CONSENSYS

REPORT

The Ethereum Proof of Work Mining Report

Analyzing the performance of Proof of Work mining on the Ethereum network, 2015 - 2020.

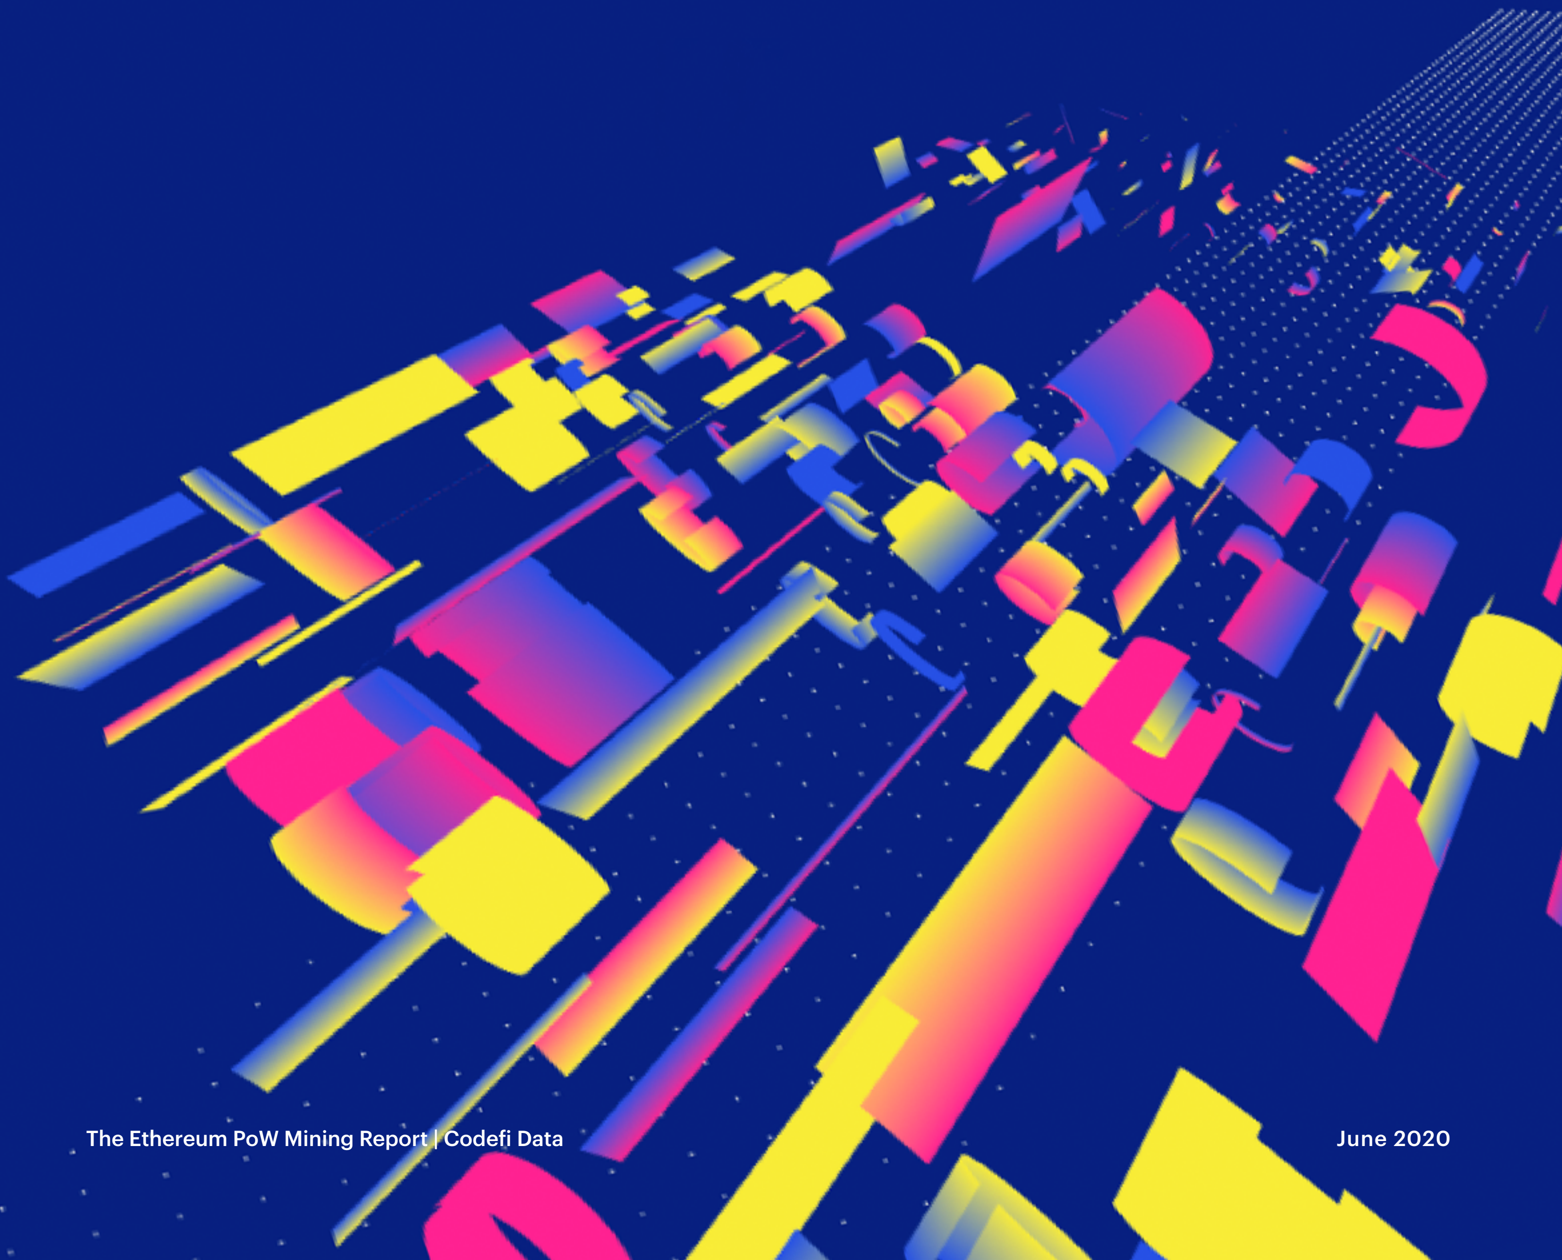

REPORT

Outline

REPORT

Executive Summary

This report is a data-driven overview of the role that mining, miners, and mining pools have played in the maintenance of Proof of Work (PoW) on the Ethereum blockchain from its launch in mid-2015 to Q2 2020. This report looks at the performance of PoW through crucial on- and off-chain ecosystem events, and identifies areas of strength and potential weakness in the past and the future.

The Proof of Work consensus mechanism has upheld the Ethereum network since its launch in 2015. Likely in 2020, Ethereum 2.0 will launch its first phase of development and the ecosystem will begin the gradual migration from Proof of Work to Proof of Stake. When Proof of Work retires, the Ethereum network will no longer have use for mining. Instead, the network will be maintained by a complex yet highly secure and scalable set of validators and staked ETH. As Ethereum 2.0 rolls out, however, the current PoW chain will continue to uphold the entire ecosystem.

Mining has traditionally been a tricky subject to analyze from a data perspective. Not all information is on-chain, and even the information that is on-chain is often difficult to understand or interpret. With robust data analytics tools offered by Codefi Data and a couple of core assumptions about the role of miners on the network (see "Research Scope"), we arrived at a set of conclusions about the Ethereum PoW mining ecosystem:

1. Mining pools are more heavily impacted by on-chain events rather than off-chain events. Payout addresses (miners) are more heavily impacted by off-chain events rather than on-chain events.
2. The overwhelming majority (>90%) of all blocks produced on Ethereum for the past two years have been mined by 56 known mining pools.
3. The overwhelming majority (>90%) of all block rewards won on Ethereum for the past two years have been attributed to 56 known mining pools.
4. The overwhelming majority (>80%) of Ethereum's hashrate for the past two years has been attributed to 56 known mining pools.
5. The second half of 2017 coincided with considerable consolidation of network activity among mining pools.
6. Twice in Ethereum's history, two mining pools have owned more than 51% of the network's hashrate.

