

The merits of using Ethereum MainNet as a Coordination Blockchain for Ethereum Private Sidechains

PETER ROBINSON^{1,2} 

¹*Protocol Engineering Group and Systems (PegaSys), ConsenSys, Brisbane, QLD, Australia*
e-mail: peter.robinson@consensys.net

²*School of Information Technology and Electrical Engineering, University of Queensland, Australia*

Abstract

A *Coordination Blockchain* is a blockchain that coordinates activities of multiple private blockchains. This paper discusses the pros and cons of using Ethereum MainNet, the public Ethereum blockchain, as a Coordination Blockchain. The requirements Ethereum MainNet needs to fulfil to perform this role are analyzed within the context of Ethereum Private Sidechains, a private blockchain technology which allows many blockchains to be operated in parallel, and allows atomic crosschain transactions to execute across blockchains. We found that Ethereum MainNet is best suited to storing long-term static data that need to be widely available, such as the Ethereum Registration Authority information. However, due to Ethereum MainNet's probabilistic finality, it is not well suited to information that needs to be available and acted upon immediately, such as the Sidechain Public Keys and Atomic Crosschain Transaction state information that need to be accessible prior to the first atomic crosschain transaction being issued on a sidechain. Although this paper examined the use of Ethereum MainNet as a Coordination Blockchain within reference to Ethereum Private Sidechains, the discussions and observations of the typical tasks a Coordination Blockchain may be expected to perform are applicable more widely to any multi-blockchain system.

1 Introduction

This paper analyses the advantages and disadvantages of using Ethereum MainNet as a *Coordination Blockchain*, by demonstrating the ways in which a Coordination Blockchain may be leveraged in a blockchain network that runs several parallel blockchains. We conduct an in-depth review of the features of Ethereum Private Sidechains, which is an example of such a blockchain system, to explore their potential usage of Coordination Blockchains as an exposition of using Coordination Blockchains more generally. The analysis builds on the Symposium on Distributed Ledger Technology paper *Future of Blockchain* (Robinson, 2018a), and other work on Ethereum Private Sidechains including the following: *Requirements for Ethereum Private Sidechains* (Robinson, 2018b), *Ethereum Registration Authorities* (Robinson, 2018c), *Anonymous Pinning* (Robinson & Brainard, 2019), and *Atomic Crosschain Transactions* (Robinson *et al.*, 2019).

Ethereum MainNet is the largest public deployment of the Ethereum platform. It is a permissionless network, allowing any node to join the network. It is said to offer good authenticity, integrity, and non-repudiation properties, along with an economic system to discourage transaction spamming (Xu *et al.*, 2017; Griffith, 2018). To date, there has been no work that has analyzed all of these assertions. This paper remedies this deficiency by carefully analyzing whether these properties are successfully delivered.

Sidechains are blockchains that rely on a separate blockchain, a Coordination Blockchain, for their overall utility. This could be to enhance security by pinning the state of the sidechain to the Coordination Blockchain (Robinson & Brainard, 2019), for addressing information (Robinson, 2018c), or for storing data that are used across all sidechains. We analyze the appropriateness of using Ethereum MainNet as a Coordination Blockchain for the various features of sidechains, using as a reference Ethereum Private Sidechains.

This paper is organized as follows: the *Background* section briefly introduces Ethereum MainNet, the platform that forms the basis for this paper. We describe the concept of private blockchains and the enterprise version of Ethereum and introduce the concept of block ‘finality’. Next, cryptanalysis of message digest and asymmetric algorithms are reviewed given classical and quantum cryptanalytical techniques. The *Ethereum MainNet Features* section analyzes whether Ethereum MainNet delivers authenticity, integrity, non-repudiation, and crypto-economic anti-spam properties. The *Ethereum Private Sidechains* section describes the features of Ethereum Private Sidechains and their usage of Coordination Blockchains. The *Pros and Cons of using Ethereum MainNet as a Coordination Blockchain* section analyzes the advantages and disadvantages of using Ethereum MainNet as the Coordination Blockchain for each of the Ethereum Private Sidechain features.

2 Background

2.1 Ethereum

2.1.1 Ethereum MainNet

Ethereum (Wood, 2016) is a blockchain platform that allows users to upload and execute computer programmes known as Smart Contracts. Ethereum Smart Contracts can be written in a variety of Turing Complete languages, the most popular being Solidity (2017). Source code is compiled into a bytecode representation. The bytecode can then be deployed using a contract creation transaction. Contracts have a special *constructor* function that only runs when the contract creation transaction is being processed. This function is used to initialize memory and call other contract code. Miners execute the bytecode inside the Ethereum Virtual Machine (EVM). At present, each miner must execute all transactions of all contracts and hold the current value of all the memory associated with all of the contracts. The Ethereum community is actively working on methodologies to scale the Ethereum network by sharding the blockchain (Sharding FAQ, 2018).

Ethereum transactions update the state of the distributed ledger but do not return values. They fall into three categories: Ether transfer, contract creation, and calling a function on a contract. Ether transfer transactions move Ether from the user’s account to another account. Contract creation transactions put code into the distributed ledger and call the constructor of the contract code, setting the contract data’s initial state. Function call transactions call a function on a contract and result in updated state. Contract creation and function call transactions also allow Ether to be transferred. All types of transactions must be signed by a private key corresponding to an account and include a nonce value that prevents replay attacks. In addition to Ethereum transactions, ‘View’ function calls can be executed on the Smart Contract code. These View function calls return a value and do not update the state of the Smart Contract.

Executing code and accessing resources, such as memory, cost certain amounts of ‘Gas’. The ‘Gas Cost’ of executing code is closely tied to the real-world cost of executing each type of instruction. Miners preferentially mine transactions that are prepared to pay a higher ‘Gas Price’. Accounts instigating transactions specify the ‘Gas Price’ they are prepared to pay for their transaction and specify the maximum amount of gas a transaction can use known as ‘Start Gas’. This commits an account holder to paying up to a certain amount of Ether for the transaction. Any unused gas is returned to the account holder at the end of the transaction. Transactions that run out of gas prior to completion are aborted, with all of the gas being expended.

In the Ethereum public network, ‘MainNet’, all contract code and data are readable by any user of any node that connects to the network. Smart Contracts on Ethereum MainNet can only perform

permissioning in contract code, limiting which accounts can update the state of a contract. However, there is no mechanism to limit which users can read contract code and data.

2.1.2 Private blockchains and Enterprise Ethereum

Private blockchains are blockchain networks that are established between nodes operated by enterprises (Robinson, 2018b). Only permissioned nodes belonging to participating enterprises are allowed to join the private blockchain's peer-to-peer network and only permissioned accounts belonging to participating enterprises are allowed to submit transactions to the nodes. These blockchains provide the privacy and permissioning required by enterprises (Enterprise Ethereum Alliance, 2018).

The need for security and permissioning features over and above what is available in standard Ethereum (Enterprise Ethereum Alliance, 2018) has led to a range of platforms being developed. J.P. Morgan (2018) developed Quorum, a fork of the Golang Ethereum implementation called Geth (Ethereum Foundation, 2018). ConsenSys's Protocol Engineering Group, PegaSys created Pantheon (PegaSysEng, 2018), an Ethereum MainNet compatible client that aims to meet the permissioning and privacy requirements of the Enterprise Ethereum Client Specification (Enterprise Ethereum Alliance, 2018). Hyperledger Fabric (Androulaki *et al.*, 2018) is a distributed ledger platform originally created by IBM and now hosted by The Linux Foundation. Similar to Quorum and Pantheon, the platform offers privacy and permissioning features. Whereas Quorum offers Ethereum-based private transactions, Pantheon offers private smart contracts that are private to a set of participants. Hyperledger Fabric offers the ability to host one or more smart contracts on a private blockchain called a 'channel'. Hyperledger Fabric allows multiple channels to be operated on the one network, thus allowing for multiple sets of private contracts between different sets of participants to operate on the one network. An analysis of the merits of Hyperledger Fabric and Quorum has been analysed elsewhere (see *Requirements for Ethereum Private Sidechains*, Robinson, 2018b).

