

White paper:

Determining the electricity consumption and carbon footprint of Proof of Stake networks

Crypto Carbon Ratings Institute (CCRI)¹

<https://carbon-ratings.com>

Last updated: March 2023

¹*Crypto Carbon Ratings Institute* is a research-driven company providing data on sustainability aspects of cryptocurrencies, blockchain and other technologies. CCRI was founded by Ulrich Gallersdörfer, Lena Klaatzen and Christian Stoll in 2021. More information: <https://carbon-ratings.com>.

Introduction

The electricity consumption and related carbon footprint of Bitcoin and other cryptocurrencies are subject to extensive discussion in public, academia, and industry. Various estimations exist, comparing Bitcoin's electricity consumption to different mid-sized countries (CCAF, 2022; Vries, Gellersdörfer, Klaaßen, & Stoll, 2022).² 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 & Nadel, 2012). While it is consensus in the broader scientific community that PoS does not exhibit the same electricity issues of PoW, the responsibility of individual PoS systems is typically less clear.

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).³ In previous industry reports, CCRI calculated both the electricity consumption and carbon footprint of a wide range of Proof of Stake networks, namely Algorand, Avalanche, Cardano, Polkadot, Tezos, and Solana (CCRI, 2022a) as well as TRON (CCRI, 2022c), Polygon (CCRI, 2022b) and Ethereum (CCRI, 2022d).⁴ There are also estimates for other PoS systems, albeit no actual hardware measurements took place (Platt et al., 2021).

This white paper aims to outline the methodology used by CCRI to derive the electricity consumption of PoS networks. By doing this, we aim to work towards more standardized approaches for assessing PoS networks and build the foundation for constant improvement based on established approaches summarized in this white paper.

² Transparency note: This article was co-authored by the founders of CCRI.

³ Transparency note: This article was authored by the founders of CCRI.

⁴ Transparency note: The industry reports have been commissioned by Avalanche, TRON, Polygon and ConsenSys.

Overview of approach

Our methodology builds upon five steps to generate data on the electricity consumption and carbon footprint of a PoS system.

- (1) In the first step, we analyze the respective PoS network and its 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.
- (2) In the second step, we estimate the electricity usage of a single node and provide upper and lower bounds as well as a best guess estimate for the network. We start by running the network's client 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, such as CPU utilization and processed blocks, 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.
- (3) In the third step, we estimate the electricity consumption of the complete network. For this, we collect information about the size of the network, as the validator node count significantly influences the amount of electricity consumed. We then multiply the electricity consumption of the best guess estimate by the number of validator nodes in the network.
- (4) In the fourth step, we analyze additional data (such as transactions and block information) to develop further metrics to explore energy efficiency in transaction throughput. We take samples of the nodes' electricity consumption periodically and examine the number of transactions that were handled by the single nodes during the respective time periods. This allows us to describe the marginal influence of the number of transactions on the electricity consumption of a node. As a result, we establish a model to estimate a node's electricity consumption based on the number of transactions. This enables us to put the electricity consumption of the analyzed network into perspective with other PoS networks and also other cryptocurrencies.
- (5) In the fifth step, we estimate the carbon emissions arising from the operation of the considered PoS network. If data on the location of the network's nodes is available, we use country-specific emission intensities to calculate a network-specific emission intensity. If no information on the regional distribution of the nodes in the network is present, we rely on the world average emission intensity.

