



Energy Efficiency and Carbon Footprint of PoS Blockchain Protocols

January 2022

© Crypto Carbon Ratings Institute, 2022

carbon-ratings.com

The Crypto Carbon Ratings Institute (CCRI) is a brand of CCRI GmbH based in Dingolfing, Germany.

Title photo by [Terry](#) on [Unsplash](#).

Energy Efficiency and Carbon Footprint of Proof of Stake Blockchain Protocols

January 2022

Ulrich Gellersdörfer, Lena Klaaßen, Christian Stoll
Crypto Carbon Ratings Institute
carbon-ratings.com

Preamble

This report is prepared by the Crypto Carbon Ratings Institute (CCRI) for
Avalanche Inc.

Executive summary

- The electricity consumption and carbon footprint of Proof of Work (PoW)-based cryptocurrencies such as Bitcoin remain significant.
- Blockchain networks based on alternative consensus mechanisms such as Proof of Stake (PoS) consume significantly less electricity.
- For the 6 PoS networks assessed in this study, we find a variance in total yearly electricity consumption of 28x – ranging from 70 MWh (Polkadot) to 1,967 MWh (Solana).
- Comparing electricity consumption per node or transaction shows different rankings of the 6 PoS networks.
- Electricity consumption can be translated into carbon emissions via emission factors of electricity generation to gauge climate impacts.
- Continuous development and evolution of networks require regular updates of measurement and analyses.

1 Introduction

The electricity consumption and related carbon footprint of Bitcoin and other cryptocurrencies are subject to extensive discussion in public, academia, and industry. For these protocols, various estimations exist, comparing Bitcoin's electricity consumption to countries like Norway (de Vries, 2021; Stoll, Klaaßen, & Gellersdörfer, 2019). The problem has been known for several years, and other systems and technologies have emerged to solve the issue. The consensus family of Proof of Stake (PoS) is deemed superior regarding the electricity requirements compared to the traditional Proof of Work (PoW) consensus mechanisms (King & Nadal, 2012). While it is consensus in the broader scientific community that PoS could solve the electricity issues of PoW, it is unclear how these PoS systems compare to each other.

Instead of requiring computational power to solve mining puzzles for securing the network in PoW, PoS requires validators to lock in funds for a specific period of time to propose or vote on new blocks. Due to the nature of the software engineering process and network architectures, different PoS systems rely on varying fundamentals regarding the hardware requirements, programming language, network size, transaction throughput, transaction complexity, and more. These factors influence the electricity consumption and, therefore, the carbon footprint of a respective network. While it is expected that the overall differences between PoS networks are minor, it is nonetheless essential to understand the absolute and relative energy efficiency of single networks (Gellersdörfer, Klaaßen, & Stoll, 2020).

In this report, we provide an analysis of the electricity consumption, carbon footprint, and influencing factors of six major Proof of Stake-based cryptocurrencies. We rank these currencies in terms of electricity consumption, carbon footprint, and other factors such as transaction throughput and energy efficiency. Table 1 summarizes all results.

Table 1: Overview of results. Green cells mark the best value for the respective column.

	Nodes [# total]	Transactions [Tx/year] ^a	Total electricity consumption [kWh/year]	Electricity per node [kWh/year]	Electricity per transaction [Wh/Tx]	Total carbon emissions [tCO ₂ e/year]
Cardano	3,002	11.9 mn	598,755	199.45	51.59	284.41
Polkadot	297	4.0 mn	70,237	236.49	17.42	33.36
Solana	1,015	11.8 bn	1,967,930	1,938.85	0.166	934.77
Tezos	375	2.5 mn	113,249	250.99	41.45	53.79
Avalanche ^b	1,084	93.9 mn	489,311	451.39	4.76	232.42
Algorand	1,190	190.0 mn	512,671	430.82	2.70	243.52

This report is outlined as follows: In chapter 2, we define the aim and scope and the selected cryptocurrencies to understand the goals and limitations of this report. We describe the methodology of how this study is conducted in Chapter 3. Chapter 4 outlines the selected hardware and the infrastructure established for electricity measurements. Chapter 5 provides the calculation of the electricity consumption and carbon footprint of the respective networks. In Chapter 6, we analyze the electricity consumptions of the networks in the context of other metrics. Chapter 7 concludes with a summary and an outlook on the future of the environmental impact of PoS-based networks.

^a We assume the amount of transactions occurred during our measurement for a daily basis.

^b These numbers represent the numbers for Avalanche network, including a major update in the node software implemented in September/October 2021. The initial measurement performed prior to the update in August 2021 yielded the following results: Total energy consumption [kWh/year]: 144,384 | Electricity per node [kWh/year]: 139.77 | Electricity per transaction [Wh/Tx]: 14.27 | Total carbon emissions [tCO₂e/year]: 68.58

2 Aim and scope

This report aims to provide insights into the electricity consumption and carbon footprint of the current state of PoS networks. To do so, we present a methodology which we describe in Chapter 3.

It is noteworthy that the methodology is a helpful tool to derive a ballpark estimate for total electricity consumption and carbon emissions as well as the relative performance. However, the networks are associated with uncertainties that impede deriving exact numbers of the electricity consumption or, respectively, of the network's carbon footprint. Numerous factors, such as the network size, varying hardware configuration, or network infrastructure, influence the overall electricity consumption. Nonetheless, we deem this report to produce the most precise electricity consumption and carbon footprint estimates for these cryptocurrencies to date, as we observe and measure the electricity consumption of single hardware components and use them as a proxy for the overall network.

The establishment of representative hardware, network sizes, and electricity measurements form the basis for future research, such as comparing different networks and their respective requirements and properties. Identical assumptions (e.g., selected hardware) for networks and adjustments for their requirements (e.g., the hardware requirements for some networks are higher than others) allow to build a valid data set for the comparison of these networks.

For our analysis, we select all Proof-of-Stake networks in the Top 40 of cryptocurrency market capitalization according to coinmarketcap.com on 7th of May 2021^c. Table 2: Overview of PoS-based cryptocurrencies by market cap and rank among all cryptocurrencies displays the respective currency with their rank, name, ticker symbol and the market capitalization. Their combined market capitalization is about \$ 117 bn USD, a share of about 4.9 % of the overall market capitalization of all cryptocurrencies.

