

# Prism Removes Consensus Bottleneck for Smart Contracts

Gerui Wang

Department of Computer Science  
University of Illinois Urbana-Champaign  
Urbana, USA  
Email: geruiw2@illinois.edu

Shuo Wang

Department of Electrical Engineering  
Stanford University  
Stanford, USA  
Email: shuow94@stanford.edu

Vivek Bagaria

Department of Electrical Engineering  
Stanford University  
Stanford, USA  
Email: vbagaria@stanford.edu

David Tse

Department of Electrical Engineering  
Stanford University  
Stanford, USA  
Email: dntse@stanford.edu

Pramod Viswanath

Department of Electrical and Computer Engineering  
University of Illinois Urbana-Champaign  
Urbana, USA  
Email: pramodv@illinois.edu

**Abstract**—The performance of existing permissionless smart contract platforms such as Ethereum is limited by the consensus layer. Prism [1] is a new proof-of-work consensus protocol that provably achieves throughput and latency up to physical limits while retaining the strong security guarantees of the longest chain protocol. This paper reports experimental results from implementations of two smart contract virtual machines, EVM and MoveVM, on top of Prism and demonstrates that the consensus bottleneck has been removed. Code can be found at <https://github.com/wgr523/prism-smart-contracts>.

**Keywords**—*smart contract; consensus; scalability;*

## I. INTRODUCTION

Existing permissionless smart contract platforms such as Ethereum is based on the longest chain consensus protocol, the original blockchain protocol invented by Nakamoto [2]. While maintaining high security against adversarial attacks, it is well known that the longest chain protocol suffers from poor throughput and latency performance. Hence, the performance of these platforms is limited by the consensus layer.

This limitation has led to practical congestion in the network; a noteworthy instance occurred when CryptoKitties made its debut on Ethereum, a spike of transactions rushed into the system, far exceeding Ethereum’s supported throughput. The pending transaction queue was growing quickly, and users had to increase transaction fees to incentivize miners to add their transactions to the chain. Decentralized Finance applications have been rapidly growing over the last few years and as it gets more popular in the near future, the demand will continue to grow, making the performance scaling of smart contract platforms an urgency.

Several promising efforts to scale the performance have been proposed. Almost every major live smart contract platform such as Ethereum, Algorand, and Tron are optimizing their existing smart contract engines to increase the throughput. A few others like Libra (led by Facebook) and Hyperledger Fabric (led by IBM) have taken the route of permissioned blockchains to obtain higher throughput. On the other hand,

Ethereum foundation has taken a sharding approach to support higher throughput. Optimistic Rollup [3], ZK-Rollup [4], and Arbitrum [5] are other off-chain scaling solutions built on top of an existing smart contract platform such as Ethereum. In these off-chain solutions, not every validator node needs to keep track of the execution of the off-chain contracts, which leads to an improved overall efficacy but at the expense of security.

Prism [1] is a recent permissionless proof-of-work (PoW) consensus protocol which naturally scales the performance of the longest chain protocol. It provably achieves throughput and latency up to computation and communication limits of the underlying physical network, while retaining the strong security guarantees of the longest chain protocol. An implementation of Prism [6] scales performance significantly in a Bitcoin-like payment system, improving the throughput of Bitcoin by about 4 orders of magnitude. The question remains as to whether Prism can successfully support a general smart contract platform and remove the consensus bottleneck. Indeed, not every blockchain consensus protocol is extensible to a smart contract platform (eg: Spectre [7]) and scalably integrating consensus with smart contract platforms is nontrivial.

This paper demonstrates that Prism can support general smart contract platforms and provide a very high level of performance. We present the design and implementation of Prism that provides a flexible interface for connecting with two common smart contract virtual machines. We report experimental results from implementation of two smart contract virtual machines, Ethereum VM (EVM) and MoveVM, on top of Prism. Fig. 1a shows throughput results for running several canonical smart contract applications on EVM on Prism, while Fig. 1b shows analogous results for MoveVM on Prism. As can be seen, the throughputs are very close to that of virtual machine execution only *without consensus*, and much larger than the throughput using the longest chain protocol. Thus, we conclude that smart contract platforms built on Prism can perform without the consensus layer bottleneck.

The rest of this paper is structured as follows. In §II




Fig. 1: Throughput of Prism clients; experimented with 100 nodes on several applications: *Native Payment*, *Do Nothing*, and *ERC20* are lightweight applications whereas others are heavyweight ones. In MoveVM, *Native Payment* is essentially the same as *ERC20*. The reader is cautioned against comparing performance across the two VM’s, as EVM is a mature technology while MoveVM is current under active development. Rather, the main point of obtaining results in two VM’s is to demonstrate the flexibility of Prism. Moreover, we did not compare with the the performance of MoveVM on Libra consensus because it is permissioned while Prism is permissionless.

we discuss smart contract scaling approaches in different dimensions. §III gives a brief overview of Prism consensus protocol. In §IV, we describe our design and implementation of Prism with EVM and MoveVM. We present evaluation results of various canonical applications in EVM and MoveVM, and discuss their implications in §V. Conclusion is in §VI.

## II. RELATED WORK

The throughput of blockchains with smart contract platform can be increased at three different points on the blockchain stack. The first approach is to improve the execution speed of the virtual machine engine. A basic approach is to optimize the execution of individual op codes (followed in EVM clients

such as Parity Ethereum and Geth) or by designing a new set of op codes from first principles (followed by Libra to arrive at MoveVM [8]). A more involved approach is to execute smart contracts in *parallel* similar to the modern design of databases such as MySQL [9] and Postgres [10]. The first technique is to run multiple smart contracts in parallel where smart contracts acquire locks on a data before editing to ensure no data is simultaneously edited by more than a single smart contract; this method is used in [11], with a 33% improvement in throughput. An alternative approach uses optimistic concurrency with rollbacks; here multiple smart contracts execute in parallel (without locks) and in the case when two smart contracts running in parallel try to edit the same data, one of them is rolled back and executed later; this approach is explored in [12]–[15] where 3-4x improvement in throughput is observed. Although the improvement in throughput is significant in these methods, it exposes the blockchain to new kinds of adversarial attacks. Moreover these methods don’t address metering which is a critical component to align incentives.