2.1.3 Finality

A block is deemed *final* when it can no longer be changed. All transactions contained within a finalized block are also deemed final.

Ethereum transactions are included in blocks. An Ethereum MainNet miner that solves the Proof of Work cryptographic puzzle can add a block to the end of the blockchain. If two or more miners solve the puzzle simultaneously, then two or more chains are created with common ancestors, and this is known as a *fork* (De, 2019). In Bitcoin, the longest chain of blocks is deemed to be the valid blockchain (Nakamoto, 2008; Courtois, 2014). In Ethereum, the fork choice is solved by means of a modified Greediest Heaviest Observed Subtree (GHOST) protocol (De, 2019) that takes into account the mining power in creating blocks that have links to the main chain, but have become stale (Gencer *et al.*, 2018). These blocks are commonly referred to as *uncle* blocks. The *weight* of a block relates to the number of previous blocks in the chain and uncle blocks. The heaviest chain of blocks is deemed to be the valid blockchain. If an Ethereum MainNet miner becomes aware of a heavier chain than it knew about, it should then only attempt to add blocks to the new chain. Blocks on the old heaviest chain that are not in common with the new longest chain are deemed *reordered*. If none of the transactions in a reordered block have been included in the blocks of the new longest chain, then the block can be included as an *uncle* block. Otherwise, the transactions that are not included in the reordered chain need to be included in a new block. There is no certainty that these transactions will be included in a new block or that transactions in a proposed uncle block will be included in the blockchain.

As more blocks are added to the end of Ethereum MainNet's blockchain, the probability of a miner finding a longer blockchain and reordering the blockchain is reduced (De, 2019). This is because a miner would need to repeatedly solve the Proof of Work cryptographic puzzle for each block faster than all other miners. As the probability of a block being reordered is reduced, the probability of the transactions included in a block being *final* increases. Hence, Ethereum MainNet is said to have *probabilistic finality* (Courtois, 2014).

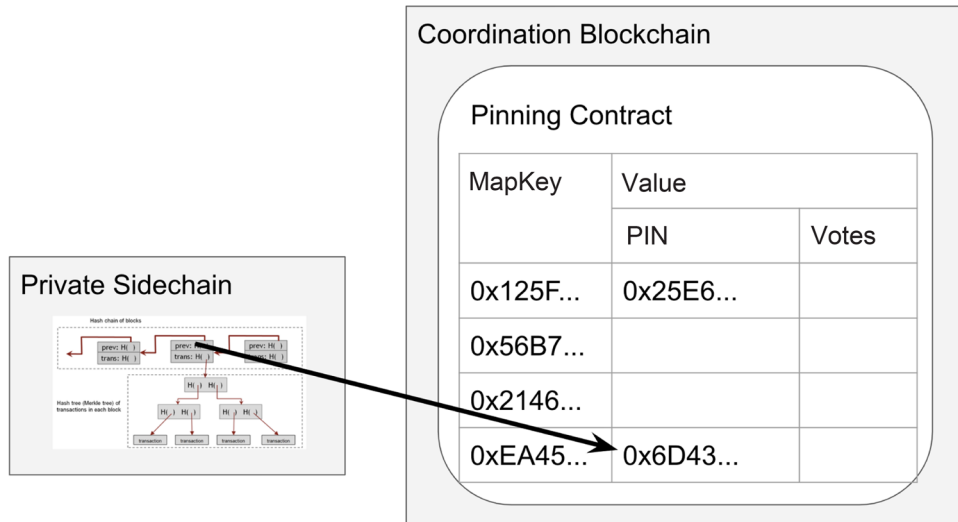


Figure 1 Pinning

Consensus algorithms such as Istanbul Fault Byzantine Tolerant (IBFT) (Lin, 2017) and Istanbul Fault Byzantine Tolerant version 2 (IBFT2) (Saltini, 2019) used in consortium blockchains give *instant* finality, where once a transaction has been included in a block minted by a validator, it can no longer be changed.

2.1.4 Pinning

The state of a blockchain or sidechain can be represented by the Block Hash of a block. The Block Hash of a final block could be submitted to a contract on a Coordination Blockchain at regular intervals (Robinson & Brainard, 2019), as shown in Figure 1. This process is known as *pinning*. Regularly, pinning sidechain state helps to protect minority sidechain participants from state reversion due to collusion by the majority of sidechain participants (Robinson & Brainard, 2019).

2.2 Cryptanalysis

This section provides background material on cryptanalysis that is needed to understand the analysis of the security properties of Ethereum MainNet.

2.2.1 Message digest algorithm cryptanalysis

Message digest algorithms have three main security properties: Preimage Resistance, Second Preimage Resistance, and Collision Resistance. Message digest algorithms are commonly called Cryptographic Hash algorithms, or simply Hash algorithms. Given a Hash algorithm h , the three security properties can be stated as follows:

- *Preimage Resistance*: Given y , it is difficult to determine x such that $y = h(x)$.
- *Second Preimage Resistance*: Given y and x_1 , it is difficult to determine x_2 such that $y = h(x_1) = h(x_2)$ and $x_1 \neq x_2$.
- *Collision Resistance*: It is difficult to determine x_1 and x_2 such that $h(x_1) = h(x_2)$ and $x_1 \neq x_2$.

2.2.2 Classical computing cryptanalysis

Gordon Moore, co-founder of Intel, stated in his 'Moore's Law' that the number of transistors on an integrated circuit doubles approximately every 2 yr Moore (1995). With the increased number of transistors has come a decrease in transistor size, which has resulted in decreased power consumption per transistor. This has resulted in an increase in computation power, while keeping the power consumption

relative static over a 50-yr period. This rate of increase of computation power and decrease of transistor size though slowing is still continuing (Cunningham, 2016). Additionally, new alternative approaches are being developed to deliver increased computational power (Simonite, 2016).

Classical computational power can be used to break algorithms such as message digest algorithms by trying all possible combinations using a ‘Brute Force’ attack. Complexity theory predicts how many attempts are likely to be needed to break an algorithm. For message digest algorithms, using classical computing power, the complexity of breaking an algorithm’s Preimage Resistance or Second Preimage Resistance property is $O(N)$, where N is the number of combinations of the digest output, whereas the complexity of breaking an algorithm’s Collision Resistance is $O(\sqrt{N})$.

The USA’s National Institute of Standards and Technology (NIST) defines *Security Strength* (Barker, 2016) as ‘A number associated with the amount of work (i.e. the number of operations) that is required to break a cryptographic algorithm or system’. Security Strength and complexity are related. The Security Strength of a message digest algorithm’s Preimage and Second Preimage Resistance properties is $\log_2 N$ and the Collision Resistance Security Strength is $\log_2 \sqrt{N}$. Recall that $\log_2 N$ corresponds to the message digest output length in bits. As such, the algorithm SHA-256’s Preimage and Second Preimage Security Strength is 256 bits and its Collision Resistance Security Strength is 128 bits, assuming classical computers (Barker, 2016).

In some instances, a message digest output is truncated. For example in Ethereum, Keccak-256 is used to generate account numbers with the output truncated from 256 to 160 bits. In this usage, the analysis of Security Strength remains unchanged: the complexity and hence Security Strength relates to the number of possible values of the digest output. If a message digest output is truncated, then the Security Strength of the overall algorithm is proportionally reduced.