Table 2: Overview of PoS-based cryptocurrencies by market cap and rank among all cryptocurrencies as of the 7th of May 2021

Rank	Name	Symbol	Market Capitalization
7	Cardano	ADA	\$ 52.8 bn
8	Polkadot	DOT	\$ 37.4 bn
17	Solana	SOL	\$ 11.8 bn
33	Tezos	XTZ	\$ 5.5 bn
37	Avalanche	AVAX	\$ 4.9 bn
38	Algorand	ALGO	\$ 4.7 bn
			\$ 117.1 bn

^c <https://coinmarketcap.com/historical/20210507/>

3 Methodology

Our methodology builds upon four steps to generate data on the electricity consumption and carbon footprint of PoS-systems and additional data to develop metrics for a valid comparison between the systems.

In the first step, we analyze the selected PoS-systems and their minimum hardware requirements. The hardware requirements are an indicator of the hardware composition of the network. We use this information and additional hardware data from PassMark to select and obtain hardware that we use to measure a single node's electricity consumption.

In the second step, we estimate the electricity usage of a single node and provide upper and lower bounds for the networks. We start by running the software on all obtained hardware devices and measure their single electricity consumption while running the network and while idling. We also measure other data points to be able to evaluate additional metrics. These values allow us to produce reasonable upper and lower bounds for running a single node, as our hardware is selected accordingly.

In the third step, we estimate the electricity consumption of the complete network. Firstly, we collect information about the size of the network, as the node count significantly influences the amount of electricity consumed. Secondly, we develop a weighting between the single hardware devices for each network. Lastly, we multiply the electricity consumption of the weighed nodes by the number of nodes in the network. In case the distribution of nodes for all networks is available, we use the respective carbon intensity factors of the regions to calculate the carbon footprint of the respective network. Otherwise, we rely on an average global carbon intensity factor.

In the fourth step, we analyze the additional data (such as transaction data) to develop additional metrics to explore energy efficiency in transaction throughput. This allows us to put the electricity consumption of single networks into perspective with other PoS networks and also other cryptocurrencies such as Bitcoin and Ethereum.

4 Proof of Stake systems, hardware requirements, and selected hardware

The six proof of stake networks we selected do not employ identical algorithms and have different prerequisites in terms of hardware, network size, transaction throughput, and other properties. For example, some of them support stake-delegation, meaning that one node can stake on behalf of other entities in the network which do not need to run a machine themselves. While this allows for some efficiency gains and allows users with lower funds to participate in the revenue generation of staking, it comes at the price of trustworthiness and decentralization. The first part of this chapter takes a closer look at the single PoS systems and hardware requirements. In the remaining two parts of the chapter, we describe the hardware selection for running the nodes and the infrastructure required to measure electricity consumption.

4.1 Overview of PoS systems and their hardware requirements

The systems previously selected in chapter 3 have different requirements for the hardware of their network participants. The hardware requirements are partly unspecific: Tezos, for instance, specifies only the amount of CPU cores required to run their software. For the consumer market, the first two-core CPU was released in 2005 (Intel® Pentium® Processor Extreme Edition 840 (Intel, 2005)), deeming almost any CPU released in the last 15 years suitable for running Tezos (although we did not test). Other networks are more precise or have higher requirements. Solana has the highest requirements in our sample, requiring 12 cores / 24 threads, 128 GB RAM, and a 2 TB NVMe SSD. Solana also recommends a specific CPU type (AMD Threadripper Zen3) and states that a node might profit from using a GPU. It is noteworthy that Polkadot does not provide a minimum hardware requirement but rather links to a recommended setup. Algorand provides two sets of requirements, one for regular and one for enterprise nodes, although it remains unclear what is meant by that.

Table 3: Hardware requirements of PoS networks

Name	Cardano	Polkadot	Solana	Tezos	Avalanche	Algorand
CPU	2x2GHz	i7-7700k	12x2.8GHz	2 cores	> 2GHz	4 / 16 cores
RAM	8 GB	64 GB	128 GB	8 GB	6 GB	4-8 / 24 GB
Storage	30 GB	80-160 GB	2TB	100 GB	200 GB	100 / 500 GB
SSD/NVMe	N/A	NVMe	NVMe	SSD	N/A	SSD
Source	(IOHK, 2021)	(Polkadot, 2021)	(Solana, 2021)	(Heng, 2020)	(Minchev, 2021)	(Algorand Foundation, 2021)

4.2 Hardware selection

Based on the hardware requirements outlined in the previous section, we define three separate categories of hardware requirements:

- **Low hardware requirements:** Cardano, Tezos, Avalanche, and Algorand specify 2 to 4 CPU cores, 4-8 GB RAM, and under 200 GB. These requirements are somewhat imprecise. However, this hints that computational power is not a concern for these systems, and users should be comfortable running the software on any system they have available. All of these networks (except for Cardano) recommend using low-energy hardware such as the well-known Raspberry Pi for running nodes within the respective networks. 4-8 GB RAM and 200 GB of storage (even an SSD) are not uncommon anymore in today's average consumer desktop PC.
- **Specific hardware requirements:** Polkadot specifies the most precise hardware requirements with the exact CPU type as well as RAM and storage. While they list it as "Requirements – Standard Hardware", the description instead clarifies that they meant "We are able to run Polkadot on this machine". While we use hardware that satisfies Polkadot's requirements, we also run Polkadot on hardware that does not meet the requirements and include them in our calculation if they are able to run a node reliably. Nonetheless, hardware requirements typically give users who intend to run a node an indication about what to expect regarding demand, influencing their final choice of hardware.
- **High hardware requirements:** Solana has surprisingly high hardware requirements. The CPU, RAM, and storage requirements are at the highest level of standard desktop computers (besides servers). They mention that the AMD Threadripper Zen3 family is famous among the community. They also mention that graphic cards could be required in the future, which hints at the immense required processing power.

Based on the hardware requirements, both an upper and a lower bound of hardware are evident. For the lower bound, as of the popularity of the Raspberry Pi computers is high within all communities (and outside), we select a Raspberry Pi 4 Model B with 8 GB RAM and 128 GB SD-card as the lower bound. We opt for an official Raspberry Pi full kit, including fan and power supply.

Solana's hardware requirements define the upper bound. We opt for an average system within the Threadripper specifications consisting of an AMD Ryzen Threadripper 3970X, 32C/64T, 256 GB RAM (DDR4-3600), and a Samsung 970 Evo Plus 2TB. As the processor does not have an onboard graphics processor, we need a graphics card. However, as a graphics card is not required at that time, we opt for a card that does not support CUDA and cannot participate in the calculations of Solana or any other network. We select an appropriate mainboard as well as a power supply.

