

Association for Information Systems

AIS Electronic Library (AISeL)

ICIS 2024 Proceedings

International Conference on Information
Systems (ICIS)

December 2024

Toward a Holistic Perspective on Blockchain Electricity Consumption

Vincent Gramlich

Frankfurt University of Applied Sciences, vincent.gramlich@fim-rc.de

Tobias Guggenberger

University of Bayreuth, tobias.guggenberger@fim-rc.de

Felix Paetzold

FIM Research Center, University of Bayreuth, felix.paetzold@fim-rc.de

Johannes Sedlmeir

University of Luxembourg, johannes.sedlmeir@uni.lu

Jens Strüker

FIM Research Center, University of Bayreuth, jens.strueker@fim-rc.de

Follow this and additional works at: <https://aisel.aisnet.org/icis2024>

Recommended Citation

Gramlich, Vincent; Guggenberger, Tobias; Paetzold, Felix; Sedlmeir, Johannes; and Strüker, Jens, "Toward a Holistic Perspective on Blockchain Electricity Consumption" (2024). *ICIS 2024 Proceedings*. 5.
<https://aisel.aisnet.org/icis2024/blockchain/blockchain/5>

This material is brought to you by the International Conference on Information Systems (ICIS) at AIS Electronic Library (AISeL). It has been accepted for inclusion in ICIS 2024 Proceedings by an authorized administrator of AIS Electronic Library (AISeL). For more information, please contact elibrary@aisnet.org.

Toward a Holistic Perspective on Blockchain Electricity Consumption

Completed Research Paper

Vincent Gramlich

Fraunhofer FIT
Bayreuth, Germany
vincent.gramlich@fit.fraunhofer.de

Tobias Guggenberger

University of Bayreuth
Bayreuth, Germany
tobias.guggenberger@fim-rc.de

Felix Paetzold

FIM Research Center
Bayreuth, Germany
felix.paetzold@fim-rc.de

Johannes Sedlmeir

SnT, University of Luxembourg
Luxembourg, Luxembourg
johannes.sedlmeir@uni.lu

Jens Strüker

Branch Business & Information Systems Engineering, Fraunhofer FIT
Bayreuth, Germany
jens.strueker@fit.fraunhofer.de

Abstract

The awareness of Bitcoin's problematic electricity consumption has carried over to the underlying technology as a whole, leading to a widespread and controversial discourse on the sustainability of blockchain networks that still reveals knowledge gaps. In this paper, we conduct a systematic analysis to identify the scientific body of knowledge on key components and factors that impact blockchain electricity consumption. We find that most research so far has focused on Bitcoin and proof-of-work-based cryptocurrencies, with less attention given to blockchain networks that operate with far less electricity-intensive consensus mechanisms or employ emerging scaling solutions. Building on a systematic literature review and additional explorative and inductive reasoning, we present a comprehensive list of determining factors of blockchain electricity consumption and discuss how they are interconnected. Our research structures methodologies and parameters for measuring the electricity consumption of blockchains and identifies important gaps and avenues for future research.

Keywords: Cryptocurrency, distributed ledger technology, energy demand, proof-of-stake, proof-of-work.

Introduction

In 2018, a Nature Climate Change publication posited that the carbon emissions attributed to Bitcoin's electricity consumption alone might push global warming beyond 2°C within 30 years (Mora et al., 2018). One of the study's foundational assumptions, a proportionality of electricity consumption and the number of transactions processed on the Bitcoin network, was quickly pointed out to be erroneous (Dittmar & Praktiknjo,

2019; Houy, 2019; Masanet et al., 2019). Nevertheless, even under more conservative assumptions, the electricity consumption of Bitcoin exceeds 100 TWh per year and, as such, matches that of medium-sized industrialized countries like Norway (CCAF, 2024; Gallersdörfer et al., 2020). On the other hand, blockchain is seen as a promising technology, projected to transform many industries. For instance, permissionless blockchains enable a new decentralized financial ecosystem or other Web3-based business models (Gramlich et al., 2023). Furthermore, employing blockchain-based infrastructures for improving sustainability is also explored, e.g., for tracing and verifying CO₂-emissions (e.g., Babel et al., 2022). Yet, the initial alarming reports about the energy needs of blockchain technology in general have led to a persistent skepticism in public discourse, which often overlooks the heterogeneity of blockchain networks regarding their electricity consumption (Schmidt, 2021; Sedlmeir et al., 2020b).¹

While a rich body of research scrutinizes the resource demands of blockchain technologies (e.g., de Vries, 2018; de Vries, 2021; Krause & Tolaymat, 2018; Sedlmeir et al., 2020b; Vranken, 2017) or also covers other environment-related aspects (e.g., Krause & Tolaymat, 2018; Wendl et al., 2023), it is increasingly evident that this scholarship suffers from certain methodological weaknesses and a lack of comprehensive vision, as highlighted by Lei et al. (2021) and Sai & Vranken (2023). The emerging consensus suggests that the field's current trajectory is fragmented, missing the interconnections and complexities of blockchain networks' electricity consumption. At the same time, there is a substantial public interest in developing a systematic and comprehensive methodology for assessing the energy consumption of activities involving crypto-assets (DG-FISMA, 2023), which provides the backbone for regulatory measures that can incentivize the use of less energy-intensive crypto-assets, e.g., through taxation. Therefore, it is apparent that future inquiries must adopt an integrative approach that encompasses the multifaceted nature of blockchain electricity consumption. This perspective should consider the intricate interplay of technical design choices – e.g., consensus mechanisms, block sizes, and block times – and economic factors – e.g., hardware affordability and electricity pricing. Only through such a comprehensive lens can the sustainability and efficacy of blockchain networks be truly evaluated, and the costs and benefits of its applications fully understood. This thorough understanding paves the way for more informed decisions regarding blockchain's role in sustainable development. Thus, the overarching question we seek to address is:

What are the key components and factors that influence blockchain electricity consumption?

To answer this research question, we conduct a systematic literature review (SLR) according to the guidelines of Kitchenham et al. (2009), supplement it with exploratory analyses, and use inductive reasoning to provide valuable insights into the electricity consumption of different blockchain designs. Our study holistically collects and organizes design options and parameters that impact electricity consumption. In particular, we highlight the importance of considering sometimes neglected economic factors alongside technical characteristics in future research on blockchain sustainability. Given its interdisciplinary nature, we argue that the IS community is uniquely positioned to lead this research.

Background

Blockchain technology refers to distributed systems where *nodes* maintain a shared and synchronized database building on a peer-to-peer (P2P) network (Butijn et al., 2020). It employs decentralized consensus mechanisms and an append-only data structure for efficient synchronization and tamper-detection (Beck et al., 2018). In addition to fault tolerance, public verifiability resulting from the replicated processing and storage of transactions is a core characteristic (Butijn et al., 2020). One dimension to distinguish blockchains is the restriction of participation in consensus (Beck et al., 2018). Permissioned blockchains require an onboarding process in which nodes' identities are verified and receive the authorization to participate, while the permissionless blockchains underpinning most cryptocurrencies allow anyone to participate in consensus.

¹Note that according to the law of conservation of energy, energy can only be transformed between different forms. Therefore, while the *electric* energy required for operating computing devices can be “consumed” by transforming it into computational work and heat, energy per se cannot be consumed. To facilitate both conciseness and correctness, we will use the term “electricity consumption” in the following. This perspective also accounts for the fact that we do not consider the primary energy needs for generating the electric energy consumed by blockchain networks.

Permissionless blockchain networks assign a node a specific voting weight in consensus by coupling its probability of being entitled to add a block with a scarce resource that is verifiable in the P2P network (Sedlmeir et al., 2020a). Participation, in turn, is incentivized by giving out rewards, mostly in cryptocurrency, to consensus participants for contributing new blocks (Stinner, 2022). In many first-mover cryptocurrencies and specifically Bitcoin, the scarce resource is computational power; thus, hardware and substantial amounts of electricity are required for “mining” operations. The corresponding consensus mechanism is called proof-of-work (PoW) (Nakamoto, 2008). To find a new block in a PoW blockchain network, nodes must brute-force a solution to a specific cryptographic puzzle by performing many hashing operations. In non-PoW blockchains, the scarce resource is different. Proof-of-stake (PoS) based blockchain networks use the cryptocurrency itself – referred to as the users’ *stake* – as the scarce resource, which can be easily verified on the public ledger (Saleh, 2021). One example of a PoS blockchain is Ethereum, which switched from PoW to PoS in September 2022 (Rieger et al., 2022). To be eligible to participate in the consensus process, users often need to lock their stake for a certain period of time. This approach addresses “nothing at stake attacks” and provides nodes with additional incentives to abide by the rules of the network, as any detected misconduct can result in a reduction or loss of the staked capital (*slashing*) (Álvarez et al., 2024). Hence, in PoS systems, the determination of voting weight is no longer coupled to computing power, thereby significantly reducing electricity consumption. Besides PoW and PoS, other scarce and digitally verifiable resources are also utilized, such as disk space in Filecoin (Fisch et al., 2018).

