ACI Research Paper #15-2023

The Energy Consumption of the Ethereum-Ecosystem

Patrick WOITSCHIG

Gazi Salah UDDIN

Taojun XIE

Wolfgang Karl HÄRDLE

August 2023

Please cite this article as:

Woitschig, Patrick, Gazi Salah Uddin, Taojun Xie and Wolfgang Karl Härdle, "The Energy Consumption of the Ethereum-Ecosystem", Research Paper #15-2023, *Asia Competitiveness Institute Research Paper Series (August 2023)*

The Energy Consumption of the Ethereum-Ecosystem

Patrick Woitschig - Humboldt-University of Berlin

Gazi Salah Uddin - Linköping University

Taojun Xie - National University of Singapore

Wolfgang Karl Härdle - National University of Singapore, Singapore Management University, Bucharest University of Economic Studies, University of Rijeka

Institute for Digital Assets

Blockchain Research Center

Abstract

Decentralized Finance (DeFi) and blockchain technology are potential game changers but come with a snag. Indeed, the cloven hoof becomes apparent in the network of cryptocurrencies: the energy consumption of cryptocurrencies presents, without any doubt, an environmental challenge. Traditional methodologies for quantifying this consumption, broadly classified into top-down and bottom-up approaches, have merits and limitations. However, discrepancies between these methods indicate an opportunity for more nuanced and precise analysis. Based on the Ethereum ecosystem, this study first presents a detailed analysis of the split into Ethereum and Ethereum Classic, and seconds the energy consumption consequences of the merge. Leveraging domain knowledge, and profitability thresholds, we estimate Ethereum's pre-Merge energy demand and demonstrate the potential of PoS consensus mechanisms to significantly reduce the energy demands of cryptocurrency networks, contributing to the broader discourse on the environmental sustainability of blockchain technologies.

1 Introduction

The rising interest surrounding blockchain technology in recent years has been nothing less than extraordinary. Once dismissed as a fringe interest confined to the crypto sphere, or regarded with skepticism as an endeavor of questionable viability, blockchain technology now commands the focus of numerous pioneering global organizations. The capital flowing into the field is nothing short of a torrent, with billions of dollars earmarked for its adoption and innovative exploration, a testament to its growing acceptance and recognition (Nagarajan (2022)). Its magnetic allure has also permeated the hallowed halls of academia, with a surge of research and development activities revolving around this transformative technology (Härdle, Harvey, and Reule (2020)). Notable academic institutions like the Blockchain Research Center at Humboldt-University of Berlin have emerged as active contributors in this arena, spearheading initiatives like CRIX - a Crypto-Currencies Index (Trimborn and Härdle (2018)). Such efforts underscore the indisputable potential of blockchain technology and the increasing academic interest in understanding the mechanics of cryptocurrencies.

However, the ascension of blockchain technology, first popularized by the advent of Bitcoin, has also spurred vigorous discourse on its energy consumption and potential ramifications on climate change (de Vries (2018); Truby (2018)). Prestigious publications, including the New York Times and the BBC, have raised concerns about the environmental repercussions of blockchain technologies, with

Bitcoin as the focal point. Journalists deduced that in 2019, the energy consumption of the Bitcoin network was on par with that of countries like the Netherlands (Huang, O'Neill, and Tabuchi (2021)) or Norway (Criddle (2021)). However, these assertions often suffer from oversimplification, lack of context, or insufficient understanding of the complex relationship between energy consumption and blockchain technology's economic or technical dimensions (Krause and Tolaymat (2018)).

Indeed, while some cryptocurrencies like Bitcoin (BTC) account for substantial energy consumption, it is imperative to acknowledge the varying environmental impacts of distinct blockchain technologies (Sedlmeir, Buhl, Fridgen, and Keller (2020a), Sedlmeir, Buhl, Fridgen, and Keller (2020b)). In the past decade, emerging blockchain platforms like Ethereum and Hyperledger Fabric have provided versatile, secure transaction networks without intermediaries. Notably, many of these platforms exhibit significantly lower energy consumption than the pioneering Proof-of-Work (PoW) technology used by BTC (Fairley (2019)).

A prime illustration of this progression is Ethereum's (ETH) landmark transition from PoW to Proof-of-Stake (PoS) via "the merge," addressing energy consumption apprehensions and exemplifying the feasibility of a sustainable blockchain technology future. This shift is projected to reduce ETH's energy consumption by an astounding 99 percent, signifying a monumental stride towards sustainability (de Vries (2023a)). This crucial milestone in ETH's journey accentuates the ongoing innovation within the blockchain domain. However, it is equally important to consider that ETH's energy savings do not equate to the net savings of the entire cryptocurrency ecosystem, as evidenced by miners transitioning to other PoW currencies. It is, therefore, necessary to create a comprehensive understanding of the energy consumption patterns of various blockchain technologies to guide public discourse.

The primary objective of this study is to foster a comprehensive understanding of the energy consumption of blockchain technologies. It focuses on ETH, exploring its energy profile before and after the transition to PoS and investigating the extent of energy retained in the cryptocurrency ecosystem as miners transition to other cryptocurrencies post-ETH transition. To address these areas, the following research questions are posed:

1. What was the extent of energy consumption by ETH before it transitioned to PoS?

2. Post-ETH's transition, how much energy was retained within the cryptocurrency ecosystem

due to miners switching to other cryptocurrencies?

The study aims to debunk misconceptions about blockchain technology's energy consumption by answering these questions and fostering a more nuanced discourse and research in this critical field. Data analysis and statistical modeling techniques are employed to achieve this.

2 Literature Review

The burgeoning prevalence and integration of blockchain technologies, predominantly Proof-of-Work (PoW) cryptocurrencies such as Bitcoin, have ignited comprehensive debates and provoked concerns about their ecological footprint, especially in energy consumption. The energy usage of Bitcoin became the focus point of most discussions, with O'Dwyer and Malone leading the research discourse as early as 2014 (O'Dwyer and Malone (2014)). Since then, a diverse range of estimates has emerged, subject to different assumptions and computational methodologies. Some have even put forth the proposition that "Bitcoin emissions alone could push global warming above 2°C" (Mora, Rollins, Taladay, et al. (2018)).