Introduction

Proof of Work is the consensus mechanism that has upheld some of the most important and widely-adopted blockchain networks to date, including Bitcoin, Ethereum, and Litecoin. Since first being tested in 2009 with the release of Bitcoin, Proof of Work (PoW) has remained extremely resilient despite the security, growth, and performance issues that affect the adoption of most emerging technologies. In particular, PoW has proven a secure mechanism to support decentralized networks, which require the proactive participation of many unconnected and competing node operators.

Though secure for most widely-adopted blockchain networks (a notable exception being the 51% attack of ETC in January 2019), PoW has not proven sustainable for networks that have more complex utility beyond the transfer of wealth, such as Ethereum. There are, broadly speaking, three issues facing PoW networks: accessibility, centralization, and scalability.

Accessibility:
The barriers to entry to becoming a PoW miner are high. A miner must purchase, set up, and maintain all the necessary hardware to run a PoW mining rig. Additionally, PoW mining is extremely energy-intensive. Not only is the underlying mechanism inefficient from an energy standpoint, but it further increases the barrier to entry. To earn significant block rewards, it is better for a miner to live in a region with lower electricity costs. Additionally, jurisdictions often offer lower electricity costs to corporations, meaning a miner who wishes to maximize their profits would need to form a company and purchase enough mining hardware to offset the effort and associated costs. Altogether, energy inefficiency, variable electricity costs, hardware costs, and corporate electricity breaks all present significant barriers to entry for most would-be miners.

Scalability:
In the current Ethereum Proof of Work chain, each block is mined consecutively. Each block can only contain a certain amount of data, known as the block size. This means that if there are more pending transactions than can fit into a block, the transactions that do not make it into the next block to be mined must "wait" for the following block for another chance to be included. On Ethereum, a block is mined once every ~14 seconds, but during particularly high transaction events, some users could wait hours for their transactions to be processed.

Centralization:
Barriers to entry for mining can have the adverse secondary effect of greater centralization of miners. As it gets more costly and less profitable to become a miner, the network naturally sees a concentration of mining into two categories. First, large mining conglomerates that operate in areas with low electricity costs and cold weather (to reduce the cost of manually cooling mining hardware) such as Mongolia and Siberia. Second, mining power is centralized in the hands of mining pools. As it becomes less profitable for most people to mine individually, they buy hash power from a mining pool, which operates as a single mining entity. By the end of 2019, over 50% of blocks on Ethereum were mined by just two mining pools.

In 2020, the Ethereum community will take the first step towards Ethereum 2.0 with the launch of Phase 0. Phase 0 will implement a Proof of Stake consensus mechanism to underpin the Ethereum 2.0 blockchain. Eventually, the current PoW network will merge with the Ethereum 2.0 PoS network, and PoW will no longer be a major component of the Ethereum blockchain.

Research Scope

It is generally agreed upon that the switch from PoW to PoS will improve issues with accessibility, scalability, and centralization. As Ethereum 2.0 goes live with PoS, however, the existing PoW chain will continue to exist for likely another few years as the new network is tested and further developed. Not until Phase 1.5 will the PoW chain merge with the PoS chain. Until then, the rapidly-growing Ethereum ecosystem will continue to operate on the existing Proof of Work chain.

There are, therefore, two reasons for this data-driven analysis into the performance of the Proof of Work Ethereum blockchain:

1. As we continue to develop the Ethereum 2.0 PoS chain, we should have a firm understanding of the benefits and the shortcomings of the PoW consensus mechanism we aim to retire.
2. Because we will be dependent on the PoW chain for at least the near future, it is important for us to have as much objective data regarding the performance of the consensus mechanism that is supporting the majority of the Ethereum ecosystem.

Objective data about mining is notoriously difficult to come by. This report has been compiled by the Codefi Data team using publicly-available APIs and data analytics tools to arrive at the most objective data possible regarding fundamental components of Ethereum PoW. Where data can be interpreted differently, we will qualify the reasons for our interpretations in the body of the report or the appendix.

Through this report we aimed to answer the following questions:

1. How has Proof of Work performed over the years since Ethereum's launch in 2015?
2. How have major planned and unplanned ecosystem events affected core elements of Proof of Work, including hashrate, miner count, payouts, and block production.
3. How has the role of miners and mining pools evolved over the years?
4. How have Proof of Work mining block rewards been distributed, and how have they affected the balance of wealth in the Ethereum ecosystem?

Additionally, due to the occasional difficulty of arriving at clean data with respect to mining, we have viewed our data by adopting a few assumptions:

1. Every significant miner in recent years is a mining pool. Many of these miners can be named and definitely identified as mining pools (see the appendix), but the current state of the network makes it sufficiently uneconomical for single miners. This economic pressure is significant enough for us to write this report under the assumption that the overwhelming statistical majority of Ethereum miners are mining pools.
2. Mining pools send the block rewards they mine directly to the miners. Some mining pools may adopt the technique of sending block rewards to proxy addresses (typically smart contracts), which then manage the distribution of rewards to the appropriate miner addresses. However, there is no comprehensible way to distinguish the few mining pools that manage rewards via proxy, so we are writing this report under the assumption of "direct miner payouts."

Miners & Mining Pools: A Historical Count

Critical Ecosystem Events

Both miners and mining pools have upheld the Ethereum blockchain by mining blocks into existence and receiving ETH rewards as incentive. The only incentive for miners and mining pools to spend hash power on a certain blockchain is the opportunity to profit off of block rewards. Their chosen network, therefore, must contain the highest opportunity for profitability. This opportunity results from a combination of on-chain factors such as block rewards, hashrate, network security, and difficulty, and off-chain factors such as crypto price, electricity rates, and confidence in particular protocols. Over the course of Ethereum's history, there have been a number of significant on- and off-chain events that have impacted the participation of miners and mining pools. In particular, we marked the following:

JUNE 2016
The DAO hack, which resulted in the unplanned hard fork and the creation of Ethereum and Ethereum Classic.

OCTOBER 2017
The Byzantium upgrade, which reduced the block rewards from 5 ETH to 3 ETH.

JANUARY 2018
Record ETH prices at ~$1,400 USD.

DECEMBER 2018
Record ETH lows (in recent history) at <$90 USD (a 1500%+ fall in price in fewer than 12 months).

FEBRUARY 2019
The Constantinople upgrade, which reduced block rewards from 3 ETH to 2 ETH.

MARCH 2020
The month of significant market events, which saw the price of ETH drop by over 43% in one day from March 11 to March 12 due to worldwide economic skepticism around the COVID-19 pandemic.

Mining Pools and Payout Addresses, 2015 - 2020

Figure 1 below shows the number of miners and the number of payout addresses from the launch of the network in mid-2015 through April 2020.

Between mid-2015 and the peak of ETH price in January 2018, the number of individual payout addresses increased, whereas the number of miners decreased. It is highly likely that in 2015, the high number of miners is due to the fact that mining pools had not yet gained significant traction or scalability among the Ethereum ecosystem, contributing to a high number of individual miners on the network. The steep decline in miners in 2015 is likely due to the emergence of mining pools - leading to our assumption listed in the introduction that we consider the overwhelming statistical majority of miners on the Ethereum network are mining pools. The continued decrease in miners following 2015, as we will discuss, has been due to the consolidation of mining pools.

Following the DAO attack, the growth of individual payout addresses slowed down slightly compared to the pace of the first six months of 2016, but the miner count decreased notably during the last six months of the year. In the weeks and months following the DAO hack, there was considerable ecosystem debate over the validity of the unplanned hard fork. Some miners may have chosen to stick with the original Ethereum Classic chain rather than join the 'new' Ethereum network after the DAO.

