbon Footprint of the Polygon Blockchain.

Crypto Carbon Ratings Institute (2022b) Energy Efficiency and Carbon Footprint of the TRON Blockchain.

Crypto Carbon Ratings Institute (2022c) The Merge: Implications on the Electricity Consumption and Carbon Footprint of the Ethereum Network.

Culotta C, Brüning S, Schulte AT, Gesmann-Nuissl D, Märkel C, Beck R (2022) Nachhaltigkeit im Kontext der Blockchain-Technologie: Anwendungsbeispiele, Herausforderungen und Handlungsfelder.

Daian P, Goldfeder S, Kell T, Li Y, Zhao X, Bentov I, Breidenbach L, Juels A (2020) Flash Boys 2.0: Frontrunning in Decentralized Exchanges, Miner Extractable Value, and Consensus Instability. In: IEEE Symposium on Security and Privacy, pp 910–927.

de Vries A (2018) Bitcoin's Growing Energy Problem. Joule 2(5):801–805. doi:10.1016/j.joule.2018.04.016.

de Vries A (2019) Renewable Energy Will Not Solve Bitcoin's Sustainability Problem. Joule 3(4):893–898. doi:10.1016/j.joule.2019.02.007.

de Vries A (2020) Bitcoin's energy consumption is underestimated: A market dynamics approach. Energy Research & Social Science 70:101721. doi:10.1016/j.erss.2020.101721.

de Vries A (2021) Bitcoin boom: What rising prices mean for the network's energy consumption. Joule 5(3):509–513. doi:10.1016/j.joule.2021.02.006.

de Vries A (2022) Cryptocurrencies on the road to sustainability: Ethereum paving the way for Bitcoin. Patterns:100633. doi:10.1016/j.patter.2022.100633.

de Vries A, Gallersdörfer U, Klaaßen L, Stoll C (2022) Revisiting Bitcoin's carbon footprint. Joule 6(3):498–502. doi:10.1016/j.joule.2022.02.005.

de Vries A, Stoll C (2021) Bitcoin's growing e-waste problem. Resources, Conservation and Recycling 175:105901. doi:10.1016/j.resconrec.2021.105901.

Dena (2019) Blockchain in der integrierten Energiewende. Deutsche Energie-Agentur GmbH, Berlin.

Dong Z, Lee YC, Zomaya AY (2019) Proofware: Proof of Useful Work Blockchain Consensus Protocol for Decentralized Applications.

Douceur JR (2002) The Sybil Attack. In: Goos G, Hartmanis J, van Leeuwen J, Druschel P, Kaashoek F, Rowstron A (eds) Peer-to-Peer Systems. Springer Berlin Heidelberg, Berlin, Heidelberg, pp 251–260.

Erdogan S, Ahmed MY, Sarkodie SA (2022) Analyzing asymmetric effects of cryptocurrency demand on environmental sustainability. Environmental science and pollution research international 29(21):31723–31733. doi:10.1007/s11356-021-17998-y.

Ethereum Foundation (2021) Shard Chains. https://ethereum.org/en/eth2/shard-chains/. Accessed 2023-08-08.

EU Blockchain Observatory (2021) Energy efficiency of blockchain technologies

Gallersdörfer U, Klaaßen L, Stoll C (2020) Energy Consumption of Cryptocurrencies Beyond Bitcoin. Joule 4(9):1843–1846. doi:10.1016/j.joule.2020.07.013.

Gallersdörfer U, Klaaßen L, Stoll C (2022) Energy Efficiency and Carbon Footprint of Proof of Stake Blockchain Protocols.

Gojka E-E, Kannengießer N, Sturm B, Bartsch J, Sunyaev A (2021) Security in Distributed Ledger Technology: An Analysis of Vulnerabilities and Attack Vectors. In: Arai K (ed) Intelligent Computing. Springer International Publishing, Cham, pp 722–742.

