12 Angry Miners

Aryaz Eghbali and Roger Wattenhofer

ETH Zurich, Switzerland {aeghbali,wattenhofer}@ethz.ch
https://disco.ethz.ch/

Abstract. In this paper we investigate the behavior of miners, in terms of which type of hardware they use, based on publicly-available macro-scale data of Bitcoin. We provide a model for the market share of mining hardware, which is then used to estimate the energy consumption, the distribution of electricity price among Bitcoin miners, and the total investment in the backbone of the Bitcoin network.

Keywords: Bitcoin · Data Analysis · Mining Hardware.

1 Introduction

Raw materials such as minerals, metals, coal, oil, gas, or gems are the foundation of the world's economy. Mining these materials is big business, with major investments. The raw material mining companies are somewhat secretive about their mining equipment and investments.

Cryptocurrency mining shares some of this secrecy. Cryptominers do not publish how much they invest in mining equipment or how much they actually mine. Mining pools advertise their hashing power to attract more miners, but their self-announced numbers should be handled with care, as there are various incentives to claim both higher or lower hash rates.

However, in contrast to raw materials, a lot of cryptomining information is available in the public cryptocurrency blockchain. In this paper, we ask to what degree we can understand the cryptomining business by analyzing public information. What mix of mining equipment is being used? How much energy is spent, and at what cost? When is it no longer profitable to mine with outdated equipment, and when does it become profitable again? In our paper, we focus on the largest cryptocurrency, Bitcoin.

Bitcoin mining became a serious business already in the early years of Bitcoin's inception. It also created businesses like ASIC manufacturing, mining pools, and cloud mining around it. Although the hashing power and mining hardware are integral parts of this ecosystem, since not much data has been publicly released by miners, hardware manufacturers, or mining pools, data-driven studies focusing on the hardware market have been difficult and rare.

On the other hand, the money invested in any cryptocurrency shows how interested people are in its success. Many have argued that people who have stake in Bitcoin have the incentive to keep it alive. The market capitalization shows a part of the total stakes, however, another important, if not the most important, group of stakeholders are the miners. Yet, we do not know how much stake they have in this business.

2 Related Work

Analyzing Bitcoin's or other altcoins' blockchain data has been the focus of various recent work. In [7] Decker et al. studied the propagation of messages and their delays on the Bitcoin network. The graph of transactions in Bitcoin was extracted by Ron et al. in [10]. Reid et al. showed that Bitcoin does not preserve anonymity [9]. Anderson et al. [3] did some empirical analysis on Ethereum, Namecoin, and Peercoin. Bartoletti et al. [4] analyzed the Ethereum blockchain to detect Ponzi schemes. Cong et al. [6] investigated the mining pools. Recently, Ma, Gans, and Tourky [8] analyzed the equilibria of the Bitcoin mining game.

Two simple approaches have been used to estimate the energy consumption of Bitcoin. These estimations were popular when the Bitcoin price exploded in 2017, with a lot of press coverage. One approach is to divide the total hashing power by the hash rate of the most efficient (or average) mining hardware. Another approach is to take the value of mined Bitcoins in a period of time, and divide it by either the cheapest (or average) electricity cost of world region where most of the mining actually happens. These approaches showed that the mining energy consumption of Bitcoin is significant ("in the order of a medium European country"). Since our approach models the share of hashing power by each device, we can have more accurate estimations of the energy used.

In a series of reports, Bendiksen et al. [5] investigated the electricity consumption of Bitcoin and its distribution and types of sources that it come from. Their work presents good insight into the energy aspect of Bitcoin, using a journalistic approach. They gather data from many different sources and verify them, while using expert knowledge when the data is not credible enough. The results that we present in this paper can be used in similar reports as another source of information about another aspect of this market.

3 Data

Our Bitcoin blockchain data is fetched from blockchain.com [1]. We use the price, hash rate, transaction fees, and difficulty, which are reported daily. The hash rate h is calculated from the number of blocks found, n, in a day by the following formula:

$$h = \frac{n}{144} \cdot d \cdot \frac{2^{32}}{600},$$

where d is the difficulty. The process of mining blocks is a Poisson random process and hence the number of blocks mined in a day might change even if the hashing power has remained constant. So, it is important to distinguish the perceived hash rate from the actual hash rate. We use the smoothed hash rate as the actual hash rate of the network.