(1) Hardware requirements and test environment

We observe three general categories of hardware requirements across PoS systems for nodes participating in a network:

1. **Low hardware requirements.** For PoS networks with rather low hardware requirements, we assume that computational power is not a concern for the systems, and users should be comfortable running the software on any system they have available. Typically, such networks recommend using low-energy hardware for running nodes, as for example the well-known Raspberry Pi. In today's average consumer desktop PC, 4-8 GB RAM and 200 GB of storage (even an SSD) are not uncommon anymore.
2. **Specific hardware requirements.** Some networks specify quite precise hardware requirements, for instance stating the exact CPU type as well as RAM and storage. For such networks, we normally aim for using hardware that satisfies the requirements, but we also test hardware that does not meet the recommendations if they are able to run a node reliably. In these cases, we include the respective results in our calculation. Nonetheless, hardware requirements typically give users who intend to run a node an indication of what to expect regarding demand, influencing their final choice of hardware.
3. **High hardware requirements.** Some few PoS systems exhibit surprisingly high hardware requirements. The CPU, RAM, and storage requirements can be at the highest level of standard desktop computers (besides servers). Graphic cards can be required in such networks, which hints at an immense processing power required.

Hardware description

We define a standardized hardware pool that covers the above-mentioned categories in order to ensure a high degree of hardware diversity for the analyses of PoS networks.

On the lower end of hardware requirements (configuration 1), we select a Raspberry Pi 4 Model B with 8 GB RAM and 128 GB SD-card given that the popularity of the Raspberry Pi computers is high within all communities. We opt for an official Raspberry Pi full kit, including a fan and power supply.

On the upper end of hardware requirements (configuration 6), we opt for an average system within the Threadripper specifications consisting of an AMD Ryzen Threadripper 3970X, 32C/64T, 256GB RAM (DDR4-3600), a MSI GeForce RTX 4090 graphics card, and a Samsung 970 Evo Plus 2TB in order to address high hardware requirements. We select an appropriate mainboard, power supply, and case.

These two configurations can be seen as upper and lower bounds which 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 networks to be analyzed. Therefore, we

decided to add four 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, the CPU also has several variables such 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 four CPUs to derive our final configurations. We thereby aim at three categories of performance (high, mid, and low) and select one or more CPUs with the average efficiency for their class.

For the high-tier (configuration 5), 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 with regard to benchmarking results. We opted for 64 GB DDR4 RAM and a Samsung 970 Evo Plus 2 TB to complement the system. Mainboard, power supply unit, and case have been selected appropriately.

Regarding the mid-tier section, we have extended our hardware selection with an additional device compared to our previous measurements conducted in our initial benchmarking study (CCRI, 2022a), as we assume that most standard users apply hardware from this range. Since the Intel NUC series is becoming increasingly popular for running blockchain nodes, we decided on an Intel NUC with medium equipment (configuration 4). We chose an Intel Core i5-1135G7 laptop processor with included graphics chip, which represents the upper mid-range of typically used devices quite well. This additional mid-tier computer is equipped with a 32 GB DDR4 RAM and a 2 TB NVMe SSD. Furthermore, we still stick to the Intel Core i5-8400T since it 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 NVMe 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 as well as 8 GB RAM.

We consider our selection as representative to provide a balanced set of hardware for electricity measurements with these six computers. As an operating system, we use Ubuntu Server 20.04 for all our devices, except for configuration 5. Due to driver issues, we have to opt for Ubuntu Server 21. *Table 1* displays an overview of the hardware configurations described above. 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 1: Overview of selected hardware configurations from lowest to highest requirement

	1	2	3	4	5	6
CPU	Broadcom BCM2711	Intel i3-8109U	Intel i5-8400T	Intel i5-1135G7	Intel i5-10400	AMD 3970X
CORES/THREADS	4/4	2/4	6/6	4/8	6/12	32/64
ARCHITECTURE	ARM	x86/x64	x86/x64	x86/x64	x86/x64	x86/x64
RAM	8 GB	8 GB	8 GB	16 GB	64 GB	256 GB
STORAGE	128 GB SD	512 GB SSD	256 GB SSD	2 TB SSD	2 TB SSD	2 TB SSD
GPU	Onboard	Onboard	Onboard	Onboard	Onboard	MSI 4090
PSU	USB-C	65 Watt	65 Watt	65 Watt	650 Watt	1000 Watt
CASE	Integrated	Integrated	Integrated	Integrated	Custom	Custom
OS	Ubuntu 20.04	Ubuntu 20.04	Ubuntu 20.04	Ubuntu 20.04	Ubuntu 21	Ubuntu 20.04

Hardware selection

For the analysis of a specific PoS network, it is important to decide on a case-by-case basis which hardware configurations to consider in the analysis. Depending on the hardware requirements of the network, we do not run the measurement on all hardware configurations as listed in *Table 1*. As we do not want to enforce the network's hardware requirements as a strict lower bound, we try to include hardware configurations that only barely satisfy the hardware requirements.