Regardless of the consensus mechanism, there are three different participant groups in a blockchain network (see Table 1). The first group, at the core of the network, are the consensus participants. They participate for instance, as “miners” by attempting to create new blocks in PoW and, thus, voting on previous blocks (Nakamoto, 2008) or “stakers” (by creating new blocks or attesting to blocks by signing them with a cryptographic key that controls a certain amount of stake in PoS (Saleh, 2021)). In the following, we term consensus participants of non-PoW blockchains “validators”. We associate the electricity consumption of mining hardware or other devices used for participation in consensus with the particular miner or validator that runs it. The second participant group comprises (full) nodes that receive, verify new blocks, and store the entire append-only ledger. They furthermore form the P2P network, responsible for broadcasting new transactions to be included in upcoming blocks as well as new blocks. In addition to full nodes, there are light nodes that store only a few recent blocks (or block headers) and the part of the state that affects them (e.g., their own balance) to accommodate resource-constrained devices such as mobile phones (Chatzigiannis et al., 2022). On the other hand, there are archive nodes with high storage capacity that extend full nodes’ information by maintaining a versioned history of the state (Ethereum, 2023).

Lastly, there is the broader group of network participants – all other stakeholders interacting with the blockchain network, such as end-users who do not run a (light) node but instead interact with the blockchain by sending new transactions through clients running on mobile devices or computers to other full nodes. Many blockchain-based applications (e.g., platforms for tokens or cryptocurrency exchanges) also involve additional back-end and front-end services (Gramlich et al., 2023).

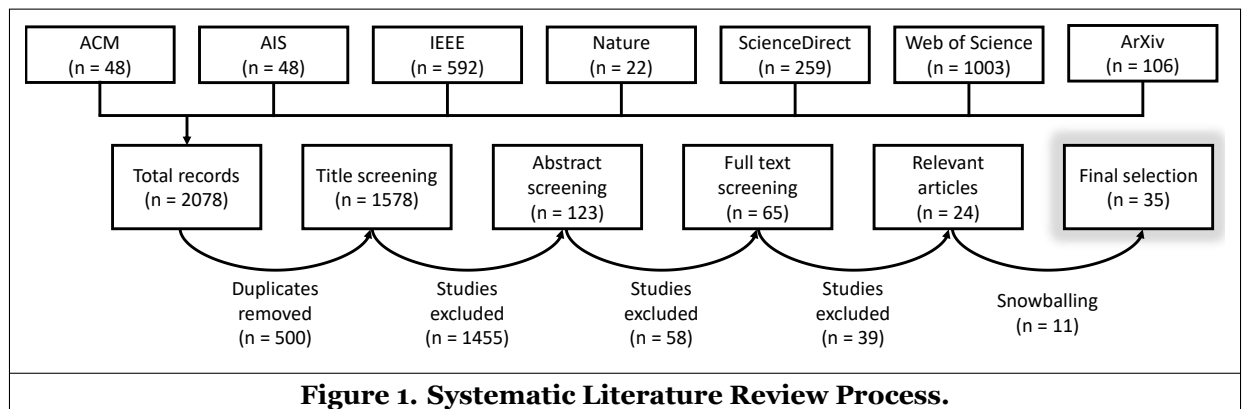
Participant group	Role in the network
Consensus participants	Attesting to or proposing new blocks and receiving the respective rewards, entitled through the provision of the scarce resource
Nodes	Storing the ledger and the current state and forming part of the P2P network responsible for the propagation of new transactions and blocks
Broader network participants	All other stakeholders interacting with the blockchain network, e.g., by using it for transactions or providing services built on top
Table 1. Different Participant Groups in a Blockchain Network and Their Roles.	

Method

We conducted an SLR to gather the existing knowledge on the electricity consumption of blockchain networks, following the research process outlined by Kitchenham et al. (2009). This approach provides a holistic and structured overview of previous work, identifies relevant knowledge gaps, and determines new research opportunities (Webster & Watson, 2002). A preliminary search of seminal publications helped us to identify relevant keywords and synonyms for our research questions, which informed the construction of a two-part search string focused on blockchain electricity consumption. The first part focuses on blockchain technology and distributed ledger technology (DLT), while the second part represents the research area on electricity consumption. We incorporated Bitcoin and Ethereum in the first part of the search string as they were the specific study objects of many early publications on blockchain electricity consumption, as well as “cryptocurrency” as a prevalent application of blockchain technology. By coupling “electricity”, “energy”, and “power” with “demand” and “consumption”, we narrowed the search to exclude publications on blockchain applications in the energy sector (Andoni et al., 2019). Similarly, we excluded “sustainability” to omit discussions of blockchain use, e.g., in the circular economy. Our final search string was

(blockchain OR cryptocurrency OR “distributed ledger” OR DLT OR Bitcoin OR Ethereum) AND [(electricity OR energy OR power) AND (consumption OR demand)].

We searched six well-established databases covering computer science and social sciences: ACM Digital Library, AiSEL, IEEE Xplore, Nature, ScienceDirect, and Web of Science. To encompass also the latest research that reflects developments like Ethereum’s switch to PoS, we furthermore added ArXiv. Applying our search string on April 4th, 2024 in a title, abstract, and keywords search yielded a total of 2078 articles. During the title screening, we only included English publications related to blockchain electricity consumption. We then reviewed the abstracts of these articles and included publications that quantify the electricity consumption of blockchain networks, determine components or factors impacting it, and conduct meta-studies on these topics. This narrowed our selection down to 65 papers. For each remaining paper, we conducted a full-text screening and included only publications where (1) authors derived their own estimates of electricity consumption for a blockchain-based network, e.g., by modeling or measuring different components’ electricity consumption, (2) refined estimates from another publication, or (3) provided meta-studies of corresponding estimates or methods to obtain them. Consequently, we excluded publications that estimate CO₂ emissions by relying on electricity consumption estimates from other sources. We also excluded publications that apply machine learning or regression methods to find correlations between trading activity and electricity consumption, as these publications do not assess the underlying factors and provide only statistical, not causal, relationships. After applying these exclusion criteria, 24 relevant publications remained. As the last step of the literature selection process, we carried out a snowballing search (i.e., forward and backward searches) (Webster & Watson, 2002). We found eleven additional relevant papers, of which nine represent gray literature. Finally, we conducted an in-depth full-text analysis of the total 35 papers. We coded the literature items using MAXQDA, systematically categorizing the methods employed, the parameters used, and the corresponding data sources for these parameters.



Most of the publications we analyzed focus on estimating the electricity consumption of individual blockchain networks. Only two publications review existing studies: Sai & Vranken (2023) provide an SLR that highlights the lack of rigor in previous works, and Lei et al. (2021) emphasize several limitations of previous publications. We found no systematic investigation of the components and factors influencing the electricity consumption of blockchain networks. As such, our SLR also justifies our research question (Müller-Bloch & Kranz, 2015).

Existing Approaches to Assessing Blockchain Electricity Consumption

As established in the Background section, a blockchain network comprises various participant groups, each contributing differently to the network's functionality and electricity consumption. Given the distributed nature of a blockchain network, the total electricity consumption is the aggregate of all the individual participants and their operational components' (Platt et al., 2021; Stoll et al., 2019). Related work accounts for two primary contributors to a blockchain network's electricity consumption: (1) participation in consensus, particularly in PoW systems where specific mining hardware is employed, and (2) node operations. According to the final selection of publications from our SLR, there are profound conceptual and numerical differences between these contributors. We structure the following discussion accordingly.

Proof-of-work networks

25 out of the 35 papers from our SLR investigate the electricity consumption of PoW networks. 12 papers analyze exclusively the Bitcoin network, while two consider Ethereum before its transition to PoS. The remaining twelve publications study Bitcoin, Ethereum, and up to 20 other PoW cryptocurrencies' electricity consumptions. Despite substantial differences among the design of PoW networks, some fundamental assumptions that apply to all PoW networks allow to derive bounds for the corresponding electricity consumption based on publicly observable data and assumptions about miners' equipment and decision-making (Sedlmeir et al., 2020a). Our SLR reveals various procedures applied by related work that can be classified into three core categories: technical, economic, and hybrid approaches.