Even though the current VMs have low throughput, the current bottleneck in today’s blockchain platform is the consensus protocol itself. Longest chain protocol and its current variants do not saturate the performance of the underlying VMs (refer Fig. 1 for details). Therefore, the second approach of designing high throughput consensus protocols is a natural avenue to scale smart contract platforms. One method is to move from permissionless to permissioned consensus protocols which can support high throughput, and Facebook’s Libra [16] and IBM’s Hyperledger Fabric [17] take this path. Libra has chosen a recent high-performing Byzantine Fault Tolerance (BFT) consensus protocol (HotStuff [18]); Hyperledger Fabric [17] proposes the execute-order-validate paradigm in order to attain both performance and extensibility, where (1) participants execute transactions and collect endorsements for the executions, (2) responsible participants order these executed transactions through a consensus protocol, and (3) transactions are validated by all participants. However, these approaches sacrifice the very important characteristic of being permissionless. In this paper we take the approach of designing and implementing a high throughput *permissionless* consensus protocol, Prism, which achieves high throughput. Protocols such as OHIE [19], Algorand [20], Bitcoin-ng [21] take a similar route. To the best of knowledge, there do not exist implementations running smart contracts on top of these protocols; hence we have not been able to make a direct comparison with Prism’s performance.

The third approach is Plasma and sharding. In 2015, Poon and Buterin proposed Plasma [22], along the lines of MapReduce, an off-chain scaling solution. Many offshoots of plasma have been proposed by different communities and refer the following webpage [23] for an overview. At a high level, Plasma is a network of secondary chains, each custom designed to serve different needs. These chains interact among each other and the main chain (on a need basis) to resolve conflicts using fraud proofs. This approach has weaker security properties and, in particular, susceptible to the “mass exit” attack. To overcome some of these security vulnerabilities, Ethereum 2.0 [24], near [25], polkadot [26], and Trifecta take the sharding approach which *horizontally* scales the throughput by running multiple instances of blockchains and pooling them




Fig. 2: Prism: Factorizing the blocks into three types of blocks: proposer blocks, transaction blocks and voter blocks.

to obtain high security. Even though this approach has better security than plasma, overall it has lower security compared to the pure consensus protocols in the previous paragraph.

### III. OVERVIEW OF PRISM

The selection of a main chain in a blockchain protocol can be viewed as electing a leader block among all the blocks at each level of the blocktree. In this light, the blocks in the longest chain protocol can be viewed as serving three distinct roles: they stand for election to be leaders; they add transactions to the main chain; they vote for ancestor blocks through parent link relationships. The latency and throughput limitations of the longest chain protocol are due to the *coupling* of the roles carried by the blocks. Prism removes these limitations by factorizing the blocks into three types of blocks: proposer blocks, transaction blocks, and voter blocks (Fig. 2). Each block mined by a miner is randomly sortitioned into one of the three types of blocks, and if it is a voter block, it will be further sortitioned into one of the voter trees.

The proposer blocktree anchors the Prism blockchain. Each proposer block contains a list of reference links to transaction blocks, which contains transactions, as well as a single reference to a parent proposer block. Honest nodes mine proposer blocks on the longest chain in the proposer tree, but the longest chain does not determine the final confirmed sequence of proposer blocks, known as the *leader sequence*. We define the *level* of a proposer block as its distance from the genesis proposer block, and the *height* of the proposer tree as the maximum level that contains any proposer blocks. The leader sequence of proposer blocks contains one block at every level up to the height of the proposer tree, and is determined by the *voter chains*.

There are  $m$  voter chains, where  $m \gg 1$  is a fixed parameter chosen by the system designer. For example, we choose  $m = 1000$  in our experiments. The  $i$ th voter chain is comprised of voter blocks that are mined on the longest chain of the  $i$ th voter trees. A voter block votes for a proposer block by containing a reference link to that proposer block, with the requirements that: 1) a vote is valid only if the voter block is in the longest chain of its voter tree; 2) each voter chain votes for one and only one proposer block at each level. The leader block at each level is the one which has the highest number of votes among all the proposer blocks at the same level (tie broken by hash of the proposer blocks.) The elected leader blocks then provide a unique ordering of the transaction blocks to form the final confirmed ledger.

By decoupling the various types of blocks, Prism can

provably achieve low latency and high throughput while maintaining high security.

#### A. Latency

The votes from the voter trees secure each leader proposer block, because changing an elected leader requires reversing enough votes to give them to a different proposer block in that level. Each vote is in turn secured by the longest chain protocol in its voter tree. If the adversary has less than 50% hash power, and the mining rate in each of the voter trees is kept small to minimize forking, then the consistency and liveness of each voter tree guarantee the consistency and liveness of the ledger maintained by the leader proposer blocks. However, this would appear to require a long latency to wait for each voter block to get sufficiently deep in its chain. What is interesting is that when there are many voter chains, the same guarantee can be achieved without requiring each and every vote to have a very low reversal probability, thus drastically improving over the latency of the longest chain protocol.

**Theorem 1** (Latency, Thm. 4.8 [1]). *For an adversary with  $\beta < 50\%$  of hash power; network propagation delay  $D$ , Prism with  $m$  chains confirms honest<sup>1</sup> transactions at reversal probability  $\epsilon$  guarantee with latency upper bounded by*

$$Dc_1(\beta) + \frac{Dc_2(\beta)}{m} \log \frac{1}{\epsilon} \text{ seconds}, \quad (1)$$

where  $c_1(\beta)$  and  $c_2(\beta)$  are  $\beta$  dependent constants.

For large number of voter chains  $m$ , the first term dominates the above equation and therefore Prism achieves near optimal latency, i.e. proportional to the propagation delay  $D$  and independent of the reversal probability.

#### B. Throughput

To keep Prism secure, the mining rate and the size of the voter blocks have to be chosen such that each voter chain has little forking. The mining rate and the size of the proposer blocks have to be also chosen such that there is very little forking in the proposer tree. Otherwise, the adversary can propose a block at each level, breaking the liveness of the system. Hence, the throughput of Prism would be as low as the longest chain protocol if transactions were carried by the proposer blocks directly.

To decouple security from throughput, transactions are instead carried by separate transaction blocks. Each proposer block when it is mined refers to the transaction blocks that have not been referred to by previous proposer blocks. This design allows throughput to be increased by increasing the mining rate of the transaction blocks, without affecting the security of the system. The throughput is only limited by the computing or communication bandwidth limit  $C$  of each node, thus potentially achieving 100% utilization.

**Theorem 2** (Throughput, Thm. 4.4 [1]). *For an adversary with  $\beta < 50\%$  fraction of hash power and network capacity  $C$ , Prism can achieve  $(1 - \beta)C$  throughput and maintain liveness in the ledger.*

<sup>1</sup>Honest transactions are ones which have no conflicting double-spent transactions broadcast in public.