Infrastructure for electricity measurements

For the measurement of the electricity consumption, we use a Mystrom WiFi Switch for each computer. These switches measure the electricity consumption as well as the room temperature and provide the values over a REST interface. The electricity measurements are made in Munich, Germany in a separate server room with near-constant room temperature.

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

A separate Raspberry Pi, equipped with a Python script, collects and monitors the systems during executing the full nodes and analyses the data generated during the runs. All computers are only connected to the power outlet and LAN. All systems share an internet connection with 350 Mbit/s download and 110 MBit/s upload.

(2) Electricity consumption of a single node

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, represented by the most and least efficient hardware, and the best guess as a weighted average between the selected computer devices. On that basis, we will then establish the electricity consumption of the overall PoS network.

Single node measurements

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

1. **Hardware Setup.** We install the node with the respective Linux version, configure Glances and configure remote access.
2. **Idle Measurement.** We run the idle measurement for the devices without any additional software installed.
3. **Node Setup.** We download and install the software necessary for executing a full node of the PoS network to be analyzed and verify the correct installation.
4. **Node Bootstrap.** On each node we include in our analysis, we start the execution of the client software necessary to run a full node of the network to be analyzed. We wait for the node fully synced since we do not want to skew the electricity consumption of the devices during the bootstrapping phase.
5. **Electricity Measurement.** We shut down the node, start the electricity measurement and then start the node again. The node runs for 24 hours executing the network's client software, as this covers an entire day cycle.

To understand what exactly we are measuring, we need to describe a standard PoS network and its common setup. It consists of nodes running a client service, either validators (participating in the consensus protocol and producing new blocks) or regular full nodes (broadcasting and verifying regular

transactions). In an ideal setup, we would differentiate between full nodes and validators, as they have slightly different roles and responsibilities within the network, however, significant stakes are commonly required to run a validator. 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 PoS network.

Idle electrical power

We measure the electricity consumption of all devices in idle mode. *Table 2* depicts the minimum, maximum, median, and the first and third quartile of the electricity consumption for 24 hours. All values are rounded to one decimal. Interestingly, configurations 2 and 3 consume less electricity than the Raspberry Pi (configuration 1), which we deemed the most energy-efficient solution beforehand.

Table 2: Electrical power in idle measured in Watt [W] – hardware selection for each of the six configurations can be found in *Table 1*

	1	2	3	4	5	6
MIN [W]	2.9	2.6	2.6	3.6	24.5	77.5
Q1 [W]	3.0	2.6	2.9	3.7	24.8	77.9
MEAN [W]	3.0	2.7	3.0	3.7	25.0	78.2
MEDIAN [W]	3.0	2.7	2.9	3.7	24.9	78.0
Q3 [W]	3.0	2.7	3.0	3.7	25.1	78.3
MAX [W]	3.9	17.8	17.3	4.4	26.6	118.1

Calculation of bounds and a best guess for electricity consumption

With the measurements of the electricity consumption for each hardware configuration when running a full node in a respective PoS network, we can provide an upper bound (the highest electricity that a node consumes), a lower bound (the least electricity a node consumes), and a best guess that captures the consumption of the average node best for the network.

Upper and lower bound. The upper and lower bound are determined by the least efficient and most efficient hardware as described in *Table 1*.

Best guess. The electricity consumption of an average node in a network is challenging to estimate. Typically, there is no empirical data on the concrete hardware that nodes are running on or indicating users' 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 barely 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 hardware selected out of our pool for the network considered, based on a regular distribution for key questions. With this distribution, we calculate the electricity consumption of an average node by taking a weighted average of the single nodes' electricity consumptions.