NIST defines algorithms with Security Strengths of 80, 112, 128, 192, and 256 bits (Barker & Roginsky, 2015). NIST have mandated the phasing out of 80-bit Security Strength algorithms in 2010 and, based on Moore’s Law, had indicated the phasing out of 112-bit Security Strength algorithms by 2030.

2.2.3 Quantum computing cryptanalysis

Quantum computers are expected to allow all currently used popular asymmetric cryptographic algorithms to be defeated and are expected to reduce the Security Strength of message digest and symmetric cipher cryptographic algorithms (Chen *et al.*, 2016). Aggarwal *et al.*, (2018) estimate that ECC 256-bit schemes will be able to be compromised with a Quantum computer using the Shor (1994) algorithm in less than 10 minutes sometime between 2027 and 2040.

Grover’s (1996) algorithm provides a speedup for database search style algorithms, such as searching for a message digest preimage or second preimage. Using Grover’s algorithm, the complexity of message digest algorithm’s Preimage or Second Preimage Resistance properties is reduced from $O(N)$ to $O(\sqrt{N})$. This means that the Security Strength assuming a sufficiently powerful quantum computer is half that when compared to the Security Strength due to classical computing power.

(Brassard *et al.*, 1997) claimed to have developed an algorithm for use with quantum computers that reduces the complexity of finding message digest collisions to $O(\sqrt[3]{N})$. Bernstein (2009) has refuted this claim, stating that there is no real advantage provided by Brassard and Tapp’s algorithm given the cost–performance analysis over classical computing power. However, Aaronson and Shi (2004) have determined a tight lower bound for the complexity of the collision problem as $O(\sqrt[3]{N})$. As such, despite Bernstein’s refutation of Brassard and Tapp’s algorithm, it can be conjectured that another algorithm may be found that meets the theoretical bound, which has a better cost–performance metric.

Despite the reduced Security Strength offered by message digest algorithms, assuming a quantum computer, they are unlikely to be a point of weakness in the near term. Developing a complex quantum computer that can defeat message digest algorithms is expected to be significantly more complex than developing one to defeat ECC 256-bit (Mosca, 2015). As such, it is likely that a quantum computer that can be used to attack message digest algorithms will not be available until at least the 2030s.

2.2.4 Algorithmic weaknesses

Researchers search for weaknesses in algorithms. These weaknesses when found can reduce the effective Security Strength offered by the algorithm. For example, various weakness have been found in the MD-5 message digest algorithm (Wang & Yu, 2005; Sasaki & Aoki, 2009). It is impossible to predict if a weakness in an algorithm such as Keccak-256 will be found, and the degree to which the algorithm would be weakened with such a compromise. Algorithmic weaknesses will not be considered in the analysis of Ethereum MainNet given the uncertainty as to whether such weakness will be found, when they will be found, and the impact such weaknesses might have.

3 Ethereum MainNet features

This section discusses in detail the features of Ethereum MainNet that are important to its usage as a Coordination Blockchain.

3.1 Authentication

The International Telecommunications Union (ITU 2003) define *authentication* in X.805 as:

...serves to confirm the identities of communicating entities. Authentication ensures the validity of the claimed identities of the entities participating in communication (e.g., person, device, service or application) and provides assurance that an entity is not attempting a masquerade or unauthorized replay of a previous communication.

In the context of Ethereum, this means ensuring Ethereum transactions are directly attributable to participants who operate Ethereum Accounts.

Ethereum transactions are signed using the private key belonging to a participant (Wood, 2016). The public key associated with the private key can be derived from the transaction signature of any transaction signed by the private key. The account number is the 20-byte truncated Keccak-256 message digest of the public key.

In Ethereum, each transaction includes a *nonce* (Wood, 2016). The initial nonce value for each account is zero. The nonce is incremented for each successfully mined transaction. Miners reject transactions with out of order or repeated nonces. Doing this protects Ethereum from transaction replay attacks.

The nonce value is represented as a 64-bit signed number in Geth (Ethereum Foundation, 2018) and Pantheon (PegaSysEng, 2018). Adding one to the maximum representable number would result in the largest negative number. If this situation was not guarded against in the code, it would lead to unexpected results, and possibly an authentication failure. However, 63 bits is large enough such that even if a single account issued every transaction on Ethereum MainNet, and could craft sufficiently small transactions and could have the gas limit increased such that they could execute 1000 transactions per second, the nonce value would not wrap around for 584 million years.

Ethereum private keys are 256 bits long. The signature algorithm ECDSA/Keccak-256 using the secp256k1 curve is used for signing transactions. The secp256k1 curve has been analyzed and found to not have any weaknesses (Mayer, 2016). This signature algorithm provides 128 bits of Security Strength (Barker & Roginsky, 2015) assuming Classical Cryptanalysis. The conversion of the public key to an account number using a 20-byte truncated Keccak-256 message digest offers 160 bits of Security Strength assuming Classical Cryptanalysis, as an attacker would need to exploit the Second Preimage Resistance property of the message digest function to determine another public key which could hash to the same value as the authentic public key. As such, overall the Ethereum signing mechanism provides 128 bits of Security Strength assuming Classical Cryptanalysis.

NIST has issued guidance that usage of algorithms offering 112-bit Security Strength assuming Classical Cryptanalysis should be phased out by 2030 (Barker, 2016). This means that Ethereum's transaction signing technique should be secure well beyond 2030, assuming Classical Cryptanalysis, given its 128-bit Security Strength.

If an attacker had access to a sufficiently powerful Quantum Computer, they could determine private keys associated with the public keys. The attacker could observe transactions that have been submitted and determine the public keys associated with each transaction using the standard ecrecover technique (Wood, 2016). Once an attacker had access to a private key, they could issue arbitrary transactions using that private key. Aggarwal's (2018) analysis indicates that the authenticity of transactions may be able to be compromised in this way some time after 2027.

The Ethereum community have recognized the threat that Quantum Cryptanalysis poses to Ethereum transaction signing. There are plans to roll-out 'Account Security Abstraction' changes that will authenticate transactions programmatically using user supplied code (Buterin, 2015, 2016, 2017). This would allow for users to choose to use Quantum Cryptanalysis resistant algorithms.

In summary, the existing transaction authentication techniques are likely to be secure until at least 2027. Prior to 2027, Ethereum is likely to be upgraded to mitigate the threat of quantum computers, thus ensuring the authenticity of transactions into the future.

3.2 Integrity

ITU (2003) defines *data integrity* as:

... ensures the correctness or accuracy of data. The data is protected against unauthorized modification, deletion, creation, and replication and provides an indication of these unauthorized activities.

In the context of Ethereum, this means ensuring that authenticated transactions and data in the distributed ledger are stored such that they cannot be modified.

Ethereum transactions are combined into blocks using Merkle Patricia trees (Wood, 2016). Similarly, data in the distributed ledger are protected using Merkle Patricia trees. Compromising values in the Merkle Patricia trees would require breaking the Second Preimage Resistance property of Keccak-256. This is unlikely to occur in foreseeable future using either Quantum or Classical Cryptanalysis techniques. However, there is always the possibility that a weakness in Keccak-256 will be found.

3.3 Non-Repudiation

ITU (2003) defines *non-repudiation* as:

...provides means for preventing an individual or entity from denying having performed a particular action related to data by making available proof of various network-related actions (such as proof of obligation, intent, or commitment; proof of data origin, proof of ownership, proof of resource use). It ensures the availability of evidence that can be presented to a third party and used to prove that some kind of event or action has taken place.

