

# Flow: Separating Consensus and Compute

Dr. Alexander Hentschel  
alex.hentschel@dapperlabs.com

Dieter Shirley  
dete@dapperlabs.com

Layne Lafrance  
layne@dapperlabs.com

## Abstract

Throughput limitations of existing blockchain architectures are well documented and are one of the most significant hurdles for their wide-spread adoption. Attempts to address this challenge include layer-2 solutions, such as Bitcoin’s Lightning or Ethereum’s Plasma network, that move work off the main chain. Another prominent technique is sharding, i.e., breaking the network into many interconnected networks. However, these scaling approaches significantly increase the complexity of the programming model by breaking ACID guarantees (Atomicity, Consistency, Isolation, and Durability), increasing the cost and time for application development.

In this paper, we describe a novel approach where we split the work traditionally assigned to cryptocurrency miners into two different node roles. Specifically, the selection and ordering of transactions are performed independently from their execution. The focus of this paper is to formalize the split of consensus and computation, and prove that this approach increases throughput without compromising security.

In contrast to most existing proposals, our approach achieves scaling via separation of concerns, i.e., better utilization of network resources, rather than sharding. This approach allows established programming paradigms for smart contracts (which generally assume transactional atomicity) to persist without introducing additional complexity. We present simulations on a proof-of-concept network of 32 globally distributed nodes. While the consensus algorithm was identical in all simulations (a 2-step-commit protocol with rotating block proposer), block computation was either included in a consensus nodes’ regular operations (conventional architecture) or delegated to specialized execution nodes (separation of concerns). Separation of concerns enables our system to achieve a throughput increase by a *factor* of 56 compared to conventional architectures without loss of safety or decentralization.

# Contents

<b>1</b>	<b>Introduction</b>	<b>3</b>
1.1	Terminology . . . . .	4
1.2	Related Work . . . . .	4
<b>2</b>	<b>Architecture Overview</b>	<b>6</b>
<b>3</b>	<b>Theoretical Performance and Security Analysis</b>	<b>8</b>
3.1	Special Role of Consensus Nodes . . . . .	11
<b>4</b>	<b>Performance Simulations</b>	<b>14</b>
4.1	Experimental Setup . . . . .	14
4.1.1	Experiment (I): Flow PoS network with the split of consensus and compute . . . . .	15
4.1.2	Experiment (II): PoS network of nodes with diverse performances . . . . .	16
4.1.3	Experiment (III): PoS network of nodes with uniform performance . . . . .	16
4.2	Experimental Results . . . . .	16
<b>5</b>	<b>Further Work</b>	<b>17</b>
<b>6</b>	<b>Conclusions</b>	<b>17</b>
	<b>Acknowledgments</b>	<b>19</b>
	<b>References</b>	<b>19</b>

# 1 Introduction

The original Bitcoin blockchain uses a consensus model referred to as Nakamoto consensus [1]. It uses a sequential model in which a block is built, mined, and verified, and consensus about it is formed by nodes building subsequent blocks on top of it. While the proof-of-work challenge (mining) that must be solved for every block provides tamper-resistance for the chain, the associated computational effort limits block-production rate and transaction throughput. The throughput limitations of existing blockchain architectures are well documented and among the most significant hurdles for the wide-spread adoption of decentralized technology [2, 3, 4].

The leading proposals for removing the overhead of proof-of-work adapt Byzantine Fault Tolerant (BFT) consensus algorithms [5, 6]. Blockchains using proof-of-work (PoW) require exceptional computational effort and, subsequently, electric power. In contrast, finalizing blocks through BFT consensus is highly efficient but requires a known set of participants, a super-majority of which must be honest.

Combining BFT consensus with proof-of-stake (PoS) [7] allows for the creation of a permissionless network with strong security properties. Under PoS, all participating nodes are required to deposit a (financial) stake that can be taken away if they violate the protocol’s rules. The amount of influence given to each node is proportional to its fraction of total stake. Then the economic pressure (i.e., the stake at risk) to follow the protocol is correlated with a node’s influence. In addition, the deposited stake has the added benefit of preventing Sybil attacks [8, 9, 10]. PoS systems promise to increase the throughput of the chain while also decreasing the total capital costs associated with maintaining the security of the chain. Even including the performance increases through adopting PoS, throughput restrictions are remaining the major challenge for wide-spread adoption [11, 12].

In this paper, we explore a novel approach to increasing the throughput of PoS blockchains. While PoS blockchain proposals remove proof-of-work as the dominant sink of computational effort, they tend to inherit most of their architecture from the proof-of-work systems of the previous generation. In particular, every full node in the network is required to examine and execute each proposed block to update their local copy of the blockchain’s state. As every transaction needs to be processed by every single node, adding nodes to the system provides no benefit in throughput or scale. Instead, adding nodes *reduces* the throughput of most BFT consensus protocols, because the message complexity to finalize a block increases super-linearly with the number of consensus nodes (see section 3.1 for details). Consequently, most PoS blockchains have to make a trade-off between a small consensus committee (weakening security) or a low block production rate (decreasing throughput).

For networks unwilling to compromise either security or decentralization, the most common approach to addressing the scaling problem has been through sharding [2] or moving work off the main chain (e.g. Lightning [13] or Plasma [14] network). Both of these approaches, unfortunately, place significant limitations on the ability of transactions to access state that is distributed throughout the network [15]. These limitations dramatically increase the complexity required for developers who wish to deploy smart contract-based applications.

Our proposal, the Flow architecture, addresses these limitations by fundamentally changing how the blockchain is formed. Flow decouples the selection and ordering of transactions from their execution so that both processes can run in parallel. The decoupling enables significantly higher transaction throughput than other blockchains architectures, without undermining security or participation rates.

Traditional blockchain architectures require a commitment to the result of each block’s state update to be included as part of the consensus process. As a result, every node must reproduce the state-update computation before it can finalize a block. Our finding is that consensus on the order<sup>1</sup> of transactions in the block is all that is required. Once that order is fixed, the resulting computation is determined even though it may not necessarily be known. Thereby, the computational effort of participating in consensus is significantly reduced, even for a very large number of transactions. Once the transaction order is determined, ensuing processes can be delegated to perform the computation itself, without affecting decentralization of the system.