The subsequent discourse aims to compare two distinct methodologies of estimating energy consumption, primarily focusing on Ethereum: the Top-Down Approach and the Bottom-Up Approach. The Top-Down Approach, pioneered by Digiconomist, can be interpreted as a ceiling estimate predicated on the assumption that miners exclusively utilize profit-generating hardware (de Vries (2023b)). This upper limit can be adjusted downwards by integrating additional plausible assumptions to yield a more realistic estimation. In juxtaposition, the Bottom-Up Approach can generate a floor estimate, assuming that miners operate with the most energy-efficient equipment available (McDonald (2021)). The estimate can also be fine-tuned towards a more pragmatic range by factoring in a more reasonable assumption regarding various hardware used.

2.1 Top-Down Energy Consumption Estimation Approach

Alex de Vries, alias Digiconomist, formulated an energy consumption estimation method for cryptocurrencies, initially applied to BTC (de Vries (2018)). This "Top-Down Approach" uses two primary factors: electricity costs and the income from mining operations. The crux of this method assumes a competitive cryptocurrency production landscape (Hayes (2015)), where miners invest resources until costs balance expected gains. Consequently, the energy consumption can be deduced by performing a break-even analysis. With the following formula, where energy consumption is in kilowatt-hours (kWh), Ethereum price is in USD per Ethereum (ETH), the reward is in ETH per second (s), and electricity price is in USD per kWh, we can bound the energy consumption:

$$\textbf{Energy Consumption}\ [kWh] \leq \frac{\textbf{Ethereum Price}\ [\frac{USD}{ETH}] \times \textbf{Mining Reward}\ [ETH]}{\textbf{Electricity Price}\ [\frac{USD}{kWh}]}$$

Ethereum miners' full recompense emanates from two principal sources: block rewards, bestowed upon successful block mining, and transaction fees, which users pay to include their transactions within blocks. Hence, the cumulative mining reward can be delineated as a synthesis of block rewards and transaction fees. Given the approximately 13-second block creation interval, known as block time, miners acquire block rewards and transaction fees for each successfully mined block. Intriguingly, transaction fees exhibit considerable variability, contrasting with the relative stability of block rewards. The elasticity of transaction fees, subject to factors like transaction complexity and network congestion, enriches the dynamic landscape of Ethereum mining.

There are wide-ranging estimates for electricity costs, which have sparked robust debates among academics and cryptocurrency enthusiasts. The Bitcoin Energy Consumption Index (BECI), developed by Digiconomist, suggests an average electricity cost of 0.05 USD per kWh, based on the premise that miners are incentivized to operate in regions where electricity is more affordable to maximize their profits. However, in (de Vries (2018)), he notes, "With an electricity rate of 10 cents per kWh (commonly used in mining profitability calculators), the power demand of miners should not exceed 9.21 GW to avoid operating at a loss (also assuming no further expenses other than electricity)." Additionally, his Ethereum Energy Consumption Index places this figure at 8.88 GW, which he puts as his realistic estimate for the Ethereum network's power demand.

Given the volatile nature of the Ethereum price, estimates can fluctuate between 8.61 GW and 9.85 GW for the first half of September, placing the 9.21 GW squarely within this range. However, according to the Bitcoin Energy Consumption Index on his website, he conjectures that 49.24% of the total miner income is dedicated to electricity. This assertion raises questions about the source of the 8.88GW used in his final calculation. Taking an electricity price of 0.05 USD per kWh and an allocation of 49.24% of the total miner income to electricity estimates for the same period range from 8.53 GW to 9.7 GW.

2.2 Bottom-Up Approach to Estimating Energy Consumption

The "Bottom-Up Approach" to computing energy consumption in cryptocurrency was initially conceived by O'Dwyer and Malone back in 2014, with a specific focus on Bitcoin (O'Dwyer and Malone (2014)). However, for our purpose, we want to discuss McDonald's adaptation and implementation of this methodology to Ethereum (McDonald (2021)). Contrasting sharply with the "Top-Down Approach" employed by Digiconomist, the bottom-up approach presents a detailed, hardware-oriented analysis.

As a point of departure from the economic assumptions embedded in the top-down model, 'mining' in the context of the bottom-up approach necessitates a thorough understanding of the hardware capabilities within the Ethereum mining landscape. This model postulates that the power consumed by the network can be meticulously estimated by cataloging the spectrum of active mining devices, each with its unique energy consumption footprints. Contrasting the breakeven analysis driving the top-down approach, the bottom-up methodology calculates energy expenditure by directly incorporating these mining machines' hash rates and energy efficiencies. At the heart of this approach is a mathematical equation that elegantly captures the power consumed in the process of mining:

$$\textbf{Power Consumption}\ [\frac{J}{s}] = \frac{\textbf{Hash Rate}\ [\frac{Hash}{s}]}{\textbf{Hash Efficiency}\ [\frac{J}{Hash}]}$$

The hash rate quantifies the total computational power committed to Ethereum mining, directly influencing the network's energy expenditure. A higher hash rate increases computational power, leading to greater energy consumption. The Ethash hash rate can be considered the rate at which miners can make 'guesses' per second to solve the Ethash puzzle, thereby successfully mining a new block. As such, the significance of the hash rate becomes clear.

Efficiency plays a crucial role in understanding the energy consumption of cryptocurrency mining. Assessing the total hash rate based on individual mining devices' hash rates allows a more accurate estimation of energy consumption patterns. This approach acknowledges mining hardware's diversity and varying efficiencies, providing valuable insights into the energy footprint. For instance, advanced ASIC miners can achieve hash rates in the tera hashes per second (TH/s) range, while less powerful GPUs may only produce mega hashes per second (MH/s). Recognizing these differences is critical for accurately assessing the energy consumption of cryptocurrency mining.