In the context of Ethereum, this means ensuring that authenticated transactions are stored such that they can not be revoked.

Ethereum blocks are linked together using Keccak-256 message digests. Compromising this linkage would require breaking the Preimage Resistance property of Keccak-256, which is unlikely to occur in foreseeable future.

As discussed in Section 2.1.3, Finality, Ethereum MainNet has *probabilistic finality*. When blocks are added to the blockchain after a block containing a transaction, the probability of a miner proposing a heavier chain that does not include the block decreases. The number of blocks added after a block is known as the number of *block confirmations*. Nakamoto (2008) showed the probability of a Bitcoin block being removed after six blocks, assuming an attacker has 10% of the mining power was 0.00024. A greater number of *block confirmations* should be observed if an attacker was assumed to have a greater percentage of the total mining power available to them, or if the user wished to have greater certainty that the block was not going to be removed.

In 2016, Gervais *et al.* (2016) determined that 37 Ethereum MainNet block confirmations were needed to offer the same level of security as six Bitcoin block confirmations, assuming Ethereum was being attacked with 30% of mining power. Since 2016, the mining power devoted to Ethereum has increased considerably such that a 30% attack now seems inconceivable. Major miners are unlikely to attack their own network as this would risk devaluing the cryptocurrency they are mining (Bitcoin Magazine, 2019; Niu & Feng, 2019). The maximum hash power which can be rented in a straightforward way is 5% (Crypto51, 2019). Purchasing hardware to generate 30% hash power (174TH/s; EthStats, 2019) would cost in excess of US\$400 million (GPU, 2019).

Scaling the results of Gervais *et al.*'s (2016) work based on the changed mining rewards of Bitcoin and Ethereum, the changed valuations, and allowing for a 10% mining power attack, indicates that eight Ethereum block confirmations correspond to six Bitcoin confirmations. Using a different methodology, Buterin (2019) determined that 6–12 confirmations were required to deem a transaction final, depending on the level of risk a user was prepared to assume.

Based on a 14 second target block time and assuming 12 confirmations, a block on Ethereum MainNet could be deemed final in approximately 3 minutes. The finality time is not a precise number as the block time is randomly distributed with an average of 14 seconds.

When Ethereum MainNet client vendors and miners agree to changes in the Ethereum protocol, the system is updated via changes known as *Hard Forks*. A Hard Fork requires all Ethereum MainNet client vendors to release updated software which will activate new functionality at a certain Ethereum MainNet block number. For the Spurious Dragon Hard Fork in November 2016 (Ogundeji, 2019), the changes were implemented slightly differently. This resulted in the Ethereum MainNet blockchain forking for some hours (CCN, 2019). The fork is resolved once the vendors' software has been corrected. However, it is possible that a transaction which was part of a block accepted into the fork which was discarded was reverted and not resubmitted to the blockchain. This type of forking and state reversion due to mismatched feature implementation is much less likely to occur now and in the future than it did in 2016 as Ethereum MainNet clients undergo significantly more review and testing than they did in 2016 (De Angelis, 2018; CoinDesk, 2019).

If an attacker could dedicate 51% of the total mining power to attacking the network, they would be able to mount a *51% Attack* (Hertig, 2018). This would allow the attacker to rewrite the history of the blockchain. The three largest Ethereum MainNet miners could collude to mount such an attack. However, these miners are disincentivized to do such an attack as this would adversely affect confidence in Ethereum MainNet. This would lead to a dramatic drop in the value of Ether (Bitcoin Magazine, 2019; Niu & Feng, 2019), substantially decreasing the value of their Ether and their Ethereum infrastructure investments.

Though the Ethereum MainNet system typically can not be modified, after a re-entrancy bug was exploited in the Decentralized Autonomous Organization attack (Atzei *et al.*, 2017), the system was modified to reverse the results of the attack. Doing this caused some to question trust in blockchain systems and Ethereum MainNet in particular (Spode, 2017). However, this type of *irregular state change* (Khoo, 2016) to reverse the results of such an attack appear unlikely to occur again in Ethereum MainNet. Despite a bug in the Parity Wallet contract that resulted in hundreds of millions of dollars of funds becoming inaccessible, proposals to alter history to restore the funds have been refused (Johnson, 2018; Schoedon, 2018).

3.4 *Crypto economic anti-spam*

As described in Section 2.1.1, each transaction on Ethereum MainNet costs Gas to execute, which participants pay for with Ether. Ethereum MainNet currently aims to produce new blocks every 14 seconds to every 14 seconds with 8 million Gas available for each block (EthStats, 2019). Each transaction has as a minimum cost, the transaction fee, that is currently 21 000 Gas. Simple balance transfers between accounts just cost the transaction fee, whereas complex function calls can cost more than a million gas.

As the block gas limit is 8 million, it means that no transaction can use more than 8 million gas. This translates to Ethereum MainNet supporting between four transactions per minute and 27 transactions per second. A typical simple transaction, *adding a Pin to a pinning contract*, costs 64972 Gas (Robinson & Brainard, 2019). Given the 8 million Gas limit, 8.8 of these transactions could execute per second.

Participants are disincentivized from flooding the network with transactions as each transaction has an economic cost. The cost of Gas depends on the block utilization (Ethereum Gas Station, 2018). Historically, the Gas price has spiked high when block utilization has been high (Polites, 2017). If many entities attempted to issue *adding a Pin to a pinning contract* transactions regularly, such that the block utilization was high, then the cost of issuing the *adding a Pin to a pinning contract* transactions would increase. This would incentivize the entities to find alternatives, such as reducing the frequency of submitting the transactions.

3.5 Summary

Based on the analysis in this section, it can be said that Ethereum MainNet contains transactions for which the authenticity and integrity are certain. Once 12 blocks have been appended to the block containing a transaction, the probability of the blockchain being reorganized such that the transaction is reverted is small. As such, Ethereum MainNet offers strong non-repudiation properties. Ethereum's Gas mechanism operates as an effective anti-spam tool.

4 Ethereum private sidechains

Ethereum Private Sidechains are Ephemeral, On-demand, Permissioned, Private, Confidential, blockchains that allow for Atomic Crosschain Transactions. They are Ephemeral in that they are created, they operate, and then they can be archived when they are no longer needed. Their On-demand nature allows them to be created when needed between parties that have no prior relationship. Permissioning ensures that only authorized nodes are able to join a sidechain. Their design is such that to the greatest extent possible, their membership and their transactions are kept Private. Confidentiality is ensured by encrypting the sidechain data when being communicated between nodes and stored on nodes. Atomic Crosschain Transactions enable transactions that update state across sidechains atomically.

Ethereum Private Sidechains have been described in terms of their requirements (Robinson, 2018b) and aspects of their technology (Robinson, 2018c; Robinson & Brainard, 2019; Robinson *et al.*, 2019). This paper is the first to present this technology holistically. Additionally, this paper introduces the idea of pinning the final state of a sidechain prior to archiving, thus allowing the sidechain to be reinstated if needed, and introduces the idea of using Ethereum MainNet gas pricing as a mechanism for rate control of Atomic Crosschain Transactions.

4.1 Ephemeral

Ethereum Private Sidechains are *Ephemeral*: they are created, they are used for a period, and then archived when they are no longer required. This limited lifespan matches many real world requirements, such as *Letters of Credit* and other business deals, which have a limited lifespan. The ability to archive the blockchain data in a sidechain is in contrast to existing blockchain technologies that are designed to be operational indefinitely.

The lifespan of a sidechain could vary widely. For usages in which sidechains are used to deploy a contract and automatically negotiate a deal, it might only be needed for some minutes, hours or days. Other usages, such as an *Oracle*, require a long or indefinite lifespan. Indefinite lifespans can be accommodated by never archiving the sidechain.