Our main body of work is section 3, where we discuss and prove our central theorem, which states that one can separate the majority of computation and communication load from consensus without compromising security. Section 4 complements the theoretical discussion by benchmarking the Flow architecture on an experimental network. We conclude the paper by outlining the future work that would be required to implement a system based on these ideas in section 5.

The focus of this paper is to formalize the split of consensus and computation and prove that this approach increases throughput while maintaining strong security guarantees. A blockchain designed around the principles outlined in this paper would need to specify a variety of additional mechanisms including:

- a full protocol for verifying computation results,
- details of the consensus algorithm,
- and adequate compensation and slashing mechanics to incentivize nodes to comply with the protocol.

Detailed formalization and analysis of these topics is reserved for follow-up publications.

## 1.1 Terminology

While most current blockchains focus solely on processing financial transactions, we consider a blockchain as a general, Turing-complete, distributed computing platform. Instead of referring to the blockchain as a ledger, we adopt the terminology of a distributed state machine wherein transactions describe transitions between computational states. Furthermore, we use the term **consensus** to refer only to linearizing the order of state transitions (but do not consider agreement about the computational result as a part of consensus).

## 1.2 Related Work

Blockchains supporting Turing-complete computation generally impose an upper limit on the computation within one block, such as Ethereum’s **gas limit**. Such a gas limit, in turn, introduces undesired throughput restrictions. One reason for imposing a gas limit in the first place is to avoid the Verifier’s Dilemma [16]. By setting the gas limit low enough, the time investment for verification is negligible compared to solving the PoW challenge for mining the block. Thereby, the gas limit ensures that performing the verification work does not introduce a pivotal disadvantage for a node to successfully mine the next block.

---

<sup>1</sup>In the full Flow architecture, Consensus Nodes work with transaction batches (**collections**). A block contains collection hashes and a source of randomness, which the Execution Nodes use to shuffle the transactions before computing them. While Consensus Nodes don’t directly compute the transaction order for a block, they *implicitly* determine the order by specifying all the inputs to the deterministic algorithm that computes the order. The detailed protocol is specified in a follow-up publication.

For PoS blockchain, the Verifier’s Dilemma persists when incentives are given for speedy operation. Especially for high-throughput blockchains, the verification of a large number of transactions consumes significant computational resources and time. By separating consensus and computation, the Verifier’s Dilemma for checking correctness of the execution result is of no concern anymore to the consensus nodes<sup>2</sup>. The maximum amount of computation in a block (the block’s gas limit) can now be increased without the consensus nodes affected by Verifier’s Dilemma or slowed down by the computation work. However, the Verifier’s Dilemma still needs to be solved for the verifiers of the computation. Potential solutions include zkSnarks [17], Bulletproofs [18], and TrueBit’s approach [19]. For Flow, we developed **Specialized Proofs of Confidential Knowledge (SPoCKs)** to overcome the Verifier’s Dilemma, which is described in detail in a follow-up paper.

Another limitation for blockchains are the resource demands to store and update the large computational state. In Ethereum, the gas limit is also used to control the rate at which the state grows [20]. By delegating the responsibility to maintain the large state to specialized nodes, hardware requirements for consensus nodes can remain moderate even for high-throughput blockchains. This design increases decentralization as it allows for high levels of participation in consensus by individuals with suitable consumer hardware on home internet connections.

The concept of separating the issue of transaction ordering from the effort of computing the results of the computations has been previously explored in the context of distributed databases [21, 22]. Here, transactions are ordered into a log through a quorum (consensus), and subsequently each node can independently resolve transaction effects. However, these systems are designed for operation in well-maintained data-centers where Byzantine faults are not concern. Specifically, the number of participating nodes is small, and node failure modes are restricted to dropouts.

Within the blockchain space, the Ekiden [23] paper describes a system where consensus is separated from computation, with the goal of preserving the privacy of the contract execution environment. The paper, which explains part of the Oasis blockchain technology [24], notes that this approach leads to improved performance. But it does not quantify the performance gain or prove that the resulting system maintains security.

---

<sup>2</sup>The Verifier’s Dilemma (checking the correctness of the execution result) is of no concern anymore to the Consensus Nodes. However, checking cryptographic signatures and proofs in PoS can still require noticeable computational work. Delay through verifying signatures might induce a Verifier’s Dilemma for the consensus nodes. Though, on a much smaller scale compared to requiring the consensus nodes to re-execute all transactions.

## 2 Architecture Overview

The Flow architecture is founded on the principle of ‘separation of concerns’. The network has specialized roles: **Consensus Nodes** and **Execution Nodes**. The core differentiation between the two node types is **Objectivity** vs. **Subjectivity**. *Objective tasks* are those for which there is an objectively correct answer. Any traditional mathematical calculation is objective; you don’t need an authority or expert to confirm the correctness of  $2 + 2 = 4$ , or to confirm that an actor that claims  $2 + 2 = 5$  is Byzantine. *Subjective tasks* have no such deterministic solution. Most human governance systems (“laws”) typically deal with subjective issues. At different times in different societies, the rules about who can do certain things can be very different. The definition of the word **consensus** means the agreement on subjective problems, where there is no single correct answer. Instead, one answer must be selected through mutual agreement.

Blockchains combine objective rules with a decentralized mechanism for resolving subjective problems. One example is if two transactions are submitted at the same time that try to spend the same coins (e.g., no double-spends), which one resolves correctly, and which one fails? Traditional blockchain architectures ask the nodes participating in the network to solve both kinds of problems at the same time. In Flow, the Consensus Nodes are tasked with all subjective questions, while the Execution Nodes are responsible solely for fully deterministic, objective problems. While we reserve a detailed discussion of the nodes’ roles, tasks, and interactions for the follow-up paper, we briefly define the Consensus Role and Execution Role for nodes. For an illustration, see Figure 1.

### Consensus Role


Consensus Nodes form blocks from transaction data digests. Essentially, Consensus Nodes maintain and extend the core Flow blockchain. An agreement to accept a proposed block needs to be reached by many nodes which requires a **Byzantine-Fault-Tolerant (BFT)** consensus algorithm [5]. It should be noted that the results in this paper hold for *any BFT consensus algorithm* with deterministic finality.

In Flow, a block references its transaction and defines their execution order. However, a block contains *no* commitment to the resulting computational state *after* block execution. Accordingly, Consensus Nodes do not need to maintain the computational state or execute transactions.