Technical approach

The technical approach estimates a PoW network's electricity consumption based on its hash rate and the efficiency factor of mining devices. Thus, the network's *power* consumption P , i.e., its electricity consumption per second, is $P = H \cdot e$, where H is the network's hash rate in hashes per second and e is the average electric energy required to compute a single hash (Coroamă, 2022; Sedlmeir et al., 2020b). Multiplying by the number of seconds in a certain time duration, like a year, determines the electric energy consumption for that period. Table 2 lists the 22 publications using a technical approach for estimating the electricity consumption of PoW networks and outlines the parameters used in these publications: Hash rate, the maximum or minimum hashing efficiency (HE) of mining hardware, and power usage effectiveness (PUE) as an indicator for the share of electric energy that can be used for powering mining hardware to holistically account for the required infrastructure, e.g., for cooling the mining devices (de Vries, 2018). From the seven studies that include PUE as an extra multiplier, only two validate their PUE in expert interviews (CCAF, 2024; Stoll et al., 2019). Furthermore, in all publications, a global PUE value is used, ignoring regional characteristics of the mining operations like the dependency of cooling on the climate zone (de Vries, 2018; Lei et al., 2021).

While a PoW blockchain network's hash rate is not directly observable, one can compute it accurately at any time based on the average number of attempts required for solving a hash-puzzle (which is publicly recorded by its difficulty) and the average number of blocks created per time period. Note that this approach does not consider orphaned blocks and neglects the mining of parallel temporary forks, so the actual hash rate may be higher, particularly for blockchains with smaller block time (Gervais et al., 2016). Many studies derived the hardware efficiency factor, or electric energy per hash, from manufacturers' specifications (e.g., Coinshare, 2022; Stoll et al., 2019; Vranken, 2017). The open, pseudonymous nature of permissionless blockchains makes direct observation of the distribution of miners' hash rates and hardware choices elusive (Sedlmeir et al., 2020a; Stinner, 2022). Nevertheless, research has proposed strategies to infer the distribution of mining hardware. 15 studies from the SLR have pinpointed the most efficient hardware to determine the

Publication	Network(s)	Hash Rate	HE Avg.	HE Max.	PUE
CCAF (2024)	Bitcoin and Ethereum	x		x	x
Coinshare (2022)	Bitcoin	x	x		x
Coroamă (2022)	Bitcoin	x		x	x
CCRI (2022d)	Ethereum	x	x		
de Vries (2018)	Bitcoin	x		x	x
de Vries (2020)	Bitcoin	x	x	x	x
de Vries (2022)	Ethereum	x		x	
Digiconomist (2024)	Bitcoin	x		x	x
Gallersdörfer et al. (2020)	Top 20 mineable currencies	x			
Jones et al. (2022)	Bitcoin	x			
Küfeoğlu & Özkuran (2019)	Bitcoin	x		x	
Li et al. (2019)	Dodgecoin	x	x		
Mora et al. (2018)	Bitcoin	x	x		
O'Dwyer & Malone (2014)	Bitcoin	x		x	
Qin et al. (2021)	Bitcoin	x	x		
Sedlmeir et al. (2020b)	Bitcoin, Ethereum, Bitcoin Cash, Bitcoin SV, Litecoin	x		x	
Sedlmeir et al. (2020a)	Bitcoin, Ethereum, Bitcoin Cash, Bitcoin SV, Litecoin	x		x	
Shi et al. (2023)	Ethereum	x		x	
Song & Aste (2020)	IOTA		x		
Stoll et al. (2019)	Bitcoin	x		x	x
Vranken (2017)	Bitcoin	x			
Zade et al. (2019)	Bitcoin and Ethereum	x		x	
HE: Hashing efficiency of mining hardware, PUE: Power usage effectiveness of mining facilities					
Table 2. Publications Using the Technical Approach.					

network's minimum power requirement (e.g., de Vries, 2022; Sedlmeir et al., 2020b). Yet, this approach underestimates the actual electricity consumption, as it presumes that all miners utilize the most efficient gear and neglects factors such as aged equipment or hardware shortages (CCAF, 2024; Sai & Vranken, 2023; Zade et al., 2019). In contrast, Li et al. (2019) and Song & Aste (2020) quantified the electricity consumption of mining devices using consumer hardware in a controlled testbed environment, though more efficient specialized mining hardware was predominantly utilized at the time.

An accurate estimate requires a nuanced understanding of the hash rate's distribution across different efficiency levels (Lei et al., 2021) and involves a weighted sum of various hardware contributions, implying that "average hardware efficiency" must be interpreted carefully. Some researchers infer hardware distribution from manufacturers' reports (de Vries, 2018; Shi et al., 2023; Zade et al., 2019), but this is feasible mainly for networks like Bitcoin with specialized hardware (e.g., application-specific integrated circuits (ASICs)) where sales data can be directly linked to mainly a single blockchain (de Vries, 2018). Other studies analyze network-wide mining activities to derive the hardware distribution. For example, Mora et al. (2018) associated a specific mining device's energy efficiency with a mined block, disregarding the higher probability of more efficient devices successfully mining a block (Houy, 2019) and, therefore, potentially underestimating efficiency. In contrast, Coinshare (2022) proposed a method to detect hardware by analyzing patterns in the solutions found to hash puzzles, although this approach has faced criticism for overestimating the corresponding device type's share in hardware distribution (CCAF, 2024).

Economic approach

The second approach to quantifying the electricity consumption of a PoW network assumes an equilibrium in the mining economy (Sedlmeir et al., 2020b). It is based on the model of rent-seeking competition among miners, with higher investments into computing power providing higher revenues (Capponi et al., 2023). The competition's prizes are block rewards and transaction fees paid in the native cryptocurrency. Miners' revenues are also influenced by the current cryptocurrency price, as cryptocurrency is arguably not broadly used for electricity and hardware procurement (Stinner, 2022). Consequently, higher exchange rates result in higher revenues and, thus, incentivize higher electricity consumption. In particular, miners' revenues define an upper limit to their investment in electricity, assuming that rational miners don't spend more than

they earn: total mining revenue \geq total mining costs. Table 3 features the nine publications we identified based on this method. Total mining revenue, derivable from historical data on block rewards and transaction fees (Stinner, 2022), is contrasted with total mining costs, which are not directly observable and depend on individual miners' cost structures. To estimate the network's electricity consumption, one must make assumptions about the cost distribution among miners. The miner's cost structure is primarily reflected by variable electricity costs, which are linked to expected revenue (Capponi et al., 2023), and the fixed hardware costs. Revenues minus fixed costs set an upper bound for a miner's electricity costs.

A challenge arises due to limited empirical data on electricity prices paid by miners, obscuring transparency in many publications (Sai & Vranken, 2023). Only Qin et al. (2021) and CCAF (2024) have validated their electricity price models through interviews with miners. Most studies presume a global electricity price ranging from \$0.05 to \$0.13 per kWh, often without citing primary sources. Moreover, they typically assume a constant electricity price over several years. Gonzalez-Barahona (2021) present sensitivity analyses for varying electricity prices, and Shi et al. (2023) consider temporal and regional price fluctuations. Nonetheless, their reliance on average household rather than industrial prices likely overestimates costs (CCAF, 2024). Hayes (2015) overlook the full revenue stream, omitting transaction fees as part of the miners' rewards and, therefore, potentially underestimating miners' total income and affecting the accuracy of the calculated upper limit for their electricity costs. The approach assumes rational miners only use their hardware as long as it is profitable (Garratt & Oordt, 2023; Houy, 2019). This assumption could lead to an underestimation of the offered hash power by ignoring miners who are not motivated solely by monetary incentives.

Seasonal impacts, such as China's rainy season's effect on electricity prices, are seldom quantitatively addressed in electric energy consumption studies, except for mentions by Sedlmeir et al. (2020b). Notably, the impact of such exogenous shocks on the economics of mining has been studied in publications on the security of PoW blockchains (e.g., Stinner, 2022). Additionally, many studies ignore long-term expenses, like new hardware and personnel, erroneously assuming that miners spend all income on electricity. To correct for this overly simplification, several publications from our SLR examine the interplay between fixed (hardware) and variable (electricity) costs in determining electricity consumption (e.g., de Vries, 2021; Gonzalez-Barahona, 2021). Among the studies that consider fixed costs or hardware depreciation, most agree that Bitcoin and Ethereum miners allocate between 50 and 70 % of their total revenue to electricity costs (de Vries, 2018; Digiconomist, 2024; Qin et al., 2021; Shi et al., 2023).