To accurately determine the energy consumption of Ethereum mining, the bottom-up approach considers various factors that contribute to overhead and efficiency losses. These factors include power supply efficiency (PSE), hardware overhead (HO), data center overhead (DO), and grid loss (GL) (McDonald (2021)). By accounting for these factors, the bottom-up approach recognizes the inefficiencies in power supply units, the varying GPU resources and configurations, the additional energy consumption within mining operations, and the losses incurred during transmission and distribution in the power grid. This approach allows for a more accurate estimation of the overall energy footprint of Ethereum mining.

$$\textbf{Realised Hash Efficiency}\ [\frac{J}{Hash}] = \textbf{Theoretical Hash Efficiency}\ [\frac{J}{Hash}] \times \frac{\textbf{PSE}}{\textbf{HO} \times \textbf{DO} \times \textbf{GL}}$$

To provide a holistic understanding of Ethereum's energy consumption, McDonald's bottom-up approach (McDonald (2021)) combines granular hardware-specific energy calculations with an in-depth understanding of Ethereum mining intricacies. This methodology considers the Ethash hash rate, the mining difficulty level, hashing efficiency, and various overhead and efficiency losses. McDonald's offers a more refined estimation approach. He uses a weighted average considering the different efficiencies of all possible mining hardware, leading to a more realistic estimate of 2.44 GW.

2.3 Discrepancy Between the Approaches

It is clear from the earlier discussion that the two primary methodologies for estimating the power consumption of PoW cryptocurrencies — the top-down and bottom-up approaches — yield vastly different estimates. While the top-down approach of Alex de Vries provides a range of 8.88 GW to 9.21 GW, the bottom-up approach of Kyle McDonald has a substantially lower estimate of 2.44 GW.

Such a wide discrepancy is due to the fundamental difference in the two approaches' premises and underlying assumptions. The top-down approach essentially focuses on the economics of mining, assuming that miners will only participate in the process if it is profitable for them. It considers miners' economic incentives and extrapolates energy usage based on the electricity costs that these incentives could cover. On the other hand, the bottom-up approach is hardware-focused, starting from the ground and working upwards. It evaluates the energy efficiency of all available mining hardware. From this basis, it estimates the total power consumption by aggregating the consumption of all the individual mining machines.

This discrepancy between the top-down and bottom-up approaches underscores the need for a more nuanced methodology that can accommodate the dynamic and rapidly evolving landscape of cryptocurrency mining. The following section will introduce a unique approach to bridge this gap, reconciling the two methodologies and addressing their inherent limitations to provide a more robust and comprehensive analysis of energy consumption in Proof-of-Work cryptocurrencies. The ultimate goal is to find a middle ground that integrates the strengths of both methodologies while mitigating their weaknesses, thereby painting a more accurate picture of energy usage in PoW-cryptocurrency mining.

3 Methodology

This section aims to provide an enhanced and accurate estimate of Ethereum's energy consumption before the transition to Proof-of-Stake (PoS), often called "the merge." This assessment is pivotal as it provides a benchmark to gauge the energy efficiency improvements brought about by the merge. Our approach is twofold: we first undertake a comprehensive analysis of the impact of the Ethereum merge on the Ethereum Classic network; secondly, we delve into the underlying principles of profitability thresholds that play a pivotal role in cryptocurrency mining.

The first part of our investigation focuses on the migration of miners from Ethereum to Ethereum Classic post-merge. This migration is critical to our analysis, as these miners constitute a distinctive group within the Ethereum mining community, distinguished by their specific hardware efficiency benchmarks. By understanding the behaviors and dynamics of this group, we can gain valuable insights into the energy usage patterns of Ethereum pre-merge. The second part of our study aims to explore the profitability thresholds in cryptocurrency mining. This is a critical aspect to consider, as it significantly influences miners' decisions to participate in a network, affecting its overall energy consumption.

Through these dual prongs of investigation, we strive to shed light on the complex dynamics of Ethereum's energy consumption before its significant transition. The insights gleaned from this study could inform discussions about the sustainability of blockchain technologies and provide valuable context for understanding the energy efficiencies introduced by the merge.

3.1 The Genesis and Journey of Ethereum Classic

Ethereum Classic (ETC) and Ethereum (ETH) were carved from the same initial blockchain, their histories intertwined up until block 1,920,000. This shared heritage laid the foundation for their parallel evolution until 2016, when a significant shift led them onto distinct trajectories. In that year, the infamous DAO (Decentralized Autonomous Organization) attack sent tremors through the Ethereum community. The DAO was a cutting-edge smart contract on the Ethereum blockchain, conceived as an innovative, investor-led venture capital fund. Regrettably, in June 2016, an attacker exploited a loophole in The DAO's code, draining around a third of its funds, equivalent to roughly 50 million dollars at the time (Mehar et al. (2019)).

This event sparked a critical debate within the Ethereum community: should the community accept the attack as an irreversible but unfortunate incident, thus maintaining the ethos of "code is law," or should they intervene by adjusting the blockchain's history to refund the pilfered funds? The majority, including the Ethereum Foundation and a significant portion of the project's developers, chose the latter option, executing a hard fork that effectively annulled the transactions that allowed the hacker to siphon The DAO's funds (Mehar et al. (2019)). This action birthed a new iteration of the Ethereum blockchain, now known as Ethereum (ETH). The original, unaltered Ethereum blockchain, where the DAO attack and its aftermath remained intact, persisted and was rebranded Ethereum Classic (ETC).

The figures below show both networks' prices and hash rates from Ethereum's inception until mid-2017. Initially, it was uncertain whether Ethereum or Ethereum Classic would become the dominant Ethash cryptocurrency. However, by early 2017, Ethereum emerged as the clear frontrunner, largely due to wider adoption. This trend is also evident in each network's number of active addresses. Ethereum's active addresses began to rise at the beginning of 2017, while the number of active addresses in Ethereum Classic remained constant.

