GoChain: Blockchain at Scale

Version 1.0

1. Introduction

Cryptocurrency is changing the world economy. What was previously impossible in most of the world – transferring money globally – is now not only possible, but it's safe, cheap and easy.

It also acts as a store of value, like gold, but for the digital age. Asset ownership is stored in a secure, decentralized database called the blockchain. This decentralization prevents any one government or company from controlling it. At the time of this writing, the total market cap of cryptocurrencies is $718 billion, up from $15 billion a year ago.

The potential of the blockchain is huge and world-changing but the lack of scalability, lack true decentralization, and excessive energy use are issues that plague existing cryptocurrencies.

1.1. Background

1.1.1. Speed and Volume. Public, decentralized cryptocurrencies suffer from slow transactions and low transaction volume. Bitcoin can only process 7 transactions per second [1], Ethereum can only process 13 per second [2]. Additionally, the time to verify transactions can range from several minutes to several hours depending on current volume [4].

In contrast, Visa, Inc. averages 150 million transactions every day and is capable of handling more than 56,000 transactions per second [3]. Public cryptocurrencies are too slow for real world processing by 4 orders of magnitude.

1.1.2. Energy Consumption. The process of mining blocks uses an enormous amount of energy because of a consensus algorithm called Proof-of-Work (PoW) [12]. PoW requires non-trivial computational work by mining nodes which, in turn, makes it cost prohibitive for a bad actor to perform malicious acts. This computational workload requires energy.

As of today, 3.5 million US households could be powered with the energy used to run the Bitcoin network, while Ethereum uses the equivalent power of 1 million households [10], [11]. This unacceptable and unsustainable.

1.1.3. Decentralization. Decentralization is a central tenant of cryptocurrencies. It ensures no one company or government can control it. However, in practice most mining has moved to China where electricity is cheapest [26], [27]. 75% of all blocks are mined by large Chinese mining companies. This is true of Bitcoin, Bitcoin Cash, Ethereum and top cryptocurrencies [5]. In the event of company collusion or government privatization, 51% attacks [6] would be possible.

1.1.4. Rigid Contracts. Ethereum provides a unique feature called *smart contracts*. These smart contracts allow users to write small, unmodifiable programs that ensure all involved parties are committed to a set of rules with absolute certainty. This has been so successful that nearly every ICO in the past year has run on Ethereum's smart contract system.

These smart contracts, however, are not smart at all. They are extremely rigid contracts that are unable to adapt to changes like real world contracts would. Parties involved in a contract have no ability to upgrade their contract if necessary to adjust terms or to fix bugs in contract code. Successful attacks have been mounted against smart contracts [7] and there are several known attacks [8], [9].

1.2. Previous Work

1.2.1. Proof-of-Work Algorithm. Proof-of-Work (PoW) [12], [13] is a consensus algorithm that is commonly used in cryptocurrencies. PoW was originally invented as a means to combat spam [15]; if you make it computationally expensive to send email then spamming would be cost prohibitive while still being almost free for a normal user to send email. The same concept is used in cryptocurrencies to prevent malicious actions by making it prohibitively expensive to modify the blockchain.

In cryptocurrency networks, "miners" are special nodes that perform the PoW calculation on a set of transactions plus the hash of the previous block to generate the next block in the blockchain. Since the block contains the hash of the previous block, changing a historical block would require regenerating all of the subsequent blocks. Regenerating all the hashes would be computationally intensive and would require a lot of energy – and energy isn't free. It would also be time consuming. The process of proving work and generating blocks is called "mining". Miners are rewarded for this work with newly minted coins adding to the total supply.

Although PoW has helped move us towards secure distributed ledger system, it suffers from poor performance, a lack of decentralization, and excessive energy consumption.

1.2.2. Proof-of-Stake Algorithm. Proof-of-Stake (PoS) [14] is another consensus algorithm which pseudo-randomly chooses validators based on their stake in the network. The idea is that those with the most coins in circulation have the most to lose so they are positioned to work in the interest of the network. This approach avoids the cost of computing hashes, however, it makes assumptions about the interests of its members being in line with the network.

Validators within the PoS network are anonymous users who are identified only by their wallet address. This provides no additional accountability over PoW for bad actors who can amass significant wealth on the network. Second, transaction fees will go to those who already have the most money within the network and large wealth requirements exclude poorer coinholders from validation. Finally, while PoS reduces energy consumption its goal is not oriented towards high performance. Initial targets for Ethereum's Casper implementation are only 100 TPS.