The upper and lower bounds highly deviate from each other in terms of computational power and electricity consumption. Further, the two computers may not capture the complete picture of the hardware used within these networks. Therefore, we decided to add three additional computers to ensure a well-balanced set of hardware for electricity consumption measurements.

As there are millions of different computer configurations, thousands of variables, and other factors that influence the electricity consumption of devices, we opt for one key variable and derive other specifications of the system from it: The central processing unit (CPU). Nonetheless, also the CPU has several variables as the number of cores, threads, speed, turbo speed, thermal design power (TDP), and others. Further, identical variables do not necessarily lead to the same computational power or electricity consumption. To get an in-depth view and understanding of the CPU landscape, we obtain a data set from PassMark. PassMark provides a software suite able to benchmark varying types of hardware, including CPUs. The obtained data set contains over 3,100 CPU models as well as over 1 million results of their benchmarking suite (Passmark Software, 2021). Based on this data set, we select three CPUs to derive our final configurations. We thereby aim at three categories of performance (high, mid, and low) and select a CPU with the average efficiency for their class. A detailed description of our approach to select CPUs can be found in Appendix A.

For the high-tier (configuration 4), we identified the Intel Core i5-10400F as being closest to the average efficiency. As Intel's F-models only have a deactivated onboard graphics chip (Intel, 2021), we decided to opt for the non-F variant, as otherwise, a dedicated GPU would add unnecessary electricity consumption to the system. The non-F variant is almost identical to the F variant regards to benchmarking results. We opted for 32 GB DDR4 RAM and a Samsung 970 Evo Plus 512 GB NVMe SSD to complement the system. Mainboard, power supply unit, and case have been selected appropriately.

The Intel Core i5-8400T has the best fit for the average electricity consumption in the mid-tier section (configuration 3). The T-model means the CPU has a "power-optimized lifestyle", resulting in lower performance and less electricity consumption. We could not directly obtain the CPU in the market and instead opted for a completed build: The Lenovo ThinkCentre M720q Tiny 10T8S3KD00. Besides the processor as mentioned above, it includes a 256 GB SSD as well as 8 GB RAM.

In the low-tier section (configuration 2), we identify the Intel Core i3-8109U as the processor with an average energy efficiency for its class. The U-label refers to a "Mobile power-efficient" CPU but is nonetheless included in MiniPCs. To our knowledge, this CPU was never sold separately on the consumer market but is available in Intel's NUC series. We obtain the Intel NUC Kit NUC8i3BEK2 Barebone and augment it with the Samsung 970 Evo Plus 512 GB NVMe SSD as well as 8 GB RAM.

We consider our selection as representative to provide a balanced set of hardware for electricity measurements with these five computers. As an operating system, we use for all our devices Ubuntu Server 20.04, except for configuration 4. Due to driver issues, we had to opt for Ubuntu Server 21. Table 4 displays an overview of the hardware configurations. Other factors than CPU are also relevant for the electricity consumption of the systems. Nonetheless, this set of hardware yields a broad overview of used hardware within such networks.

Table 4: Overview of selected hardware from lowest to highest requirement

Category	1	2	3	4	5
CPU	Broadcom BCM2711	Intel i3-8109U	Intel i5-8400T	Intel i5-10400	AMD 3970X
Architecture	ARM	x86/x64	x86/x64	x86/x64	x86/x64
RAM	8 GB	8 GB	8 GB	32 GB	256 GB
Storage	128 SD	512 SSD	256 SSD	512 SSD	2 TB SSD
GPU	Onboard	Onboard	Onboard	Onboard	AMD 6970
PSU	USB-C	65 Watt	65 Watt	650 Watt	1000 Watt
Case	Integrated	Integrated	Integrated	Custom	Custom
OS	Ubuntu 20.04	Ubuntu 20.04	Ubuntu 20.04	Ubuntu 21	Ubuntu 20.04

4.3 Infrastructure for electricity measurements

For the measurement of the electricity consumption, we use five *Mystrom WiFi Switch* which measure the electricity consumption as well the room temperature and provide them over a REST interface. The electricity measurements are made in Munich, Germany in a separate basement room with near-constant room temperature. The average room temperature during the measurement period was between 18°C and 22°C.

All devices were equipped with the same software, a new Ubuntu server 20.04/21 installation, and the monitoring tool *Glances* (Hennion, 2021) that allows us to collect additional system information such as temperature or system load.

A separate Raspberry Pi, equipped with a Python script, collected and monitored the five systems and analyzed the data generated during the runs. All computers are only connected to the power outlet and LAN. All systems share a 100 Mbit internet connection.

5 Electricity consumption of Proof of Stake networks

The definition of the to-be used hardware allows us to establish single node measurements. With these measurements, we provide upper and lower bounds for the electricity consumption of a single node and the best guess as a weighted average between the computer devices. On that basis, we establish the electricity consumption of the respective overall network and discuss additional metrics such as the electricity use per transaction.

5.1 Single node measurements

After defining and obtaining the hardware required for our analysis, we set up the hardware and install the node software for the respective network. For that, we use the following process:

- **Hardware Setup:** We install the node with the respective Linux version, configure Glances and configure remote access.
- **Idle Measurement:** We run the idle measurement for the devices without any additional software installed.
- **Node Setup:** We download and install the node software and verify the correct installation.
- **Node Bootstrap:** The single nodes are fully synced, as we do not want to skew the electricity consumption of the devices, as some might process data faster.
- **Electricity Measurement:** We shut down the node, start the electricity measurement and then start the node again. The node runs for 24 hours, as this covers an entire day cycle. Appendix B contains an overview of every electricity measurement.

To understand what exactly we are measuring, we need to describe the network and its setup. The network usually consists of nodes, either validators (participating in the consensus protocol and producing new blocks) or regular full nodes (broadcasting and verifying regular transactions). Validators are selected to propose new blocks based on their stake. We would differentiate between full nodes and validators in an ideal setup, as they have slightly different roles and responsibilities within the network. On the Main net, significant stakes are required to run a validator; on the available test nets, we were not able to generate blocks due to the low share of the overall stake. Furthermore, previous research suggests that participating in the PoS consensus mechanism has only a negligible effect on the device's electricity consumption (Sedlmeir, Buhl, Fridgen, & Keller, 2020). Therefore, we run our electricity measurement on regular full nodes running on the Main net.