Starting in 2017, the number of payout addresses begins to climb rapidly alongside the rising price of ETH during the same time period. Payout address count reaches its zenith, unsurprisingly, at the height of ETH price at ~$1,400 in January 2018. Even with the Byzantium upgrade in October 2017, which reduced block rewards from 5 ETH to 3 ETH, the impact on payout address count was small and short-lived. Immediately before Byzantium, figure 1 shows a small dip in the number of payout addresses, which quickly recovers alongside a continued increase in ETH price.

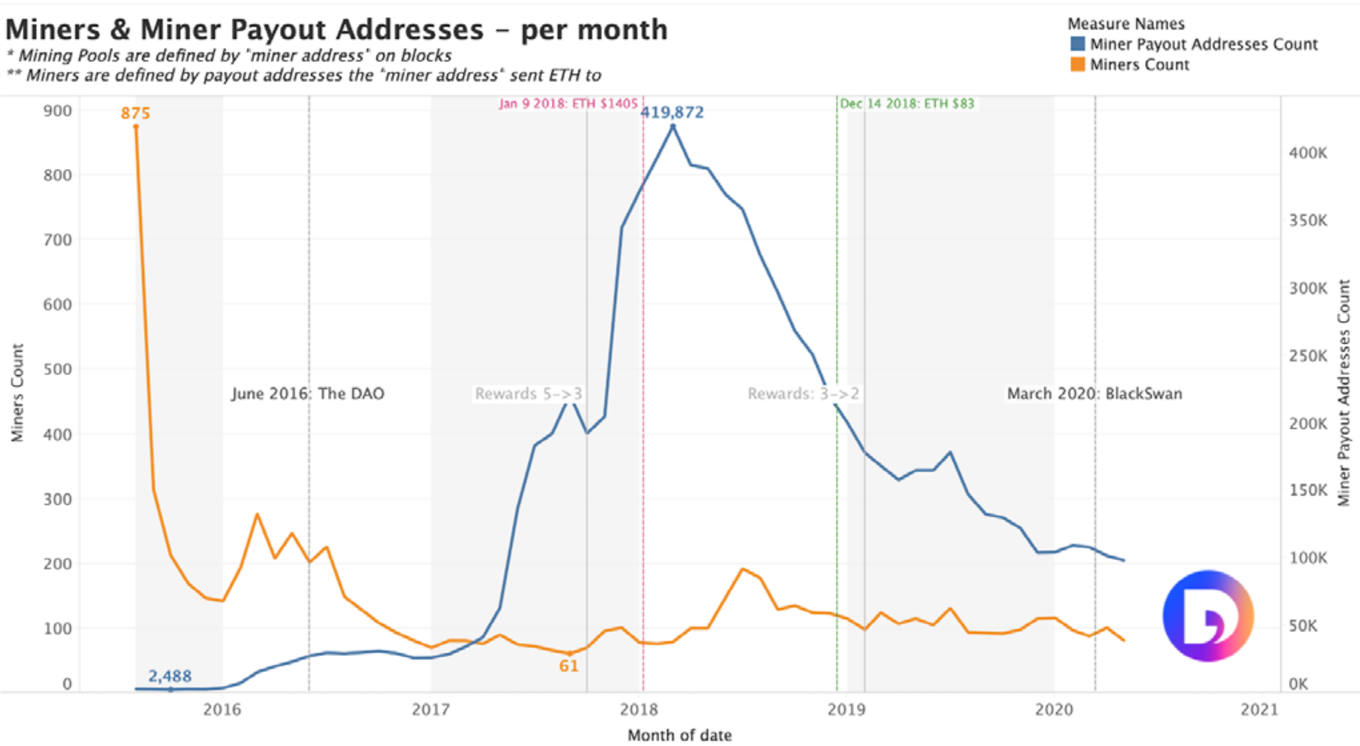

Figure 1: Miners and Payout Addresses on Ethereum, 2015 - 2020. Measured by month.

Interestingly, miner count declines in the year following the DAO during the same time frame when payout addresses were steadily and then drastically increasing. In the middle of the price soar in 2017, miner count (i.e. mining pool count) reached its all-time low of 61 active mining pools in October 2017, immediately before the Byzantium upgrade. Miner count only began to rise in mid-2018 once ETH prices began falling, suggesting either that mining pool count is inversely related to ETH price, or that mining pools simply take a longer time to set up and launch, meaning their numbers increased with a slight delay.

The former conclusion - that mining pool count might be inversely related to ETH price - is supported by the slight recovery in mining pool count starting in early 2018 despite the ETH price beginning to fall rapidly in the same time period. The reason for this observed inverse relationship could be the incentives of block rewards among miners. When block rewards decreased from 5 to 3 and the ETH price began increasing rapidly, the hashrate also began increasing (figure 2). Higher hash rate, increased difficulty, and reduced ETH rewards meant either 1) miners left the network and diverted their energy towards blockchains with a higher change of reward, or 2) the few remaining individual miners recognized the higher probability of earning block rewards if they joined a mining pool, resulting in a decrease of overall miner numbers and the further consolidation of the network in the hands of mining pools.

Starting in early 2018, the number of payout addresses steadily decreased along with the price of ETH, which reached its lowest point since mid-2017 on December 14, 2018 at ~$83 USD. During that time, the Constantinople upgrade reduced ETH rewards from 3 ETH to 2 ETH. The impact of the block reward reduction cannot be sufficiently determined by looking at payout address count, as the number of addresses continued to decrease. Though we can see a slightly accelerated decrease in the number of miners (i.e. mining pools) immediately following the block reward reduction, which then stabilizes in December 2019 once the price stopped falling.

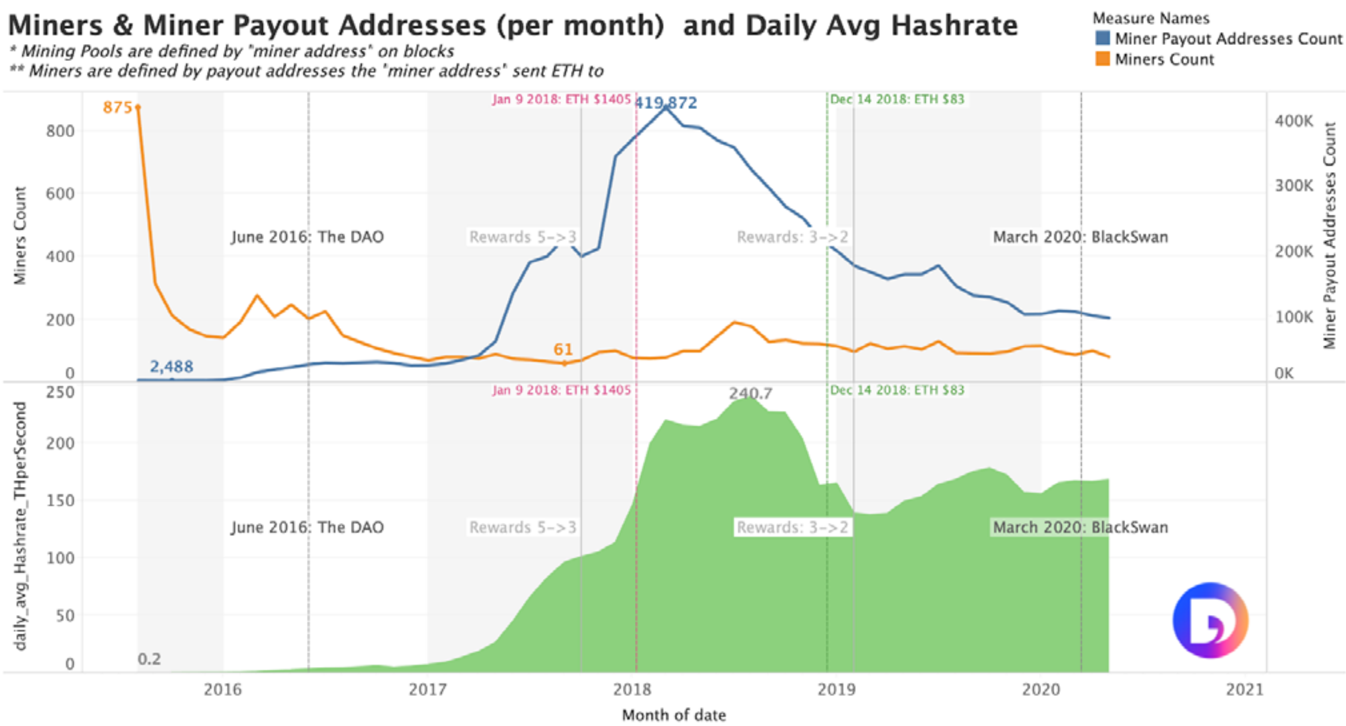

Figure 2: Miners and Payout Addresses alongside Hashrate on Ethereum, 2015 - 2020. Measured by Month.

