

# A Limitlessly Scalable Transaction System

Max Mathys, Roland Schmid, Jakub Sliwinski, and Roger Wattenhofer

ETH Zürich

**Abstract.** We present ACCEPT, a simple, asynchronous transaction system that achieves perfect horizontal scaling.

Usual blockchain-based transaction systems come with a fundamental throughput limitation as they require that all (potentially unrelated) transactions must be totally ordered. Such solutions thus require serious compromises or are outright unsuitable for large-scale applications, such as global retail payments.

ACCEPT provides efficient horizontal scaling without any limitation. To that end, ACCEPT satisfies a relaxed form of consensus and does not establish an ordering of unrelated transactions. Furthermore, ACCEPT achieves instant finality and does not depend on a source of randomness.

## 1 Introduction

The financial world is changing around the globe. With the rise of digital cryptocurrencies like Bitcoin [13] and Ethereum [16], the pressure on the traditional banking system to implement a digital currency on its own is rising. Due to the well-known limitations of permissionless blockchain systems thus far, such a digital currency is mostly envisioned based on permissioned, byzantine fault-tolerant ledger technology. Furthermore, as both central and commercial banks have no interest in a fully distributed solution that can hardly be regulated, the employment of a permissioned system constitutes a good fit to establish such a digital currency.

Previously, revolutionizing the global financial infrastructure was envisioned based on byzantine agreement protocols. Such systems are popular because they provide a reliable and robust way of transferring funds between participants by establishing a total order of all transactions. However, ordering all transactions has proven to be a throughput-limiting factor for these systems, which only achieve throughputs up to tens of thousands of transactions per second even in lab environments [15]. Despite byzantine agreement systems being optimized for high throughput, they only recently matched the demands of leading credit card providers; hence, this technology does not seem future-proof in our increasingly digital world.

*Our Contribution:* We present a system based on A Cheaper Consensus for Efficient, Parallelizable Transactions (ACCEPT), that features:

- **Limitless Scalability:** ACCEPT does not order transactions that do not depend on each other. Thus, any number of parallel transactions can be processed given a sufficient amount of hardware.

- **Instant Finality:** ACCEPT confirms transactions in 1 round-trip time to the validator nodes. Confirmed transactions are final and cannot be reversed.
- **Asynchrony:** ACCEPT does not require any network timing assumptions. An adversary having complete control over the network can halt the progress of the system (by simply disabling communication) but otherwise cannot trick the participants in any way, such as reversing a confirmed transaction or pretending that an impermissible transaction is confirmed.
- **Independent Validators:** Validator nodes only need to provide a basic API to accept new transactions, verifying them internally and returning a valid signature in case of success. This makes the nodes’ implementations independent, as they only need to interpret other nodes’ (and clients’) signatures. Validators can parallelize and balance workload as they individually see fit, without changing the way they interact with the system.
- **Simplicity:** ACCEPT does not rely on randomness and is easy to comprehend and implement. Hence, given some deployment scenario, the validators can easily implement the protocol themselves, thus achieving fault tolerance with respect to software bugs in the code. In contrast, existing (complicated) permissioned blockchain systems rely on reference implementations that are used as a black box, where a single bug can compromise the entire system.

We demonstrate ACCEPT’s horizontal scaling on common server hardware and showcase its transaction throughput to be orders of magnitudes above byzantine agreement systems.

## 2 Preliminaries

We assume that conflicting transactions can only be issued by a misbehaving party. Under this crucial assumption, the system can maintain liveness and consistency without solving consensus [8].

In contrast to orthodox blockchain systems, ACCEPT does not support consensus and does not attempt to order conflicting transactions issued simultaneously. If a misbehaving client issues two transactions spending the same funds simultaneously, it is possible that both transactions will be rejected by the system and the misbehaving client will lose the funds.

### 2.1 Model

ACCEPT is maintained by  $n$  different agents called *validators*. Similarly to other byzantine fault tolerant systems, we assume up to  $f = \frac{n-1}{3}$  of the validators are adversarial and behave arbitrarily. Any set of  $2f + 1$  validators is called a *quorum*. In addition, an arbitrary number of *clients* interact with the system by issuing and receiving transactions.