Gola C, Sedlmeir J (2022) Addressing the Sustainability of Distributed Ledger Technology. SSRN Electronic Journal. doi:10.2139/ssrn.4032837.

Gonzalez-Barahona JM (2021) Factors determining maximum energy consumption of Bitcoin miners.

Gourisetti SNG, Mylrea M, Patangia H (2020) Evaluation and Demonstration of Blockchain Applicability Framework. IEEE Transactions on Engineering Management 67(4):1142–1156. doi:10.1109/TEM.2019.2928280.

Gräbe F, Kannengießer N, Lins S, Sunyaev A (2020) Do Not Be Fooled: Toward a Holistic Comparison of Distributed Ledger Technology Designs. In: Bui T (ed) Proceedings of the 53rd Hawaii International Conference on System Sciences. Hawaii International Conference on System Sciences.

Guggenberger T, Sedlmeir J, Fridgen G, Luckow A (2022) An in-depth investigation of the performance characteristics of Hyperledger Fabric. Computers & Industrial Engineering 173:108716. doi:10.1016/j.cie.2022.108716.

Gundaboina L, Badotra S, Tanwar S, Manik (2022) Reducing Resource and Energy Consumption in Cryptocurrency Mining by using both Proof-of-Stake Algorithm and Renewable Energy. In: 2022 International Mobile and Embedded Technology Conference (MECON). IEEE, pp 605–610.

Heinonen HT, Semenov A, Veijalainen J, Hamalainen T (2022) A Survey on Technologies Which Make Bitcoin Greener or More Justified. IEEE Access 10:74792–74814. doi:10.1109/ACCESS.2022.3190891.

Hinterstocker, M, Schott, P, von Roon, S (2017). Disaggregation of household load profiles. In 10. Internationale Energiewirtschaftstagung an der TU Wien.

Hunhevicz JJ, Hall DM (2020) Do you need a blockchain in construction? Use case categories and decision framework for DLT design options. Advanced Engineering Informatics 45:101094. doi:10.1016/j.aei.2020.101094.

Jacquet P, Mans B (2019) Green mining: Toward a less energetic impact of cryptocurrencies. IEEE INFOCOM 2019-IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS):210--215. https://ieeexplore.ieee.org/servlet/opac?punumber=8831168.

Jagals M, Karger E, Ahlemann F, Brée T (2021) Enhancing Inter-Organizational Data Governance via Blockchain⌧Shaping Scopes and Research Avenues. In: Proceedings of the International Conference on Information Systems (ICIS).

Jennath HS, Asharaf S (2020) Survey on Blockchain Consensus Strategies. In: Kumar A, Paprzycki M, Gunjan VK (eds) ICDSMLA 2019. Springer Singapore, Singapore, pp 637–654.

Jiang S, Li Y, Lu Q, Hong Y, Guan D, Xiong Y, Wang S (2021) Policy assessments for the carbon emission flows and sustainability of Bitcoin blockchain operation in China. Nature communications 12(1):1938. doi:10.1038/s41467-021-22256-3.

Kannengießer N, Lins S, Dehling T, Sunyaev A (2021) Trade-offs between Distributed Ledger Technology Characteristics. ACM Computing Surveys 53(2):1–37. doi:10.1145/3379463.

Keller R, König C (2014) A Reference Model to Support Risk Identification in Cloud Networks. In: ICIS.

Khatri V, Brown CV (2010) Designing data governance. Communications of the ACM 53(1):148–152. doi:10.1145/1629175.1629210.

King S, Nadal S (2012) Ppcoin: Peer-to-peer crypto-currency with proof-of-stake. self-published paper, August 19(1).

Kitchenham B, Pearl Brereton O, Budgen D, Turner M, Bailey J, Linkman S (2009) Systematic literature reviews in software engineering – A systematic literature review. Information and Software Technology 51(1):7–15. doi:10.1016/j.infsof.2008.09.009.