Centralization of Power Among Mining Pools, 2017

An interesting conclusion from our data is that despite the growing number of miner payout addresses and the fairly stagnant number of mining pools during 2017, the impact of mining pools still outstripped that of individual payout addresses.

We will discuss this in more detail in other sections, but the percentage of blocks mined by mining pools increased from ~70% in late 2016 to >90% by the end of 2017. And the percentage of block rewards attributed to mining pools increased from ~71% in late 2016 to >97% in late 2017 (figures 3 and 4). The drastic increase of mining pool influence despite the decrease in the number of mining pools and the increase of payout addresses during 2017 further supports that considerable consolidation of power occurred during 2017, when people were rapidly joining the network through mining pools in the attempt to profit off the price spike.

In particular, the consolidation of power in the hands of mining pools seems to have occurred in the second half of the year. From Q2 to Q3, the percentage of block rewards received by mining pools decreased from 96.93% to 85.98%. During the same time period, however, the percentage of blocks mined by mining pools jumped from 78.05% to 82.56%. These inverse trends demonstrate that mining pools were growing more significant (i.e. mining more blocks), but the Byzantium upgrade reduced the overall percentage of rewards attributed to mining pools. By Q4 2017, the consolidation was complete. Mining pools accounted for 92.13% of mined blocks and 97.69% of block rewards.

Historical Count: Summary

Generally, figures 1 and 2 tell the fairly simple story that the most significant factor in the number of mining participants (payout addresses + miners) on the network is ETH price. Even during the DAO hack and two upgrades that reduced block rewards by 40% and 33%, the network responded most significantly to fluctuating ETH prices, especially during the price increase from mid-2017 to early 2018 and the price fall from early 2018 to late 2018. Individual payout address count had a direct relationship with ETH price, and mining pool count had an inverse relationship with ETH price.

If we divide the six significant network events by on- and off-chain, we can generally see that payout address numbers are susceptible to off-chain events (ETH price) and miner count is most susceptible to on-chain events (block rewards, security events, hard forks). The DAO event did not seem to noticeably affect the number of payout addresses, but catalyzed the decrease of miners on the network. The Byzantium upgrade, which reduced block rewards, resulted in only a small and temporary decrease in the otherwise-rising number of payout addresses - but it seemed to catalyze the recovery and consolidation of mining pool count, which had reached its all time low just before. The Constantinople upgrade, which also reduced block rewards, did not seem to halt or accelerate the decrease in payout address, because that statistic was so heavily affected by the then-falling ETH price instead.

In the past two months, significant on- and off-chain events have occured. The sudden price decrease in March catalyzed some large Ethereum applications such as Compound to "break" momentarily. Moreover, the first two quarters of 2020 have been marked by on-chain security events; none the size of the DAO, but significant nonetheless. We have had less than two months to observe the impact of this decrease, but further analysis in a few months will provide us with another data set to either support or contradict the divergent relationships of mining pools with ETH price.

Mining Pools: Block Production & Rewards

Data Approach: Mining Pool Centralization

As figures 1 and 2 demonstrate above, the number of miners on the Ethereum network has stagnated at less than 200 since late 2016. This number does not reflect the number of actors on the network - mining rig operators and payout addresses that buy or lend hashpower to mining pools also need to be taken into account when considering the 'size' of the Ethereum miner network. However, miners and mining pools are still fundamental on-chain components of Ethereum, and need to be taken into account when considering points of strength/weakness or (de)centralization on the network.

In the graphs below, we analyze the mining pool attribution percentages for a number of factors relevant to PoW decentralization: block production, rewards, hash rate, and diversity.

In the introduction of this report, we outlined our assumption that every miner on the Ethereum network is a mining pool. We operate on this assumption based on our conjecture that - at least recently - a statistically overwhelming number of miners on the network are mining pools, and the remaining individual miners are of a negligible enough volume to disregard.

For the following figures (3 - 6), we wanted to arrive at as close to certainty as possible when discussing the centralization of mining pool power on Ethereum. Using available off-chain data on Etherscan, we identified 56 mining pools by name. Together, these 56 mining pools control 67 addresses. In other words, though we assume every miner on the network is a mining pool - we can definitely attribute data to 56 mining pools. These 56 mining pools and their associated addresses are listed in the appendix.

See figure 3 as an example. Figure 3 shows the percentage of blocks mined by each miner quarter over quarter since mid-2015. In Q3 2015, we saw a huge diversity of miners (the narrow colored horizontal bars) and only a few recognizable mining pools (Nanopool, for example). The green line that bisects each of figures 3 - 6 is the percentage that we know is attributed to those known 56 mining pools.

So, if we look at Q4 2015, we would say: "We operate under the assumption that every miner - i.e. every colored bar - is a mining pool. However, we know for certain that 69.72% of blocks mined in that quarter were mined by known, named mining pools." By adding in this additional layer of off-chain certainty, we can further ground our conclusions in objectivity.

Mining Pool Block Production

The mining pool attribution of block production follows the trends seen in figures 1 and 2. Starting in mid-2016 following the DAO, we see that the blocks mined by mining pools decreased, from 87.26% in Q2 to 69.9% in Q4. This is in-line with the post-DAO decrease in overall miners. The green line in figure 3 suggests that a number of mining pools did not join the 'new' Ethereum network. Likely, they stayed on Ethereum Classic; they may have joined the Ethereum network later.

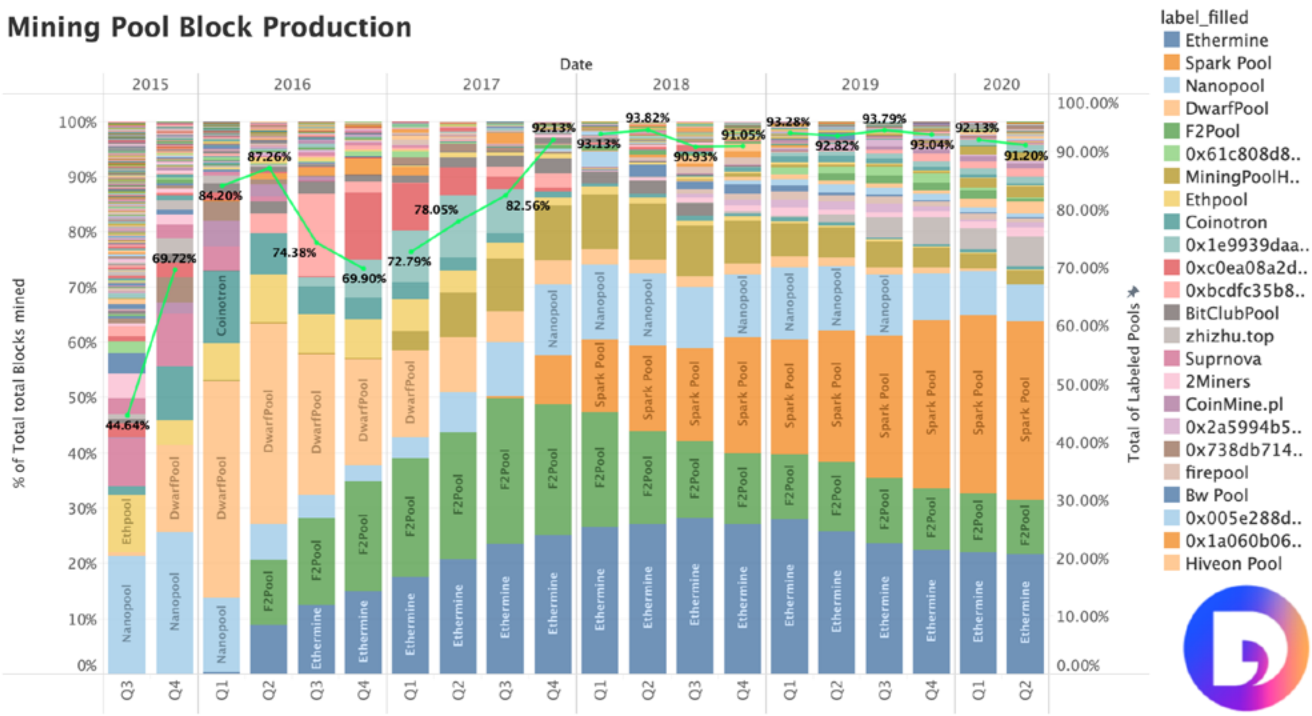

Figure 3: The % of total blocks mined by each miner quarter over quarter. The green line indicates the cumulative percentage attributed to 56 known mining pools (appendix).