Hybrid approach

To address the shortcomings of the individual approaches outlined above, some authors integrate technical and economic approaches into a hybrid approach. This approach combines the assumption of miners' rationality and the network's observed hash rate to enhance accuracy (Lei et al., 2021). As we illustrate in Table 4, a total of seven publications use this rationale for approximating the distribution of hardware, while one paper provides a longitudinal analysis of the electricity consumption of Bitcoin. The hybrid approach first determines the least efficient mining device still in use by analyzing the profitability of each device type based on its expected revenue and electricity costs. Using the observed hash rate, an upper bound on the network's electricity consumption is then derived (Lei et al., 2021). This method generally assumes fully

Publication	Network(s)	Block Reward	Transaction Fees	Electricity Costs	Fixed Costs
de Vries (2018)	Bitcoin	x	x	5 ct/kWh	x
de Vries (2022)	Ethereum	x	x	10 ct/kWh	
Digiconomist (2024)	Bitcoin	x	x	5 ct/kWh	x
Gonzalez-Barahona (2021)	Bitcoin	x	x	— *	x
Hayes (2015)	Bitcoin	x		13 ct/kWh	
Sedlmeir et al. (2020b)	Bitcoin, Ethereum, Bitcoin Cash, Bitcoin SV, Litecoin	x	x	5 ct/kWh	
Sedlmeir et al. (2020a)	Bitcoin, Ethereum, Bitcoin Cash, Bitcoin SV, Litecoin	x	x	5 ct/kWh	
Shi et al. (2023)	Ethereum	x	x	16.8 – 17.8 ct/kWh**	x
Vranken (2017)	Bitcoin	x	x	3.5 – 20 ct/kWh***	x
* Different models and sensitive analysis ** Regional prices, *** Regional and temporal prices					
Table 3. Publications Using the Economic Approach.					

Publication	Network(s)	Hash Rate	Block Reward	TX Fees	HE Min.	HE Avg.	Electricity Costs
CCAF (2024)	Bitcoin and Ethereum	x	x	x	x	x	5 ct/kWh & 10 ct/kWh
Bevand (2018)	Bitcoin	x	x		x		5 ct/kWh
Coroamă (2022)	Bitcoin	x	x	x	x		5 ct/kWh
Gallersdörfer et al. (2020)	Top 20 mineable currencies	x	x		x		5 ct/kWh
Krause & Tolaymat (2018)	BTC, ETH, Litecoin, Monero	x	x		x		5 ct/kWh
Küfeoğlu & Özkuran (2019)	Bitcoin	x	x		x		9.1 ct/kWh
Stoll et al. (2019)	Bitcoin	x	x	x	x	x	5 ct/kWh
HE: Hashing efficiency of mining hardware							
Table 4. Publications Using the Hybrid Approach.							

deprecated hardware (Bevand, 2018; CCAF, 2024; Coroamă, 2022; Gallersdörfer et al., 2020; Küfeoğlu & Özkuran, 2019; Stoll et al., 2019). However, Sai & Vranken (2023) highlight potential inaccuracies due to the reliance on strong assumptions about electricity prices and the complete rationality of miners in this approach, mirroring limitations seen in the economic approach. This could lead to an overestimation of hashing efficiency. Moreover, some studies (e.g., Bevand, 2018; Gallersdörfer et al., 2020; Krause & Tolaymat, 2018; Küfeoğlu & Özkuran, 2019) again overlook the full revenue stream, for instance, by omitting transaction fees, potentially resulting in an underestimation of miners' total income. For a best guess, rather than an upper bound, CCAF (2024), Gallersdörfer et al. (2020), and Krause & Tolaymat (2018) assumed an equally-weighted distribution of still profitable hardware in their electricity consumption estimates.

None of the publications from our SLR considers components beyond mining hardware in their electricity consumption estimates for PoW networks. This is because the electricity consumption of other network users is considered significantly lower than the electricity consumption associated with the PoW consensus (Vranken, 2017). Indeed, participating in the network as a (full) node only requires commonly available hardware, similar to what is used in non-PoW networks (see below). Notably, the economic approach covers all electricity consumption. Yet, the share of hardware- and electricity-related costs strongly differ between devices for running nodes and dedicated mining hardware. Consequently, for a holistic perspective, both components should be considered.

Non-proof-of-work networks

We identified 13 publications that investigate the electricity consumption of non-PoW blockchains, listed in Table 5. Although research on the economics of mining dates back to before 2015 (Kroll et al., 2013), and many non-PoW based consensus mechanisms were suggested and implemented before 2017, publications on the electricity consumption of non-PoW networks have only appeared since 2020. The absence of rent-seeking mining competition in non-PoW blockchains makes them much less energy-intensive (Saleh, 2021; Sedlmeir et al., 2020a). Therefore, the electricity consumption in a non-PoW network cannot be derived from observations of hash power or economic principles. Instead, it is necessary to consider each node's electricity consumption caused by its participation in the replicated transaction processing and storage, as well as communication in the P2P network. In the absence of energy-intensive mining, one can estimate that each entity within the network consumes roughly the same amount of electricity (CCRI, 2022a; Platt et al., 2021). As a starting point, the total electricity consumption of a non-PoW network is hence approximated by multiplying the number of active nodes and a node's average electricity consumption. Unfortunately, the precise number of participating nodes in public permissionless networks usually cannot be determined (Sai & Vranken, 2023). Therefore, most studies draw the number of nodes participating in consensus from online sources, without further explanation or examination (CCRI, 2022a; CCRI, 2022b; CCRI, 2022c; CCRI, 2022d; de Vries, 2022; Platt et al., 2021; Shi et al., 2023). Only CCAF (2024) used a crawler that utilizes the node discovery protocol of the core node software to identify other nodes. However, after an upgrade to the crawler, the node count doubled (CCAF, 2024), raising concerns about data reliability. A method similar to the economic approach has been applied to Ethereum after its Merge by Shi et al. (2023) to determine the number of nodes based on the expected operating cost of a node. However, they significantly overestimate the number of participating nodes, since multiple validators can be run on a single node. Their analysis also neglects the opportunity cost of staked coins.

Publication	Network(s)	# Nodes	Avg. EC per Node	EC/TX	Testbed	Avg. EC Network	Min./ Max. EC Network
CCAF (2024)	Ethereum	x	x	x	x	x	x
CCRI (2022a)	5 permissionless PoS networks	x	x	x	x	x	x
CCRI (2022d)	Ethereum (post merge)	x	x	x	x	x	x
CCRI (2022b)	Polygon	x	x	x	x	x	x
CCRI (2022c)	Tron	x	x	x	x	x	
CCRI (2023c)	4 permissionless PoS networks, 3 PoS platforms	x	x	x	x	x	x
CCRI (2023b)	Polymesh	x	x	x	x	x	x
CCRI (2023a)	Celo	x	x	x	x	x	x
de Vries (2022)	Ethereum (post merge)	x		x		x	
Platt et al. (2021)	5 PoS networks	x		x		x	
Rieger et al. (2022)	9 permissioned networks	x		x			
Roma & Hasan (2020)	Ripple				x		
Saingre et al. (2022)	Ethereum Clique				x		
Shi et al. (2023)	Ethereum (post merge)	x				x	
EC: Energy consumption, TX: Number of transactions							

Table 5. Publications Estimating Energy Consumption in Non-PoW Networks.

In a non-PoW network, all components need to support general-purpose computation, so highly optimized hardware like ASICs does not play a relevant role. For simplification, de Vries (2022) and Platt et al. (2021) assumed each node uses the same hardware and used this to determine upper and lower bounds for node electricity consumption based on minimum system requirements for node operation and plausible limitations of servers commonly used in cloud computing. To obtain higher accuracy, Roma & Hasan (2020) measured individual nodes' electricity consumption in a testbed situation. Rieger et al. (2022) used the computational and bandwidth use of cloud-hosted blockchain networks under stress as a proxy to derive the corresponding electricity consumption. Some industry studies offer a more sophisticated approach, analyzing multiple configurations and measuring the nodes' electricity consumption under varying transaction throughput to derive an estimation for the whole network (CCRI, 2022b; CCRI, 2022c; CCRI, 2022d; CCRI, 2023b). To finally derive an average electricity consumption of the nodes, the authors estimated a distribution of hardware capable of participating in the network, resulting in a projected hardware distribution with some degree of uncertainty. Additionally, they examined the electricity consumption of different configurations of Ethereum's execution and consensus clients (CCRI, 2022d). CCAF (2024) used these measurements to estimate the Ethereum network's electricity consumption, further distinguishing between various consensus client software types, such as Geth and Reth, and building on the crawler described above. In their updated report, the CCRI (2023c) distinguished between traditional PoS blockchain networks and those supporting side chains such as Avalanche, Cosmos, and Polkadot. Rather than calculating the average electricity consumption for the entire network, they established a linear relationship between the electricity consumption of side chains with the highest and lowest activity levels. Based on this relationship, they then determined the average electricity consumption of a node depending on the side chains it participates in.