Although Ethereum Classic did not garner the same adoption levels as Ethereum, it stands as a testament to the principle of immutability, upholding the belief that blockchain transactions should remain irreversible regardless of circumstances. Consequently, despite their shared genesis, Ethereum and Ethereum Classic represent differing philosophical ideologies within the broader Ethereum ecosystem (Mehar et al. (2019)). At the heart of these ideologies and their manifestations is the chosen consensus algorithm: Ethash. This algorithm, designed for its robust security and

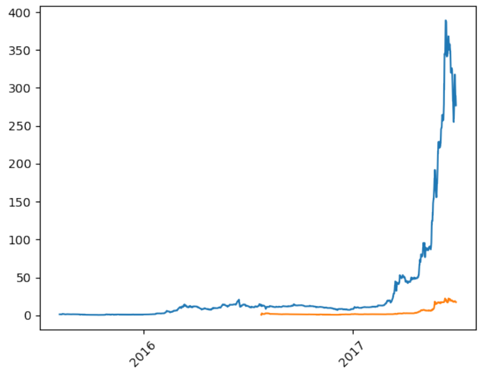

Figure 1: Hashrates [GH/s]

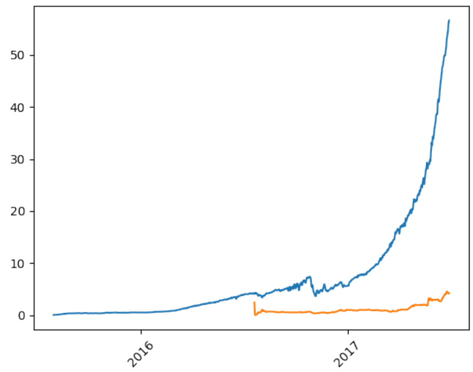

Figure 2: Prices [USD]

decentralization features, was selected at Ethereum's inception and retained by Ethereum Classic. As we navigate the narrative of the Ethereum 2.0 upgrade, it's crucial to remember this shared Ethash lineage.

3.2 Ethereum 2.0 and the Transition to Proof-of-Stake

Unveiled on September 15, 2022, Ethereum 2.0, colloquially known as "Eth2" or "Serenity," represents a significant advancement in the Ethereum blockchain's development. Ethereum's founding team initiated this progression with three core aims: improving the scalability, enhancing the security, and promoting the sustainability of the Ethereum network. This transition pivots around a critical shift in the blockchain's consensus mechanism: the transition from Proof-of-Work (PoW) to Proof-of-Stake (PoS) (Buterin (2020)).

The PoW consensus mechanism, though lauded for its security and decentralization attributes, poses several challenges. As discussed in the previous sections, PoW demands considerable computational resources and power to solve intricate mathematical problems, allowing miners to add new blocks to the blockchain. This reliance on massive energy consumption is an environmental concern and presents an economic barrier to miners, particularly with soaring energy prices globally. Moreover, PoW is said to confront limitations in scalability, a challenge that comes to the fore as the number of Ethereum transactions increases. (Buterin (2020)).

The inception of Ethereum 2.0 and its shift to the PoS model aim to mitigate these issues. Within a PoS framework, the system selects validators to create new blocks based on the quantity of Ether they own and their willingness to stake it as collateral. This change dramatically reduces the demand for computational power and energy consumption, enabling greater network scalability and environmental sustainability. It also facilitates broader participation in the Ethereum network, as the entry barriers concerning hardware investments and energy costs are significantly lower compared to the PoW paradigm (Buterin (2020)).

This paradigm shift raises a critical question: what happens to the miners, specifically those utilizing ASIC machines designed for Ethash, Ethereum's unique PoW algorithm? With Ethereum's shift to PoS, these machines can no longer participate in validating Ethereum transactions, potentially rendering them obsolete. The most feasible alternative for these miners is to migrate their operations to the Ethereum Classic network, which remains committed to the Ethash PoW consensus mechanism. This mining migration from Ethereum to Ethereum Classic has notable implications for Ethereum Classic's network dynamics and security and our understanding of Ethereum's pre-merge energy consumption. In the following section, we will delve deeper into Ethereum and Ethereum Classic mining dynamics and explore how these factors can accurately estimate Ethereum's pre-merge energy consumption.

3.3 Potential Switch From Ethereum to Ethereum Classic

The unveiling of Ethereum 2.0 profoundly impacted miners, especially those who invested in ASIC machines specifically designed for mining Ethash—the distinctive PoW algorithm of the Ethereum ecosystem. Application-Specific Integrated Circuits (ASICs) are highly specialized hardware modules crafted to execute a single function with remarkable efficiency. In the context of cryptocurrencies, these devices are purpose-built to manage a specific hashing algorithm, optimizing the transaction processing and new block generation within a blockchain operating on the corresponding algorithm. This signifies that Ethereum miners' ASIC devices are fine-tuned for the Ethash algorithm, an exclusive hashing function integral to Ethereum's PoW consensus model (Orender, Mukkamala, and Zubair (2020)).

What becomes the fate of these miners and their machinery? Among the array of cryptocurrencies still employing the Ethash algorithm, Ethereum Classic stands out as a compelling choice for miners. Being the most lucrative cryptocurrency that continues to use Ethash post-Ethereum, Ethereum Classic naturally appeals to these miners. Its unwavering dedication to PoW and the Ethash algorithm means that miners could seamlessly redirect their Ethash ASICs towards mining Ethereum Classic in the wake of Ethereum 2.0's implementation (Orender et al. (2020)). In addition, the versatility of GPU mining devices extends beyond Ethash, enabling them to be repurposed for mining any cryptocurrency.

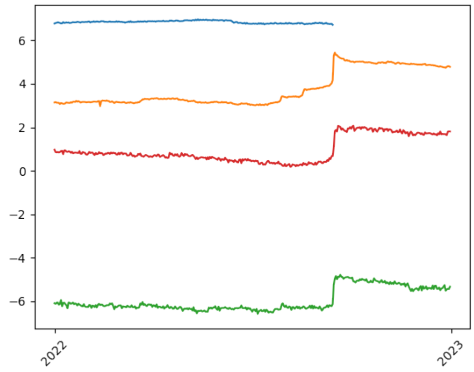

Figure 3: Log Hashrates [Terahash/s]

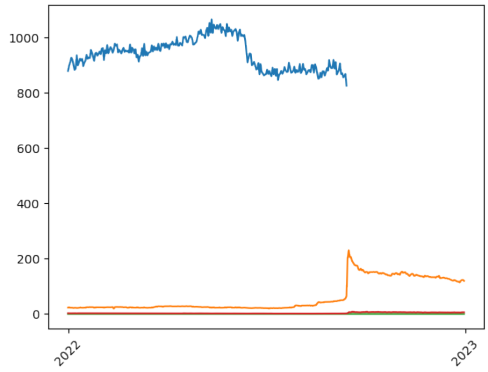

Figure 4: Hashrates [Terahash/s]

When looking at the hash rates of various cryptocurrencies, we can observe a significant move on the day of the Ethereum 2.0 merge. However, while several cryptocurrencies, including Bitcoin Gold (Green) and Vertcoin (Red), saw a substantial increase in their hash rates on a logarithmic basis, these jumps are less meaningful when assessed on standard scales. In contrast, Ethereum Classic (Yellow) showed a notable increase in hash rate, signifying a sizable migration of miners from Ethereum to Ethereum Classic.