Beginning with the rapid price increase in early 2017, the percentage of blocks mined by known mining pools increased dramatically. After reaching a low of 69.9% in Q4 2016, the percentage of blocks that mining pools mined had reached over 90% by Q4 2017, and has remained over 90% since then. Moreover, we see centralization among just a few mining pools. Ethermine, Spark Pool, Nanopool, and F2Pool alone accounted for nearly 70% of blocks mined as of Q2 2020.

Mining Pool Block Rewards

Figure 4 shows the block rewards attributed to mining pools quarter over quarter since mid-2015. Again, the green line shows the percentage attributed to known, named mining pools. The trend of block rewards over time matches that of block production over time. Mining pools lost a significant amount of block production percentage in 2016 after the DAO, and quickly recovered in 2017 with the price increase. Since Q4 2017, the 56 mining pools have definitely accounted for over 90% of block rewards quarter over quarter. Similar to block production, block rewards are further concentrated under a few 'whale' mining pools, but to a greater degree. Ethermine, Nanopool, Spark Pool, and F2Pool together accounted for ~80% of all block rewards in Q2 2020.

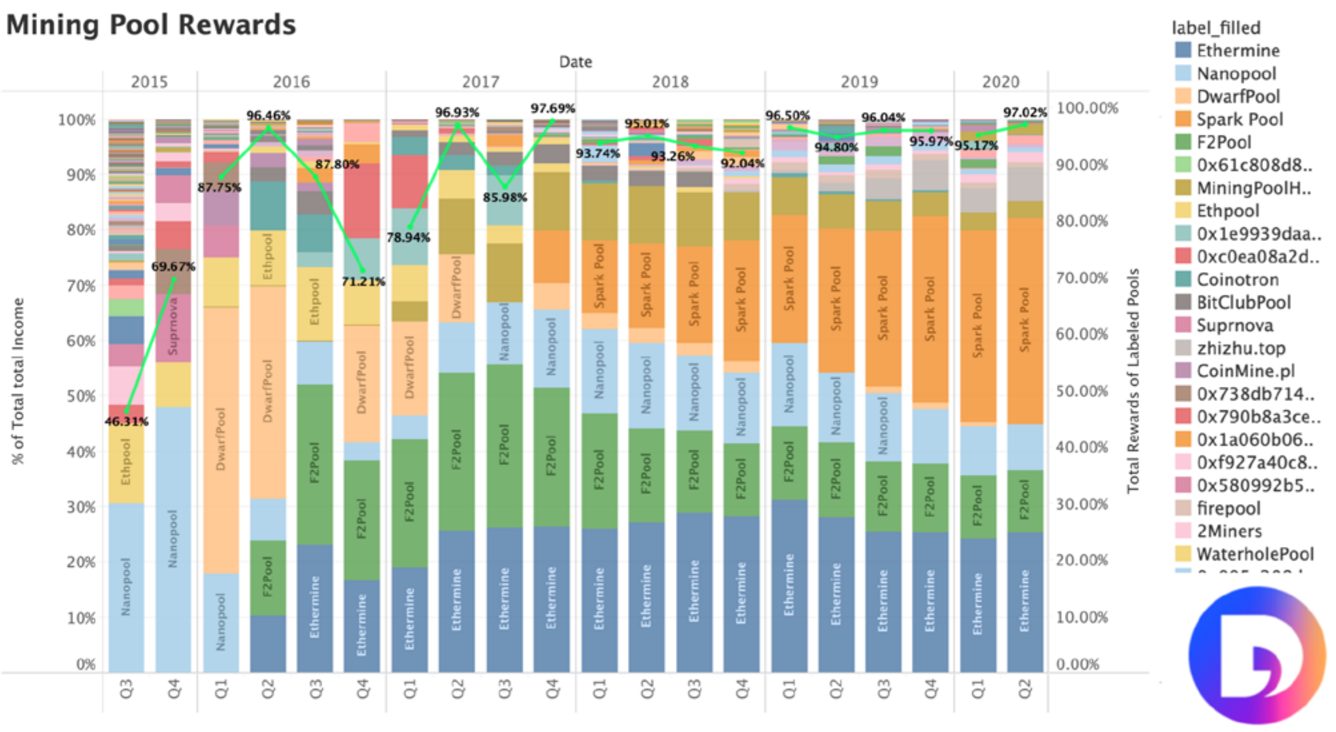

Figure 4: The % of rewards received from mining (including from blocks, uncles and fees) quarter over quarter. The green line indicates the cumulative percentage attributed to 56 known mining pools (appendix).

When compared to the block production chart (figure 3), figure 4 shows supporting evidence for the significant impact on-chain events have on miners. In Q3 2017, the percentage of block rewards attributed to known mining pools dropped to 85.98%, down from 96.63% in Q2. During this same time, however, the percentage of total blocks mined by mining pools increased, from 78.05% in Q2 to 82.56% in Q3. The decrease in block rewards during the same time period as the increase in block production is due to the Byzantium upgrade, which occurred in Q3 2017 and reduced the block rewards from 5 ETH to 3 ETH.

Interestingly, the Constantinople upgrade in Q1 2019 - which reduced block rewards from 3 ETH to 2 ETH - does not appear to have had the same impact. Block production increased from 91.05% in Q4 2018 to 93.28% in Q1 2019, and block rewards increased from 92.04% to 96.05% during the same time frame. The negligible impact of reduced block rewards on the overall rewards attributed to mining pools could be due to the smaller reduction in rewards (33% in 2019 compared to 40% in 2017) or perhaps a more stable and resilient mining pool ecosystem.

Mining Pools: Hashrate and Diversity

Mining Pool Hashrate

There has never been a moment in Ethereum's history when one miner has had control of 51% or more of the hashrate. However, we have seen two moments in Ethereum's history where two mining pools have had over 51% of hashrate - Q2 2017 and Q2 2020.

In Q2 2017, Ethermine and F2Pool together accounted for nearly 60% of the network's hashrate that quarter. It is during this same time period that ETH price was skyrocketing, perhaps suggesting that mining rigs were rushing to contribute their resources to mining pools, thus resulting in the temporary dominance of F2Pool and Ethermine.

In Q2 2020, Ethermine and Spark Pool accounted for just over 51% of the network's hashrate. This recent centralization seems to have less of an identifiable reason behind it. Whereas between Q1 and Q2 2017, F2Pool + Ethermine's hashrate jumped from ~35% to just under 60%, the recent dominance of Ethermine and Spark Pool has been the result of a gradual upward trend since Q1 2018. This gradual increase is a greater warning signal to the Ethereum community. Not only does it suggest we are seeing power concentrate in the hands of two mining pools. Gradual power concentration over the last two years (compared to sudden power concentration in Q2 2017) also means that a +51% dominance might be harder to mediate.

