Practical Smart Contract Sharding with Ownership and Commutativity Analysis

George Pirlea*
National University of Singapore
Singapore
gpirlea@comp.nus.edu.sg

Amrit Kumar
Zilliqa Research
United Kingdom
amrit@zilliqa.com

Ilya Sergey
Yale-NUS College
National University of Singapore
Singapore
ilya.sergey@yale-nus.edu.sg

Abstract

Sharding is a popular way to achieve scalability in blockchain protocols, increasing their throughput by partitioning the set of transaction validators into a number of smaller committees, splitting the workload. Existing approaches for blockchain sharding, however, do not scale well when concurrent transactions alter the *same* replicated state component—a common scenario in Ethereum-style smart contracts.

We propose a novel approach for efficiently sharding such transactions. It is based on a folklore idea: state-manipulating atomic operations that *commute* can be processed in parallel, with their cumulative result defined deterministically, while executing non-commuting operations requires one to *own* the state they alter. We present CoSPLIT—a static program analysis tool that soundly infers *ownership and commutativity summaries* for smart contracts and translates those summaries to *sharding signatures* that are used by the blockchain protocol to maximise parallelism. Our evaluation shows that using CoSPLIT introduces negligible overhead to the transaction validation cost, while the inferred signatures allow the system to achieve a significant increase in transaction processing throughput for real-world smart contracts.

CCS Concepts: • Computing methodologies → Distributed programming languages.

Keywords: Smart Contracts, Static Analysis, Parallelism

ACM Reference Format:

George Pirlea, Amrit Kumar, and Ilya Sergey. 2021. Practical Smart Contract Sharding with Ownership and Commutativity Analysis. In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation (PLDI '21)*, June 20–25, 2021, Virtual, Canada. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3453483.3454112>

*Work partially conducted while employed at Zilliqa Research.

This work is licensed under a Creative Commons Attribution International 4.0 License.

PLDI '21, June 20–25, 2021, Virtual, Canada

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8391-2/21/06.

<https://doi.org/10.1145/3453483.3454112>

1 Introduction

The idea of Nakamoto consensus (*aka* blockchain) has been instrumental for enabling decentralised digital currencies, such as Bitcoin [48]. The applications of blockchains have further expanded with the wide-spread adoption of *smart contracts* [62]—self-enforcing, self-executing protocols governing an interaction between several mutually distrusting parties. The Ethereum blockchain has provided a versatile framework for defining smart contracts as blockchain-replicated stateful objects identified by their account numbers [65].

The open and decentralised nature of Nakamoto consensus comes at the price of throughput scalability. At a high level, in order for a sequence of transactions (so-called *block*) to be agreed upon system-wide, the system's participants (so-called *miners*) have to validate those transactions, with each miner executing them individually [4]. As a result, the throughput of blockchain systems such as Bitcoin and Ethereum does not improve, and even slightly deteriorates, as more participants join the system: Bitcoin currently processes up to 7 transactions per second, while Ethereum's throughput is around 18 transactions per second. Even worse, popular smart contracts may cause high congestion, forcing protocol participants to exclusively process transactions specific to those contracts. This phenomenon has been frequent in Ethereum: in the past, multiple ICOs (Initial Coin Offering, a form of a crowdfunding contract) and games, such as CryptoKitties, have rendered the system useless for any other purposes for noticeable periods of time [14].

Sharding in Blockchains. One of the most promising approaches to increase blockchain throughput is to split the set of miners into a number of smaller committees, so they can process incoming transactions in parallel, subsequently achieving a global agreement via an additional consensus mechanism—an idea known as *sharding*. Sharding transaction executions, as well as sharding the replicated state, has been an active research topic recently, both in industry [23, 30, 39, 51, 60, 64, 67] and academia [1, 17, 36, 45, 66].

Many of those works focus exclusively on sharding the simplest kind of transactions—*user-to-user transfers of digital funds*,—which are paramount in blockchain-based cryptocurrencies, while ignoring sharding of smart contracts [36, 45, 66, 67]. Existing proposals tackling smart contracts impose

heavy restrictions on contract-manipulating transactions, for instance, requiring the accounts of both the contract and its user to be assigned the same shard, or processing all such transactions in a specialised shard [17, 30, 39]. Other solutions assume a complex cross-shard communication protocol to reconcile possible conflicts [23, 30, 51, 60, 64], or adopt a contract design very different from Ethereum [1].