5.1.1 Idle electricity consumption

We measure the electricity consumption of the devices idle. Table 5 depicts the minimum, maximum, median, and quantile 1 and 3 of the electricity consumption for 24 hours. All values are rounded to one decimal. Interestingly, the setup 2 consumes less electricity than the Raspberry Pi (configuration 1), which we deemed the most energy-efficient solution beforehand.

Table 5: Electricity consumption in Idle measured in Watt [W] – hardware selection for each of the five clusters can be found in Table 4

	1	2	3	4	5
Min [W]	2.3	2.2	2.6	23.5	75.2
Q1 [W]	2.9	2.4	2.7	24.9	76.7
Median [W]	3.1	2.5	2.8	25.4	77.8
Q3 [W]	3.1	2.5	2.8	25.1	77.1
Max [W]	3.3	2.6	2.8	25.5	78.5

5.1.2 Node electricity consumption

Due to the hardware requirements outlined in chapter 4.2, we do not run all networks on all nodes. In Table 6, we give an overview about which measurements take place on which machines. As we do not want to enforce the hardware requirements as a strict lower bound, we also test the lower-tier device if software is available (e.g., due to Raspberry Pi’s architecture, software might not be available).

Table 6: Hardware configurations used for measuring the electricity consumption of networks. (✓ (yes), ~ (test), ✗ (no), the result of the test is given in brackets)

Config.	Cardano	Polkadot	Avalanche	Algorand	Tezos	Solana
1	✗	✗	~ (✗)	~ (✓)	✓	✗
2	✓	~ (✓)	✓	✓	✓	✗
3	✓	✓	✓	✓	✓	✗
4	✓	✓	✓	✓	✓	~ (✗)
5	✓	✓	✓	✓	✓	✓

We outline challenges that occurred during the electricity measurements for some of the networks.

Due to technical restrictions, we were not able to run all nodes in **Avalanche** simultaneously throughout a period of time but had to test [2], [3,4], and [5] hardware separately. The hardware used for establishing an Internet connection crashed reproducibly when we started five nodes simultaneously. Therefore, we tested on different days accounting and controlling for the respective circumstances, such as transaction throughput. Further, PC 1 was not able to catch up to the network, therefore, we excluded this configuration from our calculations.

Cardano outlines in its requirements that an IPv4 address is necessary to connect to the network. We obtained an IPv4 address from our ISP to be able to run fully functional nodes. Due to this limitation, it can be assumed that a large portion of the network either is deployed on-premise or use hosting solutions, as anecdotal evidence suggests that a large share of private internet connections only natively support IPv6 and use Dual Stack Lite for IPv4 services. Future research could look into how this affects the hardware distribution of the Cardano network.

We also had issues with running **Tezos** on all five devices for 24 hours. During several runs, some nodes were unable to continue their work in the network and the CPU utilization dropped to zero percent. To account for this, we removed the respective timeslot (six hours are missing for computer 4) and only consider the time the nodes were fully connected and functional.

Solana poses high requirements for the internet connection. The same problems that occurred during the Avalanche measurements also occurred while running a single node on Solana. Due to that limitation, we relocated configuration number 5 to a different place with a 1 Gbit internet connection in which the node worked as expected. Noteworthy: Solana also requires an IPv4 address, but does not state so in its requirements.

In Table 7, we outline the mean consumption of all networks and nodes.

Table 7: Mean consumption of nodes for the respective network in Watt [W]

	1	2	3	4	5
Algorand	5.53	34.89	32.23	69.95	168.59
Avalanche	N/A	23.44	24.03	57.34	144.67
Cardano	N/A	3.90	3.70	27.56	84.47
Polkadot	N/A	4.31	5.35	29.25	107.86
Tezos	4.86	19.25	19.74	52.41	141.65
Solana	N/A	N/A	N/A	N/A	221.33

Two currencies are very close to each other in terms of the electricity consumption of the single nodes. Cardano and Polkadot only deviate within the range of +/- 10 % of each other throughout all hardware configurations. On the other hand, Algorand seems to consume most throughout the networks that support all devices. The Raspberry Pi (configuration 1) seems to be running at almost maximum capacity, consuming 5.5 watts. Also for configuration 2 and 3, Algorand consumes up to the factor of 9 compared to other cryptocurrencies such as Cardano and Polkadot. Tezos is in between these two groups. It consumes up to factor 5 of the low-energy blockchains, but only about 50% of Algorands consumption. Solana is different due to its high hardware requirements. It only runs on configuration

5 and consumes the most out of any network-device combination with over 220 watts while running. Avalanche consumes electricity similar to Tezos while not supporting configuration 5.

5.2 Calculation of bounds for electricity consumption

To calculate the electricity consumption of the overall networks, we need to understand the average calculation for a single node. We measured the electricity consumption for five different computers. With these measurements, we can provide upper bounds, meaning the highest electricity that a node consumes, lower bounds, the least electricity a node consumes, and a best guess that captures the consumption of the average node best for the respective network.

5.2.1 Upper and lower bound

The upper and lower bound are measured by the least efficient and most efficient hardware, respectively. The lower bound therefore is constituted by the Raspberry Pi (configuration 1) for Algorand and Tezos. Configuration 2 is used for Avalanche, Cardano, and Polkadot as a lower bound. Configuration 5 serves both as the lower bound for Solana and as an upper bound for all currencies, as for Solana only one device is available. The respective upper and lower bounds are displayed in Table 8.

Table 8: Overview of lower and upper bounds for the respective network per single node

	Lower bound [watts]	Lower bound [kWh / year]	Upper bound [watts]	Upper bound [kWh / year]
Algorand	5.53	48.47	168.59	1,476.81
Avalanche	23.44	205.33	144.67	1,267.31
Cardano	3.90	34.15	84.47	740.00
Polkadot	4,31	37.72	107.86	944.88
Tezos	4.86	42.58	141.65	1,240.89
Solana	N/A	N/A	221.33	1,938.85

5.2.2 Best guess

The electricity consumption of an average node in the networks is challenging to estimate. There is no empirical data on the concrete hardware that nodes are running on or indicating user's preferences. For node owners, several factors are relevant for their decision on which hardware to run their node on. First, owners stake tokens to receive rewards and want their revenue to be stable, aiming for hardware designed for long-term operations. Second, due to the profit structure, they do not intend to spend all their revenue on hardware and might rather opt for sufficient hardware within the hardware requirements. These thoughts might influence their decision in one way or another but might not directly translate to a hardware selection. Therefore, we opt for a binomial distribution for

the hardware selection, based on a regular distribution for key questions. The distribution for each hardware type is displayed in Table 9.