The data for mining hardware was extracted by a manual process. We consider the S series (S1-S11) of AntMiner by Bitmain, which became the dominating company to produce ASIC based mining hardware. We also consider one sample hardware from each generation of the pre-ASIC era, i.e. a CPU (Intel Core i7 920), a GPU (Nvidia GTX 460), and an FPGA (BitForce SHA256 Single). The specifications of devices used can be found in Table 1. The hashing power, and electricity usage of each device comes from what the manufacturer advertised, and the prices and release dates are either from the manufacturer or the main seller. The release date of the GPU is chosen as the first date which there is evidence of GPU mining happening in Bitcoin. We also added 30 days to the release dates of the Antminers before they can be used to account for the delivery times.

Other mining hardware, like Avalon miners, DragonMint miners, and even other CPUs and GPUs, have also been available on the market and used in mining. Nonetheless, we wanted to select a sample of devices that can be representative of all mining equipment, and avoid the factors of marketing and sales that affect the use of similar devices. So wherever our data refers to a particular mining device, the reader can imagine any device with similar specifications.

Table 1: Mining hardware that we consider and their specifications.

Device	Release date	Hash rate (TH/s)	Power consumption (KJ/TH)	Price (USD)
CPU	2008-11-01	0.0000192	4166.67	305.0
GPU	2010-07-18	0.00006831	2341.92	229.0
FPGA	2011-06-01	0.000832	96.1538	599.0
S1	2013-09-21	0.18	2.0	299.0
S2	2014-04-01	1.0	1.1	2259.0
S3	2014-07-10	0.4	0.77	382.0
S4	2014-10-16	2.0	0.7	1400.0
S5	2014-12-27	1.155	0.51	370.0
S5+	2015-08-14	7.722	0.44	2307.0
S7	2015-09-30	4.73	0.25	479.95
S9	2016-06-12	13.5	0.098	1987.95
S11	2018-11-01	20.5	0.075	1173.38

4 Model

Cryptomining is a serious industry, and almost all miners are participating because of monetary incentives, namely block rewards and transaction fees. Since buying and running mining hardware is expensive, people that mine because of other reasons are few, without considerable hashing power. So we assume all miners to be rational agents, i.e. they want to maximize their profit. To gain more profit, a miner might want to increase or decrease its hash rate based on

Bitcoin's current or future price, its electricity costs, in relation to the total hash rate of the network. The actions an agent can take are to either turn off some running hardware, buy hardware, or turn on some hardware that is currently off. We will now explain how the model works in detail.

Starting from the first day of Bitcoin we calculate the hash rate share of each device on a daily basis from the total hash rate. Each day that the hash rate has increased from the previous day, some miners have either turned on some unused hardware, or they have bought new mining equipment. If the hash rate has decreased, then they must have turned off some of their hardware.

Case 1: Hash rate has decreased. From all currently running devices we start with the lowest performing one (in terms of both the power consumption and hash rate), which in our list is the oldest one. We decrease its hash rate share until it reaches zero hash rate and then move to the next lowest performing device, or the amount of decrease is smaller than the hash rate of that device, where in this case we move to the next day after decreasing its hash rate, such that the sum of hash rates matches the hash rate of that day. As long as a device is making profit, the rational decision is to keep it running. So when a mining hardware is turned off, it means that its profitability just dropped below zero. Knowing the profit is zero and the revenue of mining from the price and hash rate data, we can calculate the running cost in this situation, which is almost completely the electricity cost. Extracting the electricity price from this model is a result of the novelty of our approach. The electricity price is calculated by dividing the revenue per hash rate of that specific device by the hash rate deducted from that device. We have the revenue per hash rate for each day in the history of Bitcoin by dividing the total value of Bitcoins mined in that day by the hash rate in the same day. The reduction in hash rate is saved as "turned off" to be used in the next case. We also save the electricity price for each hardware that has been turned off.