Furthermore, Consensus Nodes adjudicate slashing requests from other nodes, for example claims that an Execution Node has produced incorrect outputs.

### Execution Role

Execution Nodes provide the raw computational power needed to determine the result of the transactions when executed in the order determined by the Consensus Nodes. They produce cryptographic attestations declaring the result of their efforts in the form of **Execution Receipts**. These receipts



**Figure 1:** Overview of the message flow through Consensus and Execution Nodes. For brevity, only the messages during normal operation are shown. Messages that are exchanged during the adjudication of slashing requests are omitted.

can be used to challenge the claims of an Execution Node when they are shown to be incorrect. Also they are used to create proofs of the current state of the blockchain once they are known to be correct. The verification process – by which Byzantine Receipts are rejected (and the Execution Nodes which produced them are slashed), and by which valid receipts are accepted (and shared with observers of the network) – is outside the scope of this paper and will be addressed in future work.

### 3 Theoretical Performance and Security Analysis

In this section, we present a theoretical derivation that one can separate the majority of computation and communication load from consensus nodes without compromising security. Furthermore, we provide an analysis that explains the source of the experimentally observed throughput increase.

Flow is designed to guarantee<sup>3</sup> that any error introduced by the Execution Nodes maintains four critical attributes:

- **Detectable:** A deterministic process has, by definition, an objectively correct output. Therefore, even a single honest node in the network can detect deterministic faults, and prove the error to all other honest nodes by pointing out the part of the process that was executed incorrectly.
- **Attributable:** The output of all deterministic processes in Flow must be signed with the identity of the node that generated those results. As such, any error that has been detected can be clearly attributed to the node(s) that were responsible for that process.
- **Punishable:** All nodes participating in a Flow network, including Execution Nodes, must put up a stake that can be slashed if they are found to have exhibited Byzantine behaviour. Since all errors in deterministic processes are detectable and attributable, those errors can be reliably punished via slashing.
- **Recoverable:** The system must have a means to undo errors as they are detected. The property serves to deter malicious actors from inducing errors that benefit them more than the slashing penalty.

An important property of this design is that for each system-internal operation, the participants are accountable. Specifically, for all operations *except for the Consensus Nodes*, the execution of each operation is delegated to a subset of nodes, the **operation processors**. Verifying the outcome is assigned to a disjoint node set, the **operation verifiers**. Informally, the protocol works as follows:

- Both **operation processor** and **operation verifier** groups are chosen at random. The selection of nodes uses a verifiable random function [25], such that the outcome is deterministic but resistant to hash grinding.
- The inclusion probability for a node in either group is proportional to its stake. This enforces that Byzantine actors must lock up a significant amount of stake in order to have a non-negligible probability of affecting the system. Specifically, this hardens the system against Sybil attacks [26].
- The required amount of stake for both groups is set sufficiently high such that the probability of sampling *only* Byzantine actors in both groups is sufficiently small.
- As long as at least one honest node is involved in either group, the honest node will detect and flag any error.
- If a potential error is flagged, the case is adjudicated by the Consensus Nodes, malicious nodes are slashed, and the operation’s outcome is rejected if faulty.

The process above guarantees that malicious Execution Nodes are slashed with near certainty. Furthermore, it is nearly impossible for the malicious actors to succeed with introducing an error.

A question that might arise is why Flow has separate groups of **operation processors** and **operation verifiers** instead of the **operation processors** verifying each other’s results. We chose this separation of concern to address the **Verifier’s Dilemma** [16]. Without a dedicated verifier role, there is a

---

<sup>3</sup>In Flow, guarantees are probabilistic. Specifically, errors are detected and corrected with probability  $p = 1 - \epsilon$  for  $0 < \epsilon \ll 1$ . Through system parameters,  $\epsilon$  is tuned such that the desired properties hold with near certainty.



conflict of interest for an operation processor to compute the next block vs. verifying (re-computing) the last block’s result. In Flow, this dilemma is alleviated by dedicated operation verifiers who are compensated solely for verification. The technical details of our block verification protocol are presented in the follow-up paper, including a solution to the ‘freeloader problem’ (verifiers just approving any result without doing the actual computation) and ‘maliciously flagging’ (verifiers challenging correct results to congest the network).

The following theorem formalizes the security guarantees of the Flow architecture and proves that introducing an error into the system by publishing or approving faulty results is economically infeasible.

**Theorem 1 (Probabilistic security for delegation of work to small groups)**

*Introducing an error into the system by deliberately publishing or approving faulty results is economically infeasible, if the following conditions hold for any operation, except for those from the Consensus Nodes.*

1. *The operation is delegated to two sets of randomly selected nodes:*
  - (a) *set of operation processors: members of this group execute the operation and provide cryptographically secure commitments to their result*
  - (b) *set of operation verifiers: members of this group verify the operation’s result and provide cryptographically secure commitments to the result if they approve*
2. *Both groups can be relatively small as long as the probability of choosing only Byzantine actors in both groups at the same time is sufficiently low.*
3. *At the time the operation processors generate the result, the membership of the operation verifier group is unknown to them.*
4. *Consensus Nodes*
  - (a) *either verify that a significant majority have committed to the published outcome and there are no objections raised by participating nodes*
  - (b) *or adjudicate objections, determine the faulty nodes (attributable), and slash them (punishable).*

It is essential to highlight that Consensus Nodes are not required to check the correctness of the results of an operation. Instead, they ensure that other nodes with sufficient stake are accountable for the verification. Furthermore, Theorem 1 holds for any BFT consensus algorithm with deterministic finality.

**Proof of Theorem 1:**

We will show that Theorem 1 can always be satisfied under realistic conditions where

- at least one of the two groups is sampled from a large population with a super-majority of honest nodes;
- Byzantine actors cannot suppress communication between correct nodes, i.e., if there is one honest node objecting to the result and proving its faultiness, the erroneous nodes will be slashed.