Table 9: Overview of node distribution for the six networks

		1	2	3	4	5
Nodes	Algorand, Tezos	6.25 %	25.00 %	37.50 %	25.00 %	6.25 %
Nodes	Avalanche, Cardano, Polkadot	N/A	12.50 %	37.50 %	37.50 %	12.50 %
Nodes	Solana	N/A	N/A	N/A	N/A	100 %

With this distribution, we calculate the weighted electricity consumption of an average node:

$$\sum_{i \in \text{hardware}} \text{avgEnergyConsumption}_i * \text{share}_i$$

An average node in the networks consumes from 22.77 watts for Cardano to 221.33 watts for Solana. This difference of over factor 13 implies that not only the number of nodes are relevant for the electricity consumption of PoS networks, but also the underlying software and its requirements. Table 10 gives an overview about the best guess electricity consumption for every network.

Table 10: Overview of best guess estimates per single node

	Best guess [watts]	Best guess [kWh / year]
Algorand	49.18	430.82
Avalanche	51.53	451.40
Cardano	22.77	199.45
Polkadot	27.00	236.49
Tezos	34.47	302.00
Solana	221.33	1,938.85

Additionally, we display the electricity consumption estimates for the upper and lower bound as well as the best guess for each network in Figure 1.




Figure 1: Overview of lower bound, upper bound and best guess estimates for a single node of the respective network in watts

5.3 Electricity consumption of the networks

We apply our lower bound, upper bound as well as our best guess at the number of nodes in the respective networks. We obtain the number of nodes from the respective block explorers. Appendix C gives an overview of all data sources. The results are depicted in Table 11.

Table 11: Overview of electricity consumption of the respective networks applying the best guess estimate

	Node count	Electricity intensity network [W]	Consumption / day [kWh]	Consumption / year [kWh]
Algorand	1,190	58,524.08	1,404.58	512,670.92
Avalanche	1,084	55,858.52	1,340.58	489,311.19
Cardano	3,002	68,351.08	1,640.43	598,755.44
Polkadot	297	8,018.01	192.43	70,237.76
Tezos	375	12,928.06	310.27	113,249.81
Solana	1,015	224,649.62	5,391.59	1,967,930.68

The electricity consumption of the networks amounts from 70 to 1,967 MWh annually in our best guess, Polkadot with the lowest electricity consumption and Solana with the highest electricity consumption in the field. These results imply that there is a factor of more than 28 between the electricity consumptions of the different networks, hinting at the respective carbon footprint. Figure 2 gives an overview of the electricity consumptions of the networks and their respective node count in comparison.




Figure 2: Electricity consumption per year [kWh] and count of nodes of the respective networks

5.4 Electricity consumption per transaction

An often-used metric in comparing electricity consumption between systems is the electricity consumption per transaction. This allows comparing systems that have different architectures, transaction throughput, and electricity requirements. Nonetheless, companies that want to report emissions associated with cryptocurrency expose should not use a transaction-based allocation approach and should rely on other methodologies (Gallersdörfer, Klaatzen, & Stoll, 2021).

The complexity of this metric is based on the fact that some systems provide a theoretical electricity consumption per transaction, simulating the network at full speed. Other calculations are based on transaction rates measured in the networks, making comparisons skewed. Further, the definition of a transaction might vary from network to network.

An additional complexity is the attribution of the electricity consumption solely to the transactions. The system requires a base electricity consumption to keep up with the consensus without providing any transactions. Nonetheless, given the base load of a network, running a node in a "low-transaction"-period might yield higher electricity per transaction costs than usually to be expected. While this metric provides a straightforward insight into different protocols, its base assumptions need to be understood and its results must be treated with care.

As we measured the electricity consumption of our nodes in real-world scenarios, we also apply the transaction numbers that took place during the respective time periods. As we sometimes have different time slots for when we tested our hardware (for Avalanche and Tezos), we also apply the

respective transaction rates before we weigh the single nodes for the overall network. Thereby we also apply the previously described binomial distribution.

Table 12: Best guess electricity consumption of PoS networks on a per-transaction basis. The weighted average of transactions is the weighted number of transactions that took place on the respective blockchain during our measurements.

	Wh/tx per node	Wh/tx per network	Weighted avg. of tx
Algorand	0.00227	2.70	520,417
Avalanche	0.00439	4.76	276,907
Cardano	0.01718	51.59	32,559
Polkadot	0.05865	17.42	11,037
Tezos	0.11054	41.45	7,033
Solana	0.00016	0.166	32,383,318

The range for the electricity consumption per transaction goes from 0.166 watthours for Solana up to 51.59 watthours for Cardano. As expected, this metric depends on the amount of transactions taking place on the respective blockchain, also the overall electricity consumption per transaction further depends on the number of nodes connected to the respective network. Generally, these numbers are expected to go down with an increase in the transaction rate, regardless which blockchain is in use.

5.5 Carbon footprint of PoS networks

The electricity consumption of any system has no direct environmental impact, as mere usage does not cause any harm. However, the indirect impacts due to the carbon intensity of the underlying energy sources used for electricity consumption do cause damage to the environment and need to be considered for sustainable business operations.

Depending on the underlying energy sources, the respective carbon footprint of electricity consumption can vary. For a precise estimate of the carbon footprint, two pieces of data are essential: The location of the electricity consumers as well as the carbon intensity of the respective grid.

There are several ways for local electricity consumers to ensure that their energy is carbon neutral and has no displacement effect, meaning that other local electricity consumers are pushed from sustainable energy into fossil fuels. Corporate Power Purchase Agreements (PPAs), unbundled energy attribute certificates (EACs), or off-grid electricity production for self-consumption are ways to ensure that an entity is carbon-neutral from an electricity consumption perspective. As these are instruments aimed at energy-intensive industries or large corporations, we do not assume that any node relies on such methodologies and apply an average grid intensity factor.

Previous research localized nodes in other protocols by relying on internet search machines aimed at ASIC devices, IP addresses, or pool addresses. These approaches allowed for an estimate of how the nodes are distributed worldwide. Unfortunately, structured data on the location of nodes in all networks is not available.