1.2.3. Proof-of-Authority Algorithm. Proof-of-Authority [16] is a new consensus algorithm where a trusted set individuals provide all transaction processing. This trust allows transaction processing speed to improve significantly by skipping the PoW hash computation. A few networks exist but they currently only focus on private networks or do not focus on performance as a goal. Many also do not have compatibility with the Ethereum network.

One public network relies the US state-level Notary Public system to verify the identity of 12 individuals who will act as validators on the network [17]. Candidates requesting validator status submit proof of physical address, bank account, social network, and mobile phone to verify their real world identity. While PoA removes the the computational burden of mining, trusting individuals for transaction processing breaks down at scale for several reasons.

First, there is a disparity between the net worth of the network versus the market cap of the network. This is what the PoS system attempts to solve. Assuming an average net worth of an individual in the United States is $68,828 [18], the total net worth of the validators is $825,936:

$$12 * \$68,828 = \$825,936$$

Even if the number of validators increased by an order of magnitude, the total net worth of the validators is a tiny fraction of the $6.8T in transactions processed by Visa, Inc. every year [19]. This disparity introduces a strong incentive for bribery.

Second, validators must post their physical address publicly which opens the potential for intimidation or physical threats. A terrorist organization or rogue state can mount an attack on a large scale financial system by controlling half of these validators.

Finally, most individuals lack the experience and infrastructure to run a secure transaction processing system. This significantly increases the network's exposure to malicious hacking.

2. Implementation

2.1. Proof-of-Reputation

GoChain uses a Proof of Reputation (PoR) consensus model that depends on the reputation of its participants to keep the network secure. A participant must have a reputation that is important enough that they would face dire consequences if they were to cheat the system–in both financial terms and branding. Most businesses would face serious consequences if they were caught cheating a financial network. Larger companies with more to lose will be chosen over smaller companies with less to lose.

Once a company proves reputation, they may be voted into the network as an authoritative node and at this point, it operates just like a Proof of Authority network (PoA). Only authoritative nodes can sign and validate blocks.

We are building on Ethereum's network because it is much more than just a store of value. That is why we believe it's the best cryptocurrency and blockchain on the market today and why we are using it as a starting point. All Ethereum wallets and development tools will be compatible with GoChain.

2.2. Why Reputation?

Reputation is critical to a business. A business that acts in an unethical way suffers on many levels including fines, loss of revenue, decrease in valuation, branding, and public relations. Trust is a cornerstone to a successful business and once a brand loses trust with their customers, it can take years repair.

The Volkswagen emissions scandal [20] is a perfect example. They operated alone to deceive the public and their own customers. Once caught, it was a financial and public relations disaster which caused a 30% drop in stock price and a $25 billion fine [21].

GoChain uses this model to allow companies to keep each other in check and keep the network secure. Imagine if Volkswagen worked with Ford, Toyota and others to validate and verify each others' emissions tests. It is unlikely that Volkswagen could have gotten away with their fraudulent emissions tests. If they had attempted it then they would quickly be voted out of the consortium and lose their rights to be a part of the network.

PoR is more attractive to the broader business community than untrusted networks based on PoW or PoS. Risk-adverse companies will be able to rely on known brands in the same way they trust companies like Visa, Inc. or JPMorgan Chase & Co. In PoR, everyone knows exactly who they are trusting with their data.

2.3. Measuring Reputation

Reputation is impossible to measure precisely but we are weighing several important metrics in our decision:

1) Market Cap
2) Publicly Traded
3) Brand Significance

Companies with a large market cap have more to lose than small cap companies. We use this metric when evaluating companies because the value of the member companies needs to be in parity with the value of the processing network to disincentivize cheating.

We next look at whether a company is publicly or privately held. The effect of reputation affects publicly traded companies more directly and immediately than private companies. An unethical decision on the network can impact a public company's stock price in minutes.

Finally, we give preference to companies which require strong public brands for their business. For example, a loss of reputation for a company like Coca-Cola or Apple will be more impactful than for a coal mining company.

2.4. Authorized Signers

Authorized signers are trusted nodes that create blocks, sign them, and distribute them to other nodes. Similar to miners in a Proof-of-Work (PoW) system in that they create blocks and sign them but without the mining cost.

A list of authorized signers will be maintained on the blockchain. Only authorized nodes can sign blocks and all blocks are verified that this is true by checking the signer is in the authorized list. The signing algorithm is essentially the same signature algorithm as PoW but with a different set of headers. PoW-specific headers will be removed and additional headers added to enable voting.