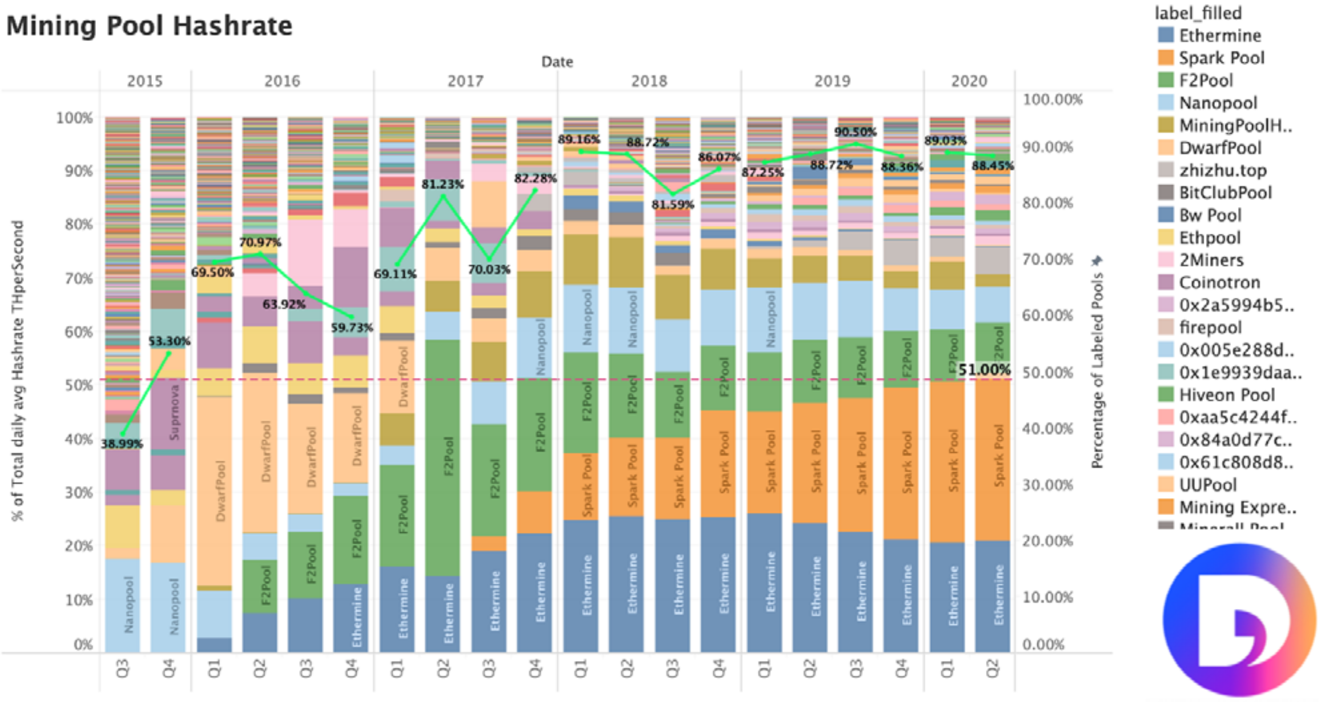

Figure 5: The % of hashrate each miner contributed to the network quarter over quarter. The green line indicates the cumulative percentage attributed to 56 known mining pools (appendix).

Mining Pool Diversity and Fund Flows

We can see from figures 3 - 5 above that mining pools are growing more centralized in their impact and dominance on the network. We see notable mining 'whale' outliers - Ethermine, Spark Pool, F2Pool, Nanopool - but overall can say that for the last two years, at most 56 mining pools have been responsible for over 90% of the block production and block rewards, and over 80% of the hashrate on Ethereum.

When we look at the diversity of mining pools - i.e. which mining pools distribute funds to the greatest number of payout addresses - we see a clearer picture of this growing centralization. Figure 6, in a way, can be seen as a graph showing which mining pool has the greatest number of "customers" quarter over quarter since the network launch in 2015.

First, it is important to note that figure 6 shows the diversity of mining pools based on our first assumption listed in the introduction - that mining pools pay rewards directly to their customer addresses. If a pool adopts a different mechanism - e.g. sending funds to a proxy smart contract that then sends funds out to recipient addresses, we have no objective way to determine how many addresses that pool is paying out to. Take, for example, Spark Pool. In figures 3 - 5, Spark Pool has had some of the highest percentages of block production, rewards, and hashrate. Yet, on figure 6 below, suggests a very small user base. Spark Pool, however, pays out to its users in a different mechanism that cannot be confidently tracked on-chain, and therefore does not show up predominantly in figure 6. Spark Pool is the only major mining pool we know of that uses an indirect payout mechanism. The other major mining pools seen below pay out directly to their miners.

The trend in figure 6 is clear. Since 2016, Ethermine has gradually become the dominant mining pool with respect to payout addresses. As of 2020, Ethermine distributes to nearly 50% of all the payout addresses on the network (again, among direct payout mining pools). Since the launch of the network in Q2 2015, Nanopool has remained a consistent player in mining pool diversity - maintaining ~20% of payout addresses even as other pools have gained and lost market share. DwarfPool seems to have an inverse relationship with F2Pool and Spark Pool, with DwarfPool's market share beginning to call just as F2Pool and Nano Pool come on the scene.

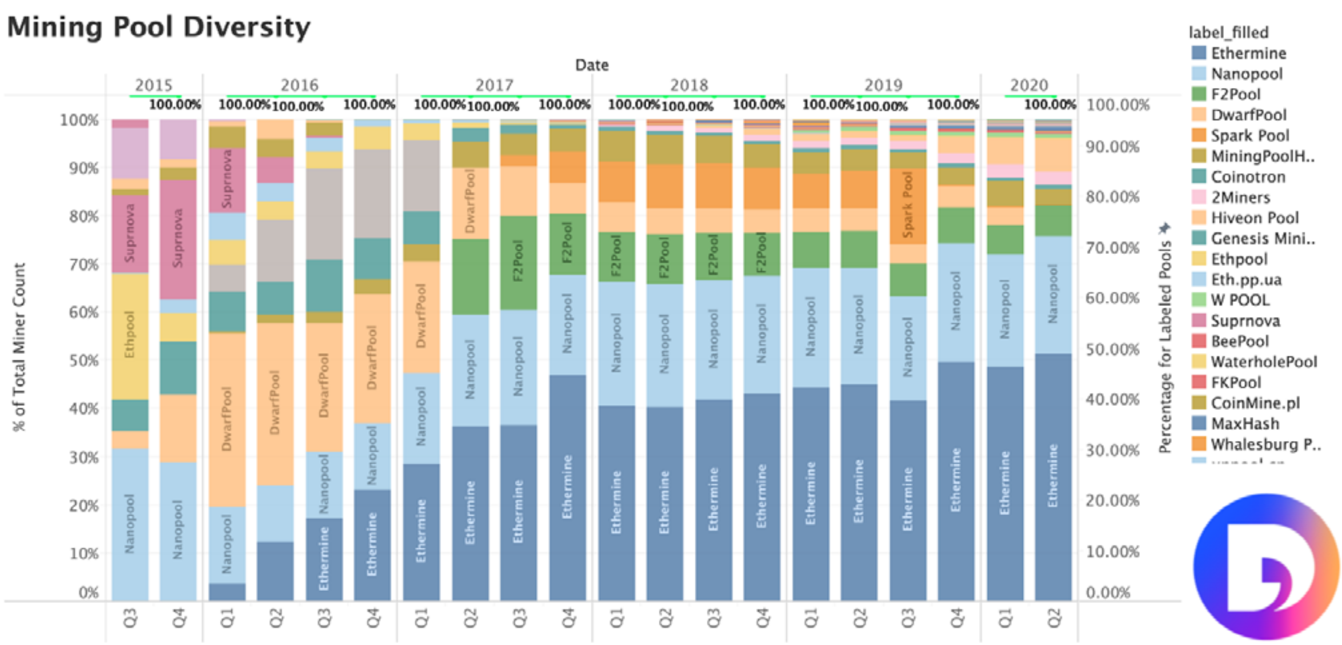

Figure 6: The % of total payout addresses attributed to each miner quarter over quarter.