To the best of our knowledge, none of these approaches allows for parallel sharded executions involving the *same* smart contract. That is, none solve the mentioned congestion problem in Ethereum, caused by highly-popular contracts.

In this work, we describe a novel approach for significantly increasing the throughput of blockchains for smart contract-manipulating transactions. To achieve this, instead of treating contract implementations as “black boxes” (as do all the works mentioned above), we design a solution based on PL techniques, specifically, on static program analysis.

Our Approach. Why can user-to-user money transfers be sharded efficiently without complex inter-shard communication, and how can we generalise (perhaps, conservatively) this logic to shard arbitrary smart contracts?

Consider a transaction tx_1 that manifests a transfer of 10 units of some digital currency from the user A to B , and a transaction tx_2 that states that A transfers 20 units to C . In order to ensure that A does not double-spend, both tx_1 and tx_2 have to be executed in the same shard—the one that *owns* A ’s account and keeps track of A ’s balance. However, neither B nor C need to be owned by A ’s shard: as long as tx_1 and tx_2 are validated within A ’s shard, the positive deltas to B and C ’s accounts can be simply broadcast through the network, so their balances are increased accordingly with no extra inter-shard interaction.

Now consider a transaction tx_3 , in which D transfers 15 units to C . Notice that it does not matter in which order tx_1 and tx_3 are going to be processed, as they *commute*: either of their relative orderings will increase C ’s balance by 35.

The notions of state *ownership* and operation *commutativity* have been central in a number of works dedicated to reasoning about deterministic parallelism and proving correctness of concurrent programs [19, 20, 22, 34, 40, 44, 50]. In those works, the ownership discipline determines what parts of the shared state need to be manipulated sequentially by the *same* thread, while commutativity allows certain actions to be executed in *concurrent* threads in parallel, with a deterministic result. The virtues of commutative operations have also been studied in the systems community for scaling concurrent software [3, 13, 52, 55] and achieving faster consensus in replication protocols [11, 37, 42, 47]. However, to the best of our knowledge, no attempts to *automatically*

leverage commutativity in *user-defined replicated computations* (e.g., smart contracts) have been made to date.

In this work, we present CoSPLIT—a static analysis tool that soundly infers both *ownership* and *commutativity* information from source code of smart contracts and translates it to *sharding signatures*. The signatures are used, upon the deployment of a contract, to define a *sharding strategy* for the contract-manipulating transactions via the following rules:

- All transactions touching parts of a contract’s state *owned* by a shard S must be executed in this shard;
- Transactions executed in different shards are guaranteed to *commute*. Their cumulative result can be obtained by means of “joining” their respective contributions in a way prescribed by the sharding signature.

These two rules allow the system to enjoy a notion of consistency for parallel transaction executions adopted from works on the semantics of concurrent revisions [7, 8, 43]:

1. Potentially conflicting contract-manipulating transactions will be executed in some *globally-agreed order*.
2. Commuting transactions can be executed in parallel, as their effect *does not depend* on their order.

As we will discuss in Sec. 2, popular Ethereum-style contracts often allow for a “logical split” of their state into disjointly owned components, which is much more fine-grained than assigning an entire contract to a single shard. This split makes it possible to process transactions affecting those contracts in parallel in different shards, thus providing a practical solution to scale up the network throughput.

Our Contributions. The contributions of this work are:

- Identifying logical state ownership and operation commutativity as enabling mechanisms for sharding Ethereum-style contracts, and demonstrating adequacy of those notions for real-world Ethereum-style contracts (Sec. 2).
- A compositional static analysis that infers ownership and commutativity signatures for contracts written in Scilla [57] and translates them to shard allocation strategies (Sec. 3).
- An implementation of the analysis and of the algorithm for deriving sharding strategies in the tool called CoSPLIT and an end-to-end integration of CoSPLIT with a production-grade sharded blockchain protocol [45, 67] (Sec. 4). The archived software artefact is publicly available [53].
- Evaluation of parallelism enabled by CoSPLIT-inferred signatures, demonstrating a consistent increase in system throughput with increasing the number of shards (Sec. 5).

2 Motivation and Key Ideas

2.1 Contract Usage in Ethereum

To motivate the design of our approach for sharding, we first present the trends for smart contract usage in Ethereum. Since there are over 700 million Ethereum transactions to date, processing all the execution traces is too computationally expensive. Therefore, we selected a random sample of

Figure 1. Left: Ethereum transaction breakdown per type; the percentage distribution is averaged over 100K block periods. Right: breakdown of single-call transactions.