#### IV. DESIGN AND IMPLEMENTATION

We implement a Prism full-node client with VMs in around 10,000 lines of RUST code. In this section, we describe the architecture of the client and highlight several design choices that are tailored to Prism consensus.

##### A. Architecture

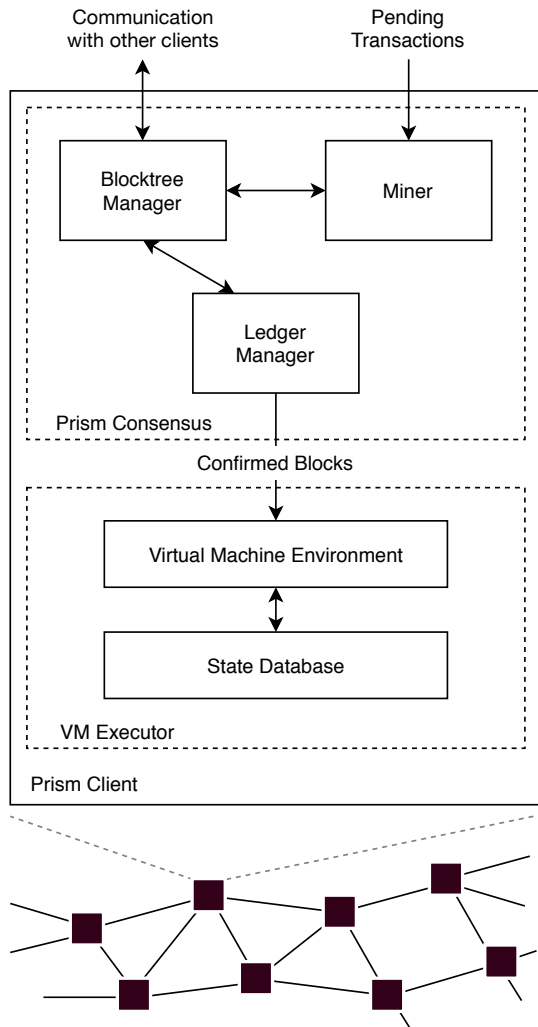


Fig. 3: Architecture of the Prism client. In the peer-to-peer network, each node is running a Prism full-node client.

Our implementation of Prism full-node client consists of two modules, Prism Consensus module and Virtual Machine Executor (VM Executor) module. Prism Consensus module is in charge of exchanging blocks with peers, following Prism consensus to confirm blocks, and push confirmed blocks to VM Executor. VM Executor maintains the *state* of the confirmed ledger, i.e., the state that results from executing transactions up to the last confirmed block. When VM Executor receives new confirmed blocks from Prism Consensus, it retrieves transactions from those blocks and updates the state accordingly. This architecture is illustrated in Fig. 3.

**Prism Consensus** module can be divided into the following three parts:

- 1) **Blocktree Manager**, which maintains the client’s view of the blockchain, and exchanges blocks with peers;
- 2) **Ledger Manager**, which confirms blocks by following Prism protocol, and pushes confirmed blocks to VM Executor;
- 3) **Miner**, which contains a transaction memory pool and assembles new blocks.

**Blocktree Manager** consists of an event loop and a thread pool. The event loop keeps listening to events such as sending/receiving blocks, and assigns a thread from the thread pool to process it. When the client receives a new block from a peer, Blocktree Manager checks its proof of work, and stores the block locally. After that, it relays the block to peers in case they have not received it. It then checks data availability, i.e., whether all the blocks referred by reference links in this block have been received. If not, it buffers the block and defers further processing until data availability is satisfied. After data availability is satisfied, Blocktree Manager checks sortition and transaction signatures. Finally the block is inserted into Prism blocktree.

**Ledger Manager** is a busy-waiting loop that queries Blocktree Manager periodically to see whether there are new confirmed blocks, following Prism’s confirmation rule. If there are, it will retrieve the blocks from local storage and push them to VM Executor via a message-passing channel. The choice of the busy-waiting loop suits the high transaction workload since the busy-waiting overhead is negligible when it takes a long time to retrieve a large number of blocks and push them to VM Executor. Both Blocktree and Ledger Managers use RocksDB as the storage backend [27], [28]; this choice is made due to its high performance and ease of integration.

**Miner module** maintains a memory pool that collects pending transactions and assembles them into new blocks. The Miner module does not actually try to solve the PoW hash inequality, instead simulating the mining process by a Poisson process (of fixed growth rate, corresponding to the mining difficulty level); the Poisson processes are statistically independent across the different nodes (matching the distributed nature of PoW mining). When a new block is mined, it is pushed to Blocktree Manager, which will broadcast the block to peers. Transactions carried by assembled or received blocks are checked for duplication in the memory pool, with duplicates being purged.

**VM Executor** is in charge of maintaining the *state database*, i.e., the persistent storage for the state of the confirmed ledger. State database stores account information such as address and balance, and manage data in a hash accumulator (Merkle Patricia tree is used in Ethereum and sparse Merkle tree is used in Libra). VM Executor receives confirmed blocks from Ledger Manager, retrieves transactions from those blocks, and executes them sequentially. To execute a transaction, VM Executor first initializes a virtual machine environment, such as program counter, stack, and memory. Then it executes the instructions coded inside the transaction and/or the smart contract, during which it may interact with the state database. The execution result of a transaction will be a success or a failure, depending on whether the transaction is valid or not. Invalid transactions with failure results should be *sanitized* out of the confirmed ledger and have no effect

on the state. Valid transactions will update the state according to the execution result. After executing all transactions in a confirmed block, VM Executor commits the updates to the state database.

We ported the VM Executors from two open source projects, Open Ethereum [29] (popularly known as Parity Ethereum) and Libra [16], and adapt the structure of transactions, the hash function, and the signature schemes to these projects respectively. The port only required us to add or modify less than 20 lines of code (LOC) for Open Ethereum and less than 160 LOC for Libra in their RUST language codebases; in addition, 2 LOC were modified in Move language for Libra codebase. The two VM Executors run single threads, with no parallel transaction execution capability. We will use the name of virtual machines, EVM and MoveVM, to refer them hereafter.

## B. Highlights

The key design and implementation challenge is in translating the high throughput, low latency and high confirmation probability that Prism provides on raw block and transaction level into an application layer programming construct via the virtual machine intermediaries. On one hand, the client must process blocks and transactions at a rate much higher than most traditional blockchains. On the other hand, low latency and high confirmation probability enables confirmation of the ledger, which the implementation can benefit from. Here, we highlight several implementation choices that are tailored for Prism consensus and distinguish our implementation from traditional blockchains.