Before the rollout of Ethereum 2.0, the network hash rate of Ethereum Classic hovered around 45 TH/s. However, this figure significantly escalated following Ethereum 2.0's unveiling, reaching a peak of 230 TH/s. This sharp increase in Ethereum Classic's hash rate signifies the influx of miners into its network. While the hash rate later subsided to around 120 TH/s by the year's end, the increase from 45 TH/s to 120 TH/s suggests a considerable retention of Ethereum's mining capacity within the broader Ethereum ecosystem via Ethereum Classic.

3.4 Revisiting ASIC Mining Profitability Thresholds

Having discerned the migration of miners from Ethereum to Ethereum Classic, we can utilize this information to refine our estimation of the energy consumption associated with Pre-Merge Ethereum. This perspective informs our return to the Digiconomist's Top-Down Approach, underpinned by the principle that rational economic incentives guide miners to avoid exceeding their mining returns through energy expenditure.

As discussed earlier, this concept allows us to perform a break-even analysis to estimate the energy involved in this process. At its foundation, we use the formula below, where energy consumption is in kilowatt-hours (kWh), Ethereum price is in USD per Ethereum (ETH), the reward is in ETH per second (s), and the electricity price is in USD per kWh:

$$\textbf{Energy Consumption}\ [\frac{kWh}{s}] \leq \frac{\textbf{Ethereum Price}\ [\frac{USD}{ETH}] \times \textbf{Reward}\ [\frac{ETH}{s}]}{\textbf{Electricity Price}\ [\frac{USD}{kWh}]}$$

While this equation abstracts from specific hardware considerations, we can restructure it to include hardware factors. Recalling our discussion on the evolution of mining hardware, we highlighted the energy efficiency improvements per hash. This is crucial as it allows us to rethink the concept of Reward from a 'per time unit' to a 'per Hash' basis.

The Reward per second (disregarding transaction fluctuations) is relatively stable in native units, making changes in hash rate significantly impactful. A miner's share of the total Reward is proportional to their share of hashing power. Hence, they should only consume as much energy per hash until it is no longer economically viable, underscoring the relevance of energy efficiency improvements per hash:

$$\textbf{Energy Efficiency}\ [\frac{kWh}{Hash}] \leq \frac{\textbf{Ethereum Price}\ [\frac{USD}{ETH}] \times \textbf{Reward}\ [\frac{ETH}{Hash}]}{\textbf{Electricity Price}\ [\frac{USD}{kWh}]}$$

3.5 The Role of Profitability Thresholds in ASIC Mining

This amount of energy consumption per hash introduces the concept of a profitability threshold. This threshold represents a breakeven point at which revenues generated from cryptocurrency mining equate to the associated costs, factoring in energy expenditures and hardware efficiency. While this principle has been primarily employed in Bitcoin's context, it holds equal relevance for other Proof-of-Work cryptocurrencies, including Ethereum and Ethereum Classic.

3.6 Implementing Profitability Thresholds to Ethereum and Ethereum Classic

In our shift from the theoretical framework of profitability thresholds to their practical applications, we spotlight Ethereum and Ethereum Classic. Before the launch of Ethereum 2.0, the profitability threshold for Ethereum Classic was higher than Ethereum's, a reflection of the economic challenges of mining Ethereum Classic given its market price.

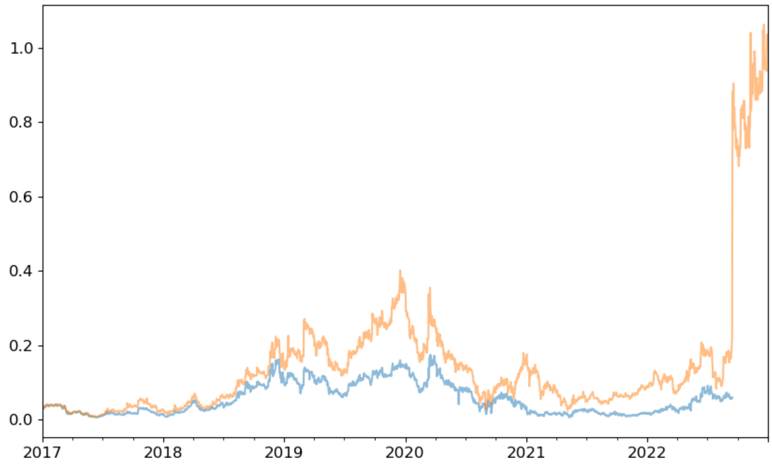

Figure 5: Profitability thresholds for Ethereum and Ethereum Classic [MH/J]

However, an exciting dynamic becomes evident when considering Ethereum mining hardware's energy efficiencies. Many mining devices were sufficiently efficient to mine Ethereum and Ethereum Classic. Despite the profitability differences being slight and at times, virtually nonexistent, the picture was more nuanced. Although Ethereum Classic's price was consistently lower than Ethereum's, the total hash rate for the Ethereum Classic network was proportionally smaller. Consequently, the price-to-hash rate ratios for both cryptocurrencies were remarkably similar. This led to instances where the mining profitability of Ethereum and Ethereum Classic could align, highlighting the fluid nature of mining profitability.