Due to the absence of such data, we rely on the average grid intensity worldwide. A formula to calculate the respective carbon footprint is shown.

$$\sum_{i \in \text{node}} \text{EnergyConsumption}_i * \text{carbonIntensity}$$

We assume the carbon intensity of the grid to be 475 g CO₂e/kWh. We derive this value from the International Energy Agency (IEA) Global Energy & CO₂e Status Report (International Energy Agency, 2018). With that, we can derive the carbon footprint of all networks. The respective values are depicted in Table 13 and in Figure 3.

Table 13: Overview of CO₂e emissions of the networks on an annual basis

CO ₂ e emissions / year [t]	Lower bound	Upper bound	Best guess
Algorand	27.40	834.77	243.52
Avalanche	105.74	652.54	232.42
Cardano	48.69	1055.20	284.41
Polkadot	5.32	133.30	33.36
Tezos	7.58	221.03	53.79
Solana	N/A	934.77	934.77




Figure 3: Carbon emissions per year [in tonnes] of each network for lower bound, upper bound and best guess estimate.

6 Discussion and comparison of results

In the previous chapters, we introduced our methodology and conducted measurements to derive the electricity consumption and carbon footprint of six Proof of Stake networks. In this chapter, we contextualize the results of our work.

6.1 Yearly electricity consumption in the context of other systems

The electricity consumption of a single network is often meaningless without a context to compare the system. In chapter 5, we outline that the yearly electricity consumption of the networks range from 70,000 kWh to roughly 1,900,000 kWh. An average US household consumes about 10,600 kWh per year and therefore, the least electricity consuming network Polkadot consumes about 6.6 times the electricity and the most electricity consuming network Solana about 200 times the electricity (U.S. Energy Information Administration, 2021). In comparison to the decentralized cryptocurrency Bitcoin, PoS networks consume less than 0.001 % of the Bitcoin network assuming 89.78 TWh on the 21st August 2021 (CBECI, 2021). Bitcoin consumes much more electricity than any Proof of Stake system due to its Proof of Work consensus mechanism, resulting in the deployment of energy-intensive hardware. In Figure 4 , we compare the Bitcoin network, the Ethereum network assuming 17.3 TWh on the 21st August 2021^d, the PoS networks, and the average US household.




Figure 4: Yearly electricity consumption for Bitcoin, Ethereum, Proof of Stake networks Polkadot, Tezos, Avalanche, Algorand, Cardano and Solana, and an average US household in kWh. Logarithmic scale.

^d We used the methodology outlined in (Gallersdörfer et al., 2020) and applied it to the respective date.

6.2 Electricity consumption per transaction

In chapter 5, we also calculate the electricity costs per transaction for each network. We compare these numbers to other approaches such as Visa, Bitcoin and Ethereum. Visa consumes about 1.5 Wh per transaction, while Bitcoin consumes about 1722.24 kWh per transaction (Statista, 2021). We obtain the figure 37 kWh per transaction from own calculations based on (Gallersdörfer et al., 2020). Figure 5 gives a comparison of these systems. Generally, it is not surprising that a centralized system like Visa is more energy-efficient than a decentralized system. Due to the difference in the amount of computational hardware and distribution worldwide, most systems must consume more electricity. Due to its Proof of Work (PoW), Bitcoin and Ethereum consume much more on an overall and on a per-transaction basis.




Figure 5: Electricity consumption [Wh] per transaction for Bitcoin, Ethereum, Visa, and all PoS systems. Logarithmic scale.

6.3 The carbon footprint of Proof of Stake networks

Overall, the emissions of the PoS networks are very low. As outlined in Chapter 5, the networks emit from 33.36 for Polkadot to 934.77 tonnes for Solana of CO₂e yearly. For example, nine round trips from Munich (MUC) to San Francisco (SFO) in business class emit about the same amount of carbon dioxide (MyClimate, 2021) produced by the Tezos network. It can be assumed that the carbon emissions of companies behind these networks are higher than the emissions from the networks itself. Figure 6 compares the carbon footprints of the six networks to each other and to a roundtrip MUC – SFO in business class.




Figure 6: Yearly carbon footprint of PoS networks compared to a roundtrip flight in business class

7 Conclusion

In this report, we outline an approach for calculating the electricity consumption and carbon footprint of six Proof of Stake networks, namely Algorand, Avalanche, Cardano, Polkadot, Tezos, and Solana. We selected hardware, made measurements of the protocols, and calculated the respective metrics. We discussed our results and introduced several other key metrics, such as the Bitcoin and Ethereum network for comparison.

Our best guess estimates the yearly electricity consumption of the Proof of Stake networks from 70 MWh for Polkadot to 1,967 MWh for Solana. This results in carbon footprints between 33 and 934 tonnes of CO₂e annually, respectively. Compared to other electricity consumers such as an average US household, these networks consumes up to 200 times more electricity, and produce up to 153 times the amount of an intercontinental roundtrip flight.

Given the continuous development and evolution of Proof of Stake networks, our results can only be taken as a snapshot of the respective timeframe. Further measurements and analyses are required to update and further enhance the validity of the metrics for electricity consumption and carbon footprint of Proof of Stake and other networks. Additionally, other networks employing different consensus mechanisms as well as second layer networks need to be taken into account to gain a holistic picture of the environmental impact of cryptocurrencies and tokens.

In recent years, Bitcoin has faced harsh criticism for its electricity demand and carbon emissions. In the public, these fears and accusations have often been applied to other blockchain protocols, regardless of their technical foundations or capabilities, harming the adoption of blockchain protocols in the industry, public sector, and private investors. Based on the total emissions calculated for these six networks, one may conclude that Proof of Stake-based blockchain protocols consume an amount of electricity that does not justify the discussions about their environmental footprints. Instead, an extensive perspective, including corporate footprints and the ecosystem, must be taken. For practitioners selecting a PoS blockchain protocol, other factors such as decentralization, network throughput or functionality (e.g., Smart Contracts) should increase in relevance as decision criteria.