Specifically, let us consider a population of  $N$  nodes from which we want to randomly draw a subset with  $n$  nodes. Furthermore, we assume that there are at most  $M < N/3$  Byzantine nodes. In the following, we focus on the case where all nodes are equally-staked<sup>4</sup>, i.e., their inclusion probabilities are identical. Drawing an  $n$ -element subset falls in the domain of simple random sampling [27] without replacement. The probability of drawing  $m \leq n$  Byzantine nodes is given by the hypergeometric distribution<sup>5</sup>

$$\mathcal{P}_{n,N,M}(m) = \frac{\binom{M}{m} \binom{N-M}{n-m}}{\binom{N}{n}}. \quad (1)$$

The probability of a successful attack,  $P(\text{successful attack})$ , requires that there is no honest node that would contradict a faulty result. Hence,

$$P(\text{successful attack}) \leq \mathcal{P}_{n,N,M}(n), \quad (2)$$

where  $\mathcal{P}_{n,N,M}(n)$  is the probability of sampling only Byzantine nodes.

$$\mathcal{P}_{k,N,M}(k) = \frac{M! (N-k)!}{N! (M-k)!} \quad \text{for } k \leq M \quad (3)$$

$$\Rightarrow \frac{\mathcal{P}_{n,N,M}(n)}{\mathcal{P}_{n+1,N,M}(n+1)} = \frac{(N-n)! (M-(n+1))!}{(N-(n+1))! (M-n)!} = \frac{N-n}{M-n} > 1 \quad (4)$$

As eq. (4) shows, the probability of sampling only Byzantine nodes is strictly monotonously decreasing with increasing  $n$ . Eq. (4) states that the larger the sample size  $n$ , the smaller the probability to sample only Byzantine node.

For a node to deliberately attack the network by publishing a faulty result or approving such, we assume the existence of some reward  $r$  which the node receives in case its attack succeeds. However, if the attack is discovered, the node is slashed by an amount  $\xi$  (by convention positive). The resulting statistically expected revenue from attacking the network is

$$\text{revenue} = P(\text{successful attack}) \cdot r - (1 - P(\text{successful attack})) \cdot \xi \quad (5)$$

$$\stackrel{(2)}{\leq} \mathcal{P}_{n,N,M}(n) \cdot r - (1 - \mathcal{P}_{n,N,M}(n)) \cdot \xi \quad (6)$$

For the attack to be economically viable, one requires  $0 \stackrel{!}{\leq} \text{revenue}$ , which yields the central result of this proof:

$$\frac{r}{\xi} \geq \frac{1}{\mathcal{P}_{n,N,M}(n)} - 1. \quad (7)$$

<sup>4</sup>The argument can be extended to nodes with different stakes. In this case, each node would have an inclusion probability equal to its fraction of total stake. However, the probability of sampling fully Byzantine groups is depending on the specific fractions of total stake for the individual nodes. A basic solution for allowing nodes with different stakes is to introduce a unit quantity  $\varrho$  of stake. For a node with the stake  $s$ , the multiplicity  $k = \lfloor s/\varrho \rfloor$  represents how many *full* staking units the node possesses. For operational purposes (including voting and node selection), the blockchain treats the node identically to  $k$  independent nodes each having stake  $\varrho$ .

<sup>5</sup>For conciseness, we only handle the case  $m \leq M$ . For  $m > M$ ,  $\mathcal{P}_{n,N,M}(m) = 0$  per definition.

Furthermore, eq. (4) implies that  $r/\xi$  increases strictly monotonously with increasing  $n$ . Using results from [28], one can show that  $r/\xi$  grows exponentially with  $n$  for  $n < N/2$ .

The left-hand side of equation (7),  $r/\xi$ , is a measure of security as it represents the statistical cost to attack the system in the scenario where an attacker bribes nodes into byzantine behavior. As an example, let us consider the case with  $N = 1000$ ,  $M = 333$ ,  $n = 10$ . For simplicity, assume that for publishing or approving a faulty result, the node’s entire stake is slashed. Then, for an attack to be economically viable the success reward  $r$  would need to be 65 343 times the node’s stake. If the operation verifiers staked \$1000 each, an attacker would have to expend \$65.3 million on average to cover all the slashing costs. It would be cheaper for the attacker to run the entire pool of operation verifiers instead of attempting to slip an error past the honest verifiers. When increasing  $n$  even further to  $n = 20$ , an attacker would need to expend  $r/\xi = 5.2 \cdot 10^9$  times the stake to slip a single error past a super-majority of honest verifiers.

In summary, we have analyzed the case where either processing an operation or verifying its result is delegated to a small, random subset of a much larger population of nodes. We have shown that under realistic assumptions, introducing an error into the system by publishing or approving faulty results is economically infeasible. Note that this result only covers node types *other* than Consensus Nodes. Hence, it is sufficient for the Consensus Nodes to check that enough nodes have participated in executing the operation as well as verifying it. However, they do not need to check the result itself to problematically guarantee its integrity.

□

Theorem 1 is a key insight, as it allows us to:

- separate the majority of computation and communication load from consensus;
- develop highly specialized nodes with distinct requirement profiles (as opposed to having one node type that has to have outstanding performance in all domains or otherwise diminish the network throughput);

While other nodes verify each others’ operations in small groups, the entire committee of Consensus Nodes audits themselves.

### 3.1 Special Role of Consensus Nodes

Consensus Nodes determine the relative time order of events through a BFT consensus algorithm. While our results hold for *any* BFT consensus algorithm with deterministic finality, HotStuff [29, 30] is the leading contender. However, we continue to assess other algorithms such as Casper CBC [31, 32, 33] or Fantôme [34].

In contrast to the operations of other nodes, which requires an auditing group to approve the result, Consensus Nodes finalize blocks without external verification. While the contents of a block can be verified and Consensus Nodes punished if they include invalid entries, blocks are not rebuilt in this scenario, unlike other verification processes. External parties can inspect the finalized blocks after the fact. However, in the event of an adversarial attack forking the chain, a double-spend attack might have already succeeded at this point. To increase the resilience of the entire system, the committee of Consensus Nodes should consist of as many staked nodes as possible.