Given N authorized signers, a signer may only sign a block every $(N/2)+1$. This ensures that someone would need to control $> 50\%$ of signers to perform a malicious attack [6].

2.4.1. Incentives / Rewards. Authorized signers will be rewarded GoChain Coins (GOC) per block signed. Initially this rate will be 5% of the total tokens which is 50,000,000 new tokens in year one. This rate will decrease over time. The amount per block will depend on the finalized block times. If block times are 10 seconds for example, then the node would be rewarded an average of 15.9 tokens per signed block in the first year.

There will also be small transaction fees that the authoritative node that signs the block containing the transactions will keep.

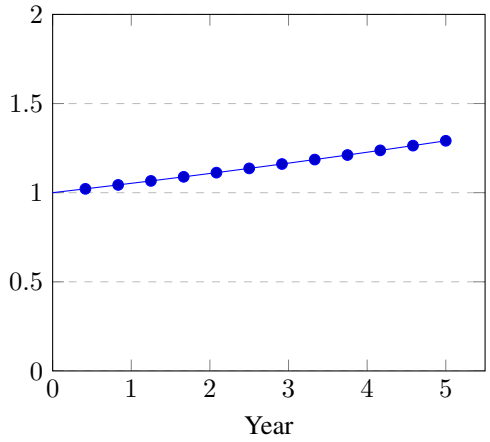

Figure 1. Total GOC (in billions)

The consensus protocol ensures fairness and liveness by incentivizing the assigned signer of a block to perform the signing but also allowing other blocks to sign if the assigned signer is unavailable. The assigned signer for a block is determined by a round-robin lookup of the authorized signer list. If the assigned signer doesn't respond then then other signers can sign at a lower block difficulty level. See Algorithm 1.

Algorithm 1: Signing Blocks

```
1  Function Sign(block)
2      if signed in last (n/2) + 1 blocks then
3          return;
4      end
5      if block.number%signer_n = signer_index then
6          block.difficulty = 2;
7      else
8          block.difficulty = 1;
9      end
10 end
```

2.5. Voting

GoChain has implemented a two-phase voting process. For the initial rollout, the GoChain Foundation will add the first 50 signers to the authorization list. These original 50 signers will be companies from multiple industries and spread out across multiple countries. This will help ensure forced decentralization and avoid interference by any single government.

Once 50 authorized signers have been established then voting control will be handed over to the signers to govern themselves. This voting process is shown in Algorithm 2. The PoA implementation repurposes several block headers to pass voting information between nodes.

2.6. Signer Verification

Companies which operate authorized signer nodes will go through a verification process to ensure their identity is correct. These validation steps will be automated through the use of smart contracts on the blockchain.

The PoA implementation provides point-in-time signer & voting state that we can build upon to provide full transparency to end users. Combined with the verification data stored within smart contracts, users can what companies are running which nodes at any given point.

2.6.1. Company Verification. The first verification step requires companies to provide their Dun & Bradstreet D-U-N-S number. This identifier allows voters obtain a D&B report on the company to view official contact information. This information will be used to establish communication. See Figure 2.

Algorithm 2: Vote to add/remove signer

```
foreach block do
    if block is epoch then
        clear all votes;
    end
    if block candidate set then
        if signer already voted then
            clear previous vote;
        end
        if block vote = 'yes' then
            cast 'yes' vote;
        else
            cast 'no' vote;
        end
        if candidate has absolute majority then
            if candidate is signer then
                remove candidate from list;
            else
                add candidate to list;
            end
        end
    end
end
```

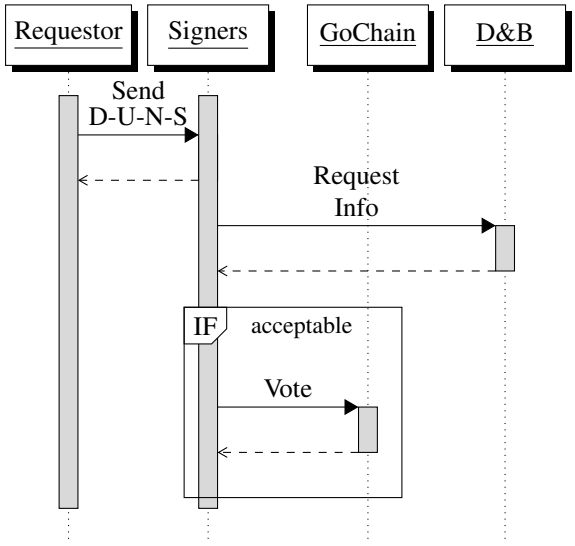

Figure 2. Company Verification Sequence

2.6.2. DNS Verification. A secondary verification step requires companies to add a TXT entry to their DNS records with a random token. This is a common practice when verifying domain ownership [28]. GoChain will host a DApp which generates tokens for requestors and validates DNS records. Authorized signing nodes can use this DApp to view the verification. See Figure 3.