In addition to the total electricity consumption of selected blockchain networks, the reports by the Crypto Carbon Rating Institute (CCRI) provide a simple calculation of electricity consumption per transaction by dividing the network's electricity consumption by the number of transactions during the monitored period. This method's reliability is, however, compromised by short observation periods, which neglect long-term changes in throughput and block capacity demand. Moreover, related work emphasizes that the direct comparison between networks based on electricity consumption metrics alone can be problematic because varying levels in decentralization and transaction complexity are not considered (CCRI, 2022a; Platt et al., 2021; Saingre et al., 2022). Platt et al. (2021) also present a model of electricity consumption as a function of network throughput, showing that an increase in throughput leads to a decrease in transactional electricity consumption, as the nodes' idle consumption is spread across a higher number of transactions. As such, the adequacy of the energy-per-transaction metric is contested not only in PoW (Rieger et al., 2022). Moreover, Platt et al. (2021) investigate Hedera, a permissioned PoS design, but only provide estimates instead of empirically measuring nodes' electricity consumption. Rieger et al. (2022) are the only authors to study permissioned blockchain networks with non-PoS consensus (e.g., Quorum and Hyperledger Fabric), demonstrating the impact of the choice of fault tolerance on the marginal electricity consumption per transaction.

Discussion and Future Research Opportunities

Our SLR finds a clear consensus among studies that the electricity consumption characteristics significantly differ between PoW and non-PoW blockchains. Literature also agrees that the electricity consumption of these networks is cumulative, encompassing the total electricity consumption of all participants (O'Dwyer & Malone, 2014; Stoll et al., 2019). Thus, we can assume that from a systemic perspective, the primary distinction in electricity consumption between PoW and non-PoW networks stems from the energy demands of the components involved in consensus. Against the backdrop, they also share common sources of electricity consumption: firstly, the energy expended in transaction processing and storage, replicated across numerous full nodes (CCAF, 2024; CCRI, 2022d); and secondly, the electric energy used by a diverse array of other network participants, such as those transacting or using smart contracts and related services in decentralized finance (Gramlich et al., 2023). Our review underscores that research on PoW blockchains has predominantly focused on the electricity consumption of consensus participants, whereas studies on non-PoW systems have primarily examined the electricity consumption of full nodes. Both strands of literature, however, tend to overlook the implications of electricity consumption for the broader network of participants, suggesting a gap in the current understanding of blockchain electricity consumption dynamics.

In the following, we reflect on the existing knowledge in the literature on blockchain electricity consumption, highlight discrepancies between and gaps in the PoW and non-PoW literature, and point towards open questions – both for the aspects already covered by the existing literature, e.g., the electricity consumption of consensus participants in PoW, and the aspects overlooked by current literature, e.g., the electricity consumption of consensus participants in non-PoW. We structure this discussion according to the three participant groups, i.e., consensus participants, blockchain nodes, and broader network participants, and according to the two main literature strands, i.e., PoW and non-PoW blockchains (see Table 6).

Consensus participants

The electricity consumption associated with consensus participants is arguably the most relevant and, thus, most studied aspect of electricity consumption of existing PoW blockchains. As outlined above, it is well established that a PoW blockchain's hash rate and an estimate of the electricity efficiency of mining devices can be used to approximate electricity consumption. However, the causal relationship stems from the monetary incentive for miners. The value of the incentive provides an upper bound for the expenses of rational miners (Stinner, 2022). As such, literature agrees that the extent of mining incentives, the allocation of expenses related to mining operations, and electricity prices determine miners' electricity consumption (Sedlmeir et al., 2020b). However, the current basket of literature only considers mining rewards and transaction fees as sources of miner revenue. Studies of blockchains with rich economic activities on the application layer suggest that miners also have additional revenue sources. For instance, maximum extractable value (MEV) denotes revenues that block proposers (miners in the case of PoW) can capture by leveraging control over the choice and order of transactions that are included in the next block. As such, they have the opportunity to front-run other transactions, i.e., inserting their transaction in front of another transaction that was submitted to the mempool earlier (Daian et al., 2020), or sell transaction positions in blocks or even ordered bundles of transactions. Qin et al. (2022) estimate that on the Ethereum blockchain, MEV accounted for over 50 million USD in May 2021. A substantial share of MEV is received by miners, i.e., increases miners' revenues. On the Ethereum blockchain, the roadmap toward "proposer-builder separation" aims to ensure that most MEV is extracted by block proposers, such that MEV can be re-distributed more fairly to avoid misaligned incentives that compromise security (Chitra & Kulkarni, 2022). If all MEV were distributed among block proposers when Ethereum was still PoW-based, it would have increased miners' income by several percent and, therefore, potentially their electricity consumption, as miners might have invested more in computational resources and energy to enhance their mining capabilities.

A second factor in the overall equation that governs miners' electricity consumption is costs. For parameters that affect costs, including electricity prices, hardware costs, and other costs, related work has used estimates based on past and current market prices. For electricity prices, predominately uniform prices have been used, which leads to inaccuracies, especially since the country and location of many miners cannot be determined with certainty (Lei et al., 2021). Another aspect influencing electricity costs is the miner's

Participant group	Proof-of-work	Non-proof-of-work
Consensus participants	Extensive knowledge (Tables 2, 3, and 4) Specific open questions: <ul style="list-style-type: none"> • How to obtain more precise estimates for electricity price and hardware distribution? • How to account for additional revenue, e.g., MEV or electricity flexibility? 	Scant knowledge (e.g., Rieger et al., 2022) Basic open questions: <ul style="list-style-type: none"> • How to determine the electricity consumption? • Can methods for node count and node electricity consumption estimations be applied? • Is there a significant difference in electricity consumption for different consensus mechanisms?
Nodes	Absent knowledge Basic open questions: <ul style="list-style-type: none"> • Can methods from non-PoW literature be transferred? • Is there a PoW blockchain network where node electricity consumption is significant? 	Extensive knowledge (Table 5) Specific open questions: <ul style="list-style-type: none"> • How to obtain more precise estimates for node count and hardware distribution? • What is the impact of current trends, e.g., sharding, rollups, or light nodes?
Broader network participants	Not considered at all Basic open questions: <ul style="list-style-type: none"> • In which groups can the broader network participants be divided? • What is the electricity consumption of these groups? • How can it be influenced by technical or economic factors? 	
Table 6. Current State of the Literature and Questions for Future Research for the Electricity Consumption of Participant Groups and Network Types.		

interaction with the electricity grid (Fridgen et al., 2021). In popular mining areas such as Texas, miners adjust their operations to benefit from fluctuating electricity prices and offer flexibility to the grid, potentially reducing overall expenses (Niaz et al., 2022). Such cost optimization strategies, involving both adaptive responses and active contributions to grid stability, add another layer of complexity to accurately estimating miners' total expenses and, ultimately, their electricity consumption. Lastly, hardware and other costs are oftentimes only considered by adding a fixed percentage to electricity costs. Some studies, like Stoll et al. (2019), survey this percentage factor for hardware and other costs based on interviews with mining business operators, but the survey still operates on a small sample size. More in-depth research that engages in empirical data collection can provide more precise estimates and, therefore, further increase the precision of the overall approximation of electricity consumption.

As opposed to PoW, the electricity consumption that originates from participating in non-PoW consensus mechanisms is poorly explored. There is a broad spectrum of consensus mechanisms that can strongly vary in the scarce resource used, the way consensus is achieved, and the security guarantees regarding the type of fault tolerance and the prioritization of integrity or availability guarantees. Moreover, the number of consensus participants can significantly affect bandwidth and processing requirements, especially in some PoS mechanisms that leverage signature aggregation. However, the current literature does not even account for the additional electricity consumption of consensus participation as it equates to the electricity consumption of validators and full nodes (Platt et al., 2021). The only work in our SLR that investigates differences in the electricity consumption between different non-PoW consensus mechanisms was carried out by Rieger et al. (2022) but is merely focused on consensus mechanisms in permissioned blockchains. An expansion of this research to the broader spectrum of non-PoW consensus mechanisms, especially for permissionless

blockchains, can thus improve our understanding of the electricity consumption of non-PoW blockchains and guide more informed decision-making when designing or choosing consensus mechanisms.

