The Energy Footprint of Blockchain Consensus Mechanisms Beyond Proof-of-Work

Workshop Paper

Moritz Platt, Johannes Sedlmeir, Daniel Platt, Jiahua Xu, Paolo Tasca, Nikhil Vadgama and Juan Ignacio Ibañez

6 December 2021

Presentation at the 21st IEEE Conference on Software Quality, Reliability, and Security

Summary

Section

Goal

Popular permissionless distributed ledger technology (DLT) systems using proof-of-work (PoW) for Sybil attack resistance have extreme energy requirements, drawing stern criticism from academia, business and the media. DLT systems building on alternative consensus mechanisms, particularly proof-of-stake (PoS), aim to address this downside. In this paper, we take an initial step towards comparing the energy requirements of such systems to understand whether they achieve this goal equally well.

Identity Inflation

Sybil attacks, which pose a critical problem for DLT systems, occur when an attacker creates an artificially large number of bogus identities [1] to skew the results of majority decisions on the admission and order of transactions.

Permissioned vs. Permissionless Systems

In *permissioned* networks, gatekeeping strategies can be applied that limit access to a network to previously vetted actors [2], thereby preventing such attacks.

However, for *permissionless* networks, in which participants can partake in consensus without any control [3], more complex mechanisms need to be applied to combat Sybil attacks.

Introduction

Proof-of-Work

Scarce Resources

Consensus mechanisms for *permissionless* networks entail aligning entitlement to participate in consensus proportionally with the possession or expenditure of resources that can be digitally verified [4]. Proof-of-work (PoW) is an example of a Sybil attack resistance scheme that has been used in most early cryptocurrencies such as Bitcoin [5]. To counteract Sybil attacks, PoW uses cryptographic puzzles of configurable difficulty with efficient verification so that it becomes computationally expensive for attackers to interfere with consensus [6].

Energy Consumption of Proof-of-Work

However, the energy consumption of PoW-based cryptocurrencies is connected to their respective market capitalisations, leading to extreme energy demand for popular implementations [7]. For instance, the electricity demand of Bitcoin is now in the same range as that of entire industrialised nations [8] and has been positioned as a dangerous contributor to global warming, producing up to 22.90 Mt CO_2 [9].

Introduction

Proof-of-Stake

Using Stake to Prevent Sybil Attacks

In proof-of-stake (PoS), participants with larger holdings of a crypto-currency have a greater influence in transaction validation. While PoS is generally understood as being more energy-efficient than PoW, the exact energy consumption characteristics of PoS-based systems, and the influence that network throughput has on them, are not widely understood.

Introduction

Approach

Previous Approach

Two main approaches to quantifying the energy consumption of a DLT system have been used in the past. One is to measure the consumption of a representative participant node and then extrapolate from this measurement. An alternative approach is to develop a mathematical model that includes the core metrics of a DLT system to calculate its energy consumption.

So far, most work has focused on PoW blockchains[1], and some research has investigated individual non-PoW systems.

[1] For the purpose of this manuscript, the term 'Blockchain' refers to any type of DLT, even if it does not make use of the 'block' concept, first described by Nakamoto [5].

Our Approach

We propose a simple energy consumption model, applicable to a broad range of DLT systems that use PoS for Sybil attack resistance. Specifically, this model considers the number of validator nodes, their energy consumption, and the network throughput, based on which the energy consumption per transaction is estimated.

Section

Related Work

We conducted an informal literature review using the 'Bielefeld Academic Search Engine'. We thereby obtained 413 results of prior studies analysing the energy demand of different DLT systems, with a significant focus on PoW blockchains in general, and specifically on Bitcoin Commonly, models take one of the following two forms: **Experimental models** [10]–[13] and **mathematical models** [7], [14]–[21].

Section

Blockchain Systems

Platform	Permissioned	Permissionless
Ethereum 2.0		•
Algorand		•
Cardano		•
Polkadot		•
Tezos		•
Hedera	•	

Table 1: Comparison of the analysed DLT systems in node permission setting.

Solution

Comparing archetypal permissioned and permissionless systems allows us to understand patterns.

Apples and Oranges?