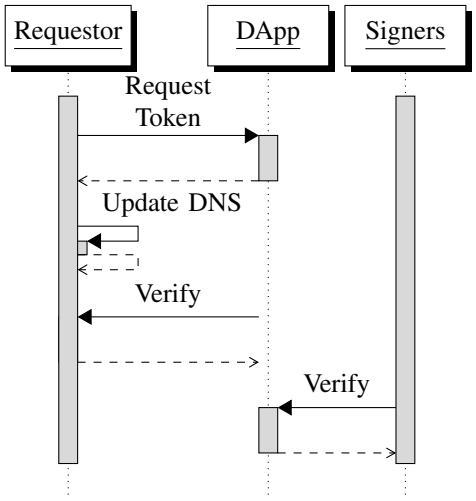

Figure 3. DNS Verification Sequence

2.7. Checkpoints

A checkpoint is a signed snapshot of the current state of the entire blockchain at a particular block number. It will contain all non-zero account balances and smart contract states. Once a checkpoint is generated, all the previous blocks and data can be removed.

When a new node is started, it will download a recent checkpoint, then continue retrieving blocks and state from that point on. This will save hours, if not days, getting in sync. This means a node can be up and running in minutes.

GoChain will provide a publicly-available, read-only API to retrieve any historical block so anyone can look up data by keys. This will be open source so anyone can run this to keep a full history. We encourage it for further verification and accountability. This will make it easier to build third party services such as block explorers.

2.8. Performance and Optimizations

2.8.1. Speed and Volume of Transactions. By using trusted nodes, transactions can be verified very quickly and the volume of transactions the network can handle increases by orders of magnitude. Similar to systems we use everyday that can handle high volumes, like a Google search or Visa payments, those systems can handle high load only because they trust the servers and the network they are running on.

Other factors such as block size and gas limits are artificially low because of the computation power required by PoW. By trusting the consensus nodes we can increase the volume of the network by **100x** more than Ethereum can currently handle. Trusted consensus is used outside cryptocurrencies in systems such as etcd which can reach 141,578 transaction per second on a 3-node cluster using modest hardware [**?**]. Improving throughput is critical as the growth rate of Ethereum is skyrocketing to an unsustainable rate. Ethereum runs at 13 tx/second right now; we are targeting 1,300 tx/second at mainnet launch.

The two major parameters we can tweak are block size (gas limit) and block times. Due to the fact that we have a relatively small set of signers (vs the number of miners in PoW) with known capabilities, we are able to increase the block size drastically and reduce the block times. This alone greatly increases the number of transactions per second. As stated above, the reason you can't increase block size in PoW is that it makes the hashing algorithm too hard and too expensive. GoChain does not have that limitation.

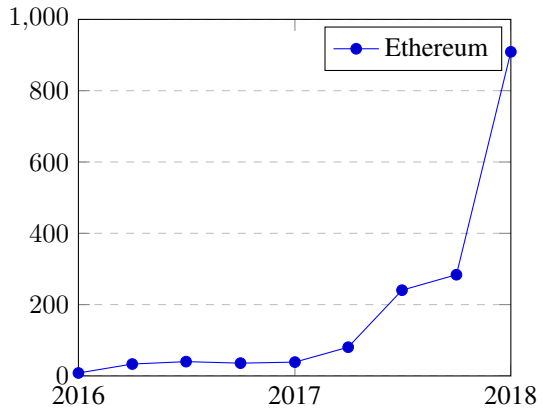

Figure 4. Transactions Per Day (in thousands) [29]

2.8.2. Energy Consumption. Using a trusted network of authoritative nodes means that there will be no mining. No mining means there will be no battle between computers to win blocks and therefore no wasted energy. Nodes will only require a small fraction of this energy to process transactions, run smart contracts, and verify blocks.

Ethereum's estimated energy consumption at the time of this writing is 14 TWh and rising [32]. Assuming 450W power usage per server [31], our 50-node cluster will only use 197.1 MWh or 0.001% of the energy of the Ethereum network.