The network is asynchronous: The adversary controls the network, dictating when messages are delivered and in what order. There is no bound on the time it might take to deliver a message. Under such weak network requirements,

an adversary delaying the delivery of messages can delay the progress of the protocol, but otherwise will not be able to interfere.

We assume the functionality of digital signatures where a public key allows the verification of a signature of the associated secret key. We also assume cryptographic hashing, where for every message a succinct, unique hash can be computed. Apart from these primitives, ACCEPT is completely deterministic.

*Security & Threat Model.* Validators and clients hold public/private key pairs. All participants know the public keys of all validators.

The adversary knows the protocol and controls all adversarial validators and any number of clients. The adversary controls the network and can delay, replay, reorder messages, etc. The adversary does not know the private key of any correct participant.

### 3 Protocol

The ACCEPT protocol differentiates two main roles:

- **Validators:** Validators are agents that *verify* and *sign* transactions. There is a fixed number of validators. Validators do not have to exchange messages with each other. Validators can be sharded across multiple servers to increase the throughput of the system.
- **Clients:** Clients are end users of the system who *issue/receive* transactions. The system supports an arbitrary number of clients in the system. Clients possess funds that can be sent to other clients via a transaction. Clients may follow the protocol correctly or not; however, it is only guaranteed that they can spend their funds if following the protocol.

Transactions are processed by the system in the UTXO model [7]. The initial state of the system, called *genesis*, consists of (number, public key) pairs, where the number represents the available funds, and the public key identifies the party able to spend them. These pairs are called unspent transaction outputs (UTXO).

If a client (sender) wishes to transfer funds to another client (recipient), the sender issues a transaction. The transaction contains input UTXOs that the sender can spend. The sender specifies output UTXOs, where inputs and outputs sum up to the same amount of funds. For example, a single transaction can specify two outputs, where one output represents the transacted amount and includes the public key of the recipient, and the second output represents the change and includes the public key of the sender. Ultimately, the sender signs the transaction with the private key(s) corresponding to the inputs.

*Transaction Pipeline.* ACCEPT performs four steps to confirm a transaction from client  $c_1$  to client  $c_2$ :

1. **Issuing a Transaction:** The client  $c_1$  composes and signs a transaction. Client  $c_1$  sends the transaction to the validators. If the inputs of the transaction are not part of genesis, the client also sends the confirmations of the inputs to the validators.

2. **Verifying a Transaction:** Each validator  $v_i$  verifies the transaction signature and inputs' confirmations (if not part of genesis). Also,  $v_i$  checks that it has not validated any transaction spending the same inputs thus far.
3. **Signing a Transaction:** If the transaction is valid, validator  $v_i$  signs the transaction and returns the signature to  $c_1$ .
4. **Finalizing a Transaction:** A set of  $2f + 1$  signatures of distinct validators constitutes a *confirmation* of the outputs. The client  $c_1$  can show the corresponding transaction output and the confirmation to  $c_2$  to prove the transfer took place. The client  $c_2$  accepts the transfer after verifying the output and the confirmation.

### 3.1 Complexity

In Section 3.5 we discuss a signature aggregation scheme where a client aggregates all signatures of a given output, such that each validator receives and verifies only one signature for each transaction input. With this improvement, the computational and communication complexity of processing one transaction are both  $O(1)$  for each validator.

Without aggregating signatures, each validator receives and verifies  $2f + 1$  signatures for each input. Thus both the computational and communication complexity of processing one transaction are  $O(n)$  for each validator. However, the batch signature scheme we discuss in Section 3.4 will be more efficient than aggregating signatures for smaller  $n$ , as discussed in more detail in Section 4.1.

### 3.2 Correctness

*Double-spending.* Suppose some execution of the protocol produced two confirmed transactions  $t_1$  and  $t_2$  that spend the same output. Each confirmed transaction is signed by a validator quorum. Since the adversary controls at most  $f$  validators, at least  $f + 1$  correct validators signed  $t_1$  and  $t_2$ . Since there are  $2f + 1$  correct validators, some correct validator  $v$  signed both  $t_1$  and  $t_2$ . However, when signing a transaction, correct validators check whether they have not signed any of the inputs previously – a contradiction.

*Finality.* Given a confirmed UTXO, it can only be invalidated if some validators observe a transaction that spends the UTXO. Only the owner of the UTXO can sign such a transaction.