*a) Confirmation:* In Ethereum and other longest chain protocols, the state of the longest chain tip is used for transaction validation. However, blocks in longest chain may be switched due to honest or adversarial forking blocks. To smoothly update state when the longest chain switch happens, Ethereum's implementation keeps a short-term journal containing actions in recent forking blocks. This makes the management of state less efficient, which is a particular impediment due to the high mining rate (and high throughput) of Prism. In our design, we find it relevant to only maintain the state of the last confirmed block; this is because of two reasons: (a) Prism guarantees confirmation with overwhelmingly high probability (e.g.  $1 - 10^{-9}$ ) so confirmed blocks are not likely to be deconfirmed. (b) Prism does not validate transactions before including them in blocks so it is unnecessary to maintain the state of the unconfirmed latest proposer block. This not only makes maintenance more efficient, but also enables the integration with VM of BFT consensus such as MoveVM.

In traditional blockchains (Bitcoin and Ethereum), blocks are mined at a relatively low rate and a newly mined block is likely to change the longest chain. Hence in their implementation, they update state when they receive a new block. In Prism, blocks are mined at a high mining rate; confirming blocks and updating state upon receipt of a new block would be onerous – we make a design choice to update the state only when blocks are confirmed and to conduct the confirmation procedure at periodic intervals.

*b) Decoupling Transaction Validation and State Update:* In most traditional blockchains, transaction validation

and state update are coupled with consensus. For example, Ethereum miners must make sure all the transactions in a block are valid, update Ethereum state accordingly, and record the result state root in that block. Prism, by design, decouples transaction validation and state update from consensus: Prism miners do not conduct transaction validation or update state. Only after a block is confirmed, transactions in it are validated, and state is updated accordingly. In this procedure, invalid transactions are *sanitized* out of the confirmed ledger. We note that invalid transactions still incur gas fees for the senders and thus a rational user has no incentive to send invalid transactions. If the transaction sender has inadequate balance to pay the gas fee, the transaction will be treated as spam and skipped. Nevertheless this type of invalid transactions could reduce the utility of network bandwidth. To mitigate this spamming attack, miners could validate transactions (by checking sender's balance is no less than gas fee) with respect to their latest confirmed state, giving the adversary only a small window to create invalid transactions and spam the system. By this method, spam traffic is reduced by 80% whereas the confirmation latency is only increased by 5 seconds [6]. We did not implement this defense against spamming attack in this paper.

Prior work Hyperledger Fabric [17] also separates transaction validation and state update from consensus by the following three-step execute-order-validate paradigm: (1) nodes execute transactions and collect endorsements for the executions, (2) responsible nodes order the executed transactions through a consensus protocol, and (3) the ordered transactions are validated by all nodes and the state is updated according to valid transactions. Prism is similar to Hyperledger Fabric in the sense that they both separate transaction validation and state update from consensus. Notably, a common feature for Prism and Hyperledger Fabric is that the ledger could possibly contain invalid transactions at first, which would be sanitized out of the ledger later. Nevertheless, they are different in two ways. (a) Prism orders transactions, then executes and validates them. In other words, it adopts the order-execute-validate paradigm in contrast to Hyperledger Fabric's execute-order-validate paradigm. This order-execute-validate paradigm is closely related to traditional consensus protocols, whereas the paradigm of Hyperledger Fabric deviates far from traditional ones. (b) In Hyperledger Fabric, the execution of a transaction only occurs on a special set of nodes, and the validation requires these nodes to sign and transmit endorsements. Whereas in Prism, validation does not require such endorsements in that a transaction is validated by checking its local execution outcome; this execution and validation are replicated on every node. These differences imply that Prism fits well with current platforms such as Ethereum and can replace traditional consensus protocols seamlessly, while Hyperledger Fabric requires some efforts to design full-node and light-node clients in order to meet its novel requirements.

*c) No Pending Transaction Exchange:* Most traditional blockchain clients exchange pending transactions in their memory pools with peers. Because the block mining rate is very low and the next block author is unpredictable, transaction exchange is necessary to ensure that pending transactions get included in the next block. This reduces network bandwidth utility since transactions are broadcast twice in the network: first as pending transactions and then as part of a block.

In Prism, pending transaction exchange can be onerous to the network bandwidth, due to the high throughput. We design our implementation to avoid exchanging pending transactions, by noting that a pending transaction can be easily included in a *new* block in a very short amount of time by any individual miner thanks to the high mining rate of Prism’s transaction blocks. Transaction blocks carrying pending transactions are broadcast to peers, in the same way blocks are broadcast in traditional blockchains. Notice that a user can still send a transaction to multiple miners for redundancy; however, miners need not exchange it. This avoids the waste of network bandwidth and contributes to the final high throughput.

*d) Signature Verification in Consensus:* Transaction signature verification is a significant fraction of total computation; this burden is only worse when the achieved throughput is higher. We design our implementation to conduct the signature verification in *parallel inside Prism Consensus* via the thread pool functionality. This is a departure from implementations in EVM (Ethereum) and MoveVM (Libra) which conduct signature verification inside the VM executor. Either sequentially or in parallel, signature verification burdens the VM executor and harms the throughput.

## V. EVALUATION

In this section we describe our experiments and performance results of our implementation of a prototype client designed based on the guidelines highlighted in the previous section. We describe experiment settings and the applications that we measure (§V-A). Then we present the throughput and confirmation latency results of Prism integrated with two virtual machines, EVM and MoveVM, from which we analyze that Prism removes the consensus bottleneck (§V-B, §V-C). In addition, we measure how our design and implementation of Prism scales with more network participants (§V-D).

### A. Experiment Setting

We evaluate our implementation of Prism by integrating it with two smart contract virtual machines: *EVM Prism* and *MoveVM Prism* respectively. The performance (upper bound) baselines are provided by *VM Executor Only* (single node, no consensus) and *Prism Consensus Only* (no smart contract platform, raw transaction throughput). *VM Executor Only* experiment feeds transactions to VM Executor running on a single node and demonstrates the optimal throughput of the VM Executor. *Prism Consensus Only* experiment runs consensus with raw blocks and transactions and measures the raw data throughput. It shows the performance that the consensus is able to support. In addition, we also implement Ethereum’s consensus protocol (essentially the longest chain protocol) and its performance provides a (lower bound) baseline.

**Applications:** We evaluate a suite of canonical applications, which can be classified into three categories.

1) *Basic applications:* We evaluate two basic applications: *Native Payment* and *Do Nothing*. *Native Payment* transactions are payments of native tokens in those smart contract platforms. *Do Nothing* is a contract with a void function, and is the simplest possible contract.