Nodes

All publications from our SLR approximate the electricity consumption of non-PoW blockchain networks by multiplying the estimated number of full nodes by their average electricity consumption. For PoW blockchains like Bitcoin, this source of electricity consumption is often ignored due to its insignificance compared to mining activities. However, as the literature outlines, miners' electricity consumption is primarily determined by economic incentives. Therefore, the principle that the electricity consumption of consensus participants is orders of magnitude larger than the electricity consumption of nodes might not hold for all PoW blockchains. In PoW networks with low mining incentives and with many full nodes with a high computational load, the electricity consumption associated with node operation could dominate. The trajectory of Bitcoin with regular halvings of block rewards and low transaction fees may even converge in this direction (Carlsten et al., 2016). Thus, the insignificance of other sources of electricity consumption besides PoW mining activities cannot be assumed for all PoW blockchains. Since the methods used to determine the electricity consumption of non-PoW blockchains are independent of the used consensus mechanism, these methods can also be applied to PoW blockchains.

A major limitation of current methods for determining the electricity consumption of blockchain full nodes is that they rely primarily on experimental measurements, which are then extrapolated to the entire network (e.g. CCRI, 2022d; de Vries, 2022; Platt et al., 2021). These quantifications are derived by setting up a small number of full nodes in a blockchain network and measuring their electricity consumption over a specific time period. The first source of uncertainty arises when the node-specific measurements are extrapolated over a longer time frame, such as an entire year, during which the computational load on the blockchain network and the corresponding electricity consumption may change systematically (e.g., due to increasing adoption and demand for transactions or due to changing computational complexity of transactions) (Rieger et al., 2022; Saingre et al., 2022). Secondly, electricity consumption is not uniform for all full nodes, as it strongly depends on the type of hardware. As a result, some studies use a sample of nodes with different hardware configurations (e.g., CCRI, 2022a). It is only feasible to test a limited number of hardware configurations, which means that many configurations likely used by some full nodes in the network remain untested. However, even the number of blockchain full nodes is often uncertain. One approach used in the literature on PoS blockchains is to use the number of validator nodes or participants who have staked some of their crypto-assets to obtain voting rights as an estimate for the number of full nodes. Yet, some non-validator participants may still operate a full node. Furthermore, in cases where validator nodes are limited to a specific threshold or amount of stake (e.g., 32 ETH per validator in Ethereum), multiple validators may be operated on the same machine or full node. Another approach derives the number of nodes and their hardware by crawling the node discovery protocol for P2P communication. This approach can also help identify the assessment of hosting characteristics of full nodes by means of the IP address of the nodes or the consensus client used (CCAF, 2024), which can also help narrow down electricity consumption estimates. On the other hand, this approach is challenged by virtual private networks and anonymization layers such as onion routing used by many blockchain nodes (Stoll et al., 2019). Consequently, the empirical distribution of hardware configurations among full nodes and even the number of full nodes itself will likely remain opaque, making precise estimations of the entire network's electricity consumption quite challenging. These uncertainties complicate accurate electricity consumption extrapolations and amplify uncertainty via the propagation of errors. There is a need for more empirical research that combines surveys, interviews, and technical and economic models to obtain more reliable and precise estimates. Such studies could quantify how factors like transaction throughput, the size and relevance of the blockchain network, the number of nodes, and the used hardware impact each other and overall electricity consumption.

Lastly, the influence of recent developments and advances in scaling solutions for blockchain technology on electricity consumption is only explored in narrow contexts, such as CCRI (2022b)'s analysis of a sidechain, or totally unknown, as it is the case for sharding, succinct blockchains, layer-2 solutions such as rollups (Rossi et al., 2022), or data availability layers. Both optimistic and validity ("zk-") rollups can reduce the computational effort for nodes' transaction validation and state updates per transaction, therefore increas-

ing the number and complexity of transactions that can be processed. On the other hand, they also add components that cause additional electricity consumption, such as external data availability layers or cryptographic provers (Principato et al., 2023; Sedlmeir et al., 2020b). Additionally, the emergence of the concept of light nodes is adding more nuances to blockchain nodes. In contrast to full nodes, light nodes only store and process a subset of all blockchain transactions and block information, which also should reduce their electricity consumption (Chatzigiannis et al., 2022). In general, blockchain scalability solutions influence the number or complexity of computations handled by individual nodes and lead to a wider palette of nodes and their respective roles, which in turn also strongly influences their electricity consumption and complicates its determination. The order of magnitude of electricity savings and the influence on the electricity consumption of the overall system is still unknown. To keep up with the emergence of these more complex blockchain solutions, the literature needs to develop an understanding of their influence on the blockchain network and its electricity consumption as a whole and design models and methods to quantify them.

Broader network participants

The broad spectrum of network participants, beyond the small subset of blockchain nodes and consensus participants, is largely unexplored in current literature (e.g., Platt et al., 2021). A more detailed discussion is found only in Lei et al. (2021), who propose considering blockchains as a set of subsystems, i.e., client devices, access networks, computing centers, distributed storage, validating nodes, and the core network, each contributing to the overall electricity consumption. However, they present this abstract model without data on the electricity consumption of these subsystems or guidance on how to determine the electricity consumption. Furthermore, their model takes the perspective of the technical components, e.g., a miner is part of the core network, has a validating node and distributed storage, and runs a computing center, while the other existing literature on blockchain electricity consumption was able to establish estimates has taken a more stakeholder-focused perspective. Obtaining an understanding of the different groups the broader network participants are comprised of could enable future research to determine the electricity consumption of the individual groups and the technical and economic factors driving it.

As outlined in Table 6, future research should start with the identification and precise definition of the different groups inside the broader network participants. Subsequent steps would be to determine the electricity consumption of the individual groups, including refinement of electricity consumption estimates through targeted methods or models and extending the scope of research on technical or economic aspects. One example requiring more sophisticated technical analysis is the development of models to approximate the number of (full) nodes in a blockchain network. An intuitive approach involves determining the number of blockchain addresses with a non-zero balance and the different IP addresses that participate in the P2P network. Furthermore, surveys on the devices or wallet software used by these participants can be used to gain insights into the average electricity consumption of participants. Building on this, future research could explore causal relationships. These relationships could be technical aspects, such as the influence of participant numbers and their devices, or economic aspects, e.g., the impact of the blockchain's application area or the extent of economic activity on the electricity consumption of different stakeholders in the network. This causal understanding can then help to find trade-offs and design choices to reduce the electricity consumption of the different participant groups.

Conclusion

This study presents a review of blockchain electricity consumption through an SLR. Our analysis reveals that previous studies have primarily distinguished between PoW and non-PoW blockchains. For PoW blockchains, the focus has been on the electricity consumption of miners, as they cause the largest share of overall consumption. The widely-accepted methods for determining miners' electricity consumption are based on observed technical parameters (hash rate and mining hardware distribution), the assumption of rational economic decision-making (where electricity expenses are limited by miners' revenues from mining rewards and transaction fees), and hybrid approaches combining these two methods. On the other hand, the electricity consumption of non-PoW blockchains is purely technical in nature and has been approximated using the electric energy consumed by full nodes. Some important parameters, such as PUE, have not been considered in all approaches. More generally, we observe that the prevailing narrow focus on a subset of

sources and an incomplete perspective on the incentives that drive mining activities are insufficient for a comprehensive understanding of electricity consumption in blockchain networks.

We categorize the various actors involved in blockchain ecosystems into three primary participants: network participants, blockchain nodes, and consensus participants. This classification consolidates insights from the separate research areas of PoW and non-PoW electricity consumption. We also identify gaps in the current literature arising from the neglect of the larger set of participants and their associated electricity consumption. A critical first step would be to establish a framework for this broader set of participants, defining its overall scope and the various subgroups within it. While the existing literature provides a solid understanding of nodes and consensus participants, the growing adoption of light nodes and scaling solutions (e.g., sidechains and rollups) challenge previous findings. Moreover, for PoW blockchains, we highlight influential factors developed or revealed by other research areas, such as blockchain scalability (Principato et al., 2023; Rossi et al., 2022) and the economics of mining (e.g., Stinner, 2022), that are not yet fully accounted for in research on blockchain electricity consumption. A more comprehensive understanding of the different roles, technical equipment, and economic decisions of these participants would provide deeper insights into industries and their revenue structures that can help refine electricity consumption estimates. In addition, this categorization could facilitate the identification of novel aspects, such as additional revenues from MEV or providing services to the grid, that have not yet been adequately explored.