(3) Electricity consumption of the entire network

To calculate the electricity consumption of the overall network, we take the number of validator nodes in the considered PoS network and multiply it by our best guess estimate of the electricity consumption of an average node. We only consider validator nodes since regular full nodes do not participate in the consensus mechanism. By using the lower and upper bound of the electricity consumption of a single node, we may also derive a respective lower and upper bound for the network's electricity consumption. Usually, we obtain the number of nodes of the network from a publicly accessible source.

(4) Analysis of efficiency in transaction throughput

Discussion on electricity consumption per transaction metric

An often-used metric in comparing electricity consumption between systems is the electricity consumption per transaction, which is frequently calculated by dividing the total electricity consumption by the total number of transactions within a specific time period. The underlying idea is to compare systems that have different architectures, transaction throughput, and electricity requirements. However, such comparisons are associated with several challenges and complexities.

First, underlying definitions may influence the results significantly. For instance, 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.

Second, there is additional complexity caused by 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 and vice versa. This may also distort comparisons across networks when electricity consumption per transaction is used as the central metric. Consequently, albeit the metric may provide straightforward insights into different protocols, its base assumptions need to be understood and its results must be treated with care and in context of other metrics.

Third, the electricity consumption per transaction is only one metric describing the sustainability of a network. It is of utmost importance to understand that this metric needs to be seen in the context of other metrics such as decentralization, security, transaction complexity, state size, and others. This metric alone is not sufficient to decide whether a cryptocurrency is sustainable or if a cryptocurrency is worth investing in; in an extreme case, a network consisting of a single, high-performance computer, would be the most sustainable cryptocurrency, however making nonsense of the decentralization idea.

Marginal electricity consumption per transaction

To address the above-mentioned challenges, we measure the power demand of our nodes in a real-world scenario and consider the transactions per second throughput that took place during the respective time period. We then derive the marginal electricity consumption per transaction based on a statistical regression model.

With this approach, we aim to obtain more accurate results than merely considering average values. Furthermore, it has the advantage that constant power demand, which is independent of the transaction count, can be differentiated from the power demand that is driven by transactions. The regression model is set up completely on the basis of our own measurements. We filter outliers and construct a regression line for each of the hardware configurations from *Table 1* selected for the considered PoS network, based on our periodically taken measurement samples that consist of the current power demand and the transactions per second throughput at that time.

Based on this, we can establish a linear equation for a regression line to predict the power demand of a best guess node operating in a PoS network (P_{BG}) for a given transactions per second (tps) throughput. To determine a general slope (m_{BG}) for a network's best guess node, which represents the electricity consumption per transaction in Watt seconds, we weigh the slopes of the individual regression lines of all hardware configurations included in our measurements. Likewise, we derive a weighted y-axis intercept ($base_{BG}$), which represents the node's power demand while executing no transactions but running the client software of the network. As a result, we obtain a linear regression equation to determine the power demand of a best guess validator node in the considered network depending on the transactions per second the best guess node processes:

$$P_{BG}(tps) [W] = m_{BG} * tps + base_{BG}$$

Multiplying a network's validator node count by the resulting power demand of a best guess node allows us to estimate the overall network's power demand for a specific transaction per second throughput based on our model:

$$P_{network} [W] = P_{BG}(tps) * validatorCount$$

(5) Estimation of the network’s carbon emissions

The electricity consumption of any system has no direct climate impact, as mere usage does not necessarily cause any emissions. However, the impacts due to the potential emissions of the underlying energy sources may cause damage to the climate and need to be considered for sustainable business operations.