Commonalities between the Protocols

1. Participants can act as validators. (*In permissioned networks, the set of participants that can act as validators is limited*)
2. To act as a validator, a participant needs to be operating a computer that can send and receive data across the Internet
3. Must be able to perform the computations required to establish the correctness of proposed transactions
4. Operating such a validator node is opt-in (*Within limits of being permissioned*)
5. Validator nodes need to remain 'always on'

Section

Improvements over Previous Work

Our model differs from previous work in that we also consider energy consumption *per transaction*, as opposed to only the overall energy consumption of an entire DLT system.

We use existing models combined with additional data arising from the scientific literature, reports, and public ledger information to form a baseline that can be used to avoid time-consuming experimental validation.

The Model of Powell *et al.*

Powell *et al.* [21] define an elementary mathematical model for the energy consumption of the Polkadot blockchain that can be generalised as:

$$p_t = p \cdot n_{\mathrm{val}}, \tag{1}$$

where p_t is the overall average power the DLT system consumes, p is the average power consumed by a validator node, and n_{val} is the number of validator nodes.

Due to the low computational effort associated with PoS and the low throughput of permissionless blockchains it is assumed that validating nodes run on similar types of commodity server hardware, irrespective of the network load.

Energy consumption per Validator Node

Since it is nearly impossible to determine which type of hardware is used by validators, we use an approximation derived from industry recommendations. Dramatically different hardware recommendations are put forward for permissionless systems and permissioned systems.

Configuration	Hardware Type	Demand (W)
Minimum	Small single-board computer	5.5
Medium	General purpose server	168.1
Maximum	High-performance server	328

Table 2: Conceivable upper and lower bounds for the power demand of a validator machine.

Number of Validator nodes

Platform	# Validators	TPS Cont. (tx/s)	TPS Max. (tx/s)
Ethereum 2.0	2649[☆]	15.40[☆]	3000
Algorand	1126	9.85	1000
Cardano	8874	0.36	257
Polkadot	297	0.12	1000
Tezos	399	1.70	40
Hedera	21	48.20	10 000

[☆] Ethereum Mainnet measurements used as approximation

Table 3: The current number of validators, contemporary throughput, and the upper bound of throughput postulated.

Energy consumption per transaction

To arrive at an energy consumption per transaction metric (c_{tx}), the number of transactions per unit of time needs to be considered. The actual numbers are dynamic and fluctuate over time. The contemporary network throughput (*Cont.*) is defined as the actual throughput recently experienced by a system.

Treating the average power consumed by a validator node (p, measured in W) as a constant means that an inverse relationship between consumption per transaction (c_{tx}) and system throughput (l) can be established within the bounds of $(0, l_{\max}]$:

$$f_{c_{\mathrm{tx}}}(l) = \frac{n_{\mathrm{val}} \cdot p}{l}. \tag{2}$$

Modelling c_{tx} as a function of the number of transactions per second

Can we develop a model for c_{tx} that depends on one variable, namely l, only? This is plausible as the total number of users in a **permissionless** system increases, of the new users, a share becomes validators and another non-disjoint share executes transactions. This suggests that n_{val} and l are positively correlated.

Modelling c_{tx} as a function of the number of transactions per second

Equation (2) depends on two variables: n_{val} and l.

Data from the Cardano blockchain[2] suggest that the number of validators n_{val} and the number of transactions per second l are positively correlated. Namely, Pearson's correlation coefficient[3] for n_{val} and l for 375 data points from 29 July 2020 to 7 August 2021 is 0.80. The correlation coefficient for n_{val} delayed by 28 days and l (not delayed) for the same data is 0.87.

We will now present a model for c_{tx} that depends on one variable, namely l, only.

[2] https://data.mendeley.com/datasets/4jv2wmwrc5/1

[3] The correlation coefficient takes values in $[-1, 1]$ and a value of ± 1 would imply that n_{val} is an affine function in l.