2) *Benchmark applications:* To test Prism client with standard computation or storage read/write, we propose two

applications: *CPU Heavy* and *IO Heavy*. *CPU Heavy* runs a worst case of quick sort for an integer array of length 255. *IO Heavy* does key-value pair write 255 times followed by key-value pair read 255 times for both forward and backward order (thus total 510 times). The value type is bytes32 in EVM and bytearray in MoveVM, which are both 256-bit data type.

3) *Realistic applications:* As a counterpoint to the above applications, we evaluate here the performance with respect to two real world applications: *ERC20* and *CryptoKitties*. *ERC20* is an Ethereum token standard [30], and we implement it by using the reference implementation in [31]. *CryptoKitties* is a game that allows users to breed virtual pets. The genes of offspring are determined by a function named *mixGenes* that mixes the genes of its parents [32]. We adopt *mixGenes* function in our experiments, and feed random parent genes to it. This function is significantly computational heavy compared to basic applications.

Applications for EVM are developed in Solidity programming language. We use the official Solidity compiler v0.6.3 to compile all smart contracts to bytecode except for *CryptoKitties*, which we follow the version v0.4.18 in the contract. We set the compiler to Constantinople version and enable the default optimization. When creating a smart contract in EVM, an account address is created and bytecode is stored under the address. Applications for MoveVM are developed in Move IR. The smart contracts are first published as modules under the sender’s address and then are called via scripts. We use Move IR compiler to compile the modules and scripts to bytecode. We have basic applications and benchmark applications and they have the same functionality as corresponding applications for EVM. Native tokens in MoveVM have essentially the same function as ERC20 tokens in EVM, hence *ERC20* experiment for MoveVM is unnecessary. As the Move language is in rapid development and not yet mature at the moment of our experiments, it is not straightforward to implement *CryptoKitties* in MoveVM.

Table I presents the statistics of applications. Transaction sizes differ because we pass different input parameters to these applications. Number of instruction and gas are indicators of the complexity in terms of both computation and storage read/write. MoveVM does not provide the statistics for number of instruction.

To generate the workloads for our evaluations, we implement a transaction generator that periodically generates transactions and push them into the mempool, generating different transaction types for different applications. We cap the generation rate according to the throughput of *VM Executor Only* experiment, in order not to exhaust the virtual machine.

We acquire data from the first 100 million transactions on Ethereum to derive a distribution on the number of transactions sent and received by an account. We sample our transactions using this distribution to mimic the usage of Ethereum in our experiments. In our experiments we use 10,000 accounts in total for both sender and receiver. The transaction generator of each node is initialized with 10,000 key pairs; one key pair for each account. In order to mimic the usage of Ethereum

---

<sup>2</sup>Since we pass random inputs to *CryptoKitties*, the number is also random and we present an approximation in the table.

TABLE I: EVM and MoveVM application statistics.

	Native Payment	Do Nothing	CPU Heavy	IO Heavy	ERC20	Crypto-Kitties
EVM Tx Size (Bytes)	533	536	567	567	601	631
EVM Gas	21000	21394	334390	435244	26602	140000 <sup>2</sup>
Num of Instruction	0	32	88417	25364	309	25000 <sup>2</sup>
MoveVM Tx Size (Bytes)	424	329	365	366	N/A	N/A
MoveVM Gas	43076	629	2275420	2956846	N/A	N/A

for *Native Payment* and *ERC20*, each node randomly and independently draws a sender and a receiver address from the aforementioned distribution. Other applications like *Do Nothing*, *CryptoKitties*, *CPU Heavy*, and *IO Heavy* have a fixed receiver (EVM) or no receiver (MoveVM) and hence we only sample the sender address.

**Experiment environment.** We perform our experiments on Amazon EC2’s 100 c5d.4xlarge instances. Each instance has 16 CPU cores, 32 GB memory, and NVMe SSD storage. Each instance hosts one Prism client and they are connected to form a random 4-regular topology; the diameter of the network is 6. To emulate a realistic peer-to-peer network, we introduce a propagation delay of 120 ms on each link to match the typical delay in Ethereum’s network [33], and a rate limiter of 300 Mbps for both ingress and egress traffic, except for *Prism Consensus Only* experiment where the rate limiter is 600 Mbps in order to show the performance upper bound that the consensus can reach.

**Parameters.** For *EVM Prism* and *MoveVM Prism*, we choose a high adversarial hash power capability of  $\beta = 0.4$  and a very low deconfirmation probability  $\epsilon = 2 \times 10^{-9}$ . We use  $m = 1000$  voter chains and cap the size of transaction blocks to be 200 tx/block. Given the testbed with 120 ms peer-to-peer delay, we tune the mining rate of Prism’s proposer and voter blocks to be 0.08 block/s, at which the empirical forking rate<sup>3</sup> is less than 0.11 in all experiments, and thus it ensures the security of Prism. We tune the mining rate of transaction blocks differently for different applications to match the throughput of *VM Executor Only* experiment: In *EVM Prism*, *Native Payment* 108; *Do Nothing* 180; *ERC20* 70; *CryptoKitties* 3.78; *CPU Heavy* 1.08; *IO Heavy* 2.34 block/s. In *MoveVM Prism*, *Native Payment* 12.6; *Do Nothing* 7.2; *CPU Heavy* 1.44; *IO Heavy* 3.06 block/s.

For the *Prism Consensus Only* experiment, we increase the size of transaction blocks to 400 tx/block and the mining rate to 200 block/s in order to show the performance upper bound that the consensus can reach. As for the *Ethereum* experiment, we use a mining rate of 0.1 block/s and a block size of 200

tx/block, which resemble the live Ethereum parameters.




Fig. 4: Time series plot of the throughput for *Do Nothing* application in *EVM Prism* for 10 minutes. Around the first 30 seconds there are very few processed transactions, since the clients are just started and have not extended the ledger significantly. After the ledger starts to be extended significantly, the throughput soon increases and becomes stable. This phenomenon occurs in all Prism experiments.

All experiments are run for at least 10 minutes. As we see in the time series plot of the throughput (Fig. 4), in the first several seconds, the nodes don’t process any transaction because they just started mining blocks and there are not enough blocks to extend the confirmed ledger. This phenomenon only happens at the beginning and does not affect the performance afterwards. Hence, the final throughput calculation involves the average performance over the last 9 minutes of the experiment.