For a simple BFT algorithm, the message complexity  $\eta$  per block (i.e., the total number of messages sent by all  $N$  nodes) is  $O(N^2)$  [7]. More advanced protocols achieve  $\eta \in O(N \log N)$  [35] or  $\eta \in O(c \cdot N)$  [36, 37], for  $c \ll N$  an approximately constant value for large  $N$ . The overall

bandwidth load  $\mathcal{B}$  [MB/s] for the entire consensus committee is

$$\mathcal{B} = \beta \cdot b \cdot \eta, \tag{8}$$

for  $\beta$  the block rate [ $\text{s}^{-1}$ ],  $b$  the message size [MB/s]. For a node receiving a message  $m$  and processing it, imposes a latency

$$\mathcal{L}(m) = f(b) + \text{process}(m), \tag{9}$$

where  $f$  denotes the network transmission time for receiving  $m$  and  $\text{process}(m)$  represents the computation time for processing  $m$ . It is apparent that both  $\mathcal{B}$  and  $\mathcal{L}$  strongly impact the throughput of the consensus algorithm. The specific details are highly dependent on the chosen BFT protocol, as well as the employed gossip topology [38]. Other factors include delays or bandwidth limitations in the underlying network hardware.

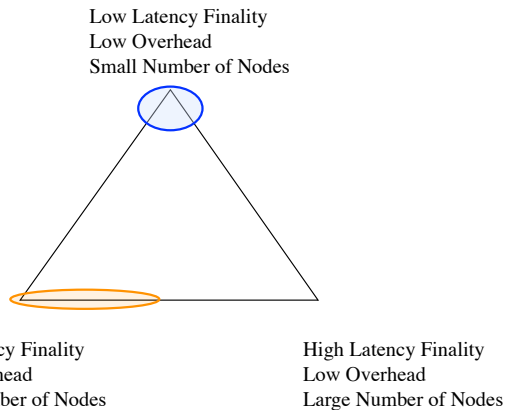
Currently, we are targeting a consensus committee on the order of several thousand nodes. To support decentralization and transparency, hardware requirements for Consensus Nodes should be limited such that private groups can still afford to run a node and participate in consensus. Hence, given a desired minimal block rate (e.g.,  $\beta = 1 \frac{\text{block}}{\text{s}}$ ) and an environment-determined function  $f$ , the consensus committee can be increased only by decreasing message size  $b$  or  $\text{process}(m)$ .

For completeness, we provide a brief outlook on how we simultaneously reduce  $b$  and  $\text{process}(m)$  in Flow. For the detailed mechanics of the respective processes, the reader is referred to subsequent papers.

- We delegate the computation of the transactions to the specialized Execution Nodes. The delegation removes the need for Consensus Nodes to maintain and compute state, which significantly reduces  $\text{process}(m)$ .
- Consensus Nodes do not require the full transaction texts during normal operation. Instead, specialized Collector Nodes receive transactions from external clients and prepackage them into batches, called collections. Consensus Nodes only order collection references (hashes), which substantially reduces the messages size  $b$  compared to handing individual transaction texts.

We conclude this section by comparing our approach with the ‘triangle of compromises’ proposed by Vlad Zamfir [39], which we re-create in Figure 2. The triangle illustrates Zamfir’s impossibility conjecture for proof of stake systems. While the conjecture has not been proven formally, we concur that the compromises are correctly identified. However, Flow optimizes where these compromises are made:

- *Consensus Nodes* work as part of a large consensus committee for maximal security. To ensure security and fast generation of finalized blocks, we accept a higher communication overhead (the bottom-left corner of the triangle). However, unlike other blockchains, this consensus only determines the order of transactions within a block, but not the resulting state. The architecture compensates for the resulting bandwidth higher



**Figure 2:** Zamfir’s triangle of compromises; orange: Consensus Nodes; blue: Execution Nodes.

overhead by minimizing message size through the use of collection references (references to batches of transactions) instead of full individual transactions. To further increase throughput and reduce communication latency, all possible computation is delegated to other node types.

- The *Execution Nodes* unpack the collections and process individual transactions. Without a large number of Byzantine actors involved, this can be handled directly between the Collector and Execution Nodes without involving Consensus Nodes. Furthermore, tasks can be parallelized and processed by small groups which hold each other accountable for correct results. Only during Byzantine attacks, malicious actions would be reported to Consensus Nodes to adjudicate slashing.

## 4 Performance Simulations

The theoretical analysis presented in section 3.1 suggests that transaction throughput can be increased by separating consensus about the transaction order from their execution. However, the theoretical analysis makes no assertion as to what the realistically achievable speedup is, as throughput heavily depends on a variety of environmental parameters such as message round-trip time, CPU performance, etc. Therefore, we have implemented a simplified benchmark network that solely focuses on transaction ordering (consensus) and transaction execution.

### 4.1 Experimental Setup


In a 2015 study analyzing the distribution of computational power in the Bitcoin network [40], the authors estimated that 75% of the mining power was provided by roughly 2% of the nodes. We simulated a system whose centralization metrics are roughly half of the Bitcoin scores. In our simulations, roughly 38% computational power is provided by the **fast nodes**, which represent approximately 6% of the nodes. For the remaining 62% of the network’s total computational power, we have applied a less-extreme ratio: two-third of the nodes (**slow nodes**) hold one-third of the remaining computational power (i.e.,  $62\%/3 \simeq 20\%$  of the total computational power). The remainder is assigned to **medium nodes**.

To assign the Execution role to nodes with the most computation power requires incentive mechanisms that compensate nodes for the resources used by the network. Assuming the existence of such incentive mechanisms, it is economically rational for a **fast node** to stake as an Execution Node. In any other role, its resources would not be utilized to the maximum potential leading to diminished revenue. Hence, we assumed that the most powerful nodes would stake specifically to become Execution Nodes.

We conducted three different experiments. The common characteristics of all simulations are described in the following. Section 4.1.1 to 4.1.3 present the specific details for each individual experiment. Figure 3 illustrates the different setups. For each experiment, the network resources (node types) and the assignment of responsibilities are given in Table 1.

		Slow Nodes	Medium Nodes	Fast Nodes
<b>Experiment (I):</b>	number of nodes	20	10	2
	role of nodes	consensus	consensus	compute
<b>Experiment (II)</b>	number of nodes	20	10	2
	role of nodes	consensus and compute	consensus and compute	consensus and compute
<b>Experiment (III)</b>	number of nodes	32		
	role of nodes	consensus and compute		

**Table 1:** Network configuration for each experiment.



**Figure 3: Illustration of experiments.** (I) Flow network with the division of consensus and compute; (II) Conventional PoS network containing nodes with different performance levels; (III) Conventional PoS network containing only slow nodes.

### Common Characteristics

*Transactions:* For simplicity, our network was processing benchmark transactions, which all had identical complexity (number of instructions). Batches of benchmark transactions were directly generated by the Consensus Nodes instead of integrating a dedicated submission process into the simulation.