*Liveness.* Any  $2f + 1$  validators can confirm any transaction. Since at least  $2f + 1$  validators are correct, the correct validators can confirm transactions if the adversary refrains from participating.

### 3.3 Signature Protocol

Naively, validators can sign each transaction separately and verify separate signatures for each output. In addition to this naive approach, we design two different protocols for batch-processing the transactions, thereby vastly improving the system's performance: the Merkle scheme and the BLS scheme.

### 3.4 Merkle Scheme

In the Merkle scheme, validators combine many transactions into Merkle trees and only sign the root, effectively signing many transactions at once.

*Signing.* By *pooling* multiple signing requests, the validator collects a large number  $p$  of unsigned outputs. The hashes of the outputs  $h_0(i) = h(o_i)$  are hashed in pairs  $h_1(i) = h(h_0(i), h_0(i+1))$ , the resulting hashes are hashed in pairs  $h_2(i) = h(h_1(i), h_1(i+2))$  and so on, to create a complete binary tree of hashes, where the leaves are the hashes of the UTXOs to be signed.

The validator signs the root of the tree. For each output  $o_i$ , the validator will return to the issuer of  $o_i$  the hashes needed to compute the path from  $h(o_i)$  to the root: the hash  $x_1$  to compute  $h_1(i) = h(h_0(i), x_1)$ , the hash  $x_2$  to compute  $h_2(i) = h(h_1(i), x_2)$ , and so on. These hashes  $x_j$  together with the signature of the tree root constitute  $o_i$ 's signature in the Merkle scheme (see Appendix A.1).

When a validator signs  $p$  outputs, only one signing operation is executed (compared to  $p$  signing operations with the naive scheme). However, many hash operations must be performed for both signing and verification in this scheme.

*Signature verification.* As in the naive scheme, for an output to be confirmed, there must be  $> \frac{2}{3}n$  signatures from different validators. We verify each Merkle signature  $s_i = (\{x_j\}, s_i^{root})$ : we reconstruct the hash path to the root using the  $x_j$ 's and verify the signature of the root. There are  $p$  signatures with the same Merkle root; hence, the verification result can be cached in memory by the validator such that the validator only verifies the root signature the first time it is observed. For all  $p - 1$  subsequent encounters of the root, it suffices that the validator performs  $\log(p)$  hash operations.

*Optimal Merkle Tree Size.* If the Merkle tree is very large, the hashing time dominates the verification and signing process. However, if the Merkle tree size is small, the cryptographic operations dominate. We omit the work not related to cryptographic operations or hashing and estimate the optimal tree size.

Let  $q$  be the quorum size of the system,  $N$  the number of leaves in the Merkle tree,  $c_h$ ,  $c_s$ ,  $c_v$  the costs of hashing, signing, and verification. Let  $C_{naive}$  and  $C_{merkle}$  denote the average time cost incurred by a validator to process one UTXO under the naive and Merkle signature protocols. Each validator signs a UTXO once and later has to verify the signatures constituting that UTXO's confirmation. Hence, the expected cost for the naive scheme is  $C_{naive} = q \cdot c_v + c_s$ .

The expected cost of the Merkle scheme is

$$C_{merkle} = q \left( \underbrace{c_h \log N + \frac{c_v}{N}}_{\text{Verification}} \right) + \underbrace{\frac{2Nc_h + c_s}{N}}_{\text{Signing}} = \log N (qc_h) + \frac{1}{N} (qc_v + c_s) + 2c_h.$$

The expression is minimized by  $N = \frac{qc_v + c_s}{qc_h} \ln 2$ . For example, if the relative operation costs are around  $c_h = 1$ ,  $c_s = 63$ ,  $c_v = 107$  and there are 10 validators (quorum size 7), we estimate the optimal number of leaves in the Merkle tree at

$N^{optimal} \approx 80.4$ . For binary trees ( $\log N$  being an integer), either a Merkle tree with 64 or 128 leaves is optimal.

### 3.5 BLS Scheme

BLS [6] is a signature scheme where signatures can be aggregated. One verification operation on an aggregated signature can be used to verify all constituent signatures at once. Combining the properties of BLS with Shamir’s secret sharing, it is possible to construct a threshold signature scheme. In this scheme, each validator possesses a different private key and signs transactions for clients individually. Once a client obtains at least  $2f + 1$  signatures for their transaction, they can use the signatures to compute a unique *master* signature of their transaction. The master signature is unique and the same, irrespective of which  $2f + 1$  validators’ signatures were used to compute it.