Case 2: Hash rate has increased. We start by looking at the currently turned off hardware ordered from highest to lowest performing. We perform the following routine until the increase in total hash rate is justified in our model. If buying the newest available hardware on the market breaks even on electricity costs alone in less than three months, then we assume the miner buys the new device (when considering the break-even time based on the electricity costs alone for various mining hardware there is a large gap, which justifies this assumption). Otherwise, if the new hardware breaks even in a longer period of time, the miner turns on the highest performing devices that she has. Calculating the break-even point requires the knowledge of electricity price of the miner. We know the electricity prices by saving the average electricity cost of miners that use this particular device when they are turned off in the previous case. If the increase in hash rate is still higher than what we already considered, then we assume new hardware was bought. Similar to the previous case, the rational reason to turn on or buy hardware is that it is profitable. So we can again compute the electricity prices. However, these electricity prices that we calculate are not exact but bounds. So when we calculate electricity prices in Case 1, they are the lower bounds for the electricity prices of those miners, and in Case 2 they are the upper bounds. The reason that these are not exact prices is that in assuming agents to be rational we only consider current profits but not price and market speculations.

The biggest limitation of our model is that since the granularity of our data is limited to a day, we cannot capture changes that happen with higher frequency. However, as the nature of Bitcoin mining has randomness (the number of blocks mined in a day follows a Poisson distribution), shorter time frames cannot provide data with good accuracy. Even the daily data has fluctuations in hash rate caused by randomness. To avoid the error caused by these random fluctuations, we smooth the hash rate by taking a weighted average of neighboring points.

Using the total hash rate we cannot detect some changes that happen at the same time, such as when some mining device breaks down and somewhere else a device is turned on. These two events would cancel out each other and not appear in our analysis.

Our model inherently considers cases that are not mentioned above. For example, in recent years there have been some other cryptocurrencies that could be mined using the same hardware as Bitcoin's and some miners "coin hop", which is mining the most profitable coin and frequently switching between them. We incorporate this behaviour in our model by taking the Bitcoin Cash (BCH) data (total hash rate and price) and assuming when it is more profitable to mine Bitcoin Cash, the hash rate on Bitcoin decreases, and the hash rate on Bitcoin Cash increases, miners are moving from Bitcoin to Bitcoin Cash, and the other way around. The reason that we chose Bitcoin Cash is that it is the most popular cryptocurrency that uses the same proof-of-work as Bitcoin. Moreover, other altcoins with the same proof-of-work do not have noticeable hash rates and we do not have as much data as we need for them. Another case is when a wealthy miner which already has a mining farm running on miner x decides to buy miner y because it is more profitable. This miner might want to sell his x miners and buy y miners. Note that since xs are still profitable, not using them is not rational. Hence, this scenario is also captured by our model, as those x miners are considered to still be running (by another miner).

5 Analysis

Figure 1 shows our main result, which is the market share of each mining hardware through time, in logarithmic scale.

From the hash rate share we immediately get the power consumption by device, which is shown in Fig. 2 and compared to [2]. Although both estimates follow the same trend, the difference is the methods used. We use the hashing power of running devices to calculate how much electricity they use. In [2] the price of Bitcoin is used to estimate the energy costs, and by assuming an average of 5 cents per KWh, the total energy consumption is achieved. It is important to note that the hash rate and price of Bitcoin are not always synchronized. There have been times where the hash rate lagged behind price, and there have

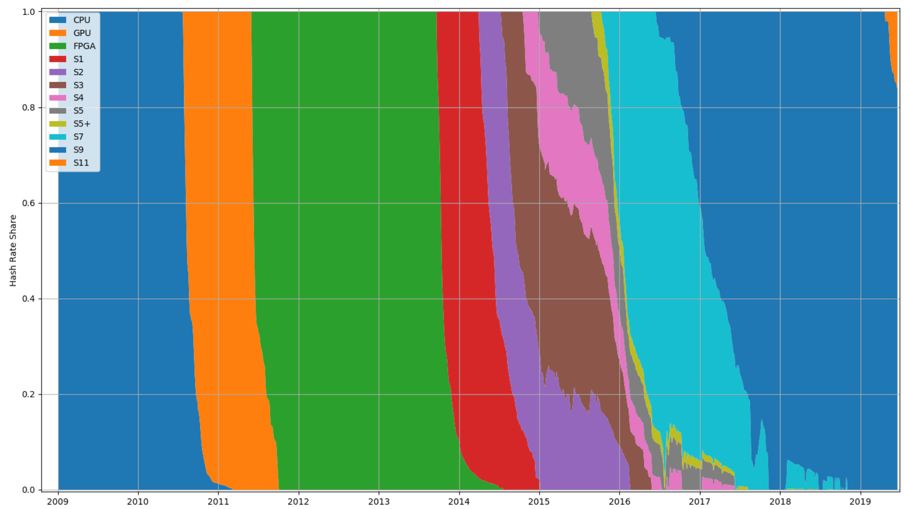

Fig. 1: Market share for each mining hardware.

been times where the price lagged behind the hash rate. We refer the reader to Bitcoin's charts on [1].