Kohli V, Chakravarty S, Chamola V, Sangwan KS, Zeadally S (2022) An Analysis of Energy Consumption and Carbon Footprints of Cryptocurrencies and Possible Solutions.

Koomey J (2019) Estimating Bitcoin Electricity Use: A Beginner's Guide. https://www.coincenter.org/app/uploads/2020/05/estimating-bitcoin-electricity-use.pdf. Accessed 2023-08-08.

Kostal K, Krupa T, Gembec M, Veres I, Ries M, Kotuliak I (2018) On Transition between PoW and PoS. In: 2018 International Symposium ELMAR. IEEE, pp 207–210.

Krause MJ, Tolaymat T (2018) Quantification of energy and carbon costs for mining cryptocurrencies. Nature Sustainability 1(11):711–718. doi:10.1038/s41893-018-0152-7.

Król M, Sonnino A, Al-Bassam M, Tasiopoulos A, Psaras I (2019) Proof-of-Prestige: A Useful Work Reward System for Unverifiable Tasks.

Kubler S, Renard M, Ghatpande S, Georges J-P, Le Traon Y (2023) Decision support system for blockchain (DLT) platform selection based on ITU recommendations: A systematic literature review approach. Expert Systems with Applications 211:118704. doi:10.1016/j.eswa.2022.118704.

Küfeoğlu S, Özkuran M (2019) Bitcoin mining: A global review of energy and power demand. Energy Research & Social Science 58:101273. doi:10.1016/j.erss.2019.101273.

Labazova O (2019) Towards a framework for evaluation of blockchain implementations.

Lasla N, Alsahan L, Abdallah M, Younis M (2020) Green-PoW: An Energy-Efficient Blockchain Proof-of-Work Consensus Algorithm.

Lee SU, Zhu L, Jeffery R (2018) Designing data governance in platform ecosystems.

Lei N, Masanet E, Koomey J (2021) Best practices for analyzing the direct energy use of blockchain technology systems: Review and policy recommendations. Energy Policy 156:112422. doi:10.1016/j.enpol.2021.112422.

Li B, Chenli C, Xu X, Shi Y, Jung T (2019) DLBC: A Deep Learning-Based Consensus in Blockchains for Deep Learning Services.

Loe AF, Quaglia EA (2018) Conquering Generals: an NP-Hard Proof of Useful Work. In: Proceedings of the 1st Workshop on Cryptocurrencies and Blockchains for Distributed Systems. ACM, New York, NY, USA, pp 54–59.

Lundbæk L-N, Janes Beutel D, Huth M, Jackson S, Kirk L, Steiner R (2018) Proof of Kernel Work: a democratic low-energy consensus for distributed access-control protocols. Royal Society open science 5(8):180422. doi:10.1098/rsos.180422.

Marangappanavar RK, Kiran M (2021) Proof-of-Equality: Fairness Ensured Consensus Mechanism for Blockchain Technology. In: Kumar R, Dohare RK, Dubey H, Singh VP (eds) Applications of Advanced Computing in Systems. Springer Singapore, Singapore, pp 153–161.

McDonald K (2021) Ethereum Emissions: A Bottom-up Estimate.

Milutinovic M, He W, Wu H, Kanwal M (2016) Proof of Luck: an Efficient Blockchain Consensus Protocol. In: Proceedings of the 1st Workshop on System Software for Trusted Execution. ACM, New York, NY, USA, pp 1–6.

Mohsin M, Naseem S, Zia‐ur‐Rehman M, Baig SA, Salamat S (2020) The crypto‐trade volume, GDP, energy use, and environmental degradation sustainability: An analysis of the top 20 crypto‐trader countries. International Journal of Finance & Economics. doi:10.1002/ijfe.2442.

Mora C, Rollins RL, Taladay K, Kantar MB, Chock MK, Shimada M, Franklin EC (2018) Bitcoin emissions alone could push global warming above 2°C. Nature Climate Change 8(11):931–933. doi:10.1038/s41558-018-0321-8.