While a sidechain is operational, the sidechain could be pinned to a Coordination Blockchain at regular intervals (Robinson & Brainard, 2019). Regularly pinning sidechain state helps to protect minority sidechain participants from state reversion due to collusion by the majority of sidechain participants (Robinson & Brainard, 2019).

A key aspect of Ephemeral sidechains is the requirement to be able to restart the sidechain after archiving. This can be achieved by pinning the last block of the sidechain to a Coordination Blockchain. Now that the Block Hash of the last block has been securely stored in the Coordination Blockchain, the state of the sidechain can then be stored offline. To restart the sidechain, the stored data are compared against the final Block Hash to confirm the correct state is being used to restart the sidechain.

4.2 On-demand between parties with no prior relationship

Ethereum Private Sidechains need to be able to be deployed between parties that have no prior relationship. That is, the parties need to be able to establish a sidechain without knowing each others' node IP addresses, cryptographic keys, or other information required to set-up a secure connection. Establishing sidechains in this dynamic way is in contrast to existing permissioned blockchains that are largely static systems that require complex set-up. For example, set-up of a Quorum (Morgan, 2018) network requires enode addresses (IP addresses and Ethereum account numbers) for each node to be shared out of band with all other nodes. Adding new nodes to the network requires this sharing and manual intervention on each node.

The on-demand sidechain establishment is analogous to a user of a Web browser establishing a secure connection with a Web server by simply entering in a URL such as <https://example.com/>. The user does not know the IP address of the computer corresponding to *example.com* or the public key that can be used to verify the communications emanating from *example.com*. However, using the domain name, some initial trust, and the Domain Name Service (DNS) and Transport Layer Security (TLS) protocols, they are able to establish a secure connection. Similarly, Ethereum Private Sidechains need to be able to establish a secure sidechain using just domain names.

The Ethereum Registration Authorities system is a set of smart contracts that can be used to provide discoverable information to enable establishment of sidechains between organizations with no prior relationship (Robinson, 2018c). A Coordination Blockchain could be used to locate the information using domain names that can be grouped according to different trust levels and different trust relationships. Moreover, a Coordination Blockchain that provides organizations with a secure, decentralized, censorship-resistant mechanism for storing information that can be located using domain names and grouped according to different trust levels and different trust relationships would overcome the limitations of previous technologies that did not provide the security and censorship resistance properties that users of blockchain technologies expect.

4.3 Permissioned

Ethereum Private Sidechains need to be operated by authorized nodes using authorized Ethereum accounts. These requirements match those of the Enterprise Ethereum Client Specification (Enterprise Ethereum Alliance, 2018). The implementation of these requirements do not use a Coordination Blockchain.

4.4 Private

Ethereum Private Sidechains should, to the greatest extent possible, keep their membership private from other sidechains they interact with and from any Coordination Blockchains they use to facilitate their actions.

4.5 Confidential

Ethereum Private Sidechains should encrypt their blockchain and state data such that the transaction information is kept confidential, both when it is communicated between nodes on a sidechain and when it is stored in a node's local data store. The implementation of this feature does not use a Coordination Blockchain.

4.6 Atomic crosschain transactions

Ethereum Private Sidechains technology needs to enable transactions that update state across sidechains atomically (Robinson *et al.*, 2019). That is, if an Atomic Crosschain Transaction is across sidechains A, B, and C, then the state updates related to the transaction are either applied on all sidechains or ignored on all sidechains. A Coordination Blockchain holds a Crosschain Coordination Contract. This contract is used to indicate that an Atomic Crosschain Transaction has commenced, has been committed, or should be ignored. The contract acts as a common time-out reference for all sidechains and helps prevent denial of service attacks. The data in the Crosschain Coordination Contract need to be available until the last sidechain using it has been archived.

The Atomic Crosschain Transaction system uses threshold signatures to prove values across sidechains. The public key that corresponds to the private key shares held by each of the sidechain validators is known as a Sidechain Public Key. This key needs to be available to all sidechains that need to verify values coming from a sidechain. As such, this value should be stored on a Coordination Blockchain. The Sidechain Public Key needs to be re-generated and uploaded to the Coordination Blockchain each time a validator is added or removed from the sidechain. Assuming that sidechain membership is largely static, this regeneration and upload is likely to be a rare event.

5 Pros and cons of using Ethereum MainNet as a Coordination Blockchain

The subsections below analyze the advantages and disadvantages of using Ethereum MainNet as the Coordination Blockchain for the operations of an Ethereum Private Sidechain. The findings of the subsections are summarized in Table 1.

5.1 Private node discovery—Ethereum Registration Authorities

The Ethereum Registration Authorities system (Robinson, 2018c) uses smart contracts on a Coordination Blockchain to enable discovery of sidechain node address and cryptographic key information, as described in Section 4.2. As the information is used to bootstrap a sidechain, it is fundamental to the entire Ethereum Private Sidechain system that this information is authentic.

The data in the Ethereum Registration Authority smart contracts are largely static. That is, the IP address and cryptographic key information, once set, change rarely. Given this largely static data, the economic cost of storing information on Ethereum MainNet would only be incurred rarely. It is likely to cost less than US\$1.00 to set-up an enterprise in the Ethereum Registration Authority system on Ethereum MainNet, based on current prices (Robinson, 2018c).

Sidechain users who wish to establish a sidechain need to be able to access the bootstrap information stored in Ethereum Registration Authority smart contracts for the system to be useful. The information needs to be stored on a permissionless network or a permissioned network that has a black list of banned nodes. Doing this allows users who have no prior relationship with the operators of the Coordination Blockchain to access the information.

5.2 State pinning

A private blockchain state pinning approach should be used to prevent state reversion as described in Section 4.2. Posting Pins to Ethereum MainNet leverages the authenticity, integrity, and non-repudiation properties of Ethereum MainNet. However, submitting transactions costs money. Pinning once per hour for a year would cost US\$508 (Robinson & Brainard, 2019). Additionally, if many sidechains pinned to Ethereum MainNet simultaneously, it would cause transaction congestion. Another issue with pinning directly to Ethereum MainNet is that any disputes that occur would need to occur on Ethereum MainNet, thus making the participant list of the sidechain public.

Pins could be posted directly to a smart contract on Ethereum MainNet or could be posted via a smart contract on an intermediate blockchain using a hierarchical pinning approach (Robinson & Brainard, 2019). Using a hierarchical pinning approach, many private blockchains could treat another private

Table 1 Advantages and disadvantages of using Ethereum MainNet as Coordination Blockchain for Ethereum Private Sidechains

Ethereum Private Sidechain operation	Ethereum MainNet as Coordination Blockchain	
	Advantages	Disadvantages
Discover using Ethereum Registration Authorities	Good authenticity, integrity, and non-repudiation properties Permissionless, public network, enables discovery	
State Pinning & Final State Pinning	Good authenticity, integrity, and non-repudiation properties	Economic cost Increased congestion on Ethereum MainNet Pinning and disputes are public
State Pinning & Final State Pinning via an intermediate private blockchain	Leverage Ethereum MainNet security properties while minimizing cost and congestion Pinning and disputes are not public	Pins take more time to become <i>final</i> on Ethereum MainNet than if pinned directly Sidechain participants must observe all levels of pinning
Sidechain Public Key	Public keys widely available	Significantly delays when first Atomic Crosschain Transactions can be issued
Atomic Crosschain Transaction State	Leverages Ethereum MainNet anti-spam capabilities	Significantly delays when first Atomic Crosschain Transactions can be issued Economic gas cost Increased congestion on Ethereum MainNet

blockchain as a Coordination Blockchain posting Pins to it. This private blockchain could in turn post Pins to another private blockchain or to Ethereum MainNet. This is shown diagrammatically in Figure 2. Pinning to a hierarchy of Coordination Blockchains in this way means that only a small number of Pins on Ethereum MainNet could be used to secure a large number of private blockchains. The cost of submitting Pins to the private blockchain could be either free or significantly less than Ethereum MainNet.