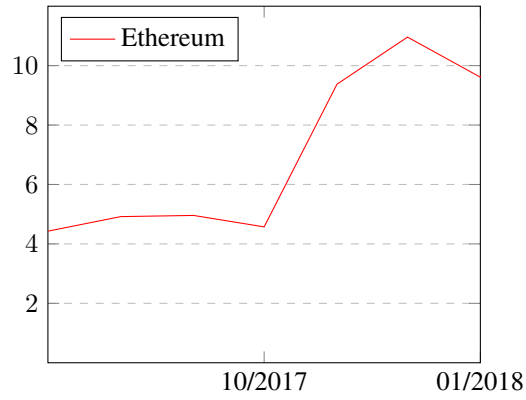

Figure 5. Estimated TWh per Year [32]

2.8.3. Networking. Signers will communicate directly with each other. This means that the node who just finished signing will send the just signed block to other signers in the authorized signers list before sending to a replication node. This ensures the authorized signers get the information they need as fast as possible while offloading blockchain and API queries for the rest of the network to dedicated replica nodes.

The replication layer exists for non-signer nodes (everyone else) to request blocks and query the state using a read-only API. Because the replication layer is read-only, we can horizontal scale to meet the needs of a global scale set of users. Figure 6 shows an example of this 3-tier network strategy.

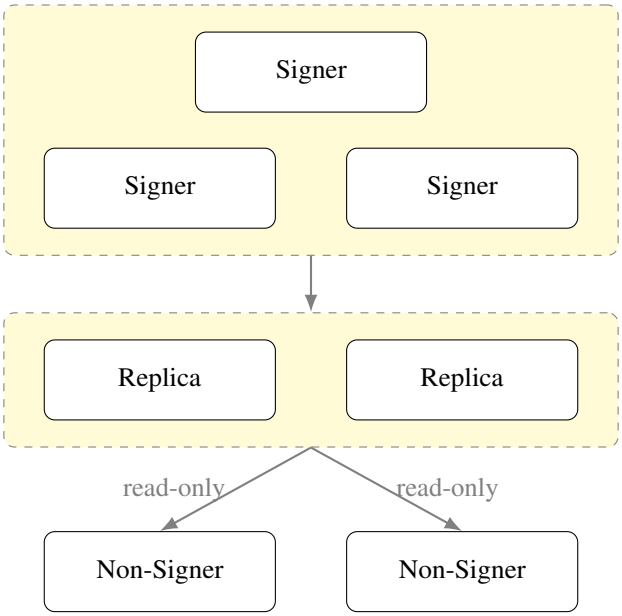

Figure 6. GoChain Network Topology

2.8.4. Storage. The storage requirements to store the entire blockchain is quite large–Ethereum size is hundreds of gigabytes and it's growing rapidly. It can take hours or days to synchronize to a new node which makes it impractical for the average user that just wants to send a transaction. There are newer modes you can run to reduce the size such as *fast* and *light* mode, which reduce the size drastically and that is a good step in the right direction.

Since GoChain will be handling 100x more transaction volume, storage becomes much bigger–potentially 100x bigger. As of this writing, Ethereum transactions average 174 bytes and 700,000 transactions occur per day [29] which produce 120.4MB of block data per day or 43.9GB per year. Increasing throughput by 2 orders of magnitude will generate 4.4TB of block data per year. Propagating this across all 23,000 nodes in the Ethereum network [30] would require 101 petabytes.

By limiting the set of nodes operating on the dataset we reduce the network traffic and storage requirements. Checkpointing allows nodes to only store the small fraction of the total blockchain that is required for current process-

ing. Current cloud pricing of $0.022 USD per GB/month [33] makes storing a copy of the blockchain history only $1,161.60 USD per year of block data.

2.8.5. The Future. Beyond our initial goals described above, we have a plan to upgrade the smart contract system to make it easier and less error prone. Software almost always contains bugs that are unknown at the time of release and developers need a way to fix those bugs. Ethereum does not allow you to upgrade your contracts and that results in $100's of millions of value being stolen [7]. We intend to make writing smart contracts easier to write and easier to deploy, as well as making them safer to prevent the massive amounts of theft that is happening. Smart contracts need to be more like the real world, where they can be amended, paused, and/or terminated.