We hope that our research stimulates further discourse on the sustainability of IT systems beyond blockchain within the IS community. Other emerging digital technologies, such as artificial intelligence and, in particular, large language models, have already raised similar controversial conversations about sustainability and initiated nascent research efforts to quantify the associated electricity consumption (García-Martín et al., 2019; Rillig et al., 2023), yet also suffer from a lack of economic perspectives to refine estimates. Our findings suggest that to accurately assess the electricity consumption of information systems, it is critical to consider various factors, including the diversity of stakeholders, the economic incentives influencing decision-making, and technological advances. Our research highlights the strong connection between these factors and the electricity consumption of blockchains and the need to integrate literature from computer science and economics more broadly. The interdisciplinary nature and rigorous methods of the IS discipline appear well-suited for meaningful contributions in this domain.

Acknowledgements

We thank Jens-Christian Stoetzer for his valuable feedback on the manuscript.

This research was funded in part by the Bavarian Ministry of Economic Affairs, Regional Development and Energy through the project “Fraunhofer Blockchain Center (20-3066-2-6-14)” and the Luxembourg National Research Fund (FNR) in the FiReSPaR-X (grant reference 14783405) and PABLO (grant reference 16326754) projects, as well as Banque et Caisse d’Épargne de l’État, Luxembourg (Spuerkeess) and the Directorate-General for Financial Stability, Financial Services and Capital Markets Union (DG FISMA). For the purpose of open access, and in fulfilment of the obligations arising from the grant agreements, the authors have applied a Creative Commons Attribution 4.0 International (CC BY 4.0) license to any Author Accepted Manuscript version arising from this submission.

References

- Álvarez, I. A., Gramlich, V., & Sedlmeir, J. (2024). Unsealing the secrets of blockchain consensus: A systematic comparison of the formal security of proof-of-work and proof-of-stake. *Proceedings of the 39th ACM/SIGAPP Symposium on Applied Computing*, 278–287. <https://doi.org/10.1145/3605098.3635970>
- Andoni, M., Robu, V., Flynn, D., Abram, S., Geach, D., Jenkins, D., McCallum, P., & Peacock, A. (2019). Blockchain technology in the energy sector: A systematic review of challenges and opportunities. *Renewable and Sustainable Energy Reviews*, 100, 143–174. <https://doi.org/10.1016/j.rser.2018.10.014>
- Babel, M., Gramlich, V., Körner, M.-F., Sedlmeir, J., Strüker, J., & Zwede, T. (2022). Enabling end-to-end digital carbon emission tracing with shielded NFTs. *Energy Informatics*, 5(1). <https://doi.org/10.1186/s42162-022-00199-3>

- Beck, R., Müller-Bloch, C., & King, J. L. (2018). Governance in the blockchain economy: A framework and research agenda. *Journal of the Association for Information Systems*, 19(10), 1020–1034. <https://doi.org/10.17705/1jais.00518>
- Bevand, M. (2018). Electricity consumption of Bitcoin: A market-based and technical analysis.
- Butijn, B.-J., Tamburri, D. A., & Heuvel, W.-J. v. d. (2020). Blockchains: A systematic multivocal literature review. *ACM Computing Surveys*, 53(3). <https://doi.org/10.1145/3369052>
- Capponi, A., Olafsson, S., & Alsabah, H. (2023). Proof-of-work cryptocurrencies: Does mining technology undermine decentralization? *Management Science*, 69(1), 6417–7150. <https://doi.org/10.1287/mnsc.2023.4840>
- Carlsten, M., Kalodner, H., Weinberg, S. M., & Narayanan, A. (2016). On the instability of Bitcoin without the block reward. *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security*, 154–167. <https://doi.org/10.1145/2976749.2978408>
- CCAF (2024). Cambridge Bitcoin Electricity Consumption Index. Cambridge Centre for Alternative Finance
- CCRI (2022a). Energy efficiency and carbon emissions of PoS networks.
- CCRI (2022b). Energy efficiency and carbon footprint of the Polygon blockchain.
- CCRI (2022c). Energy efficiency and carbon footprint of the TRON blockchain.
- CCRI (2022d). The Merge: Implications on the electricity consumption and carbon footprint of the Ethereum network.
- CCRI (2023a). CELO: Electricity consumption and carbon footprint of the CELO Network.
- CCRI (2023b). Polymesh: Electricity consumption and carbon footprint.
- CCRI (2023c). PoS benchmark study 2023: Energy efficiency and carbon footprint of PoS blockchain networks and platforms.
- Chatzigiannis, P., Baldimtsi, F., & Chalkias, K. (2022). SoK: Blockchain light clients. *Financial Cryptography and Data Security: 26th International Conference*, 615–641. https://doi.org/10.1007/978-3-031-18283-9_31
- Chitra, T., & Kulkarni, K. (2022). Improving proof of stake economic security via MEV redistribution. *Proceedings of the ACM CCS Workshop on Decentralized Finance and Security*. <https://doi.org/10.1145/3560832.3564259>
- Coinshare (2022). The Bitcoin mining network: Energy and carbon impact.
- Coroamă, V. C. (2022). Exploring the energy consumption of blockchains through an economic threshold approach. *11th International Conference on Energy Efficiency in Domestic Appliances and Lighting & 17th International Symposium on the Science and Technology of Lighting*. <https://ieeexplore.ieee.org/document/9932250>
- 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. *IEEE Symposium on Security and Privacy*, 910–927. <https://doi.org/10.1109/SP40000.2020.00040>
- de Vries, A. (2018). Bitcoin’s growing energy problem. *Joule*, 2(5), 801–805. <https://doi.org/10.1016/j.joule.2018.04.016>
- de Vries, A. (2020). Bitcoin’s energy consumption is underestimated: A market dynamics approach. *Energy Research & Social Science*, 70, 101721. <https://doi.org/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. <https://doi.org/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. <https://doi.org/10.1016/j.patter.2022.100633>
- Digiconomist (2024). Bitcoin energy consumption index.
- Dittmar, L., & Praktiknjo, A. (2019). Could Bitcoin emissions push global warming above 2°C? *Nature Climate Change*, 9(9), 656–657. <https://doi.org/10.1038/s41558-018-0321-8>
- Ethereum (2023). Ethereum archive node.
- Fisch, B., Bonneau, J., Greco, N., & Benet, J. (2018). Scaling proof-of-replication for Filecoin mining.
- DG-FISMA (2023). Developing a methodology and sustainability standards for mitigating the environmental impact of crypto-assets.
- Fridgen, G., Körner, M., Walters, S., & Weibelzahl, M. (2021). Not all doom and gloom: How energy-intensive and temporally flexible data center applications may actually promote renewable energy sources. *Business Information Systems Engineering*, 63, 243–256. <https://doi.org/10.1007/s12599-021-00686-z>