Mulligan C, Scott JZ, Warren S, Rangaswami JP (2018) Blockchain Beyond the Hype. A Practical Framework for Business Leaders. World Econmic Forum.

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

O'Dwyer KJ, Malone D (2014) Bitcoin Mining and its Energy Footprint. In: 25th IET Irish Signals & Systems Conference 2014 and 2014 China-Ireland International Conference on Information and Communities Technologies (ISSC 2014/CIICT 2014). Institution of Engineering and Technology, pp 280–285.

Otto B, Jarke M (2019) Designing a multi-sided data platform: findings from the International Data Spaces case. Electronic Markets 29(4):561–580.

Ouaili L, Banerjee S, Kornyshova E (2022) Towards Possibilities of Energy Minimization in Consensus and Mining Paradigm. In: 2022 9th International Conference on Future Internet of Things and Cloud (FiCloud). IEEE, pp 272–276.

Panian Z (2010) Some practical experiences in data governance. World Academy of Science, Engineering and Technology(62.1):939–946.

Pedersen AB, Risius M, Beck R (2019) A Ten-Step Decision Path to Determine When to Use Blockchain Technologies. MIS Quarterly Executive:99–115. doi:10.17705/2msqe.00010.

Platt M, Sedlmeir J, Platt D, Xu J, Tasca P, Vadgama N, Ibanez JI (2021) The Energy Footprint of Blockchain Consensus Mechanisms Beyond Proof-of-Work. In: 2021 IEEE 21st International Conference on Software Quality, Reliability and Security Companion (QRS-C). IEEE, pp 1135–1144.

Polemis ML, Tsionas MG (2021) The environmental consequences of blockchain technology: A Bayesian quantile cointegration analysis for Bitcoin. International Journal of Finance & Economics. doi:10.1002/ijfe.2496.

Polygon Labs (2022) Polygon Announces The World's First Zero-Knowledge (ZK) Scaling Solution Fully Compatible with Ethereum. https://polygon.technology/blog/polygon-announces-the-worlds-first-zero-knowledge-zk-scaling-solution-fully-compatible-with-ethereum, Accessed 2023-08-08.

Qin S, Klaaßen L, Gallersdörfer U, Stoll C, Zhang D (2021) Bitcoin's future carbon footprint.

Qu X, Wang S, Hu Q, Cheng X (2021) Proof of Federated Learning: A Novel Energy-Recycling Consensus Algorithm. IEEE Transactions on Parallel and Distributed Systems 32(8):2074–2085. doi:10.1109/TPDS.2021.3056773.

Reetz, F. (2019) Herausforderungen und Förderstrategien für die Blockchain-Technologie, Studien zum deutschen Innovationssystem, No. 10-2019, Expertenkommission Forschung und Innovation (EFI).

Rieger A, Roth T, Sedlmeir J, Fridgen G (2022) We Need a Broader Debate on the Sustainability of Blockchain. Joule 6(6):1137–1141. doi:10.1016/j.joule.2022.04.013.

Roth T, Stohr A, Amend J, Fridgen G, Rieger A (2022) Blockchain as a driving force for federalism: A theory of cross-organizational task-technology fit. International Journal of Information Management:102476. doi:10.1016/j.ijinfomgt.2022.102476.

Sai AR, Vranken H (2022) Promoting Rigour in Blockchains Energy & Environmental Footprint Research: A Systematic Literature Review.

Samonas S, Coss D (2014) The CIA strikes back: Redefining confidentiality, integrity and availability in security. Journal of Information System Security 10(3).

Schellinger B, Völter F, Urbach N, Sedlmeir J (2022) Yes, I Do: Marrying Blockchain Applications with GDPR. In: Proceedings of the 55th Hawaii International Conference on System Sciences, pp 4631–4640.