8 References

- Algorand Foundation. (2021). Algorand Network Architecture. Retrieved from <https://algorand.foundation/algorand-protocol/network>
- CBECI. (2021). Cambridge Bitcoin Electricity Consumption Index. Retrieved from <https://cbeci.org>
- de Vries, A. (2021). Bitcoin boom: What rising prices mean for the network's energy consumption. *Joule*, 5(3), 509-513.
- Gallersdörfer, U., Klaaßen, L., & Stoll, C. (2020). Energy consumption of cryptocurrencies beyond bitcoin. *Joule*, 4(9), 1843-1846.
- Gallersdörfer, U., Klaaßen, L., & Stoll, C. (2021). Accounting for carbon emissions caused by cryptocurrency and token systems. *arXiv preprint arXiv:2111.06477*.
- Heng, C. (2020). Guide to Using the Tezos-node Part I: Setting Up a Node. Retrieved from <https://www.tezos.org.sg/guide-to-setting-up-a-tezos-node/>
- Hennion, N. (2021). Glances - An eye on your system. Retrieved from <https://github.com/nicolargo/glances>
- Intel. (2005). Intel® Pentium® Processor Extreme Edition 840 Retrieved from <https://ark.intel.com/content/www/us/en/ark/products/27613/intel-pentium-processor-extreme-edition-840-2m-cache-3-20-ghz-800-mhz-fsb.html>
- Intel. (2021). Intel® Processor Names and Numbers. Retrieved from <https://www.intel.com/content/www/us/en/processors/processor-numbers.html>
- International Energy Agency. (2018). Global Energy & CO2 Status Report. Retrieved from https://iea.blob.core.windows.net/assets/23f9eb39-7493-4722-aced-61433cbffe10/Global_Energy_and_CO2_Status_Report_2018.pdf
- IOHK. (2021). Cardano Mainnet - Stake Pool Minimum System Requirements Retrieved from <https://iohk.zendesk.com/hc/en-us/articles/900001951586-Stake-Pool-Minimum-System-Requirements>
- King, S., & Nadal, S. (2012). Ppcoin: Peer-to-peer crypto-currency with proof-of-stake. *self-published paper*, August, 19(1).
- Minchev, R. (2021). Technical Requirements for Running a Validator Node on Avalanche. Retrieved from <https://support.avax.network/en/articles/4064879-technical-requirements-for-running-a-validator-node-on-avalanche>
- MyClimate. (2021). Carbon Footprint calculator. Retrieved from https://co2.myclimate.org/en/flight_calculators/new
- Passmark Software. (2021). Hardware & Software Market Trends. Retrieved from <https://www.passmark.com/services/market-analysis.php>
- Polkadot. (2021). Run a Validator. Retrieved from <https://wiki.polkadot.network/docs/maintain-guides-how-to-validate-polkadot>
- Sedlmeir, J., Buhl, H. U., Fridgen, G., & Keller, R. (2020). The energy consumption of blockchain technology: beyond myth. *Business & Information Systems Engineering*, 62(6), 599-608.
- Solana. (2021). Validator Requirements. Retrieved from <https://docs.solana.com/running-validator/validator-reqs>
- Statista. (2021). Bitcoin average energy consumption per transaction compared to that of VISA as of August 16, 2021 Retrieved from <https://www.statista.com/statistics/881541/bitcoin-energy-consumption-transaction-comparison-visa/>
- Stoll, C., Klaaßen, L., & Gallersdörfer, U. (2019). The carbon footprint of bitcoin. *Joule*, 3(7), 1647-1661.
- U.S. Energy Information Administration. (2021). How much electricity does an American home use? Retrieved from <https://www.eia.gov/tools/faqs/faq.php?id=97&t=3>

Appendix A: Hardware Selection

We use the Passmark CPU Benchmark Dataset. Our methodology to select three CPUs consists of the following steps:

1. The data set contains many processor types that are not relevant to us. We filter out:
 1. CPUs with less than 50 benchmarking results, as we expect that they are not relevant for the validator community.
 2. CPUs that were released before 1/1/2015, as we consider less usage of outdated hardware and a practical reason: We cannot buy these CPUs in the market.
 3. CPUs with missing or incomplete data.
 4. CPUs of AMD. Intel is the dominating manufacturer of CPUs with over 80 % market share over the last years. Not all values in the data set are consistent between both producers, and already one AMD system is included in our data set. Therefore we decided not to consider AMD processors.
 5. CPUs intended for servers or notebooks. We think that the share of server hardware is low and notebooks nonexistent. Some CPUs are marked as "Laptop only" in our dataset; however, we find them included in MiniPCs, e.g., the Intel NUC. To account for these CPUs, we consulted geizhals.de as a source of CPU models sold within MiniPCs and did not remove them from the data set.
2. After obtaining a cleaned data set, we can separate the data set into three equally large categories for later selection: High-level, mid-level, and low-level. While the hardware within the networks might not be equally distributed among these three categories, this approach allows us to shift the allocation for single networks between the devices depending on their hardware requirements.
3. We are confronted with the fact that older, high-level CPU models might have the same computational power as recent low-level CPU models but different energy efficiencies, leading to entirely different results. Therefore, we introduce an additional variable in our data set called *energy efficiency*. The energy efficiency of a processor is the average benchmarking result divided by the TDP. The TDP serves as a proxy for a processor's energy demand capabilities, as it describes the maximum amount of heat measured in Watts the CPU cooling system has to deal with.
4. This variable allows us to calculate the average energy efficiency for each category of CPUs (4-high/3-mid/2-low) and select an average processor from the respective tier. This approach ensures that we a) cover three different performance categories and b) select an average energy efficiency for their respective class.

Appendix B: Electricity Measurements of Single Nodes

All electricity measurements are conducted in Watt.