Modelling c_{tx} as a function of the number of transactions per second

For simplicity we assume that the correlation is perfect, i.e., $n_{val} = \kappa + \lambda \cdot l$ for some $\kappa, \lambda \in \mathbb{R}, \lambda > 0$, and using (2) we obtain

$$f_{c_{\mathrm{tx}}}(l) = \frac{(\kappa + \lambda l) \cdot p}{l}. \tag{3}$$

Modelling c_{tx} as a function of the number of transactions per second

For Algorand, Polkadot, Tezos, and Hedera, we compute κ, λ based on two data points. For Cardano, we use linear regression implemented as ordinary least squares regression to compute κ, λ that have the maximum likelihood of modelling $f_{c_{\mathrm{tx}}}(l)$ under the assumption that $f_{c_{\mathrm{tx}}}(l)$ is an affine function with Gaussian noise:

Platform	κ	λ
Algorand	102.8	103.9
Cardano	3803.4	8877.6
Polkadot	297	0
Tezos	440.7	−24.6
Hedera	7.6	0.3

Section

Energy Consumption per Transaction

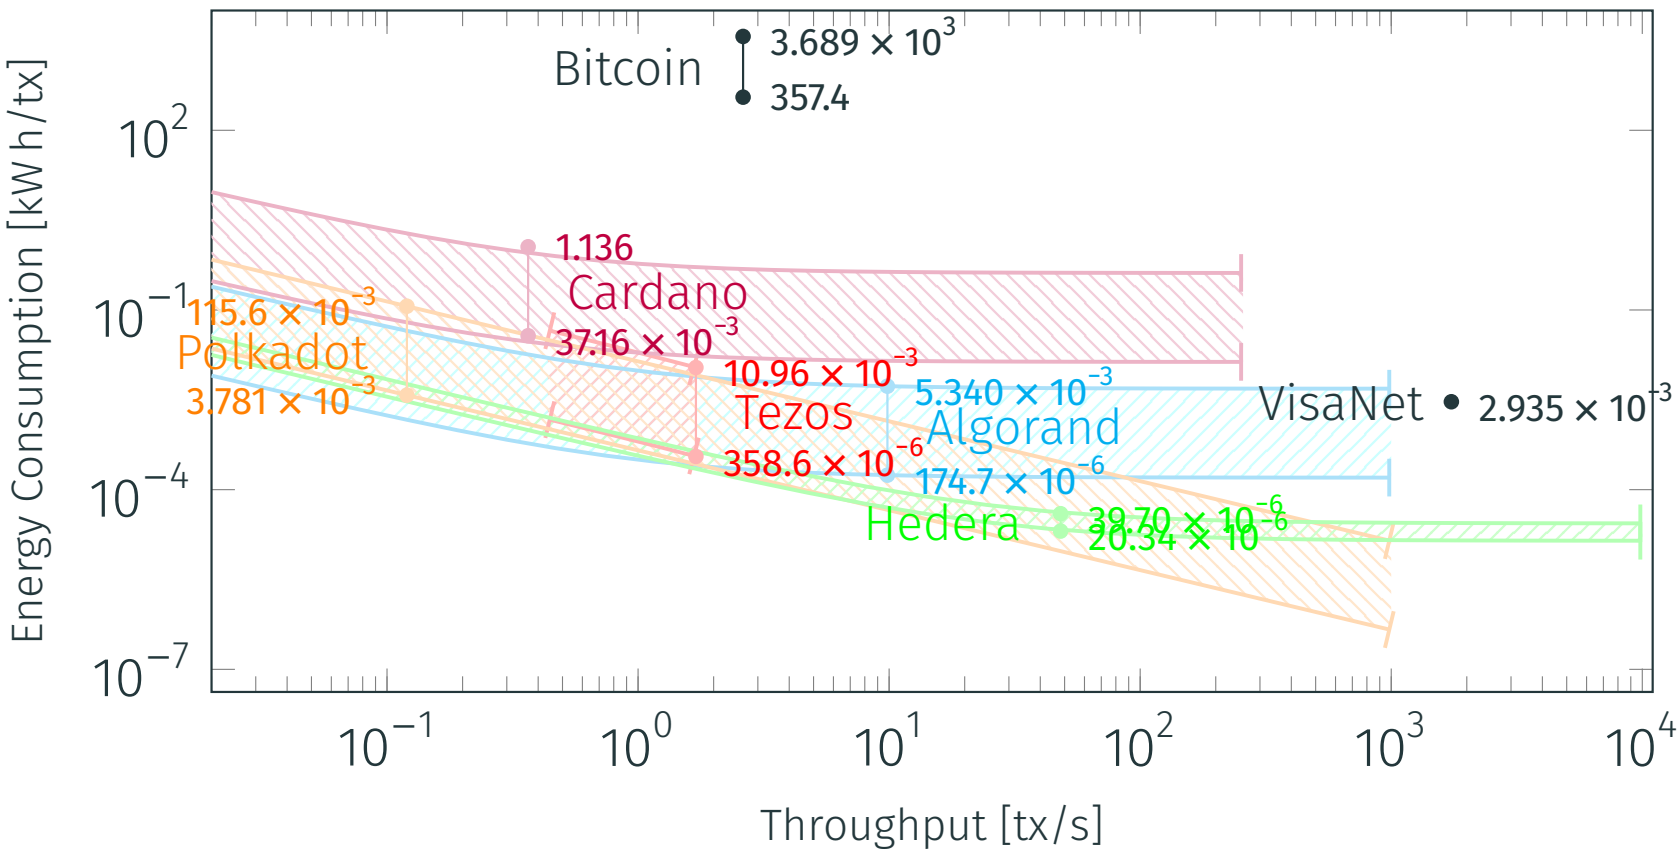

Energy Consumption per Transaction

Platform	Global (kW)			Per transaction (kW h/tx)		
Eth2[☆]	14.6	–	445.3	0.000 26	–	0.008 03
Algorand	6.2	–	189.3	0.000 17	–	0.005 34
Cardano	48.8	–	1491.7	0.037 16	–	1.135 62
Polkadot	1.6	–	49.9	0.003 78	–	0.115 56
Tezos	2.2	–	67.1	0.000 36	–	0.010 96
Hedera	3.5	–	6.9	0.000 02	–	0.000 04
Bitcoin	3 373 287.7	–	34 817 351.6	360.393 00	–	3691.407 00
VisaNet			22 387.1			0.003 58

[☆] Ethereum Mainnet measurements used as approximation

Table 4: Global power consumption ranges

Section

Conclusion

1. The energy footprint of PoW is significant: Bitcoin's energy consumption exceeds the energy consumption of all PoS-based systems analysed by at least three orders of magnitude.
2. There are significant differences in energy consumption among the PoS-based systems analysed, with permissionless systems having a larger energy footprint overall owing to their higher replication factor.