The emergence of Ethereum 2.0 and the resulting shift of miners to Ethereum Classic significantly elevated this profitability threshold. The wave of miners increased competition within the Ethereum Classic network, underscoring the importance of hardware efficiency to remain profitable and subsequently driving up the profitability threshold. This shift illustrates that mining profitability is not solely dictated by the cryptocurrency's price - the network's hash rate plays a crucial role. When two cryptocurrencies offer identical mining rewards but differ in network hash rates, the miner's potential share in the network hash rate changes. This results in a dynamic perception of cryptocurrency mining profitability, emphasizing the complex interplay between network hash rate and mining reward.

This narrative indicates that the migration of miners was not a random event but a strategic decision.

Switching to Ethereum Classic was feasible only for miners whose hardware efficiency matched or exceeded the heightened profitability threshold. The shift was mainly driven by miners who could mine Ethereum Classic at costs equal to or lower than the anticipated mining revenue. For miners with less efficient hardware or higher operating costs, the economic feasibility of this transition was low, prompting them to seek alternative solutions.

Intriguingly, the newly elevated profitability threshold for Ethereum Classic has risen to such an extent that GPU mining devices no longer possess the capacity to mine it profitably. Only ASIC machines, designed for specialized tasks and unable to perform others without significant recalibrations, have enough efficiency to undertake this job. Despite the eventual utility of such specialization, it's critical to underscore that an ASIC machine designed for Etash with an energy efficiency lower than 1GH/J would be deemed worthless. Such a device would have no other practical use without undergoing substantial adjustments, highlighting the precarious balance between efficiency and versatility in cryptocurrency mining hardware.

3.7 Bringing It All Together

Having embarked on this in-depth exploration of the dynamics and intricacies of cryptocurrency mining, we can now synthesize our findings to form a coherent narrative. Firstly, our understanding of consensus mechanisms, specifically Ethash, enabled us to see the potential for miners to transition from one network to another. This understanding laid the foundation for observing the actual transition that ensued after the launch of Ethereum 2.0. It became evident that miners overwhelmingly shifted their operations to Ethereum Classic, while other cryptocurrencies remained relatively untouched by this migration, highlighting their marginal role in the mining landscape.

Our discussion then ventured into profitability thresholds, crucial in determining who could make this transition. We found that on the day of the merge, out of the initial 825TH/s of Ethereum, 75TH/s migrated to Ethereum Classic. Rationality dictates that all miners involved in this shift had to possess hardware that was at least as efficient as the newly elevated profitability threshold for Ethereum Classic. This leads us to the final piece of our analytical puzzle: the pre-merge energy consumption of Ethereum. Our preceding exploration allows us to leverage our understanding of the profitability threshold, miner behavior, and network dynamics to make well-founded estimates of Ethereum's energy consumption before the merge.

4 Empirical Results

4.1 Ethereum Pre-Merge

4.1.1 Top-Down Approach

In this section, we hypothesize that 75 TH/s of mining capacity meets the new profitability threshold for Ethereum Classic in terms of efficiency. The remaining 750 TH/s of mining capacity is operating at the efficiency level of the older Ethereum profitability threshold. This stipulation sets an upper boundary on energy consumption, proceeding from the assumption that most miners employ less efficient hardware and consume more energy per unit of hash rate. The equation for calculating energy consumption under these conditions can be articulated as follows:

$$\textbf{Power Consumption} = (75[\frac{TH}{s}] \div 1[\frac{MH}{J}] + 750[\frac{TH}{s}] \div 0.1[\frac{MH}{J}]) = 7.58[GW]$$

Despite deploying the same core methodology, this calculation offers a reduced estimate compared to de Vries's upper limit of 9.21 GW. Like de Vries, we employ a break-even analysis to estimate the maximum energy consumption that would not result in a financial loss. However, our calculation results in a lower figure, contributing to reconciling the differing estimates. Assuming that miners allocate roughly 49.53 percent of their mining revenue to cover electricity costs, they must operate at approximately double the efficiency. This adjustment alters the energy consumption estimate as follows:

$$\textbf{Power Consumption} = (75[\frac{TH}{s}] \div 2[\frac{MH}{J}] + 750[\frac{TH}{s}] \div 0.2[\frac{MH}{J}]) = 3.83[GW]$$

It's crucial to acknowledge that variations in electricity prices can significantly influence profitability thresholds, thereby leading to altered power consumption estimates. Inherent uncertainty exists concerning electricity costs and the percentage of mining revenue allocated toward electricity expenses. Moreover, during periods of substantial volatility in the price of the mined cryptocurrency, this percentage can undergo dramatic shifts within mere days. Hence, it's paramount to consider these variables' fluidity when evaluating power consumption estimates in cryptocurrency mining.

4.1.2 Bottom-Up Approach

Switching to a refined bottom-up perspective, we aim to ascertain the average efficiencies within these two subgroups. We employ data from the Cambridge Center for Alternative Finance and incorporate overheads and grid losses as calculated by Kyle McDonald (McDonald (2021)). This approach yields an average mining device efficiency of 1.62 MH/J for the group that migrated to Ethereum Classic and 0.35 MH/J for the remaining group. Consequently, we derive the following energy consumption estimate:

$$\textbf{Energy Consumption} = (55[\frac{TH}{s}] \div 1.62[\frac{MH}{J}] + 770[\frac{TH}{s}] \div 0.35[\frac{MH}{J}]) \times 1.33 = 2.96[GW]$$

Intriguingly, our estimate surpasses that of Kyle McDonald. This upward adjustment primarily stems from adopting a lower efficiency value for the remainder of the Ethereum network. Our calculation relies on an average efficiency derived from the potential hardware list the Cambridge Centre for Alternative Finance provided. Given the absence of concrete data on how the hash rate is distributed across different hardware types, this method offers the most reasonable approach. This reasoned, yet different approach to hardware efficiency leads to our estimate diverging from McDonald's.

In summary, both the top-down and bottom-up methods provide valuable perspectives on the energy consumption of Ethereum. The top-down method offers an upper limit for consumption, highlighting the influence of hardware efficiency on overall power usage. In contrast, the bottom-up approach provides a more granular view by accounting for individual efficiencies of different miner groups. Ultimately, both approaches point to the pivotal role of efficiency in determining energy consumption.