### B. Throughput and Latency of EVM

In this experiment, we measure the transaction throughput and confirmation latency of various applications in *EVM Prism* and analyze the difference in throughput for different applications. We also compare the throughput with *EVM Executor Only* experiment, the optimal throughput of EVM on a single node. If the former is able to reach the latter, then the throughput of our Prism client is very close to the optimal throughput of the virtual machine and we can conclude that Prism removes the consensus bottleneck for smart contracts. Finally we compare *EVM Prism* with *Prism Consensus Only* to study whether Prism is able to support even higher throughput without the limitation of the virtual machine. This experiment would also indicate whether *EVM Prism*’s performance can be further improved if the underlying virtual machine becomes faster.

TABLE II: Throughput in terms of tx/s on EVM applications.

	Native Payment	Do Nothing	CPU Heavy	IO Heavy	ERC20	Crypto-Kitties
EVM Executor Only	21535	35723	207	467	13095	710
EVM Prism	18660	35329	197	447	11210	661
Prism Consensus Only <sup>4</sup>	98022	97473	92144	92144	86931	82798
Ethereum				21		

<sup>3</sup>Forking rate is calculated by  $1 - \frac{\# \text{ blocks in longest chain}}{\# \text{ blocks}}$ .

**Throughput:** As shown in Table II, for *EVM Prism*, the throughput of two basic applications is able to reach 18K and 35K tx/s respectively. For *ERC20*, *EVM Prism* gets 11K tx/s. The throughput of these three applications shows that we have a good chance to get above ten thousand tx/s for those applications that do not involve heavy computation or storage read/write. The reason that *Do Nothing* is almost twice as fast as *Native Payment* is that for *Do Nothing*, the VM Executor module updates the account information of a random sender per transaction and a fixed receiver contract account, whereas for *Native Payment* it updates a random sender and a random receiver account information. As a result, the VM Executor needs to maintain the state database and hash accumulator for half account information updates in *Do Nothing* as that in *Native Payment*.

For the *CryptoKitties* application, *EVM Prism* achieves 661 tx/s due to its computational heavy nature. Similar things happen for *CPU Heavy* and *IO Heavy*, which get 197 and 447 tx/s respectively. According to the statistics in Table I, these applications require more than 25K instructions in the virtual machine, which explains their low throughput. However, the low throughput in both *EVM Prism* and *EVM Executor Only* also indicates that EVM has a large opportunity to improve the efficiency of execution. We write exactly the same *CPU Heavy* application in Java and run in JVM, and we get a throughput over 90K tx/s. Considering the large gap between 197 and 90K, we believe that EVM has the potential to execute instructions more efficiently. We don't compare *IO Heavy* or *CryptoKitties* since they are not as straightforward to implement as a standalone program in Java.

**Is Prism consensus the bottleneck?** For all EVM applications, *EVM Prism* reaches 85% of *EVM Executor Only* throughput. This high percentage indicates that *EVM Prism* is able to reach the optimal EVM throughput very closely. As for *Prism Consensus Only*, we can see the high throughput over 80K tx/s for all applications. This high number illustrates the ability of supporting a high throughput without the limitation of the virtual machine. It also shows that if the virtual machine becomes faster in the future, Prism is able to support its performance as well. Hence, Prism consensus is not the throughput bottleneck; the virtual machine itself is the bottleneck.

Compared to 21 tx/s in *Ethereum* experiment, which adopts Nakamoto's longest chain consensus, it is clear that the current Ethereum is limited by consensus.

**Latency:** The end to end latency of a transaction consists of two parts: confirmation latency and execution latency. Confirmation latency is the time between a transaction is generated and the corresponding block is confirmed. This latency is decided by Prism's confirmation rule and has a proved bound [1]. Execution latency is the time that a transaction waits in a queue to be executed and the time of execution. As long as we cap the transaction generation rate below the optimal virtual machine throughput in experiments, the time in the queue is negligible. Also the execution time is less than ten milliseconds since all applications have over one hundred tx/s throughput. Hence, execution latency is negligible compared to confirmation latency.

Prism's confirmation rule guarantees a confirmation latency regardless of its throughput. In all Prism experiments including *EVM Prism*, *MoveVM Prism*, and *Prism Consensus Only*, the confirmation latency is no more than 130 seconds. Notice that this latency is achieved with adversarial ratio  $\beta = 0.4$  and reversal probability  $\epsilon = 2 \times 10^{-9}$ . To provide the same latency under the same condition in Ethereum, it needs to wait for ( $k = 267$ )-deep [2] and it translates to 2670 seconds if a block is mined in 10 seconds on average.

**Resource utility:** In a Prism client, the Prism Consensus module uses multiple threads to process messages from/to peers efficiently. The VM Executor module, on the contrary, runs in a single thread. In addition, RocksDB uses a few threads in the background. In total, a Prism client should only use no more than 50 threads. In our experiments, the live usage of CPU never exceeds 50% per core on average (notice that one instance has 16 CPU cores). Though, there are possible optimization to do in the future. For example, give a high priority to the VM Executor thread to prevent competing CPUs with the Prism Consensus module.

By profiling the CPU usage of a client in *EVM Prism Do Nothing* experiment, we find that transaction signature verification takes up to 39.2% of total CPU time (excluding mining), a relatively high percentage. When the experiment is running at a high throughput, the requirement of a large amount of signature verification is a major bottleneck; this emphasizes the importance of removing signature verification from the VM Executor module. In our design, we have moved signature verification into the Prism Consensus module, thus freeing the VM Executor from this heavy burden.

The VM Executor of EVM is implemented efficiently with abundant number of in-memory cache. However, it levies a heavy memory burden on the VM Executor; our port of EVM does not include the whole client-level cache management, as a result, the VM Executor does not free memory efficiently, and the memory usage increases along with the workload. This is one possible future optimization for our port of EVM.

Table III provides a breakdown statistics for three Prism block types in *EVM Prism Do Nothing* experiment. We can see that transaction blocks take up to 71.2% of total generated block data, other two blocks only 28.8%. This indicates that the majority of utilized bandwidth contributes to the high throughput (transaction blocks), whereas Prism overhead takes up only a small fraction (proposer and voter blocks). For other *EVM Prism* (and *MoveVM Prism*) experiments, this breakdown statistics will remain similar except for transaction blocks. The higher the throughput, the higher the transaction block data and percentage. Thus, we do not analyze the breakdown statistics for other experiments.

TABLE III: Statistics for three block types in *EVM Prism Do Nothing* experiment in 10 minute duration.

	# Mined Block	Block Data	Data Percentage
Proposer Block	44	4.0 MB	4.0%
Voter Block	47166	25.2 MB	24.8%
Transaction Block	107514	72.2 MB	71.2%

<sup>4</sup>In *Prism Consensus Only*, the consensus throughput is 418 Mbps and is converted to tx/s based on transaction size.