A benefit of pinning directly to Ethereum MainNet, rather than via an intermediate blockchain, is that the pinned state becomes *final* faster. That is, if a Pin is posted to a private blockchain, whose state is in turn pinned to Ethereum MainNet, then the sidechain Pin could be deemed to become *final* only once the private blockchain is pinned to Ethereum MainNet.

Posting Pins via a private blockchain significantly reduces the cost of pinning, as only one blockchain needs to submit transactions to pin its state to Ethereum MainNet, and sidechains can pin to that private blockchain. Doing this reduces the number of transactions on Ethereum MainNet, thus reducing congestion, and means that the cost of submitting transactions is only incurred once for the private blockchain, rather than once for each sidechain.

A disadvantage of posting Pins via a private blockchain is that participants of the sidechain need to observe and be ready to challenge Pins being posted at each level of the hierarchy. If sidechain state Pins are posted directly to Ethereum MainNet, then the sidechain participants only need to observe the pinning contract on Ethereum MainNet.

An additional benefit of pinning to a private blockchain is that the chain's permissioning could be set such that only certain nodes could view the blockchain and only certain accounts could submit

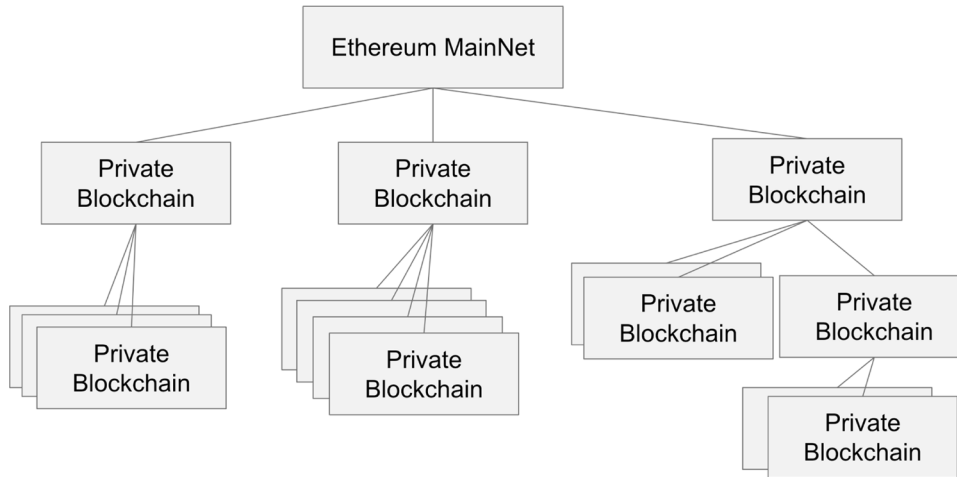


Figure 2 Hierarchical pinning

transactions to the blockchain. Pinning directly to Ethereum MainNet means that the organization pinning to the contract is public. If there is a dispute, then masked participants will need to unmask themselves, and thus link themselves to the sidechain and the other organizations on the sidechain. If an intermediate blockchain was used, then the pinning and any disputes could happen in a more private setting.

5.3 Final state pinning for archiving

Final State Pinning is the same as State Pinning, with the exception that rather than the pinning being on an ongoing basis, it is just to pin the final state of a sidechain prior to archiving, as described in Section 4.2. As such, the advantages and disadvantages are similar to those described in the previous section. As only one pin is posted, the concerns over having to observe pins on a private blockchain in addition to Ethereum MainNet are not significant as the observation is for a single event. Similarly, concerns over cost of posting pins to Ethereum MainNet and congestion are reduced. As such, the advantages are reduced to the pin becoming final sooner and the disadvantages are reduced to any dispute over the value of the pin being public.

5.4 Sidechain public keys

As described in Section 4.6, the Atomic Crosschain Transactions feature needs Sidechain Public Keys to be stored on a Coordination Blockchain. The Sidechain Public Keys need to be stored in a contract (Robinson *et al.*, 2019) that allows voting on new public keys and allows masked and unmasked participants. Given the participants are the same as those for the pinning scheme, it makes sense for these to be stored in the same contract as the pinning information. Keeping the logic in the same contract for pinning and holding the Sidechain Public Keys is useful as it means that membership changes need to only occur in one contract. However, the Sidechain Public Keys need to be visible by all sidechains that wish to verify information coming from the sidechain, whereas the pinning information need only be visible by sidechain participants and government regulators who would be appealed to in case of dispute.

Given the Sidechain Public Key is likely to be set once only, the economic cost of storing the key is likely to only be incurred once. No analysis of the gas cost of setting a Sidechain Public Key has been undertaken yet. However, given the small size of the public keys, 48 bytes, the incremental gas cost of storing the public key is likely to be in the order of 60 000 Gas, assuming the voting infrastructure has already been set-up. However, if the voting infrastructure did need to be set-up, the gas cost could be much larger.

If a sidechain was short lived, then incurring the cost of setting up the voting infrastructure and posting the Sidechain Public Key to Ethereum MainNet could be deemed considerable. However, if the sidechain was long lived, then this relative cost might not be deemed as significant.

A disadvantage of using Ethereum MainNet to hold Sidechain Public Keys is transactions take at least 12 blocks before they should be deemed *final* (see Section 3.3). This means that, given a target block time of fourteen seconds, users could not use the Sidechain Public Keys for Atomic Crosschain Transactions for 3 minutes after the transaction that posts the Sidechain Public Key is included in a block on Ethereum MainNet.

5.5 Atomic crosschain transaction state

The Atomic Crosschain Transactions capability described in Section 4.6 uses a Crosschain Coordination Contract to control when a crosschain transaction has started, been committed, or should be ignored. This information need to be available to all validators on all sidechains involved in the crosschain transaction. The information in the contract needs to be available until the last sidechain using the contract is archived.

Storing the Atomic Crosschain State on Ethereum MainNet means that each Atomic Crosschain Transaction costs money to execute. This economic cost could be seen as an advantage, as it provides an anti-spam control external to the sidechain system. However, forcing enterprises to incur a cost for each crosschain transaction is likely to be viewed as an unnecessary cost.

Additional issues with storing the Atomic Crosschain State on Ethereum MainNet is that this would leak the participants of a sidechain, as a transaction would need to be submitted linking the sidechain and the participant. Furthermore, this would leak the rate that the participant was issuing crosschain transactions.

In a similar way that storing Sidechain Public Keys on Ethereum MainNet delays when the first Atomic Crosschain Transaction can be issued, as discussed in Section 5.4, storing Atomic Crosschain Transaction State could delay the effective start of each transaction. This is because sidechain participants might want to wait for blocks that contain transactions that indicate the Atomic Crosschain Transaction start to be final prior to acting on the start indication.

6 Conclusion

Coordination Blockchains perform various coordination tasks in private blockchain systems. We used Ethereum Private Sidechains as an exposition of such a system, highlighting the features of Ethereum Private Sidechains and discussing each feature's need to leverage a Coordination Blockchain. Based on the unique requirements of each feature and coordination activity, we examine whether public Ethereum MainNet would be a suitable platform for each of those tasks.