4.2 Ethereum Classic Pre- and Post-Merge

Ethereum Classic provides fascinating insights when examined through the lenses of both top-down and bottom-up methodologies. Let's first delve into the top-down perspective. By assuming an electricity price of 0.1 USD per kWh, we can trace the evolution of power consumption, represented by the blue line in the chart. Our calculations yield an estimated power consumption of approximately 0.125 GW with these parameters. It's worth noting that energy consumption is intrinsically tied

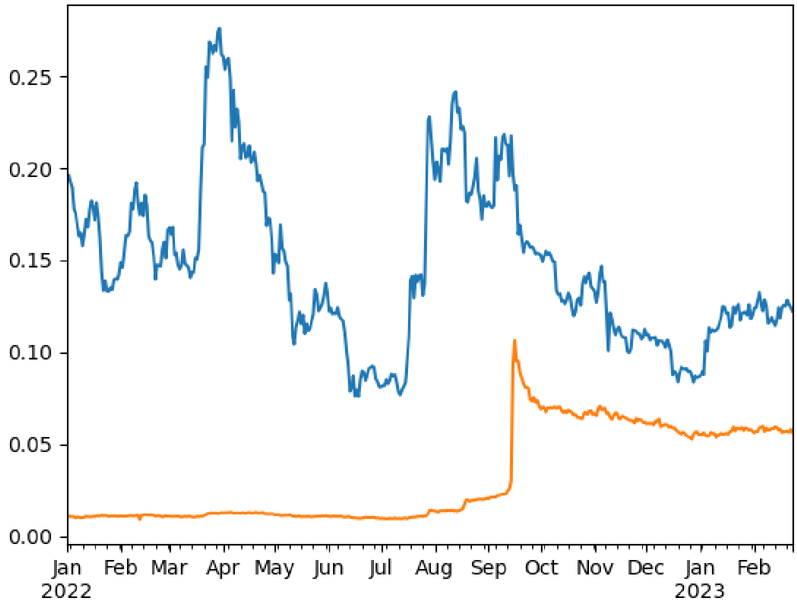

Figure 6: Ethereum Classic's Power Consumption

to price fluctuations. For instance, the decline in price towards the end of 2022 is paralleled by a corresponding decrease in energy consumption, as inferred from this calculation method.

Contrastingly, the lower limit, indicated by the blue line, is based on calculations incorporating the hash rate and the efficiency of the most advanced hardware available. This approach generates an estimate of around 0.065 GW. Interestingly, there appears to be a convergence between the upper and lower limits, suggesting a narrowing range for potential power consumption. It becomes clear that the hash rate directly influences the estimate in this scenario. Consequently, when the hash rate experienced a significant surge in September 2022, there was a corresponding uptick in energy consumption.

5 Policy Discussions

Cryptocurrency mining, due to the implementation of Proof-of-Work (PoW) algorithms, has come under scrutiny due to the environmental impact of its high energy consumption. These concerns are compounded by the popularity of cryptocurrencies such as Bitcoin and Ethereum, which have led to increased global mining activities. This issue is relevant to the discourse on the sustainability of the PoW mechanism and has significant implications for climate policy, energy policy, and the future of the digital economy.

On the one hand, cryptocurrencies and blockchain technology have brought forth innovative solutions for building trust in decentralized networks and optimizing processes across various industries. The transformative potential of blockchain technology is demonstrated in areas like supply chain finance, smart contracts, international trade, and manufacturing operations. Additionally, national digital currencies based on blockchain technology, such as the Digital Currency Electronic Payment (DCEP) in China, are being explored.

On the other hand, the environmental impact of cryptocurrencies, especially those using the PoW mechanism, poses a significant challenge. The energy consumption and carbon emission patterns of PoW mining activities are high and are increasingly contributing to global carbon emissions. This is concerning, given the urgent need to reduce carbon emissions and mitigate the impacts of climate change. We need to carefully evaluate the trade-off between the benefits of blockchain technology and the environmental costs associated with its current predominant consensus mechanism, PoW.

Furthermore, alternative consensus mechanisms such as Proof-of-Stake (PoS) or Proof-of-Authority (PoA) can significantly reduce energy consumption but also bring challenges. For instance, PoS and PoA, while being more energy-efficient, are often criticized for their potential to lead to increased centralization of network control. This is because they allocate more power to entities that hold a larger stake in the network, which could potentially undermine the decentralization principle that underlies cryptocurrencies. Therefore, any shift towards these alternatives must be handled carefully to maintain the decentralizing ethos of cryptocurrencies while mitigating their environmental impact.

Policy discussions should therefore focus not only on encouraging the shift towards these less energy-intensive consensus mechanisms but also on finding ways to mitigate potential centralization issues associated with them. This could be achieved through regulatory measures, incentives for using greener alternatives, and stringent environmental standards for mining activities. Regulations that guide the site of mining operations towards regions with cleaner energy sources can also be effective. As the energy mix differs significantly among regions, policy measures can be used to discourage mining in areas heavily reliant on fossil fuels, thus reducing the overall carbon emissions associated with mining activities.

The introduction of carbon pricing mechanisms can also be considered to account for the environmental costs of mining activities. However, this measure should be accompanied by the development of renewable energy sources to reduce the dependence of miners on fossil fuels. It's essential to consider that, while policy recommendations can be made, the decentralized nature of blockchain-based systems and their global distribution can pose challenges to policy enforcement. The regulatory environment needs to be coordinated internationally to ensure that these policy measures are effective and can curb the environmental impact of cryptocurrency mining.

In conclusion, as the adoption of blockchain technology continues to grow, it's paramount to ensure its development is sustainable. Future research should aim to continue exploring and developing energy-efficient consensus mechanisms, while policy measures should encourage their adoption and the transition toward cleaner energy sources for mining activities.

6 Conclusion

The advent of cryptocurrencies and their associated technologies have presented us with novel challenges in understanding and quantifying their energy consumption. For Proof-of-Work (PoW) cryptocurrencies, the methodologies currently employed can be categorized broadly into two approaches - top-down and bottom-up. Each of these methodologies has its merits and demerits. The top-down approach provides a broader perspective, considering the entire network, while the bottom-up approach enables a granular assessment based on individual miners. Notwithstanding their advantages, a significant discrepancy persists between these two methodologies, underscoring the complexity inherent in accurately estimating the energy demands of PoW cryptocurrencies.