### C. Throughput and Latency of MoveVM

In this experiment, we measure the transaction throughput and confirmation latency for *MoveVM Prism*. We observe similar bottleneck and latency between this experiment and EVM experiment, whereas there are also discrepancies in terms of throughput.

**Throughput:** As shown in Table IV, the throughput of two basic applications is only 1.1K and 2.2K tx/s respectively; this is an order of magnitude smaller than that of *EVM Prism*. In private communication [34], core Libra developers have indicated to us that improving the performance of MoveVM is work under progress – when this improvement transpires, our Prism client can fully utilize that as well. Benchmark applications get 249 and 512 tx/s and are higher than those of *EVM Prism*, indicating that MoveVM is more efficient at executing instructions. The *CPU Heavy* throughput number, however, is still far below that of JVM (over 90K tx/s), so we believe MoveVM has the potential to execute instructions more efficiently as well.

TABLE IV: Throughput in terms of tx/s on MoveVM applications.

	Native Payment	Do Nothing	CPU Heavy	IO Heavy
MoveVM Executor Only	1441	2501	269	585
MoveVM Prism	1172	2243	249	512
Prism Consensus Only <sup>4</sup>	123222	158802	143140	142749

**Is Prism consensus the bottleneck?** For all MoveVM applications, *MoveVM Prism* reaches 81% of *MoveVM Executor Only* throughput. This phenomenon is similar to EVM and indicates that *MoveVM Prism* is able to reach the optimal MoveVM throughput. As for *Prism Consensus Only*, we can see the high throughput over 120K tx/s as well. Similar to the case of EVM, we conclude that Prism removes the consensus bottleneck for MoveVM, and the virtual machine itself is the bottleneck.

**Latency:** Prism guarantees a confirmation latency regardless of the throughput, and we do observe that in all Prism experiments including *MoveVM Prism*, the average confirmation latency is no more than 130 seconds.

**Resource utility:** *MoveVM Prism* maintains a good memory usage, which is kept under 3.2 GB in all experiments. The live usage of CPU never exceeds 32% per core on average; compared to *EVM Prism* experiment, this CPU utility reduction is due to smaller throughput and more efficient signature verification. MoveVM adopts Ed25519 signature [35] which is faster than ECDSA [36] adopted by EVM.

### D. Scalability

In this experiment, we evaluate Prism’s ability to scale with more network participants. We use a larger number, 300, EC2 instances and use the same propagation delay and rate limiter. We use a random 5-regular topology for 300 nodes, keeping diameter the same with that of 100 nodes. We also keep the same Prism parameter, including the overall mining rate, thus the individual mining rate is modified. By our design, only

the Prism Consensus module is related to scaling with more network participants, since only it communicates with peers. In addition, Prism Consensus module’s performance is not affected by which application it is running. Hence, it suffices to experiment with one VM and application to demonstrate Prism’s scalability and we use *EVM Prism* and *Native Payment* in the experiment.

The experiment for 300 nodes also runs for 10 minutes. However, it is hard to collect the fine-grained metrics for such a high number of nodes. So we calculate the overall metrics at the end of the experiment (all 10 minutes), in contrast to previous calculation (last 9 minutes).

Table V compares the performance between 100 and 300 nodes. The throughput and latency are very similar; the difference is due to the randomness of the experiments. The forking rate 0.113 in 300 nodes is a little larger than that in 100 nodes, and is mainly due to more hops and higher delay to propagate blocks throughout the peer-to-peer network, as we can see that the average path length is higher in the 300-node topology. This forking rate 0.113 is small enough to ensure the security of Prism consensus as well. In addition, the block propagation delay, as well as the forking rate, can be reduced by increasing the degree (the number of peers per node) of the peer-to-peer network; Geth [37] and Parity Ethereum [29] client have a default maximum degree of 50, which can sustain a low forking rate for peer-to-peer networks with a larger number of nodes.

TABLE V: Performance of *EVM Prism Native Payment*, with different network topologies.

#Node	Degree	Average Path Length	Diameter	Throughput	Confirmation Latency (s)	Forking
100	4	3.55	6	17268	96	0.102
300	5	3.84	6	17417	80	0.113

Resource utility on each node is also similar. In the 300-node experiment, the live usage of CPU never exceeds 50% per core on average. The heavy memory burden of the VM Executor module is also similar to that in 100-node experiment.

We conclude that Prism is able to scale to a large number of network participants, as long as the underlying peer-to-peer network provides a topology with reasonable block propagation delay. We can achieve similar throughput, latency, and security in those cases.

## VI. CONCLUSION

Blockchain research thus far has progressed in a compartmentalized manner: algorithms and protocols (many focused on consensus) are designed and studied separately from the upper layer wrappers (virtual machine, application programming) they will interact with. This is in contrast with Nakamoto’s Bitcoin design that was envisioned and designed as a complete system. This layering philosophy works well when the consensus layer is the bottleneck and much work can be expended to improve the performance (indeed, this is the case with many blockchains, including Ethereum). Prism is a recent consensus algorithm, closely inspired by Nakamoto’s longest chain

protocol, with *theoretically optimal* throughput and latency. In this paper we explore how Prism fits with two smart contract virtual machines, EVM and MoveVM, by implementing Prism underneath these virtual machines. We demonstrate that Prism seamlessly merges with both these VMs: our implementation approaches the *optimal* virtual machine throughput for a large variety of applications. This result means that Prism not only removes the consensus bottleneck of bare metal throughput and latency but also when interacting with two popular smart contract platforms. Further improvement of the smart contract performance would have to come from new designs of virtual machines and compilers and architectures capable of parallel execution of smart contracts. The early research in this area [13], [38] now takes on added urgency.

#### ACKNOWLEDGEMENT

We thank Lei Yang for his core role in the development of the Prism codebase, which we have integrated with EVM and MoveVM in this paper. We thank Sam Blackshear and Andrew Miller for helpful discussions. This research is supported in part by NSF under grants CCF-1705007, NeTS-1718270 and Army Research Office under grant W911NF-17-S-0002.