Depending on the underlying energy sources, the respective carbon footprint of the 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 claim that their electricity consumption is carbon neutral. This includes corporate Power Purchase Agreements (PPAs), unbundled energy attribute certificates (EACs) – also often referred to as Renewable Energy Certificates (RECs) –, or off-grid electricity production for self-consumption. As we commonly do not have any information on whether or to what extent the electricity consumption of a PoS network is backed by such instruments, we rely on the average grid intensity factor. As these instruments are also often aimed at energy-intensive industries or large corporations, we find the application of the average grid intensity factor to be plausible for a solid estimate of the carbon footprint of a PoS network.

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. It depends on the specific network to be analyzed whether information on the node location is available or not. If data on where the single nodes are located is present, we calculate a network-specific emission intensity utilizing location-specific emission factors. If there is no data on node locations, we assume the latest available world average emission intensity (as of 2021: 459 gCO₂/kWh (IEA, 2022)).

With that, we derive the carbon footprint of the network by multiplying the determined carbon intensity with the sum of the best guess electricity consumption over all validators in the network:

$$\text{CarbonIntensity} * \sum_{i \in \text{node}}^{\text{NetworkValidatorCount}} \text{ElectricityConsumption}_i$$

Likewise, by applying the lower and the upper electricity consumption bounds, a lower and an upper bound for the overall network’s carbon emissions can be estimated.

Conclusion

In this white paper, we outlined the methodology proposed by CCRI to derive the electricity consumption and carbon footprint of PoS networks. By relying on actual measurements instead of a purely model-based approach, this methodology allows to base the estimate on empirical data. The described approach is very well suited to provide a detailed snapshot of the respective timeframe. By monitoring the validator node and transaction count, the assessment can serve as basis to provide dynamic estimates of the electricity consumption and carbon footprint of PoS networks.

Given the continuous development and evolution of Proof of Stake networks, further measurements and analyses may help to update and further enhance the validity of the metrics for electricity consumption and carbon footprint as well as to understand the underlying drivers. As an avenue for further research, the influence of the blockchain size and its resulting state size on the network's power demand needs to be better understood. In our data, we see first indications on heightened power demands which might not be fully explained by the transaction throughput alone, thus requiring the exploration of other potential causes of this heightened demand.

References

- CCAF (2022). Cambridge Bitcoin Electricity Consumption Index. Retrieved from <https://ccaf.io/cbeci/index>
- CCRI (2022a). Energy efficiency and carbon emissions of PoS Networks. Retrieved from <https://carbon-ratings.com/>
- CCRI (2022b). Energy Efficiency and Carbon Footprint of the Polygon Blockchain. Retrieved from <https://carbon-ratings.com/>
- CCRI (2022c). Energy Efficiency and Carbon Footprint of the TRON Blockchain. Retrieved from <https://carbon-ratings.com/>
- CCRI (2022d). The Merge – Implications on the Environmental Sustainability of Ethereum. Retrieved from <https://carbon-ratings.com/>
- 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>
- Hennion, N. (2021). Glances - An eye on your system. Retrieved from <https://github.com/nicolargo/glances>
- IEA (2022). World Energy Outlook 2022. Retrieved from <https://www.iea.org/reports/world-energy-outlook-2022>
- King, S., & Nadel, S. (2012). Ppcoin: Peer-to-peer crypto-currency with proof-of-stake. *Self-Published Paper*, 19(1).
- Passmark Software (2021). Hardware & Software Market Trends. Retrieved from <https://www.passmark.com/services/market-analysis.php>
- Platt, M., Sedlmeir, J., Platt, D., Xu, J., Tasca, P., Vadgama, N., & Ibañez, J. I. (2021). Energy Footprint of Blockchain Consensus Mechanisms Beyond Proof-of-Work. Retrieved from http://blockchain.cs.ucl.ac.uk/wp-content/uploads/2021/11/UCL_CBT_DPS_Q32021_updated-2.pdf
- 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. <https://doi.org/10.1007/s12599-020-00656-x>
- Vries, A. de, Gallersdörfer, U., Klaaßen, L., & Stoll, C. (2022). Revisiting Bitcoin’s carbon footprint. *Joule*, 6(3), 498–502. <https://doi.org/10.1016/j.joule.2022.02.005>