Algorand

Table 14: Electricity Consumption while running Algorand measured in Watt [W]

	1	2	3	4	5
Min [W]	4.4	6.9	10.8	27	87.8
Q1 [W]	5.3	35.4	23	55.5	150.5
Mean [W]	5.5	34.9	32.2	70	168.6
Median [W]	5.5	38	29.7	70.4	169.9
Q3 [W]	5.8	38.5	43.3	85.2	186.1
Max [W]	7.5	40.2	51.6	118.9	232.2

Avalanche

Table 15: Electricity Consumption while running Avalanche measured in Watt [W]

	1	2	3	4	5
Min [W]	N/A	3.6	5.0	24.4	80.1
Q1 [W]	N/A	14.8	14.5	38.6	128.1
Mean [W]	N/A	20.4	19.6	49.4	138.1
Median [W]	N/A	23.4	24.0	57.3	144.7
Q3 [W]	N/A	36.8	35.5	81.6	166.7
Max [W]	N/A	45.1	48.7	94.7	196.1

Cardano

Table 16: Electricity Consumption while running Cardano measured in Watt [W]

	1	2	3	4	5
Min [W]	N/A	2.5	2.7	24.3	75.9
Q1 [W]	N/A	2.9	2.9	26.4	80
Mean [W]	N/A	3.9	3.7	27.6	84.5
Median [W]	N/A	3	3	26.8	81.2
Q3 [W]	N/A	3.3	3.5	27.1	83.4
Max [W]	N/A	36.7	31.5	69.6	150.6

Polkadot

Table 17: Electricity Consumption while running Polkadot measured in Watt [W]

	1	2	3	4	5
Min [W]	N/A	2.36	3	23.4	77.9
Q1 [W]	N/A	3.5	4.4	26.1	91
Mean [W]	N/A	4.3	5.3	29.3	107.9
Median [W]	N/A	4	5.1	27	116.5
Q3 [W]	N/A	4.7	6	30.6	118.1
Max [W]	N/A	28.2	19.9	55.2	144

Tezos

Table 18: Electricity Consumption while running Tezos measured in Watt [W]

	1	2	3	4	5
Min [W]	3.5	8.7	11.6	28	123
Q1 [W]	4.5	16.8	17.7	48.5	140.3
Mean [W]	4.9	19.3	19.7	52.4	141.7
Median [W]	4.8	19.2	19.9	55.4	141.8
Q3 [W]	5.2	21.8	22.2	57.2	143.3
Max [W]	6.8	33.2	29.4	69.4	158.9

Solana

Table 19: Electricity Consumption while running Solana measured in Watt [W]

	1	2	3	4	5
Min [W]	N/A	N/A	N/A	N/A	164.3
Q1 [W]	N/A	N/A	N/A	N/A	204.1
Mean [W]	N/A	N/A	N/A	N/A	221.3
Median [W]	N/A	N/A	N/A	N/A	217.9
Q3 [W]	N/A	N/A	N/A	N/A	233.5
Max [W]	N/A	N/A	N/A	N/A	329.5

Appendix C: Data sources for single networks

All market capitalizations are taken from <https://coinmarketcap.com>.

Algorand	Information
Measurement period	2021-08-16 08:28:19 to 2021-08-17 08:28:10
Number of nodes	https://metrics.algorand.org/ → “Decentralization”
Transaction Count	https://algoexplorer.io/top-statistics → “Transaction Metrics” → “Total Transactions”
Software version	<i>go-algorand 2.8.0</i> (https://github.com/algorand/go-algorand)
Avalanche	Information
Measurement period	[PC2]: 2021-10-19 21:54:52 to 2021-10-20 21:54:52 [PC3,PC4]: 2021-10-20 23:24:13 to 2021-10-21 23:24:52 [PC5]: 2021-10-22 00:16:11 to 2021-10-23 00:15:59
Number of nodes	https://explorer.avax.network/validators → “Validators”
Transaction count	P-Chain / X-Chain : https://explorer.avax.network/tx → counted tx with API P-Chain: https://snowtrace.io/chart/tx → Weighted count
Software version	<i>avalanchego 1.6.2</i> (https://github.com/ava-labs/avalanchego)
Cardano	Information
Measurement period	2021-08-07 00:01:02 to 2021-08-08 00:35:26
Number of nodes	https://cardanoscan.io/ → “Total stake pools”
Transaction count	https://explorer.bitquery.io/cardano/transactions?from=2021-08-01&till=2021-08-31 → “Transactions By Date”
Software version	<i>cardano-node 1.27.0</i> (https://github.com/input-output-hk/cardano-node)
Polkadot	Information
Measurement period	2021-08-28 12:59:26 to 2021-08-29 12:58:06
Number of nodes	https://polkadot.subscan.io/validator → “Validators”
Transaction count	https://polkadot.subscan.io/transfer → “Transaction history”
Software version	<i>polkadot 0.9.9</i> (https://github.com/paritytech/polkadot)
Tezos	Information
Measurement period	[PC1,PC2,PC3,PC5]: 2021-08-22 17:17:44 to 2021-08-23 17:17:24 [PC4]: 2021-08-24 22:52:57 to 2021-08-25 17:19:51 (missing hours accounted for)
Number of nodes	https://tzstats.com/bakers → “Bakers”
Transaction count	https://api.tzstats.com/series/block.json for respective time frames, see https://tzstats.com/docs/api#time-series-endpoints , used “n_tx”
Software version	<i>tezos 10.0.0</i> (https://gitlab.com/tezos/tezos)
Solana	Information
Measurement period	2021-09-09 16:58:18 to 2021-09-10 16:57:15
Number of nodes	https://solanabeach.io → “Validators”
Transaction count	https://explorer.solana.com/ → “Live transaction stats” → “TPS history” averaged over respective period. Additionally, subtracted for “votes” operations by crawling 202 random blocks from the respective period using https://solanabeach.io/ and subtracted the average share of votes (78.2 %).
Software version	<i>solana 1.6.21</i> (https://github.com/solana-labs/solana)

About CCRI

The **Crypto Carbon Ratings Institute (CCRI)** provides carbon estimates for investments in cryptocurrencies and technologies such as Blockchain and distributed ledger technologies (DLT). We have built a multi-year research track record with a specific focus on Bitcoin and its environmental impacts. We published comprehensive and formerly peer-reviewed studies on Bitcoin's carbon footprint in the renowned scientific journals. Our research has been covered by major media outlets, such as CNN and The New York Times, and has been appraised as very good estimate by major organizations, such as the IEA.

Ulrich Gellersdörfer

Ulrich Gellersdörfer is a research associate in the Department of Informatics at the Technical University of Munich. His research focuses on identity management in blockchains. His interest extends to further aspects of the technology, ranging from environmental implications to data analytics applications.

Lena Klaaßen

Lena Klaaßen has a background in Management and Technology with a particular focus on finance, power engineering and energy markets. She has conducted research on carbon accounting in the corporate and cryptocurrency space at TUM and MIT. She now focuses on research in the field of climate finance at ETH Zurich.

Dr. Christian Stoll

Christian Stoll conducts research at the Center for Energy and Environmental Policy Research at the Massachusetts Institute of Technology and at the Center for Energy Markets of the Technical University of Munich. His research focuses on the implications of climate change from an economic point of view.