Schlatt V, Sedlmeir J, Traue J, Völter F (2023) Harmonizing Sensitive Data Exchange and Double-spending Prevention Through Blockchain and Digital Wallets: The Case of E-prescription Management. Distributed Ledger Technologies: Research and Practice 2(1):1–31. doi:10.1145/3571509.

Scriber BA (2018) A Framework for Determining Blockchain Applicability. IEEE Software 35(4):70–77. doi:10.1109/MS.2018.2801552.

Sedlmeir J, Buhl HU, Fridgen G, Keller R (2020a) Recent Developments in Blockchain Technology and their Impact on Energy Consumption. Informatik Spektrum 43(6):391–404. doi:10.1007/s00287-020-01321-z. http://arxiv.org/pdf/2102.07886v1.

Sedlmeir J, Buhl HU, Fridgen G, Keller R (2020b) The Energy Consumption of Blockchain Technology: Beyond Myth. Business & Information Systems Engineering 62(6):599–608. doi:10.1007/s12599-020-00656-x.

Sedlmeir J, Lautenschlager J, Fridgen G, Urbach N (2022a) The Transparency Challenge of Blockchain in Organizations. Electronic Markets:1–16. doi:10.1007/s12525-022-00536-0.

Sedlmeir J, Ross P, Luckow A, Lockl J, Miehle D, Fridgen G (2021a) The DLPS: A Framework for Benchmarking Blockchains. In: Proceedings of the 54th Hawaii International Conference on System Sciences, pp 6855–6864.

Sedlmeir J, Smethurst R, Rieger A, Fridgen G (2021b) Digital Identities and Verifiable Credentials. Business & Information Systems Engineering 63(5):603–613. doi:10.1007/s12599-021-00722-y.

Sedlmeir J, Wagner T, Djerekarov E, Green R, Klepsch J, Rao S (2022b) A Serverless Distributed Ledger for Enterprises. In: Bui T (ed) Proceedings of the 55th Hawaii International Conference on System Sciences. Hawaii International Conference on System Sciences.

Shi X, Xiao H, Liu, Weifeng M Chen, Xi M Lackner, Klaus. S. M Buterin, Vitalik M Stocker, Thomas F. (2021) Confronting the Carbon-footprint Challenge of Blockchain.

Shibata N (2019) Proof-of-Search: Combining Blockchain Consensus Formation With Solving Optimization Problems. IEEE Access 7:172994–173006. doi:10.1109/ACCESS.2019.2956698.

Shoker A (2018) Brief Announcement: Sustainable Blockchains through Proof of eXercise. In: Newport C, Keidar I (eds) Proceedings of the 2018 ACM Symposium on Principles of Distributed Computing. ACM, New York, NY, USA, pp 269–271.

Six N, Herbaut N, Salinesi C (2022) Blockchain software patterns for the design of decentralized applications: A systematic literature review. Blockchain: Research and Applications:100061.

Solat S (2017) RDV: An Alternative To Proof-of-Work And A Real Decentralized Consensus For Blockchain.

Song Y-D, Aste T (2020) The Cost of Bitcoin Mining Has Never Really Increased. Frontiers in Blockchain 3. doi:10.3389/fbloc.2020.565497.

Stoll C, Klaaßen L, Gallersdörfer U (2019) The Carbon Footprint of Bitcoin. Joule 3(7):1647–1661. doi:10.1016/j.joule.2019.05.012.

Syafruddin WA, Dadkhah S, Koppen M (2019) Blockchain Scheme Based on Evolutionary Proof of Work. In: 2019 IEEE Congress on Evolutionary Computation (CEC). IEEE, pp 771–776.

Talukder S, Vaughn R (2021) A Template for Alternative Proof of Work for Cryptocurrencies. In: 2021 International Conference on Smart Generation Computing, Communication and Networking (SMART GENCON). IEEE, pp 1–6.