#### REFERENCES

- [1] V. Bagaria, S. Kannan, D. Tse, G. Fanti, and P. Viswanath, "Prism: Deconstructing the blockchain to approach physical limits," in *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '19, 2019, p. 585–602.
- [2] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," 2008.
- [3] *Optimism*. [Online]. Available: <https://optimism.io/>
- [4] *ZK-Rollups*. [Online]. Available: <https://docs.ethhub.io/ethereum-roadmap/layer-2-scaling/zk-rollups/>
- [5] H. Kalodner, S. Goldfeder, X. Chen, S. M. Weinberg, and E. W. Felten, "Arbitrum: Scalable, private smart contracts," in *27th {USENIX} Security Symposium ({USENIX} Security 18)*, 2018, pp. 1353–1370.
- [6] L. Yang, V. Bagaria, G. Wang, M. Alizadeh, G. Fanti, D. Tse, , and P. Viswanath, "Prism: Scaling bitcoin by 10,000 ×," *arXiv:1909.11261*, 2019.
- [7] Y. Sompolinsky, Y. Lewenberg, and A. Zohar, "Spectre: A fast and scalable cryptocurrency protocol." *IACR Cryptology ePrint Archive*, vol. 2016, p. 1159, 2016.
- [8] S. Blackshear, E. Cheng, D. L. Dill, V. Gao, B. Maurer, T. Nowacki, A. Pott, S. Qadeer, D. R. Rain, S. Sezer *et al.*, "Move: A language with programmable resources," 2019.
- [9] *My Sql*. [Online]. Available: <https://www.mysql.com/>
- [10] *PostGres*. [Online]. Available: <https://www.postgresql.org/>
- [11] T. Dickerson, P. Gazzillo, M. Herlihy, and E. Koskinen, "Adding concurrency to smart contracts," in *Proceedings of the ACM Symposium on Principles of Distributed Computing*, 2017, pp. 303–312.
- [12] P. S. Anjana, S. Kumari, S. Peri, S. Rathor, and A. Somani, "An efficient framework for optimistic concurrent execution of smart contracts," in *2019 27th Euromicro International Conference on Parallel, Distributed and Network-Based Processing (PDP)*. IEEE, 2019, pp. 83–92.
- [13] V. Saraph and M. Herlihy, "An empirical study of speculative concurrency in ethereum smart contracts," *arXiv preprint arXiv:1901.01376*, 2019.
- [14] S. Pang, X. Qi, Z. Zhang, C. Jin, and A. Zhou, "Concurrency protocol aiming at high performance of execution and replay for smart contracts," *arXiv preprint arXiv:1905.07169*, 2019.
- [15] M. Bartoletti, L. Galletta, and M. Murgia, "A true concurrent model of smart contracts executions," *arXiv preprint arXiv:1905.04366*, 2019.
- [16] *Libra*, accessed February 16, 2020. [Online]. Available: <https://github.com/libra/libra>
- [17] E. Androulaki, A. Barger, V. Bortnikov, C. Cachin, K. Christidis, A. De Caro, D. Eneyart, C. Ferris, G. Laventman, Y. Manevich *et al.*, "Hyperledger fabric: a distributed operating system for permissioned blockchains," in *Proceedings of the Thirteenth EuroSys Conference*, 2018, pp. 1–15.
- [18] M. Yin, D. Malkhi, M. K. Reiter, G. G. Gueta, and I. Abraham, "Hot-stuff: Bft consensus with linearity and responsiveness," in *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*, ser. PODC '19, 2019, p. 347–356.
- [19] H. Yu, I. Nikolic, R. Hou, and P. Saxena, "Ohie: Blockchain scaling made simple," in *2020 IEEE Symposium on Security and Privacy (SP)*. IEEE Computer Society, may 2020, pp. 112–127.
- [20] Y. Gilad, R. Hemo, S. Micali, G. Vlachos, and N. Zeldovich, "Algorand: Scaling byzantine agreements for cryptocurrencies," in *Proceedings of the 26th Symposium on Operating Systems Principles*, 2017, pp. 51–68.
- [21] I. Eyal, A. E. Gencer, E. G. Sirer, and R. Van Renesse, "Bitcoinng: A scalable blockchain protocol," in *13th {USENIX} symposium on networked systems design and implementation ({NSDI} 16)*, 2016, pp. 45–59.
- [22] J. Poon and V. Buterin, "Plasma: Scalable autonomous smart contracts," *White paper*, pp. 1–47, 2017.
- [23] *Plasma*. [Online]. Available: <https://ethresear.ch/t/plasma-world-map-the-hitchhiker-s-guide-to-the-plasma/4333>
- [24] *Ethereum 2.0 Specifications*. [Online]. Available: <https://github.com/ethereum/eth2.0-specs>
- [25] *NEAR Protocol — A sharded, developer-friendly, proof-of-stake public blockchain*. [Online]. Available: <https://nearprotocol.com/>
- [26] *Polkadot: Decentralized Web 3.0 Blockchain Interoperability Platform*. [Online]. Available: <https://polkadot.network/>
- [27] *RocksDB*. [Online]. Available: <https://rocksdb.org/>
- [28] *rust-rocksdb*, accessed February 19, 2020. [Online]. Available: <https://github.com/rust-rocksdb/rust-rocksdb>
- [29] *OpenEthereum*, accessed January 30, 2020. [Online]. Available: <https://github.com/openethereum/openethereum>
- [30] *ERC-20 Token Standard*. [Online]. Available: <https://github.com/ethereum/EIPs/blob/master/EIPS/eip-20.md>
- [31] *OpenZeppelin Contracts*. [Online]. Available: <https://github.com/OpenZeppelin/openzeppelin-contracts/blob/master/contracts/token/ERC20/ERC20.sol>
- [32] *CryptoKitties GeneScience*. [Online]. Available: <https://etherscan.io/address/0xf97e0a5b616dff913e72455fde9ea8bbe946a2b#code>
- [33] A. E. Gencer, S. Basu, I. Eyal, R. Van Renesse, and E. G. Sirer, "Decentralization in bitcoin and ethereum networks," in *International Conference on Financial Cryptography and Data Security*. Springer, 2018, pp. 439–457.
- [34] S. Blackshear, Private Communication, 2020.
- [35] D. J. Bernstein, N. Duif, T. Lange, P. Schwabe, and B.-Y. Yang, "High-speed high-security signatures," *Journal of cryptographic engineering*, vol. 2, no. 2, pp. 77–89, 2012.
- [36] D. Johnson, A. Menezes, and S. Vanstone, "The elliptic curve digital signature algorithm (ecdsa)," *International journal of information security*, vol. 1, no. 1, pp. 36–63, 2001.
- [37] *Go Ethereum*, accessed May 31, 2020. [Online]. Available: <https://github.com/ethereum/go-ethereum>
- [38] T. Dickerson, P. Gazzillo, M. Herlihy, and E. Koskinen, "Adding concurrency to smart contracts," *Distributed Computing*, pp. 1–17, 2019.