- Gallersdörfer, U., Klaaßen, L., & Stoll, C. (2020). Energy consumption of cryptocurrencies beyond Bitcoin. *Joule*, 4(9), 1843–1846. <https://doi.org/10.1016/j.joule.2020.07.013>
- García-Martín, E., Rodrigues, C. F., Riley, G., & Grahm, H. (2019). Estimation of energy consumption in machine learning. *Journal of Parallel and Distributed Computing*, 134, 75–88. <https://doi.org/10.1016/j.jpdc.2019.07.007>
- Garratt, R. J., & Oordt, M. R. van (2023). Why fixed costs matter for proof-of-work-based cryptocurrencies. *Management Science*, 69(11), 6482–6507. <https://doi.org/10.1287/mnsc.2023.4901>
- Gervais, A., Karame, G. O., Wüst, K., Glykantzis, V., Ritzdorf, H., & Capkun, S. (2016). On the security and performance of proof of work blockchains. *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security*, 3–16. <https://doi.org/10.1145/2976749.2978341>
- Gonzalez-Barahona, J. M. (2021). Factors determining maximum energy consumption of Bitcoin miners.
- Gramlich, V., Guggenberger, T., Principato, M., Schellinger, B., & Urbach, N. (2023). A multivocal literature review of decentralized finance. *Electronic Markets*, 33. <https://doi.org/10.1007/s12525-023-00637-4>
- Hayes, A. (2015). A cost of production model for Bitcoin.
- Houy, N. (2019). Rational mining limits Bitcoin emissions. *Nature Climate Change*, 9(9), 655–655. <https://doi.org/10.1038/s41558-019-0533-6>
- Jones, B. A., Goodkind, A. L., & Berrens, R. P. (2022). Economic estimation of Bitcoin mining’s climate damages demonstrates closer resemblance to digital crude than digital gold. *Scientific reports*, 12(1). <https://doi.org/10.1038/s41598-022-18686-8>
- Kitchenham, B., Brereton, O. P., 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. <https://doi.org/10.1016/j.infsof.2008.09.009>
- Krause, M. J., & Tolaymat, T. (2018). Quantification of energy and carbon costs for mining cryptocurrencies. *Nature Sustainability*, 1(11), 711–718. <https://doi.org/10.1038/s41893-018-0152-7>
- Kroll, J. A., Davey, I. C., & Felten, E. W. (2013). The economics of Bitcoin mining, or Bitcoin in the presence of adversaries. *Proceedings of the 12th Workshop on the Economics of Information Security*. <https://w.econinfosec.org/archive/weis2013/papers/KrollDaveyFeltenWEIS2013.pdf>
- Küfeoğlu, S., & Özkuran, M. (2019). Bitcoin mining: A global review of energy and power demand. *Energy Research & Social Science*, 58, 101273. <https://doi.org/10.1016/j.erss.2019.101273>
- 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. <https://doi.org/10.1016/j.enpol.2021.112422>
- Li, J., Li, N., Peng, J., Cui, H., & Wu, Z. (2019). Energy consumption of cryptocurrency mining: A study of electricity consumption in mining cryptocurrencies. *Energy*, 168, 160–168. <https://doi.org/10.1016/j.energy.2018.11.046>
- Masanet, E., Shehabi, A., Lei, N., Vranken, H., Koomey, J., & Malmodin, J. (2019). Implausible projections overestimate near-term Bitcoin CO₂ emissions. *Nature Climate Change*, 9(9), 653–654. <https://doi.org/10.1038/s41558-019-0535-4>
- Mora, C., Rollins, R. L., Taladay, K., Kantar, M. B., Chock, M. K., Shimada, M., & Franklin, E. C. (2018). Bitcoin emissions alone could push global warming above 2°C. *Nature Climate Change*, 8(11), 931–933. <https://doi.org/10.1038/s41558-018-0321-8>
- Müller-Bloch, C., & Kranz, J. (2015). A framework for rigorously identifying research gaps in qualitative literature reviews. *Proceedings of the 36th International Conference on Information Systems*. <https://aisel.aisnet.org/icis2015/proceedings/ResearchMethods/2/>
- Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system.
- Niaz, H., Liu, J. J., & You, F. (2022). Can Texas mitigate wind and solar curtailments by leveraging Bitcoin mining? *Journal of Cleaner Production*, 364, 132700. <https://doi.org/10.1016/j.jclepro.2022.132700>
- O’Dwyer, K. J., & Malone, D. (2014). Bitcoin mining and its energy footprint. *25th IET Irish Signals & Systems Conference and China-Ireland International Conference on Information and Communities Technologies*, 280–285. <https://doi.org/10.1049/cp.2014.0699>
- Platt, M., Sedlmeir, J., Platt, D., Xu, J., Tasca, P., Vadgama, N., & Ibanez, J. I. (2021). The energy footprint of blockchain consensus mechanisms beyond proof-of-work. *21st International Conference on Software Quality, Reliability and Security Companion*, 1135–1144. <https://doi.org/10.1109/QRS-C55045.2021.00168>

- Principato, M., Babel, M., Guggenberger, T., Kropp, J., & Mertel, S. (2023). Towards solving the blockchain trilemma: An exploration of zero-knowledge proofs. *Proceedings of the 43th International Conference on Information Systems*. <https://aisel.aisnet.org/icis2023/blockchain/blockchain/5/>
- Qin, K., Zhou, L., & Gervais, A. (2022). Quantifying blockchain extractable value: How dark is the forest? *Symposium on Security and Privacy*, 198–214. <https://doi.org/10.1109/SP46214.2022.9833734>
- Qin, S., Klaaßen, L., Gallersdörfer, U., Stoll, C., & Zhang, D. (2021). Bitcoin's future carbon footprint. *Applied Energy Symposium: Low carbon cities and urban energy system*. <https://doi.org/10.46855/energy-proceedings-8232>
- Rieger, A., Roth, T., Sedlmeir, J., & Fridgen, G. (2022). We need a broader debate on the sustainability of blockchain. *Joule*, 6(6), 1137–1141. <https://doi.org/10.1016/j.joule.2022.04.013>
- Rillig, M. C., Ågerstrand, M., Bi, M., Gould, K. A., & Sauerland, U. (2023). Risks and benefits of large language models for the environment. *Environmental Science & Technology*, 57(9), 3464–3466. <https://doi.org/10.1021/acs.est.3c01106>
- Roma, C. A., & Hasan, M. A. (2020). Energy consumption analysis of XRP validator. *International Conference on Blockchain and Cryptocurrency*. IEEE. <https://doi.org/10.1109/ICBC48266.2020.9169427>
- Rossi, L. M. de, Avital, M., Gleasure, R., Solcia, R., et al. (2022). The impact of layer 2 technologies on the adoption and security of blockchain. *Proceedings of the 43rd International Conference on Information Systems*. <https://aisel.aisnet.org/icis2022/blockchain/blockchain/6/>
- Sai, A. R., & Vranken, H. (2023). Promoting rigour in blockchains energy & environmental footprint research: A systematic literature review. *Blockchain: Research and Applications*, 5(1). <https://doi.org/10.1016/j.bcr.2023.100169>
- Saingre, D., Ledoux, T., & Menaud, J.-M. (2022). Measuring performances and footprint of blockchains with BCTMark: a case study on Ethereum smart contracts energy consumption. *Cluster Computing*, 25(4), 2819–2837. <https://doi.org/10.1007/s10586-021-03441-x>
- Saleh, F. (2021). Blockchain without waste: Proof-of-stake. *The Review of Financial Studies*, 34(3), 1156–1190. <https://doi.org/10.1093/rfs/hhaa075>
- Schmidt, J. (2021). Bitcoin's energy usage, explained. *Forbes*.
- Sedlmeir, J., Buhl, H. U., Fridgen, G., & Keller, R. (2020a). Recent developments in blockchain technology and their impact on energy consumption. *Informatik Spektrum*, 43(6), 391–404. <https://doi.org/10.1007/s00287-020-01321-z>
- Sedlmeir, J., Buhl, H. U., Fridgen, G., & Keller, R. (2020b). The energy consumption of blockchain technology: Beyond myth. *Business & Information Systems Engineering*, 62(6), 599–608. <https://doi.org/10.1007/s12599-020-00656-x>
- Shi, X., Xiao, H., Liu, W., Lackner, K. S., Buterin, V., & Stocker, T. F. (2023). Confronting the carbon-footprint challenge of blockchain. *Environmental Science & Technology*, 57(3). <https://doi.org/10.1021/acs.est.2c05165>
- Song, Y.-D., & Aste, T. (2020). The cost of Bitcoin mining has never really increased. *Frontiers in Blockchain*, 3. <https://doi.org/10.3389/fbloc.2020.565497>
- Stinner, J. (2022). On the economics of Bitcoin mining: A theoretical framework and simulation evidence. *Proceedings of the 43rd International Conference on Information Systems*. <https://aisel.aisnet.org/icis2022/blockchain/blockchain/7/>
- Stoll, C., Klaaßen, L., & Gallersdörfer, U. (2019). The carbon footprint of Bitcoin. *Joule*, 3(7), 1647–1661. <https://doi.org/10.1016/j.joule.2019.05.012>
- Vranken, H. (2017). Sustainability of Bitcoin and blockchains. *Current Opinion in Environmental Sustainability*, 28. <https://doi.org/10.1016/j.cosust.2017.04.011>
- Webster, J., & Watson, R. T. (2002). Analyzing the past to prepare for the future: Writing a literature review. *MIS Quarterly*, 26(3), 13–23. <https://www.jstor.org/stable/4132319>
- Wendl, M., Doan, M. H., & Sassen, R. (2023). The environmental impact of cryptocurrencies using proof of work and proof of stake consensus algorithms: A systematic review. *Journal of Environmental Management*, 326, 116530. <https://doi.org/10.1016/j.jenvman.2022.116530>
- Zade, M., Myklebost, J., Tzscheuschler, P., & Wagner, U. (2019). Is Bitcoin the only problem? A scenario model for the power demand of blockchains. *Frontiers in Energy Research*, 7. <https://doi.org/10.3389/fenrg.2019.00021>