Taylor, A; Kugler, A; Marella, P, Babu; D, Gaby G. (2022) VigilRx: A Scalable and Interoperable Prescription Management System Using Blockchain. In: IEEE Access 10, S. 25973–25986. DOI: 10.1109/ACCESS.2022.3156015.

Toulemonde A, Besson L, Goubin L, Patarin J (2022) Useful work: a new protocol to ensure usefulness of PoW-based consensus for blockchain. In: Conference on Information Technology for Social Good. ACM, New York, NY, USA, pp 308–314.

Truby J (2018) Decarbonizing Bitcoin: Law and policy choices for reducing the energy consumption of Blockchain technologies and digital currencies. Energy Research & Social Science 44:399–410. doi:10.1016/j.erss.2018.06.009.

Truby J, Brown RD, Dahdal A, Ibrahim I (2022) Blockchain, climate damage, and death: Policy interventions to reduce the carbon emissions, mortality, and net-zero implications of non-fungible tokens and Bitcoin. Energy Research & Social Science 88:102499. doi:10.1016/j.erss.2022.102499.

Tschorsch F, Scheuermann B (2016) Bitcoin and beyond: A technical survey on decentralized digital currencies. IEEE Communications Surveys & Tutorials 18(3):2084–2123.

van den Broek T, van Veenstra AF (2015) Modes of governance in inter-organizational data collaborations.

Vranken H (2017) Sustainability of bitcoin and blockchains. Current Opinion in Environmental Sustainability 28:1–9. doi:10.1016/j.cosust.2017.04.011.

Wang J, Gem Lina (2022) Consensus Algorithm of Proof-of-Stake Based on Credit Model. In: The 2022 4th International Conference on Blockchain Technology. ACM, New York, NY, USA, pp 88–94.

Webster J, Watson RT (2002) Analyzing the past to prepare for the future: Writing a literature review. MIS quarterly:xiii–xxiii.

Wei Y, An Z, Leng S, Yang K (2022) Evolved PoW: Integrating the Matrix Computation in Machine Learning into Blockchain Mining. IEEE Internet of Things Journal:1. doi:10.1109/JIOT.2022.3165973.

Weill P (2004) IT Governance: How Top Performers Manage IT Decision Rights for Superior Results. Harvard Business Review Press, Boston.

Wen Y, Lu F, Liu Y, Cong P, Huang X (2020) Blockchain Consensus Mechanisms and Their Applications in IoT: A Literature Survey. In: Qiu M (ed) Algorithms and Architectures for Parallel Processing. Springer International Publishing, Cham, pp 564–579.

Wüst K, Gervais A (2018) Do you Need a Blockchain? In: 2018 Crypto Valley Conference on Blockchain Technology (CVCBT). IEEE, pp 45–54.

Xu X, Dilum Bandara H, Lu Q, Weber I, Bass L, Zhu L (2021) A Decision Model for Choosing Patterns in Blockchain-Based Applications. In: 2021 IEEE 18th International Conference on Software Architecture (ICSA). IEEE, pp 47–57.

Xue T, Yuan Y, Ahmed Z, Moniz K, Cao G, Wang C (2018) Proof of Contribution: A Modification of Proof of Work to Increase Mining Efficiency. In: 2018 IEEE 42nd Annual Computer Software and Applications Conference (COMPSAC). IEEE, pp 636–644.

Yu L, Zhao X, Jin Y, Cai H, Wei B, Hu B (2019) Low powered blockchain consensus protocols based on consistent hash. Frontiers of Information Technology & Electronic Engineering 20(10):1361–1377. doi:10.1631/FITEE.1800119.

Zade M, Myklebost J, Tzscheutschler P, Wagner U (2019) Is Bitcoin the Only Problem? A Scenario Model for the Power Demand of Blockchains. Frontiers in Energy Research 7. doi:10.3389/fenrg.2019.00021.

www.dena.de

future-energy-lab.de

German Energy Agency