3. The type of hardware that validators use has a considerable impact on whether the energy consumption of PoS blockchains is comparable with or considerably larger than that of centralised systems.

Conclusion

The results should not be misinterpreted as an argument for increased centralisation or for permissioned networks over permissionless ones. Permissioned networks pose a risk of centralisation, which may offer minuscule advantages in terms of energy consumption but may negate the functional advantages of blockchain.

Implications

1. Urgent call for the modernisation of PoW systems and a shift towards PoS
2. A recommendation to practitioners to consider energy-saving hardware which aligns with minimal supported configurations

References i

J. R. Douceur, 'The Sybil attack,' in *Proceedings of the 1st International Workshop on Peer-to-Peer Systems*, P. Druschel, F. Kaashoek and A. Rowstron, Eds., ser. Lecture Notes in Computer Science, vol. 2429, Cambridge, MA, USA: Springer, 2002, pp. 251–260.

M. Platt, R. J. Bandara, A.-E. Drăgnoiu and S. Krishnamoorthy, 'Information privacy in decentralized applications,' in *Trust Models for Next-Generation Blockchain Ecosystems*, ser. EAI/Springer Innovations in Communication and Computing, M. Rehman, D. Svetinovic, K. Salah and E. Damiani, Eds., Springer, 2021, pp. 85–104.

References ii

P. Tasca and C. J. Tessone, 'A taxonomy of blockchain technologies: Principles of identification and classification,' *Ledger*, vol. 4, Feb. 2019.

Y. Xiao, N. Zhang, W. Lou and Y. T. Hou, 'A survey of distributed consensus protocols for blockchain networks,' *IEEE Communications Surveys & Tutorials*, vol. 22, no. 2, pp. 1432–1465, 2020.

S. Nakamoto. (2008), 'Bitcoin: A peer-to-peer electronic cash system,' [Online]. Available: *https://bitcoin.org/bitcoin.pdf* (visited on 22/07/2021).