Due to space constraints, we do not describe the working of this threshold scheme in detail (the reader can find an instructive description at [5]).

The BLS signatures are relatively costly to produce and verify. However, the scheme comes with the great advantage that aggregated BLS threshold signatures have a constant verification time, irrespective of the number of validators in the system.

Each validator receives a BLS private key, and the corresponding public keys are publicly known. The master public key is publicly known (or can be computed given the validator public keys). The validators sign the outputs of the clients with their BLS keys, similarly to the naive scheme. The validators’ signatures function on their own as usual, so clients can verify that they receive the correct signatures. After receiving  $2f + 1$  signatures, a client can compute the unique master signature for their outputs. Most importantly, when validators receive transactions to be signed, they only need to verify one master signature for each input of the transaction.

## 4 Implementation

The validator node and client are written in Go due to its performance and ease of parallelization. The implementation features the three different signature protocols described in Sections 3.3-3.5: the naive scheme, the Merkle scheme, and the BLS scheme. For signing and verification, the naive, and Merkle schemes use EdDSA with Curve25519 (Ed25519). A Go library [1] provides bindings to ed25519-donna [4]. Ed25519-donna is written in C++ and provides a fast implementation of the Ed25519 public-key signature system [3]. Batch verification can be used for greater throughput. The BLS scheme is implemented using the `herumi/bls` [11] with Go-bindings [12].

*Benchmarking Merkle tree size.* The signing and verification times for the Merkle signature scheme have been measured on the AWS reference instance for different Merkle tree sizes and 10 validators. The observed global minimum of around 64-128 leaves matched the conclusion from Section 3.4.

#### 4.1 Cryptographic Scheme Comparison

The benchmark of the cryptographic operations executed on a single core on the reference instance is denoted in Table 1.

Scheme	Operation	ns per signature
Naive	Signing	29,967
	Verifying, single	100,663
	Verifying, batch of 64	51,247
Merkle	Signing	2709
	Verifying, no caching	106,771
	Verifying, cached	6473
BLS	Signing	640,205
	Verifying	1,918,578

Table 1: Benchmarks for relevant cryptographic operations for each scheme. For the Merkle scheme, the number of leaves is 64.

If multiple signatures are verified in one batch, EdDSA can take advantage of x86 SIMD instructions. This gives verification a speedup of up to  $\approx 2$ . Assuming that the Merkle tree has 64 leaves, the Merkle scheme is about a magnitude faster than the naive scheme. Since verification of threshold BLS signatures is constant for any number of validators, we find BLS to be faster than the naive and Merkle schemes if the number of validators is  $> 37$  and  $> 475$ , respectively. Moreover, for larger quorum sizes, the naive and Merkle schemes yield larger transaction confirmations, whereas BLS confirmations have a constant size.

#### 4.2 Storage

Each validator keeps track of outputs they signed as spent. The spent outputs are stored in a thread-safe and efficient hash map. Golang’s built-in thread-safe hash map, `sync/map`, exhibits excessive lock usage and coroutine blocking; hence, we implemented a purpose-built hash map (see Appendix A.2). Running on the AWS reference instance with one coroutine per processor, Golang’s implementation reaches  $7.04 \cdot 10^5$  inserts per second, and our implementation about  $3.30 \cdot 10^7$  inserts per second.

#### 4.3 Sharding

Crucially, validators can easily shard their workload among multiple machines. Clients are assigned to different shards based on their public keys; in this implementation, inputs corresponding to different public keys cannot be mixed in one

transaction. In our implementation, the assignment of clients to shards is publicly known and clients request the machines they are assigned to. Alternatively, validators could use some load balancing approach.

## 5 Evaluation

We performed end-to-end benchmarks of the system using the Merkle scheme with servers rented at AWS. We tested with 4, 10, and 28 validators. For each number of validators, we experimented with 1, 2, and 4 shards (machines) for each validator. The server instances used in these benchmarks were c3.8xlarge with 32 virtual CPU threads and 60 GiB of RAM. Each test consisted of a preset number of client servers (corresponding to the expected throughput) generating transactions and sending them to appropriate shards of the validators. The duration of each test was ten minutes. The results of the experiment are presented in Table 2. The average CPU utilization observed was 80.9%.