We found that Ethereum Registration Authority smart contracts of Ethereum Private Sidechains need to store long term data that have to be available in a permissionless blockchain. Ethereum MainNet would therefore be well suited for this task, as it is a permissionless blockchain that incentivizes good behaviour using crypto economics, and provides good authenticity, integrity, and non-repudiation properties. Ethereum MainNet's strong security properties are also useful for State Pinning and in particular Final State Pinning, where the data need to be stored securely for long periods of time. However, pinning directly to Ethereum MainNet could lead to congestion on Ethereum MainNet, would incur high costs, and would lead to the membership of a sidechain becoming public in the case of a dispute over the value of a Pin. These issues are significantly reduced by pinning via an intermediate private blockchain. However, doing this introduces other issues, such as participants having to observe pinned values at multiple levels in the pinning hierarchy and the pinned values taking longer to become final. Ethereum MainNet is not an appropriate location for Coordination Blockchain information that needs to be final quickly, such as Sidechain Public Keys and Atomic Crosschain Transaction State.

Acknowledgements

This research has been undertaken while I have been employed full-time at ConsenSys and have been completing my PhD part-time at University of Queensland. The author acknowledge the support of his

PhD supervisor Dr Marius Portmann. The author thank Dr Catherine Jones, Horacio Mijail Anton Quiles, David Hyland-Wood, and Sandra Johnson for reviewing this paper and providing astute feedback.

References

- Aaronson, S. & Shi, Y. 2004. Quantum lower bounds for the collision and the element distinctness problems. *Journal of ACM* **51**(4), 595–605. doi: [10.1145/1008731.1008735](https://doi.org/10.1145/1008731.1008735).
- Aggarwal, A., Brennen, G., Lee, T., Santha, M. & Tomamichel, M. 2018. Quantum attacks on bitcoin, and how to protect against them. *Ledger* **3**. issn: 2379–5980. doi: [10.5195/ledger.2018.127](https://doi.org/10.5195/ledger.2018.127). <https://www.ledgerjournal.org/ojs/index.php/ledger/article/view/127>.
- Enterprise Ethereum Alliance. 2018. *Enterprise Ethereum Alliance - Enterprise Ethereum Client Specification V2*. http://entethalliance.org/wp-content/uploads/2018/10/EEA_Enterprise_Ethereum_Client_Specification_V2.pdf.
- Androulaki, E., Barger, A., Bortnikov, V., Cachin, C., Christidis, K., De Caro, A., Enyeart, D., Ferris, C., Laventman, G., Manevich, Y., Muralidharan, S., Murthy, C., Nguyen, B., Sethi, M., Singh, G., Smith, K., Sorniotti, A., Stathakopoulou, C., Vukolić, M., Cocco, S. W. & Yellick, J. 2018. Hyperledger fabric: a distributed operating system for permissioned blockchains. In *EuroSys '18: Thirteenth EuroSys Conference, 2018*. <http://vukolic.com/fabric.pdf>.
- Atzei, N., Bartoletti, M. & Cimoli, T. 2017. *A Survey of Attacks on Ethereum Smart Contracts (SoK)*. Springer, 164–186. ISBN: 978-3-662-54455-6.
- Barker, E. *NIST Special Publication 800-57 Part 1, Revision 4, Recommendation for Key Management, Part 1: General*. <http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-57pt1r4.pdf>.
- Barker, E. & Roginsky, A. 2015. *NIST Special Publication 800-131A, Revision 1, Transitions: Recommendation for Transitioning the Use of Cryptographic Algorithms and Key Lengths*. <http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-131Ar1.pdf>.
- Bernstein, D. J. 2009. *Cost analysis of hash collisions: will quantum computers make SHARCS obsolete?* Lausanne, Switzerland.
- Bitcoin Magazine – Binance Hacked for 40M, CEO Backpedals on Recoup Via Block Reorganization*. <https://bitcoinmagazine.com/articles/binance-hacked-40mceo-backpedals-recoup-block-reorganization/> (accessed 07 June 2019).
- Brassard, G., Hoyer, P. & Tapp, A. 1997. Quantum cryptanalysis of hash and claw-free functions. *SIGACT News* **28**(2), 14–19. doi: [10.1145/261342.261346](https://doi.org/10.1145/261342.261346).
- Buterin, V. *On Settlement Finality*. <https://blog.ethereum.org/2016/05/09/on-settlement-finality/> (accessed 08 June 2019).
- Buterin, V. Proposed Initial Abstraction Changes for Metropolis 86. Edited by Ethereum Improvement Proposal. github.com. <https://github.com/ethereum/EIPs/issues/86>.
- Buterin, V. 2017. *Tradeoffs in Account Abstraction Proposals - Architecture - Ethereum Research*. <https://ethresearch.ch/t/tradeoffs-in-account-abstractionproposals/263> (accessed 14 May 2019).
- Buterin, V. 2015. *Understanding Serenity, Part I: Abstraction*. Ethereum Blog. <https://blog.ethereum.org/2015/12/24/understanding-serenity-part-i-abstraction/>.
- CCN. *Ethereum's Blockchain Accidentally Splits*. <https://www.ccn.com/ethereums-blockchain-accidentally-splits> (accessed 07 June 2019).
- Chen, L., Jordan, S., Liu, Y.-K., Moody, D., Peralta, R., Perlner, R. & Smith-Tone, D. 2016. *NISTIR 8105 DRAFT, Report on Post-Quantum Cryptography*. http://csrc.nist.gov/publications/drafts/nistir-8105/nistir_8105_draft.pdf.
- Crypto51 – Cost of a 51% Attack for Different Cryptocurrencies*. <https://www.crypto51.app/> (accessed 07 June 2019).
- CoinDesk – Ethereum Developers Start Search for New Hard Fork Coordinator*. <https://www.coindesk.com/ethereum-developers-start-search-for-new-hard-forkcoordinator> (accessed 11 June 2019).
- Courtois, N. T. 2014. *On the Longest Chain Rule and Programmed Self-destruction of Crypto Currencies*. [arXiv:1405.0534](https://arxiv.org/abs/1405.0534). <http://arxiv.org/abs/1405.0534> (accessed 06/11/2019).
- Cunningham, A. 2016. *Intel Unveils Kaby Lake, Its First Post-“tick-tock” CPU Architecture*. <https://arstechnica.com/gadgets/2016/08/intel-unveils-kaby-lakeits-first-post-tick-tock-cpu-architecture/>.
- De, N. *Ethereum Developers Delay Mining Algorithm Change for Code Audit – CoinDesk*. <https://www.coindesk.com/ethereum-developers-delay-miningalgorithm-change-for-code-audit> (accessed 11 June 2019).
- De Angelis, S. 2018. *Assessing Security and Performances of Consensus Algorithms for Permissioned Blockchains*. [arXiv:1805.03490](https://arxiv.org/abs/1805.03490). <http://arxiv.org/abs/1805.03490> (accessed 11 June 2019).
- Ethereum Foundation. 2018. *Go Ethereum Github Repository*. <https://github.com/ethereum/go-ethereum>.
- Ethereum Gas Station. 2018. <https://ethgasstation.info/>.
- EthStats: Network Statistics - Instant Ethereum Blockchain Monitoring*. 2019. <https://ethstats.io/> (accessed 27 May 2019).