We are also adding standardized rulesets to contracts to define how and when contracts can be modified. We expect this will help the adoption of smart contracts by the broader business community by using familiar contract terms. For example, a co-op organization may require a quorum of members to change a contract while other organizations may require all participants in a contract to agree to a change. GoChain will continue to default contracts to be immutable by default for compatibility with Ethereum. GoChain will also add additional security features such as whitelists to protect access to contracts to minimize attack risk.

3. Conclusion

GoChain is being built to solve fundamental issues that plague existing cryptocurrencies. The GoChain team has decades of experience building high scale, distributed applications and cloud infrastructure. We will apply that to GoChain to ensure it's one of the best and most scalable blockchains on the market. GoChain will be faster, more scalable, more decentralized and greener.

References

[1] Bitcoin Scalability Problem, https://en.wikipedia.org/wiki/Bitcoin_scalability_problem

[2] Blockchains Don't Scale, https://hackernoon.com/2cb43946551a

[3] VisaNet handling 100,000 transactions per minute, https://mybroadband.co.za/news/security/190348.html

[4] Average Blockchain Confirmation Time, https://blockchain.info/charts/avg-confirmation-time, https://etherscan.io/chart/pendingtx

[5] Stop Calling Bitcoin Decentralized, https://medium.com/@homakov/cb703d69dc27

[6] 51% Attack, https://www.investopedia.com/terms/1/51-attack.asp

[7] The DAO Attacked: Code Issue Leads to $60 Million Ether Theft, https://www.coindesk.com/dao-attacked-code-issue-leads-60-million-ether-theft/

[8] A survey of attacks on Ethereum smart contracts, https://eprint.iacr.org/2016/1007.pdf

[9] Known Attacks, https://consensys.github.io/smart-contract-best-practices/known_attacks/

[10] Bitcoin Energy Consumption, https://digiconomist.net/bitcoin-energy-consumption

[11] Ethereum Energy Consumption, https://digiconomist.net/ethereum-energy-consumption

[12] Proof-of-Work, https://en.wikipedia.org/wiki/Proof-of-work_system

[13] Proof-of-Work, https://en.bitcoin.it/wiki/Proof_of_work

[14] Proof-of-Stake, https://en.wikipedia.org/wiki/Proof-of-stake

[15] Pricing via Processing, http://www.wisdom.weizmann.ac.il/~naor/PAPERS/pvp.ps

[16] Proof of Authority, https://en.wikipedia.org/wiki/Proof-of-authority

[17] POA Network Governance Overview, https://github.com/poanetwork/wiki/wiki/Governance-Overview

[18] Average net worth of Americans, http://www.businessinsider.com/heres-the-average-net-worth-of-americans-at-every-age-2017-6

[19] Visa, Inc., https://en.wikipedia.org/wiki/Visa_Inc.

[20] How VW Paid $25 Billion for 'Dieselgate' and Got Off Easy http://fortune.com/2018/02/06/volkswagen-vw-emissions-scandal-penalties/

[21] VW Scandal: How Has It Impacted Volkswagens Stock?, https://www.investopedia.com/news/vw-scandal-how-has-it-impacted-volkswagens-stock-vlkay/

[22] Block Size and Transactions per Second, https://www.bitcoinplus.org/blog/block-size-and-transactions-second

[23] Ethereum Block Size, https://www.reddit.com/r/ethereum/comments/4a3kqo/what_is_ethereums_block_size/

[24] Can we please increase the gas limit?, https://www.reddit.com/r/ethereum/comments/7hmlm4/can_we_please_increase_the_gas_limit/

[25] Block Size Limit Controversy, https://en.bitcoin.it/wiki/Block_size_limit_controversy#Arguments_in_opposition_to_increasing_the_blocksize

[26] Electricity Sector in China, https://en.wikipedia.org/wiki/Electricity_sector_in_China

[27] Why China Mines More Bitcoin Than Any Other Country, http://www.businessinsider.com/why-china-mines-more-bitcoin-than-any-other-country-2017-12

[28] Let's Encrypt, How it Works, https://letsencrypt.org/how-it-works/

[29] Ethereum Transactions Historical Chaart, https://bitinfocharts.com/comparison/ethereum-transactions.html

[30] Ethereum Mainnet Node Explorer, https://www.ethernodes.org/network/1

[31] Google Green Computing, https://static.googleusercontent.com/media/www.google.com/en//green/pdfs/google-green-computing.pdf

[32] Ethereum Energy Consumption Index, https://digiconomist.net/ethereum-energy-consumption

[33] AWS S3 Pricing, https://aws.amazon.com/s3/pricing/