Though we can see that a small handful of mining pools hold on to a significant portion of the market share diversity, figure 6 alone does not demonstrate any clear centralization. Though Ethermine pays out to ~50% of addresses on the network, each of those addresses could be a distinct person in a distinct region, resulting in a significantly diversified mining pool payout community. When we look at the flow of funds from the largest mining pools over time, however, we begin to see evidence that even within mining pools, customer centralization could be significant.

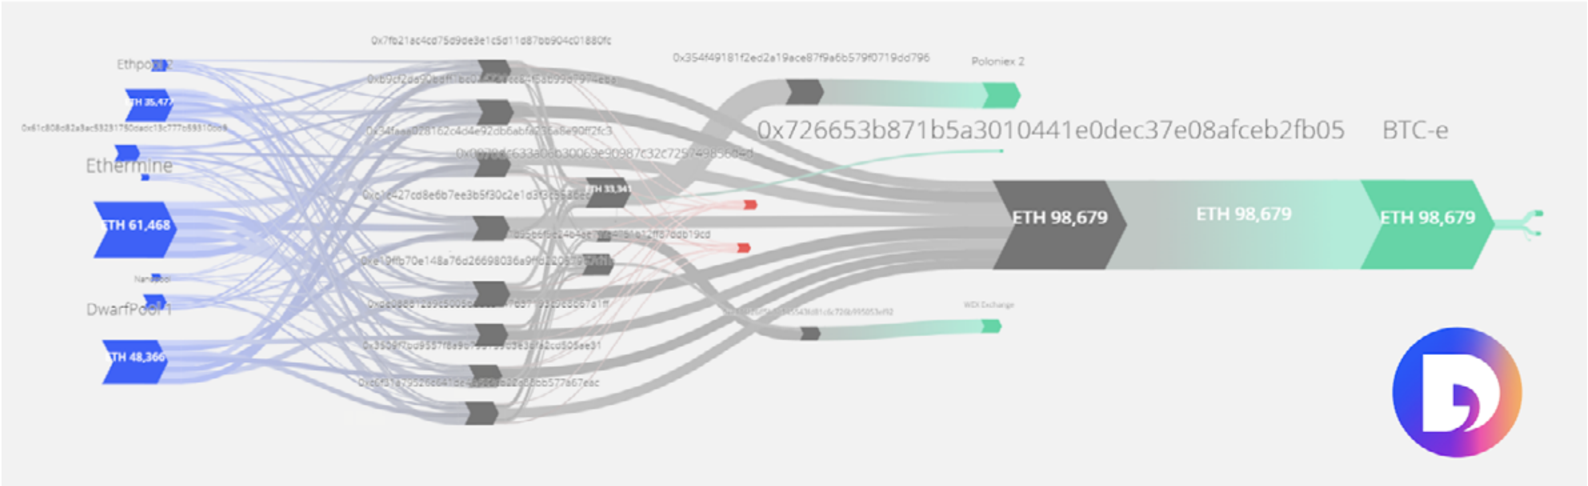

Figure 7: The flow of funds from the 8 largest mining pools, 2017 - 2020.

Figure 7 is a “Sankey” diagram visualizing the flow of funds originating from the largest mining pools between January 2017 and January 2020 (largest meaning they mined the greatest number of blocks). Only the eight largest mining pools are visible on the chart as blue nodes on the far left; the rest are too small to be seen. Of these eight pools, the most widely-known ones are Ethpool2, F2Pool, Ethermine, Nanopool, Dwarfpool and SparkPool - many of which were discussed during analysis of centralization in figures 3 through 6.

Beginning with the blue nodes on the far left, the Sankey diagram shows how funds flowed over time. Each “node” (i.e. colored arrow) is an address, and we see the “layers” that show each time funds were sent to or from an address.

From all the outgoing transactions that these mining pools made from 2017 - 2020, we can see a pattern. Despite years of payout transactions, these mining pools all share a subset of eight addresses (the first vertical “layer” made up of dark gray nodes) to which they sent over 200k ETH.

Eventually, the 200k ETH was centralized again in a few intermediary addresses (some of those addresses still hold the funds, marked by the red nodes). Other funds were either sent directly to exchanges (green nodes) or passed on to another layer of proxies (usually smart contracts) from which it also took a path towards an exchange. ~150k ETH followed this pattern and ended up in exchanges. The rest of the 200k ETH is still held by the red nodes. Granted, most of it ended up in the now defunct BTC-e platform, which was seized by US authorities in 2017. One of the proxies was 0x...5bdd, but was excluded from the chart because it was too large of a visual outlier (it sent over 21m ETH to Poloniex2).

Looking Ahead: PoW to PoS

The exact launch dates of Phase 0, 1, and 2 of Ethereum 2.0 cannot be exactly known - the consequence of having a community-led open source project. With the roll-out of Ethereum 2.0, however, the community will begin moving away from Proof of Work and testing the opportunities provided by Proof of Stake. In the introduction of this report, we discussed the three barriers to PoW: accessibility, scalability, and centralization.

Accessibility: Proof of Stake blockchains do not require validators to worry about the initial hardware costs or pay attention to electricity rates in the same way miners on PoW chains must. It is, therefore, a significantly lower barrier to entry for an individual to run a validator node on a PoS chain than run a mining node on a PoW chain. There is, however, a notable barrier to accessible entry for PoS. Validators must stake a minimum amount of crypto to run a full validator node. For Ethereum 2.0, for example, this amount is 32 ETH ($6,500 at the time of writing). For many, that is a significant amount of money and a deterrent to active participation. In the same way PoW chains have mining pools, however, there will be staking pools that aggregate the funds of participants unable or unwilling to stake 32 ETH. The pool will stake on their behalf, and they will receive rewards as a percentage of their stake.

Scalability: Proof of Stake alone does not improve scalability. However, PoS architectures allow the implementation of a scalability solution known as sharding without reducing security. Sharding is a database scaling mechanism in which a blockchain is partitioned into multiple shard chains, each of which is capable of processing blocks. This relieves the blockchain from having to process each block simultaneously, and instead allows multiple blocks (and, in other words, more sets of data) to be processed all at once. With Ethereum 2.0, for example, sharding will partition the blockchain into 64 separate shard chains, meaning the network will process transactions at minimum 64x the transaction throughput rate of the original PoW chain.

Centralization: With reduced barriers to entry and the elimination of concerns about minimizing electricity costs, PoS networks are significantly more decentralized at the node level than PoW networks. Participation in a PoS chain requires only a non-zero amount of crypto, an internet connection, and a computer (or phone/tablet). That opens up the doors of participation and revenue generation to a much larger group of people. Additionally, economies of scale are far lower in PoS economics than PoW. In PoW systems, the more hash power a miner controls, the greater the % of rewards he would be able to receive. In PoS, a validator's percent return stays constant whether she manages 1 node or 1,000.

Scalability and accessibility are often the most discussed issues with Proof of Work. The data in this report has demonstrated, however, that perhaps centralization is the silent yet more concerning factor we should be paying attention to when discussing the downsides of PoW. Since the launch of the network in 2015, through significant on- and off-chain events, we have seen power and influence concentrate in the hands not just of mining pools, but of a very select few mining pools. It can assuredly be said that the majority of activity on the Ethereum blockchain is maintained by at most 56 miners, and is overwhelmingly influenced by fewer than 8.

This centralization alone is not an issue. Malicious intent, collusion, and a successful attack must also all occur before anything detrimental can occur to the network. Ethereum was built off the promise of decentralization and the removal of vulnerabilities that arise from centralization. Proof of Work has proven tremendously successful for the Ethereum blockchain - and it will remain successful in the years to come as it continues to be necessary for the continuation of the network. From the data available to us about mining pools and PoW, however, we also see that Ethereum 2.0 and Proof of Stake will be a welcome improvement to the centralization of Ethereum.

Appendix

From Etherscan, 56 known mining pools, which together comprise 67 addresses. These 56 mining pools were used to establish the green line of certainty in figures x - 7.