*Network:* We implemented a relatively small network of 32 nodes, where each node ran on a dedicated Google Cloud instance. The nodes were spread over 8 data centres across the world to provide somewhat realistic latencies.

*Nodes:* Depending on the experiment (see Table 1), transactions were executed on nodes with different hardware:

- **slow nodes:** process a benchmark transaction in approximately 10ms
- **medium nodes:** five times as fast as slow nodes, i.e., process a benchmark transaction in 5ms
- **fast nodes:** 25 times as fast as slow nodes, i.e., process a benchmark transaction in 2.5ms

To facilitate decentralization, an ideal network should allow any participant to join, requiring only a minimal node performance. Consequently, a realistic network will contain a majority of slow nodes, some medium nodes and very few fast nodes.

*Consensus:* We implemented a Tendermint-inspired consensus algorithm with a rotating block proposer. As our goal was to benchmark achievable throughput in the absence of a large-scale Byzantine attack, our benchmark network only consists of honest nodes. The proposed blocks contain a variable number of  $t$  benchmark transactions, where  $t$  is drawn uniform randomly from the integer interval  $[240, 480]$ . However, for repeatability, we seeded the random number generator such that in all experiments, the same sequence of 20 blocks was proposed and finalized.

#### 4.1.1 Experiment (I): Flow PoS network with the split of consensus and compute

This experiment simulates a network with the division of consensus and compute:

- 20 slow nodes and 10 medium nodes form the consensus committee. They agree on the order of transactions within a block but don't store or update the chain's computational state.
- Two fast nodes execute the transactions in the order they are finalized in the blocks. They do not participate in consensus.

For further illustration, see Figure 3(I) and Table 1.

In the Flow network, blocks only specify the transaction order, but there is no information about the resulting state included. As consensus in this model only covers the transaction order, consensus nodes are oblivious about the computational state.

### 4.1.2 Experiment (II): PoS network of nodes with diverse performances

Experiment (II) closely models conventional BFT consensus systems such as Tendermint [41, 42, 43] or Hot-Stuff [29]. The network is identical to Experiment (I), i.e., it consists of exactly the same types and numbers of nodes (see Table 1 and Fig. 3(II) for details and illustration). Though, blocks contain the transaction order and a hash commitment to the resulting computational state. Due to this result commitment, each consensus node must repeat the computation of all transactions for a proposed block to validate its correctness. In essence, there is only one role for a node: to participate in the consensus algorithm, which includes the task of updating the computational state.

### 4.1.3 Experiment (III): PoS network of nodes with uniform performance

Experiment (III), illustrated in Figure 3(III), simulates a network of 32 nodes with uniform computational performance. As in Experiment (II), all nodes execute the same algorithm which combines consensus about transaction ordering with their computation.

## 4.2 Experimental Results

Our simulations aim at benchmarking the transaction throughput. For each experiment, we sent 7995 benchmark transactions through the network and measured the corresponding processing time. The results are summarized in Table 2.

Experiment (I) and (II) were executed with the same network configuration. The only difference was the separation of consensus and compute in the experiment (I), while both were combined in the experiment (II). In our moderately simplified model, separating compute from consensus increased the throughput approximately by a *factor* of 56.

Comparing experiment (II) and (III) illustrates the limited impact of increasing network resources compared. In terms of instructions per seconds, the network in the experiment (II) is 3.75 times more powerful<sup>6</sup> than in the experiment (III). However, the throughput of (II) increased only by 0.7% compared to the experiment (III).

As the results show, separation of consensus and compute allows utilizing network resources more efficiently. In a PoS network with combined consensus and compute, deterministic block finalization requires a super-majority of nodes to vote in favor of a candidate block to be finalized. For may

<sup>6</sup>Let a slow node process  $x$  instructions per second. Hence, under ideal resource utilization, the network in the experiment (III) can process  $32x$  instructions. In contrast, the network in the experiment (II) processes  $20x + 10 \cdot 5x + 2 \cdot 25x = 120x$ .

	Processing Time [s]	Throughput [TX/s]
Experiment (I)	5.14	1555.4
Experiment (II)	291	27.5
Experiment (III)	293	27.3

**Table 2: Network performances.** Processing time for 7995 transactions in seconds [s] and the resulting transaction throughput [TX/s]. While we only conducted one-shot experiments, we have repeatedly observed throughout implementation that processing times fluctuate only on the sub-second scale.



BFT protocols, such as [29, 42, 44, 45], finalization requires supporting votes with an accumulated fraction of stake  $\mathcal{S} > 2/3$ . Though, some protocols have other limits, e.g.,  $\mathcal{S} > 80\%$  for [46, 47]. All of these protocols have in common that consensus nodes are obliged to execute the computation as a sub-task to verifying the proposed block. Therefore, the time for finalizing a block is bound by the fastest  $\mathcal{S}^{\text{th}}$  percentile of nodes. Less formally, the slowest nodes determine the throughput of the entire system. Consequently, running the network on stronger nodes leaves throughput unchanged as long as the slowest  $(1 - \mathcal{S})$ -fraction of nodes do not receive performance upgrades. In contrast, separation of consensus and compute significantly shifts computational load from the consensus nodes to the fastest nodes in the network.

## 5 Further Work

A blockchain designed around the principles outlined in this paper would need to address additional problems, the most notable being the mechanism that verifies computation outputs. On the one hand, splitting consensus and compute work boosts the throughput of Flow. On the other hand, special care has to be taken that the resulting states are correct as consensus nodes do not repeat the computation.

Furthermore, in Flow, blocks no longer contain a hash commitment to the resulting state after computing the block. Therefore, a node that receives data from a block state cannot verify the validity of the received data based on the information published in the block. Nevertheless, a hash commitment for the result of a previous block can be published in a later block after passing verification. We will present the technical details of the **block verification** and commitment to the computation results (referred to as **block sealing**) in the follow-up papers.

The presented simulations provide experimental evidence to support the theoretical work of this paper. While the theoretical results (section 3) stand on their own without experimental validation, the experiments could be extended significantly. For example, we have not accounted for the extra steps required to verify computational states and commit them into the chain. Another aspect is the size of the consensus committee. It would be interesting to study the scaling of transaction throughput with different committees sizes of consensus and execution nodes. However, we have decided to prioritize implementing the Flow architecture over benchmarking a simplified model system. Throughput and other performance characteristics will be measured and published as soon as a full-fledged implementation is completed.