In the present study, we tried reconciling this divergence by introducing a fresh estimation strategy for the Ethereum network. Our methodology incorporates domain expertise, and profitability thresholds, providing a nuanced perspective to refine our estimates. Before the merge, our estimate for the Ethereum network's energy consumption ranged between 2.96 and 3.83 GW. Following the merge, the potential energy consumption for Ethereum Classic is projected to be between 0.065 and 0.125 GW. Interestingly, when comparing pre and post-merge states, Ethereum Classic's energy consumption has experienced a modest increase, not exceeding 0.05 GW. This minor increase reaffirms the complexities inherent in accurately estimating the energy demands of PoW cryptocurrencies while underscoring the potentially beneficial impact of shifting to more energy-efficient consensus mechanisms.

Future research could consider a range of fascinating topics. One potential area of exploration could be an analysis of the hash rate's convergence rate to the maximum hash rate predicted by profitability threshold analysis. This could provide insights into how swiftly the mining community adjusts to changing profitability landscapes. Moreover, a hypothetical investigation into the implications of Bitcoin switching to PoS could also yield significant findings. Given Bitcoin's position as the leading cryptocurrency, such a shift would substantially affect the cryptocurrency industry's energy consumption patterns.

References

Buterin, V. (2020, November). *Why proof of stake.* Blog post. Retrieved from https://vitalik.ca/general/2020/11/06/pos2020.html

Criddle, C. (2021, 2 10). *Bitcoin consumes 'more electricity than argentina'.* BBC. Retrieved from https://www.bbc.com/news/technology-56012952

de Vries, A. (2018). Bitcoin's growing energy problem. *Joule*, *2*(5), 801-805. Retrieved from https://www.sciencedirect.com/science/article/pii/S2542435118301776

de Vries, A. (2023a). Cryptocurrencies on the road to sustainability: Ethereum paving the way for bitcoin. *Patterns*, *4*(1), 100633. Retrieved from https://www.sciencedirect.com/science/article/pii/S2666389922002653

de Vries, A. (2023b). *Ethereum energy consumption index.* Retrieved from https://digiconomist.net/ethereum-energy-consumption (May 11, 2023)

Fairley, P. (2019, 01). Ethereum will cut back its absurd energy use. *IEEE Spectrum*, *56*, 29-32. Retrieved from https://ieeexplore.ieee.org/document/8594790

Hayes, A. (2015). Cryptocurrency value formation: An empirical analysis leading to a cost of production model for valuing bitcoin. In *Mcis 2015 proceedings.* Retrieved from http://aisel.aisnet.org/mcis2015/4

Huang, J., O'Neill, C., & Tabuchi, H. (2021, 9 3). *Bitcoin uses more electricity than many countries. how is that possible?* NYT. Retrieved from https://www.nytimes.com/interactive/2021/09/03/climate/bitcoin-carbon-footprint-electricity.html

Härdle, W. K., Harvey, C. R., & Reule, R. C. G. (2020, 02). Understanding Cryptocurrencies. *Journal of Financial Econometrics*, *18*(2), 181-208. Retrieved from https://doi.org/10.1093/jjfinec/nbz033

Krause, M. J., & Tolaymat, T. (2018). Quantification of energy and carbon costs for mining cryptocurrencies. *Nature Sustainability*, *1*, 711–718. Retrieved from https://doi.org/10.1038/s41893-018-0152-7

McDonald, K. (2021). Ethereum emissions: A bottom-up estimate. *CoRR*, *abs/2112.01238*. Retrieved from https://arxiv.org/abs/2112.01238

Mehar, M. I., Shier, C. L., Giambattista, A., Gong, E., Fletcher, G., Sanayhie, R., . . . Laskowski, M. (2019). Understanding a revolutionary and flawed grand experiment in blockchain: The dao attack. *Journal of Cases on Information Technology (JCIT)*, *21*(1). Retrieved from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3014782

Mora, C., Rollins, R. L., Taladay, K., et al. (2018). Bitcoin emissions alone could push global warming above 2°c. *Nature Climate Change*, *8*, 931–933. Retrieved from https://doi.org/10.1038/s41558-018-0321-8

Nagarajan, S. (2022, August 19). Alphabet, samsung among 40 big companies investing in blockchain, study finds. Retrieved from https://blockworks.co/news/alphabet-poured-1-5b-into-blockchain-companies-over-10-months-study-finds

O'Dwyer, K. J., & Malone, D. (2014). Bitcoin mining and its energy footprint. In *25th iet irish signals systems conference 2014 and 2014 china-ireland international conference on information and communications technologies* (p. 280-285). Retrieved from https://ieeexplore.ieee.org/document/6912770

Orender, J., Mukkamala, R., & Zubair, M. (2020). Is ethereum's progpow asic resistant? In *Proceedings of the 6th international conference on information systems security and privacy (icissp 2020)* (pp. 310–316). (https://www.scitepress.org/Papers/2020/89092/89092.pdf)

Sedlmeir, J., Buhl, H. U., Fridgen, G., & Keller, R. (2020a, nov). Ein blick auf aktuelle entwicklungen bei blockchains und deren auswirkungen auf den energieverbrauch. *Informatik Spektrum*, *43*(6), 391–404. Retrieved from https://doi.org/10.1007/s00287-020-01321-z

Sedlmeir, J., Buhl, H. U., Fridgen, G., & Keller, R. (2020b). The energy consumption of blockchain technology: Beyond myth. *Bus Inf Syst Eng*, *62*(6), 599–608. Retrieved from https://doi.org/10.1007/s12599-020-00656-x

Trimborn, S., & Härdle, W. K. (2018). Crix an index for cryptocurrencies. *Journal of Empirical Finance*, *49*, 107-122. Retrieved from https://www.sciencedirect.com/science/article/pii/S0927539818300616

Truby, J. (2018). Decarbonizing bitcoin: Law and policy choices for reducing the energy consumption of blockchain technologies and digital currencies. *Energy Research Social Science*, *44*, 399-410. Retrieved from https://www.sciencedirect.com/science/article/pii/S2214629618301750