Since we have the price of Bitcoin and the total hash rate, we can compute how much revenue each unit of hashing power makes in each day. As mentioned in Section 4, turning on a device or buying new hardware shows that it just became profitable to mine using that device or hardware, and turning a device off shows that it just became unprofitable to mine. Hence, we can calculate the points in time where the profitability changed, and what the electricity price was for those points. The distribution of electricity prices for miners who turned off their devices can be seen in Fig. 3.

We can see when each of the 12 miners were profitable and when they became unprofitable in Fig. 4.

6 Conclusion

In this paper we present an empirical model for the market share of mining hardware. Using this model we inferred information about the energy consumption, which we found is 0.14% of the global electricity usage, the profitable times for each mining equipment, and the distribution of electricity prices for Bitcoin miners, that matches many reports and claims. We further derived the amount of money invested in the mining sector of Bitcoin, which today is more than 7

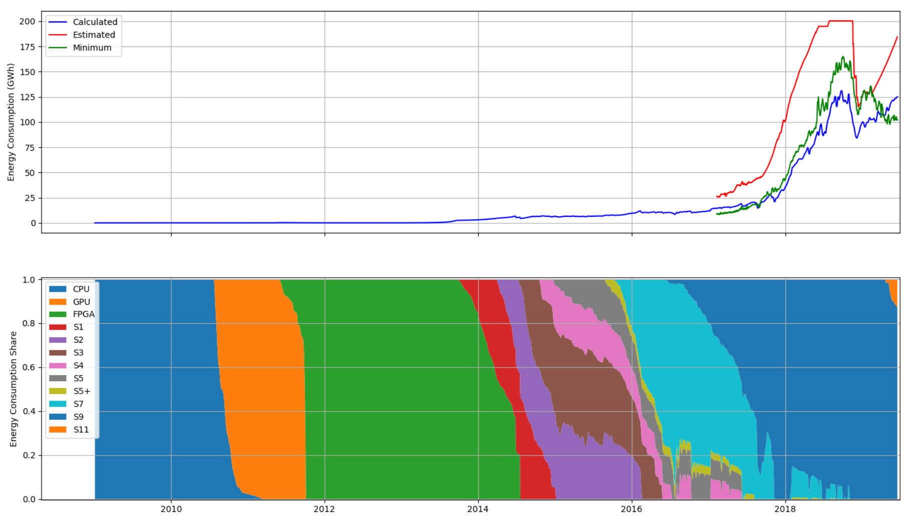

Fig. 2: The top plot shows the daily total energy consumption of Bitcoin mining in comparison to [2]. The "Calculated" line is our result, and the other two are from [2]. The bottom plot shows the daily electricity consumption share by mining device.

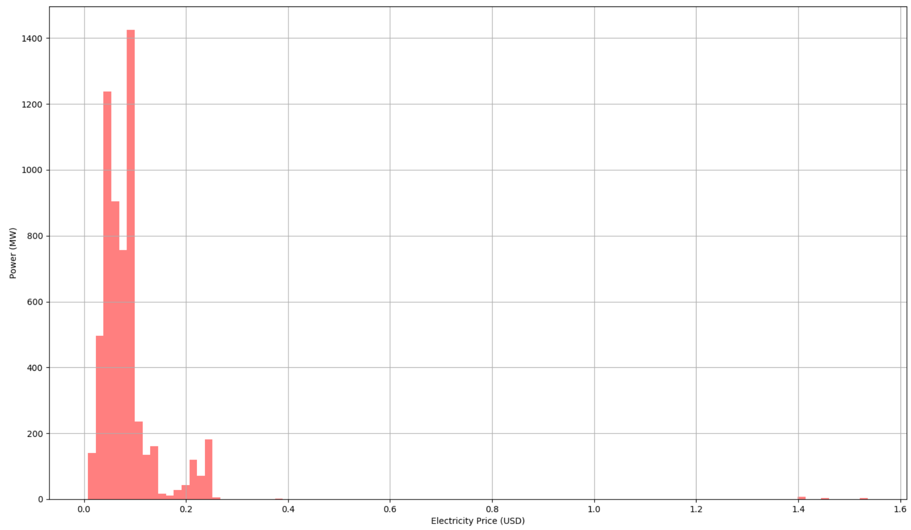

Fig. 3: The distribution of electricity prices for miners who turned off their device at some point. The distribution is weighted by the amount of energy used.