We performed an additional experiment with 4 validators and 37 shards per validator (and otherwise the same setup), yielding a throughput of 1,449,847 transactions per second on average, with an average CPU utilization of 71.9%. The results of this experiment are presented with logarithmic scales in Figure 1.

#validators	4	10	28
1 shard	48,667	34,451	22,015
2 shards	92,115	72,079	44,764
4 shards	183,380	141,767	80,917

Table 2: Average transactions per second in the experiment.

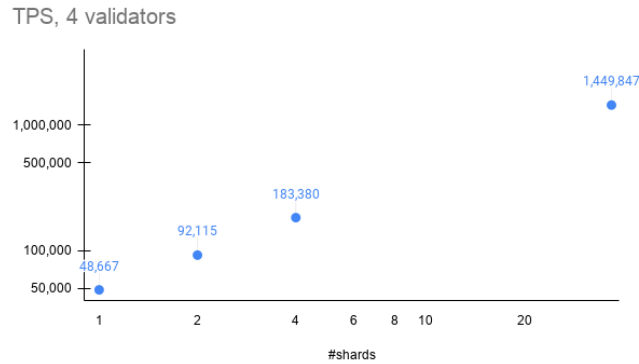


Fig. 1: TPS scalability results for 4 validators, log scales.



## 6 Related Work

The first work to suggest a simpler, consensus-free approach to processing transactions in a permissioned system was Gupta [9]. Gupta focuses on providing verifiable audit trails. We employ a similar transaction confirmation principle and focus on designing and implementing a system around horizontal scaling without loss of efficiency.

Guerraoui et al. [8] prove that the consensus number of a payment system is indeed 1 in Herlihy’s hierarchy [10].

FastPay [2] might be closest related to our work. FastPay provides an implementation of a permissioned settlement system and focuses on interfacing with a preceding, primary system. FastPay is similar in spirit to ACCEPT but more complicated in crucial aspects; for example, FastPay employs a two-phase confirmation protocol that complicates the interaction between the validators. However, FastPay does not implement the efficiency improvements described in Sections 3.4 and 3.5. Parallel processing in FastPay’s implementation is only process-based and hence does not exemplify multi-machine sharding. The performance is also not reported clearly: benchmarks are presented for the two confirmation phases separately, whereas one exhibits a bottleneck in the process-based implementation.

Cascade [14] describes how a consensus-free system can be managed similarly to proof-of-stake blockchains, thus extending the approach to the permissionless setting. Cascade also contributes some features of the protocol, such as pruning redundant contents from the blockchain and discusses some economic aspects of the permissionless setting. However, Cascade does not provide an implementation.

## References

1. ed25519 - optimized ed25519 for go. <https://github.com/oasisprotocol/ed25519> (2020)
2. Baudet, M., Danezis, G., Sonnino, A.: Fastpay: High-performance byzantine fault tolerant settlement. arXiv preprint arXiv:2003.11506 (2020)
3. Bernstein, D., Duif, N., Lange, T., Schwabe, P., Yang, B.Y.: High-speed high-security signatures. *Journal of Cryptographic Engineering* 2 (2012)
4. Bernstein, D., Duif, N., Lange, T., Schwabe, P., Yang, B.Y.: ed25519 donna. <https://github.com/floodyberry/ed25519-donna> (2020)
5. Block, A.: <https://blog.dash.org/secret-sharing-and-threshold-signatures-with-bls-954d1587b5f> (2018)
6. Boneh, D., Lynn, B., Shacham, H.: Short signatures from the weil pairing. *Journal of cryptology* 17(4), 297–319 (2004)
7. Delgado-Segura, S., Pérez-Sola, C., Navarro-Arribas, G., Herrera-Joancomartí, J.: Analysis of the bitcoin UTXO set. In: *International Conference on Financial Cryptography and Data Security*. pp. 78–91. Springer (2018)
8. Guerraoui, R., Kuznetsov, P., Monti, M., Pavlovič, M., Serebinschi, D.A.: The consensus number of a cryptocurrency. *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing - PODC '19* (2019)