- Gencer, A. E., Basu, S., Eyal, I., van Renesse, R. & Sirer, E. G. 2018. *Decentralization in Bitcoin and Ethereum Networks*. [arXiv:1801.03998](https://arxiv.org/abs/1801.03998). [http://arxiv.org/abs/1801.03998](https://arxiv.org/abs/1801.03998) (accessed 11 June 2019).
- Gervais, A., Karame, G. O., Wüst, K., Glykantzis, V., Ritzdorf, H. & Capkun, S. 2016. On the security and performance of proof of work blockchains. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, CCS '16*, 3–16. ACM. ISBN: 978-1-4503-4139-4. doi: [10.1145/2976749.2978341](https://doi.org/10.1145/2976749.2978341). [http://doi.acm.org/10.1145/2976749.2978341](https://doi.org/10.1145/2976749.2978341).
- GPU Mining Hashrates. <https://www.buriedone.com/gpu-hashrates> (accessed 07 June 2019).
- Griffith, A. 2018. *Ethereum Meta Transactions - Kauri*. <https://kauri.io/article/8b1e4a4cddb240c892925cd4779e34ca/ethereum-meta-transactions> (accessed 01 June 2019).
- Grover, L. K. 1996. *A Fast Quantum Mechanical Algorithm for Database Search*, Philadelphia, Pennsylvania, USA. doi: [10.1145/237814.237866](https://doi.org/10.1145/237814.237866).
- Hertig, A. 2018. *Blockchain's Once-Feared 51% Attack Is Now Becoming Regular*. <https://www.coindesk.com/blockchains-feared-51-attack-now-becoming-regular/>.
- International Telecommunications Union. 2003. *X.805 Security Architecture for Systems Providing End-to-End Communications*. <https://www.itu.int/rec/T-REC-X.805-200310-I/en>.
- Johnson, N. 2018. *On Parity's Proposed Changes to SELFDESTRUCT Behaviour*. <https://medium.com/@weka/on-paritys-proposed-changes-to-selfdestruct-behaviour-c3f0e5bc0f49>.
- Khoo, B. 2016. *The DAO - Give a Summary of the Fork State Changes in Block 1920000 - Ethereum Stack Exchange*. <https://ethereum.stackexchange.com/questions/7832/give-a-summary-of-the-fork-state-changes-in-block-1920000> (accessed 14 June 2019).
- Lin, Y.-T. 2017. *EIP 650: Istanbul Byzantine Fault Tolerance*. <https://github.com/ethereum/EIPs/issues/650>.
- Mayer, H. 2016. *ECDSA Security in Bitcoin and Ethereum: a Research Survey*. <https://blog.coinfabrik.com/wp-content/uploads/2016/06/ECDSA-Security-in-Bitcoin-and-Ethereum-a-Research-Survey.pdf> (accessed 07 June 2019).
- Moore, G. 1995. *Lithography and the Future of Moore's Law*. <http://www.lithoguru.com/scientist/CHE323/Moore1995.pdf>.
- Morgan, J. P. 2018. *Quorum Source Code*. <https://github.com/jpmorganchase/quorum>.
- Mosca, M. 2015. *Cybersecurity in an Era with Quantum Computers: Will We Be Ready?* <https://eprint.iacr.org/2015/1075.pdf>.
- Nakamoto, S. 2008. *Bitcoin: A Peer-to-Peer Electronic Cash System*. <https://bitcoin.org/bitcoin.pdf>.
- Niu, J. & Feng C. 2019. *Selfish Mining in Ethereum*. [arXiv:1901.04620v3](https://arxiv.org/pdf/1901.04620v3). <https://arxiv.org/pdf/1901.04620.pdf> (accessed 07 June 2019).
- Ogundeji, O. *Ethereum Hard Fork No. 4 Has Arrived as DOS Attacks Intensify*. <https://cointelegraph.com/news/ethereum-hard-fork-no-4-has-arrived-as-dos-attacks-intensify> (accessed 11 June 2019).
- PegaSysEng, 2018. *Pantheon Ethereum Client Github Repository*. <https://github.com/PegaSysEng/pantheon>.
- Polites, J. 2017. *What CryptoKitties Reveals About Ethereum Scalability Issues*. <https://www.ccn.com/cryptokitties-reveals-ethereum-scalability-issues/>.
- Robinson, P. 2018a. *Future of blockchain - an Ethereum perspective*. In *Proceedings of the 3rd Symposium on Distributed Ledger Technology*. <https://symposiumdlt.org/sdlt2018.zip>.
- Robinson, P. 2018b. *Requirements for Ethereum Private Sidechains*. <http://adsabs.harvard.edu/abs/2018arXiv180609834R>.
- Robinson, P. 2018c. *Using Ethereum registration authorities to establish trust for Ethereum Private Sidechains*. *The Journal of the British Blockchain Association*. doi: [10.31585/jbba-1-2-\(6\)2018](https://doi.org/10.31585/jbba-1-2-(6)2018). [https://doi.org/10.31585/jbba-1-2-\(6\)2018](https://doi.org/10.31585/jbba-1-2-(6)2018).
- Robinson, P. & Brainard, J. 2019. *Anonymous State Pinning for Private Blockchains*. <https://arxiv.org/abs/1903.02752>.
- Robinson, P., Hyland-Wood, D., Saltini, R., Johnson, S. & Brainard, J. 2019. *Atomic Crosschain Transactions for Ethereum Private Sidechains*. <https://arxiv.org/abs/1904.12079>.
- Saltini, R. 2019. *IBFT 2.0 Gray Paper Version 0.2*. https://github.com/PegaSysEng/ibft2.x/raw/master/gray_paper/ibft2_gray_paper.pdf.
- Sasaki, Y. & Aoki, K. 2009. *Finding preimages in full MD5 faster than exhaustive search*. In *Advances in Cryptology - EUROCRYPT 2009: 28th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cologne, Germany, April 26–30, 2009. Proceedings*, Joux, A. (ed.). Springer, 134–152. isbn: 978-3-642-01001-9. doi: [10.1007/978-3-642-01001-9_8](https://doi.org/10.1007/978-3-642-01001-9_8). http://dx.doi.org/10.1007/978-3-642-01001-9_8.
- Schoedon, A. 2018. *EIP 999*. <https://github.com/ethereum/EIPs/blob/master/EIPS/eip-999.md>.
- Sharding FAQ. 2018. *On Sharding Blockchains*. <https://github.com/ethereum/wiki/wiki/Sharding-FAQ>.
- Shor, P. W. 1994. *Algorithms for quantum computation: discrete logarithms and factoring*. In *Proceedings 35th Annual Symposium on Foundations of Computer Science*, 124–134. doi: [10.1109/SFCS.1994.365700](https://doi.org/10.1109/SFCS.1994.365700).
- Simonite, T. 2016. *Moore's Law Is Dead. Now What?* MIT Technology Review. <https://www.technologyreview.com/s/601441/moores-law-is-dead-now-what/>.
- Solidity. 2017. <http://solidity.readthedocs.io/en/latest/>.

- Spode, E. 2017. *The Great Cryptocurrency Heist*. <https://aeon.co/essays/trustthe-inside-story-of-the-rise-and-fall-of-ethereum>.
- Wang, X. & Yu, H. 2005. How to break MD5 and other hash functions. In *Advances in Cryptology - EUROCRYPT 2005: 24th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Aarhus, Denmark, May 22–26, 2005. Proceedings*, Cramer, R. (ed.). Springer, 19–35. isbn: 978-3-540-32055-5. doi: [10.1007/11426639_2](https://doi.org/10.1007/11426639_2). http://dx.doi.org/10.1007/11426639_2.
- Wood, G. 2016. *Ethereum: A Secure Decentralized Generalised Transaction Ledger*. Github. <https://github.com/ethereum/yellowpaper>.
- Xu, X., Weber, I., Staples, M., Zhu, L., Bosch, J., Bass, L., Pautasso, C., Rimba, P. 2017. *A Taxonomy of Blockchain-Based Systems for Architecture Design*. Gothenburg. doi: [10.1109/ICSA.2017.33](https://doi.org/10.1109/ICSA.2017.33). <https://ieeexplore.ieee.org/document/7930224>.