billion US dollars, and throughout these past 10 years has summed up to more than 9 billion US dollars.

Our model is implemented in Python and open sourced on Github[1], so that interested researchers can replicate our results, and use them to gain further insight into this market and potentially others. A nice extension of this work would be to implement a web-based tool that allows users to modify parameters and assumptions of our model and see the effects.

References

1. https://www.blockchain.com/charts/, accessed on June 20, 2019
2. https://digiconomist.net/bitcoin-energy-consumption, accessed on June 20, 2019
3. Anderson, L., Holz, R., Ponomarev, A., Rimba, P., Weber, I.: New kids on the block: an analysis of modern blockchains. arXiv preprint arXiv:1606.06530 (2016)
4. Bartoletti, M., Carta, S., Cimoli, T., Saia, R.: Dissecting ponzi schemes on ethereum: identification, analysis, and impact. arXiv preprint arXiv:1703.03779 (2017)
5. Bendiksen, C., Gibbons, S.: The bitcoin mining network-trends, marginal creation cost, electricity consumption & sources. CoinShares Research **21** (2018)

[1] https://github.com/EAAL/Bitcoin

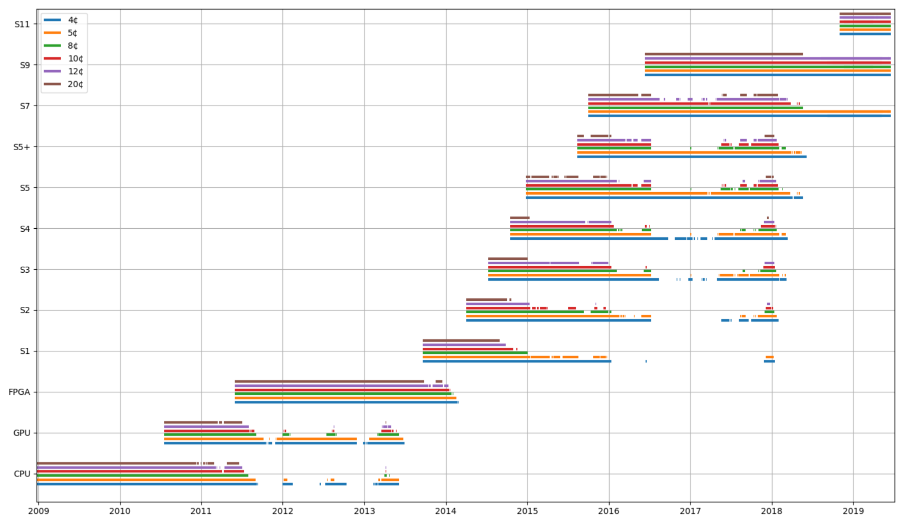

Fig. 4: The profitable period for each mining hardware is shown by colored lines (each cluster of lines represents one mining device). Different colors show different electricity prices.

6. Cong, L.W., He, Z., Li, J.: Decentralized mining in centralized pools. Tech. rep., National Bureau of Economic Research (2019)
7. Decker, C., Wattenhofer, R.: Information propagation in the bitcoin network. In: IEEE P2P 2013 Proceedings. pp. 1–10. IEEE (2013)
8. Ma, J., Gans, J.S., Tourky, R.: Market structure in bitcoin mining. Tech. rep., National Bureau of Economic Research (2018)
9. Reid, F., Harrigan, M.: An analysis of anonymity in the bitcoin system. In: Security and privacy in social networks, pp. 197–223. Springer (2013)
10. Ron, D., Shamir, A.: Quantitative analysis of the full bitcoin transaction graph. In: International Conference on Financial Cryptography and Data Security. pp. 6–24. Springer (2013)