## 6 Conclusions

In this proof-of-concept work, we have demonstrated that a separation of consensus and compute can lead to significantly increased resource utilization within PoS blockchain networks. For conventional round-based PoS networks, where one block is finalized before the subsequent block is proposed, the throughput is limited by a small fraction of the slowest nodes. In contrast, separation of consensus and compute significantly shifts computational load from the consensus nodes to the fastest nodes in the network. We have shown in Theorem 1 that such separation of concern does not compromise the network’s security. First experiments suggest that the throughput improvements enabled by such a separation of concerns are drastic. In a moderately simplified model, our simulations show

a throughput increase by a factor of 56 compared to architectures with combined consensus and block computation.

One way to substantially increase the throughput of existing blockchains, such as Ethereum, could be to increase the gas limit. However, this would accelerate the rate at which the state grows making it harder for new nodes to join the system. While in conventional proof-of-work blockchains the computational load to maintain and update the state is uniform across all (full) nodes, the large majority of the computation resources are concentrated in a small fraction of mining nodes [40].

The Flow architecture utilizes the resource imbalance naturally occurring within a network ecosystem. The few data-center-scale nodes with massive computational and bandwidth capacities can stake to become Execution Nodes to contribute their resources most efficiently. In contrast, Consensus Nodes do not store or maintain the state and, therefore, can be run on off-the-shelf consumer hardware. With such separation of concerns, sharing a large state with new Execution Nodes joining the system should not pose a substantial challenge given the operational resources available to nodes with this role.

## Acknowledgments

We thank Dan Boneh for many insightful discussions, J. Ross Nicoll for contributions to an earlier draft, and Nick Johnson, Alex Bulkin, Karim Helmy, Teemu Paivinen, Travis Scher, Chris Dixon, Jesse Walden, Ali Yahya, Ash Egan, Casey Taylor, Joey Krug, Arianna Simpson, as well as Lydia Hentschel for reviews.

## References

- [1] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system, 2009.
- [2] Kyle Croman, Christian Decker, Ittay Eyal, Adem Efe Gencer, Ari Juels, Ahmed Kosba, Andrew Miller, Prateek Saxena, Elaine Shi, Emin Gün Sirer, Dawn Song, and Roger Wattenhofer. On scaling decentralized blockchains. In Jeremy Clark, Sarah Meiklejohn, Peter Y.A. Ryan, Dan Wallach, Michael Brenner, and Kurt Rohloff, editors, *Financial Cryptography and Data Security*, pages 106–125, Berlin, Heidelberg, 2016. Springer Berlin Heidelberg.
- [3] BBC News. Cryptokitties craze slows down transactions on ethereum. 2017. <https://www.bbc.com/news/technology-42237162>.
- [4] Richard Dennis and Jules Pagna Diss. *An Analysis into the Scalability of Bitcoin and Ethereum*, pages 619–627. 01 2019.
- [5] M. Pease, R. Shostak, and L. Lamport. Reaching agreement in the presence of faults. *J. ACM*, 27(2):228–234, April 1980.
- [6] Leslie Lamport, Robert Shostak, and Marshall Pease. The byzantine generals problem. *ACM Trans. Program. Lang. Syst.*, 4(3):382–401, July 1982.
- [7] Miguel Castro and Barbara Liskov. Practical byzantine fault tolerance. In *Proceedings of the Third Symposium on Operating Systems Design and Implementation*, OSDI '99, pages 173–186, Berkeley, CA, USA, 1999. USENIX Association.
- [8] Scott Nadal Sunny King. Ppcoin: Peer-to-peer crypto-currency with proof-of-stake. 2012. <http://www.peercoin.net/>.
- [9] Ittai Abraham and Dahlia Malkhi. The blockchain consensus layer and BFT. *Bulletin of the EATCS*, 123, 2017.
- [10] Diksha Gupta, Jared Saia, and Maxwell Young. Proof of work without all the work. In *Proceedings of the 19th International Conference on Distributed Computing and Networking*, ICDCN '18, pages 6:1–6:10, New York, NY, USA, 2018. ACM.
- [11] Shehar Bano, Alberto Sonnino, Mustafa Al-Bassam, Sarah Azouvi, Patrick McCorry, Sarah Meiklejohn, and George Danezis. Consensus in the age of blockchains. arXiv:1711.03936, 2017. <https://arxiv.org/abs/1711.03936>.
- [12] Jason Spasovski and Peter Eklund. Proof of stake blockchain: Performance and scalability for groupware communications. In *Proceedings of the 9th International Conference on Management of Digital EcoSystems*, MEDES '17, pages 251–258, New York, NY, USA, 2017. ACM.
- [13] Joseph Poon and Thaddeus Dryja. The bitcoin lightning network: Scalable off-chain instant payments. 2016. <https://lightning.network/lightning-network-paper.pdf>.
- [14] Joseph Poon and Vitalik Buterin. Plasma: Scalable autonomous smart contracts. *White paper*, pages 1–47, 2017.
- [15] Mahdi Zamani, Mahnush Movahedi, and Mariana Raykova. Rapidchain: A fast blockchain protocol via full sharding. *IACR Cryptology ePrint Archive*, 2018:460, 2018.
- [16] Loi Luu, Jason Teutsch, Raghav Kulkarni, and Prateek Saxena. Demystifying incentives in the consensus computer. In *Proceedings of the 22Nd ACM SIGSAC Conference on Computer and Communications Security*, CCS '15, pages 706–719, New York, NY, USA, 2015. ACM.
- [17] Nir Bitansky, Ran Canetti, Alessandro Chiesa, and Eran Tromer. From extractable collision resistance to succinct non-interactive arguments of knowledge, and back again. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, ITCS '12, pages 326–349, New York, NY, USA, 2012. ACM.
- [18] B. Bünz, J. Bootle, D. Boneh, A. Poelstra, P. Wuille, and G. Maxwell. Bulletproofs: Short Proofs for Confidential Transactions and More. In *2018 IEEE Symposium on Security and Privacy (SP)*, pages 315–334, May 2018.