9. Gupta, S.: A Non-Consensus Based Decentralized Financial Transaction Processing Model with Support for Efficient Auditing. Master’s thesis, Arizona State University (Jun 2016)
10. Herlihy, M.: Wait-free synchronization. *ACM Transactions on Programming Languages and Systems (TOPLAS)* **13**(1), 124–149 (1991)
11. Mitsunari, S.: BLS threshold signature. <https://github.com/herumi/bls> (2020)
12. Mitsunari, S.: BLS threshold signature for ETH with compiled static library. <https://github.com/herumi/bls-eth-go-binary> (2020)
13. Nakamoto, S.: Bitcoin: A peer-to-peer electronic cash system (2008)
14. Sliwinski, J., Wattenhofer, R.: Asynchronous proof-of-stake. In: *International Symposium on Stabilization, Safety, and Security of Distributed Systems*. pp. 194–208. Springer (2021)
15. Stathakopoulou, C., David, T., Vukolić, M.: Mir-bft: High-throughput bft for blockchains. *arXiv preprint arXiv:1906.05552* (2019)
16. Wood, G., et al.: Ethereum: A secure decentralised generalised transaction ledger. *Ethereum project yellow paper* **151**(2014), 1–32 (2014)

## A Concepts

This section will clarify the concepts of Merkle signatures and hash maps. Merkle signatures are used in Section 3.4 to improve the efficiency of validators. A custom hash map is developed to improve the performance of the UTXO store, as described in Section 4.2.

### A.1 Merkle Signatures

A signature scheme based on Merkle trees is used to optimize the performance of validators when creating and verifying signatures.

*Signing.* A validator  $S_i$  collects  $n$  hashes to sign where  $n = 2^k, k \in \mathbb{N}$ . The hashes are combined into a Merkle tree. The validator  $S_i$  signs the Merkle root  $m$  using EdDSA, denoted as  $m_{S_i}$ . For each hash  $h_i, i \in \{1, \dots, n\}$ , the validator calculates the Merkle path and outputs  $path_{h_i}$ .  $path_{h_i}$  consists of the hashes and side (left or right) of the nodes from  $h_i$  to the Merkle root (in blue). The resulting signature for  $h_i$  by  $S_i$  is the tuple  $(path_{h_i}, m_{S_i})$ . An example can be seen in Figure 2.

Signing  $n$  hashes using Merkle signatures is more efficient than signing  $n$  hashes separately: the cost of an EdDSA signing operation  $c_s$  is much higher than the cost of a hash operation  $c_h$ . The cost of signing  $n$  hashes with Merkle signatures is  $2n \cdot c_h + c_v$  whereas the cost of signing  $n$  hashes separately (without Merkle signatures) is  $n \cdot c_v$ .




Fig. 2: Merkle tree and Merkle signature for hash  $h_3$ .

*Verifying.* First, the Merkle root  $m'$  is reconstructed from the path  $path_{h_i}$ . If  $path_{h_i}$  is a valid path,  $m'$  matches  $m$ . Finally,  $m_{S_i}$  is verified using EdDSA.

Verifying  $n$  hashes using Merkle signatures is also more efficient than verifying  $n$  hashes separately because the cost of an EdDSA verification operation is high, even higher than an EdDSA signing operation (and thus also higher than the cost of a hashing operation). The Merkle root can be reconstructed with a cost of  $\log_2(n) \cdot c_h$  where  $c_h$  is the cost of hashing. The Merkle root  $m$  only has to be verified for the first encountered Merkle signature, after that, the result of the verification can be cached, making this signature perform very well.

### A.2 Hash Maps

Hash maps are used to efficiently calculate set membership in the context of spent UTXO identifiers. Elements  $e \in E$  (in our case  $id_i$ ) are hashed into the hash range  $H = \{0, \dots, l\}$  and used as an index to an array of linked lists. A linked list at index  $i \in \{0, \dots, l\}$  contains all items  $e$  where  $\text{hash}(e) = i$ . The linked list is then traversed.

Hashing and indexing are implemented as a lock-free operation; traversing and modifying a linked list are protected by a mutex. If  $l$  is large enough, the probability of hash collisions is small, minimizing lock contention and the length of the linked list.