A. Back. (Mar. 1997), 'A partial hash collision based postage scheme,' [Online]. Available: *http://www.hashcash.org/papers/announce.txt* (visited on 27/07/2021).

References iii

J. Sedlmeir, H. U. Buhl, G. Fridgen and R. Keller, 'The energy consumption of blockchain technology: Beyond myth,' *Business & Information Systems Engineering*, vol. 62, no. 6, pp. 599–608, Jun. 2020.

A. de Vries, 'Bitcoin's growing energy problem,' *Joule*, vol. 2, no. 5, pp. 801–805, May 2018.

C. Stoll, L. Klaaßen and U. Gallersdörfer, 'The carbon footprint of Bitcoin,' *Joule*, vol. 3, no. 7, pp. 1647–1661, Jul. 2019.

A. Igumenov, E. Filatovas and R. Paulavičius, 'Experimental investigation of energy consumption for cryptocurrency mining,' in *Proceedings of the 11th International Workshop on Data Analysis Methods for Software Systems*, J. Bernatavičienė, Ed., Druskininkai, Lithuania: Vilnius University Press, Nov. 2019, p. 31.

References iv

D. Saingre, T. Ledoux and J.-M. Menaud, 'BCTMark: A framework for benchmarking blockchain technologies,' in *Proceedings of the 17th International Conference on Computer Systems and Applications*, Antalya, Turkey: IEEE, Nov. 2020, pp. 1–8.

C. A. Roma and M. A. Hasan, 'Energy consumption analysis of XRP validator,' in *Proceedings of the 2020 International Conference on Blockchain and Cryptocurrency*, IEEE, May 2020, pp. 1–3.

R. Cole and L. Cheng, 'Modeling the energy consumption of blockchain consensus algorithms,' in *Proceedings of the 2018 International Conference on Internet of Things and Green Computing and Communications and Cyber, Physical and Social Computing and Smart Data*, Halifax, NS, Canada: IEEE, Jul. 2018, pp. 1691–1696.

References v

- N. Lei, E. Masanet and J. Koomey, 'Best practices for analyzing the direct energy use of blockchain technology systems: Review and policy recommendations,' *Energy Policy*, vol. 156, 2021.
- U. Gallersdörfer, L. Klaaßen and C. Stoll, 'Energy consumption of cryptocurrencies beyond Bitcoin,' *Joule*, vol. 4, no. 9, pp. 1843–1846, Sep. 2020.
- S. Küfeoglu and M. Özkuran, 'Energy consumption of Bitcoin mining,' University of Cambridge, Cambridge Working Paper in Economics 1948, 2019.
- M. Zade, J. Myklebost, P. Tzscheutschler and U. Wagner, 'Is bitcoin the only problem? A scenario model for the power demand of blockchains,' *Frontiers in Energy Research*, vol. 7, p. 21, Mar. 2019.

References vi

J. Sedlmeir, H. U. Buhl, G. Fridgen and R. Keller, 'Ein Blick auf aktuelle Entwicklungen bei Blockchains und deren Auswirkungen auf den Energieverbrauch,' German, *Informatik Spektrum*, vol. 43, no. 6, pp. 391–404, Nov. 2020.

H. Vranken, 'Sustainability of Bitcoin and blockchains,' *Current Opinion in Environmental Sustainability*, vol. 28, pp. 1–9, Oct. 2017.

References vii

G. Eshani, D. Rajdeep, R. Shubhankar and D. Baisakhi, 'An analysis of energy consumption of blockchain mining and techniques to overcome it,' in *Proceedings of the International Conference on Computational Intelligence, Data Science and Cloud Computing*, V. E. Balas, A. E. Hassanien, S. Chakrabarti and L. Mandal, Eds., ser. Lecture Notes on Data Engineering and Communications Technologies, vol. 62, Kolkata, India: Springer, 2021, pp. 783–792.

L. M. Powell, M. Hendon, A. Mangle and H. Wimmer, 'Awareness of blockchain usage, structure, & generation of platform's energy consumption: Working towards a greener blockchain,' *Issues In Information Systems*, vol. 22, no. 1, pp. 114–123, 2021.