- [19] Jason Teutsch and Christian Reitwießner. A scalable verification solution for blockchains, March 2017. Accessed:2017-10-06.
- [20] Vitalik Buterin’s comment on medium article “The Ethereum-blockchain size has exceeded 1TB, and yes, it’s an issue”. 2018. <https://medium.com/hackernoon/the-ethereum-blockchain-size-has-exceeded-1tb-and-yes-its-an-issue-2b650b5f4f62>.
- [21] Matt Freels. FaunaDB: An Architectural Overview. 2018. [https://www2.fauna.com/FaunaDB\\_Tech\\_WP](https://www2.fauna.com/FaunaDB_Tech_WP).
- [22] Alexandre Verbitski, Tengiz Kharatishvilli, Xiaofeng Bao, Anurag Gupta, Debanjan Saha, James Corey, Kamal Gupta, Murali Brahmadesam, Raman Mittal, Sailesh Krishnamurthy, and Sandor Maurice. Amazon aurora: On avoiding distributed consensus for i/os, commits, and membership changes. pages 789–796, 05 2018.
- [23] Raymond Cheng, Fan Zhang, Jernej Kos, Warren He, Nicholas Hynes, Noah M. Johnson, Ari Juels, Andrew Miller, and Dawn Song. Ekiden: A platform for confidentiality-preserving, trustworthy, and performant smart contract execution. *CoRR*, abs/1804.05141, 2018.
- [24] Dawn Song. Oasis: Privacy-preserving smart contracts at scale, 2018. Microsoft Research Faculty Summit.
- [25] Silvio Micali, Salil Vadhan, and Michael Rabin. Verifiable Random Functions. In *Proceedings of the 40th Annual Symposium on Foundations of Computer Science*, FOCS ’99, pages 120–, Washington, DC, USA, 1999. IEEE Computer Society.
- [26] John R. Douceur. The sybil attack. In Peter Druschel, Frans Kaashoek, and Antony Rowstron, editors, *Peer-to-Peer Systems*, pages 251–260, Berlin, Heidelberg, 2002. Springer Berlin Heidelberg.
- [27] Frank Olken. *Random sampling from databases*. PhD thesis, University of California, Berkeley, 1993. <http://db.cs.berkeley.edu/papers/UCB-PhD-olken.pdf>.
- [28] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, March 1963.
- [29] Maofan Yin, Dahlia Malkhi, Michael K. Reiter, Guy Golan Gueta, and Ittai Abraham. HotStuff: BFT Consensus in the Lens of Blockchain. 2018. <http://arxiv.org/abs/1803.05069>.
- [30] Maofan Yin, Dahlia Malkhi, Michael K. Reiter, Guy Golan Gueta, and Ittai Abraham. HotStuff: BFT Consensus with Linearity and Responsiveness. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*, PODC ’19, pages 347–356, New York, NY, USA, 2019. ACM.
- [31] Vlad Zamfir. A Template for Correct-By-Construction Consensus Protocols. 2017. <https://github.com/ethereum/research/tree/master/papers/cbc-consensus>.
- [32] Vlad Zamfir. Casper the Friendly Ghost: A ‘Correct By Construction’ Blockchain Consensus Protocol. 2017. <https://github.com/ethereum/research/blob/master/papers/CasperTFG>.
- [33] Vlad Zamfir, Nate Rush, Aditya Asgaonkar, and Georgios Piliouras. Introducing the “Minimal CBC Casper” Family of Consensus Protocols. 2018. <https://github.com/cbc-casper/cbc-casper-paper>.
- [34] Sarah Azouvi, Patrick McCorry, and Sarah Meiklejohn. Betting on Blockchain Consensus with Fantôme. 2018. <http://arxiv.org/abs/1805.06786>.
- [35] Jing Liu, Wenting Li, Ghassan O. Karame, and N. Asokan. Scalable byzantine consensus via hardware-assisted secret sharing. *IEEE Transactions on Computers*, 68:139–151, 2019.
- [36] Guy Golan-Gueta, Ittai Abraham, Shelly Grossman, Dahlia Malkhi, Benny Pinkas, Michael K. Reiter, Dragos-Adrian Seredinschi, Orr Tamir, and Alin Tomescu. SBFT: a scalable decentralized trust infrastructure for blockchains. *CoRR*, abs/1804.01626, 2018. <https://arxiv.org/abs/1804.01626>.
- [37] Mohammad M. Jalalzai, Costas Busch, and Golden Richard III. Proteus: A scalable BFT consensus protocol for blockchains. 2019. <https://arxiv.org/abs/1903.04134>.
- [38] Stephen Boyd, Arpita Ghosh, Balaji Prabhakar, and Devavrat Shah. Randomized gossip algorithms. *IEEE/ACM Trans. Netw.*, 14(SI):2508–2530, June 2006.
- [39] Vlad Zamfir. Fundamental tradeoff in fault tolerant consensus protocols. <https://twitter.com/vladzamfir/status/942271978798534657>, December 2017.
- [40] Ann Elizabeth Miller, James Litton, Andrew Pachulski, Neal Gupta, Dave Levin, Neil Spring, and Bobby Bhattacharjee. Discovering Bitcoin’s Public Topology and Influential Nodes. 2015.
- [41] Jae Kyun Kwon. Tendermint : Consensus without mining. 2014. <https://github.com/cosmos/cosmos/tree/master/tendermint>.
- [42] Ethan Buchman. Tendermint: Byzantine Fault Tolerance in the Age of Blockchains, Jun 2016.

- [43] Jae Kwon and Ethan Buchman. Cosmos - A Network of Distributed Ledgers. 2016. <https://cosmos.network/cosmos-whitepaper.pdf>.
- [44] Cynthia Dwork, Nancy Lynch, and Larry Stockmeyer. Consensus in the presence of partial synchrony. *J. ACM*, 35(2):288–323, April 1988.
- [45] Miguel Castro and Barbara Liskov. Practical byzantine fault tolerance and proactive recovery. *ACM Trans. Comput. Syst.*, 20(4):398–461, November 2002.
- [46] Jean-Philippe Martin and Lorenzo Alvisi. Fast byzantine consensus. *IEEE Trans. Dependable Secur. Comput.*, 3(3):202–215, July 2006.
- [47] Dillkötter David. The ripple protocol consensus algorithm. *Ripple Labs Inc.*, 2014.