MINING POOL	NAME TAG	ADDRESSES
2Miners	2Miners: PPLNS	0x00192fb10df37c9fb26829eb2cc623cd1bf599e8
2Miners	2Miners: SOLO	0x002e08000acbbae2155fab7ac01929564949070d
AlphaPool	AlphaPool	0xc839ee5542b4e8413246b3634c5c739fea949562
ALTpool	ALTpool.pro	0x433022c4066558e7a32d850f02d2da5ca782174d
AntPool	AntPool	0xa855c20a1351acd2690c716e2709c7dff3978d12
BaikalMine	BaikalMine 1	0xff1b891969773159366ab6310ff63a69ac4acffd
BeePool	BeePool	0x99c85bb64564d9ef9a99621301f22c9993cb89e3
BitClubPool	BitClubPool	0xf3b9d2c81f2b24b0fa0acaaa865b7d9ced5fc2fb
BTC.com Pool	BTC.com Pool	0xeea5b82b61424df8020f5fedd81767f2d0d25bfb
Bw Pool	Bw Pool	0x52e44f279f4203dcf680395379e5f9990a69f13c
CoinMine.pl	CoinMine.pl	0x68795c4aa09d6f4ed3e5deddf8c2ad3049a601da
Coinotron	Coinotron 1	0xf8b483dba2c3b7176a3da549ad41a48bb3121069
Coinotron	Coinotron 2	0xa42af2c70d316684e57aefcc6e393fecb1c7e84e
Coinotron	Coinotron 3	0x6a7a43be33ba930fe58f34e07d0ad6ba7adb9b1f
CoolPool	CoolPool.Top: SOLO	0xe5a349fc4ff853dfdd0b7eaaa9dcd8918e768f49
Cruxpool	Cruxpool	0x249bdb4499bd7c683664c149276c1d86108e2137
DwarfPool	DwarfPool 1	0x2a65aca4d5fc5b5c859090a6c34d164135398226
DwarfPool	DwarfPool 2	0x151255dd9e38e44db38ea06ec66d0d113d6cbe37
Eth.pp.ua	Eth.pp.ua	0xa027231f42c80ca4125b5cb962a21cd4f812e88f
ETH.SoloPool.org	ETH.SoloPool.org	0xf35074bbd0a9aee46f4ea137971feec024ab704e
EthashPool	EthashPool 1	0x8fce1ef27f3add1411c7a99be402de598ad38389
EthashPool	EthashPool 2	0x52f13e25754d822a3550d0b68fdefe9304d27ae8
EtherDig	EtherDig	0x8d35067233605bef6069191ae0922d134ff80d48
EthereumPool	EthereumPool	0x9d551f41fed6fc27b719777c224dfecce170004d
Ethermine	Ethermine	0xea674fdde714fd979de3edf0f56aa9716b898ec8
Ethpool	Ethpool 1	0xe6a7a1d47ff21b6321162aea7c6cb457d5476bca
Ethpool	Ethpool 2	0x4bb96091ee9d802ed039c4d1a5f6216f90f81b01
ExtremeHash	ExtremeHash	0x6537b65a50a862391515455272f9b6c7168afe94
EzilPool	EzilPool 1	0xcc22cb1b6625b64e81909456111d76be6158dfbc
EzilPool	EzilPool 2	0x8595dd9e0438640b5e1254f9df579ac12a86865f
F2Pool	F2Pool	0x829bd824b016326a401d083b33d092293333a830
F2Pool	F2Pool Old Address	0x61c808d82a3ac53231750dadc13c777b59310bd9
firepool	firepool	0x35f61dfb08ada13eba64bf156b80df3d5b3a738d
FKPool	FKPool	0x464b0b37db1ee1b5fbe27300acfbf172fd5e4f53
Flexpool.io	Flexpool.io	0x7f101fe45e6649a6fb8f3f8b43ed03d353f2b90c
Genesis Mining	Genesis Mining	0xd34da389374caad1a048fbdc4569aae33fd5a375
HashON Pool	HashON Pool	0xd0db3c9cf4029bac5a9ed216cd174cba5dbf047c
Hiveon Pool	Hiveon Pool	0x4c549990a7ef3fea8784406c1eecc98bf4211fa5
Huixingpool.com	Huixingpool.com	0x14b30f257c2737370203a15aa343c2b600dfb675

MINING POOL	NAME TAG	ADDRESSES
Huobi Mining Pool	Huobi Mining Pool	0x9d6d492bd500da5b33cf95a5d610a73360fcaaa0
ICanMining.ru	ICanMining.ru	0xf64f9720cfcb59ca4f5f45e6fdb3f68b875b7295
KuveraPool	KuveraPool	0x4e4e23ac3c11789e23169025503ea4373b01417b
MATPool	MATPool	0x7f3b29ae0d5edae9bb148537d4ed2b12beddf8b3
MaxHash	MaxHash: EthPool	0x6c3183792fbb4a4dd276451af6baf5c66d5f5e48
MaxHash	MaxHash: Solo Mining	0xcf6ce585cb4a78a6f96e6c8722927161a696f337
Minerall Pool	Minerall Pool	0x09ab1303d3ccaf5f018cd511146b07a240c70294
Mining Express	Mining Express	0x06b8c5883ec71bc3f4b332081519f23834c8706e
MiningPoolHub	MiningPoolHub	0xda466bf1ce3c69dbef918817305cf989a6353423
MiningPoolHub	MiningPoolHub: Old Address	0xb2930b35844a230f00e51431acae96fe543a0347
myminers.org	myminers.org: Solo	0x2a98776c7e13ed1c240858bd241dcf95fc1928b4
Nanopool	Nanopool	0x52bc44d5378309ee2abf1539bf71de1b7d7be3b5
NoobPool	NoobPool	0xd5bbb4264b70ca4f58c45d27b9d7e11190754a54
PandaPool	PandaPool	0x6b7d50bb8fab584e54251a10e1c6cfa51dd7b618
PoolHub	PoolHub	0x47c439c8784b44366735fc2cfe08228cb91d5b8e
Poolin	Poolin	0xa7b0536fb02c593b0dfd82bd65aacbdd19ae4777
SaturnPool	SaturnPool	0xe16263ee79b0ee32c242c99f02559e92abaea9eb
Spark Pool	Spark Pool	0x5a0b54d5dc17e0aadc383d2db43b0a0d3e029c4c
Suprnova	Suprnova 1	0x1dcb8d1f0fcc8cbc8c2d76528e877f915e299fbe
Suprnova	Suprnova 2	0x63a9975ba31b0b9626b34300f7f627147df1f526
Uleypool	Uleypool	0xa3c084ae80a3f03963017669bc696e961d3ae5d5
UUPool	UUPool	0xd224ca0c819e8e97ba0136b3b95ceff503b79f53
W POOL	W POOL	0x44fd3ab8381cc3d14afa7c4af7fd13cdc65026e1
WaterholePool	WaterholePool	0x9435d50503aee35c8757ae4933f7a0ab56597805
Weipool	Weipool	0xd1e56c2e765180aa0371928fd4d1e41fbcda34d4
Whalesburg Pool	Whalesburg Pool: Old Address	0x7c6694032b4db11ac485e1cff0f7509d58b41569
xnpool.cn	xnpool.cn	0xe4bdced60430a90f31dba03524dd5d15a2670649
zhizhu.top	zhizhu.top	0x04668ec2f57cc15c381b461b9fedab5d451c8f7f

Authors

Everett Muzzy

Everett Muzzy is a researcher, writer, and product marketing manager at ConsenSys focusing on blockchain theory, data studies, and enterprise adoption. His past research has included quantifying decentralization and system interoperability. Get in touch with Everett.

Bogdan Gheorghe

Bogdan Gheorghe is a data scientist with a focus on emerging blockchain technology, and Ethereum Decentralized Finance. His data analyses have provided insights into DeFi protocol mechanisms, including MakerDAO's DAI migration, Compound illiquidity risks, and stablecoin functionality. Get in touch with Bogdan.

Danning Sui

Danning Sui is a data scientist specializing in emerging decentralized systems, including Ethereum. Her analyses have previously included ICO investigations, non-fungible token mechanisms, and decentralized finance protocol relationships. Get in touch with Danning.

CONSENSYS

The Ethereum PoW Mining Report | Codefi Data