16,611 blocks, containing 1.1M transactions. The sample was collected in January 2020 and includes transactions up to block 9.25M, representing 0.17% of the total number of transactions up to that point and giving our measurements a 1% margin of error at a 99% confidence level.¹ As the left plot in Fig. 1 shows, ordinary user-to-user transfers are on a solid downward trend. Moreover, single-contract transactions take up to 55% of the recent blocks in our sample.

In this work we focus on sharding single-contract transactions. The right plot in Fig. 1 shows the dominance of a specific type of such transactions that represent token transfers in a special kind of contract—ERC20 token contracts [24]. ERC20 and other similar standardised contracts pose a big bottleneck to the network throughput: each of them requires *sequential processing* of all transactions that affect it.

2.2 Towards Sharding an ERC20 Contract

Fig. 2 shows a fragment of the implementation of an ERC20 token contract [24] in Ethereum’s high-level language Solidity [25]. The contract’s replicated state is represented by two mutable fields: the mapping `balances` that contains data about the amount of tokens owned by token holders; and the mapping `allowances` that captures the amounts of tokens authorised for third-party transfers by their holders. The state manipulations are done by transactions initiated by users (*aka senders*) calling one of the functions: `transfer` for transferring tokens, `approve` for granting the transfer rights for a certain amount of tokens to a third party, and `transferFrom` for transferring tokens on behalf of the user identified as sender. The subtractions in lines 16 and 21 will fail if the approved spender (*resp.* the sender) does not have enough allowance, thus preventing double-spends.

The design of the ERC20 contract provides ample opportunities for shard-based parallelism. Consider the left part of Fig. 3 that shows a fragment of the *mutable* ERC20 state: the `balances` field mapping account addresses A–E (top part of each box) to the respective token balances (bottom part of the box), and `allowances`, which is a mapping

```
1 contract ERC20 {
2   mapping (address => uint256) balances;
3   mapping (address => mapping (address => uint256)) allowances;
4   /* Main public functions */
5   function transfer(recipient, amount) {
6     _transfer(_msgSender(), recipient, amount);
7     return true;
8   }
9   function approve(spender, amount) {
10    _approve(_msgSender(), spender, amount);
11    return true;
12  }
13  function transferFrom(sender, recipient, amount) {
14    _transfer(sender, recipient, amount);
15    _approve(sender, _msgSender(),
16             allowances[sender][_msgSender()] - amount);
17    return true;
18  }
19  /* Auxiliary internal functions */
20  function _transfer(sender, recipient, amount) {
21    balances[sender] = balances[sender] - amount;
22    balances[recipient] = balances[recipient] + amount;
23  }
24  function _approve(owner, spender, amount) {
25    allowances[owner][spender] = amount;
26  } /* More functions */ }
```

Figure 2. A fragment of an ERC20 contract in Solidity.

Figure 3. Two ways to identify footprints of four ERC20 transactions: via full affected state (left) and via state affected by non-commuting operations (right).

from addresses (*e.g.*, A) to mappings of amounts allowed to transfer by third parties (*e.g.*, D and E) on their behalf. Now consider the following four single-contract transactions accessing that state concurrently by invoking functions from Fig. 2: $tx_1 = \text{transfer}_A(B, v_1)$; $tx_2 = \text{transfer}_C(A, v_2)$; $tx_3 = \text{transferFrom}_D(A, C, v_3)$; $tx_4 = \text{approve}_A(E, v_4)$. Here, a subscript denotes the transaction sender’s address (accessed via `_msgSender()` in Fig. 2), while v_i stand for various non-negative amounts, whose exact value is not important. All those transactions alter the contract state; the left part of Fig. 3 shows their corresponding *footprints*, *i.e.*, components of the state that they interact with. It is easy to see that the footprints of, *e.g.*, tx_1 and tx_4 are *disjoint*, thus, their effects on the contract’s state commute. Therefore, assuming the system provides an operation to *join* (*i.e.*, merge) updates

¹The Ethereum dataset and analysis are part of the archived artefact [53].

on the “logically disjoint” state components, it should be possible to execute, *e.g.*, tx_1 and tx_4 in different shards.

Sharding Strategy 1: Disjoint State Ownership. Let us formulate the constraints for parallel execution of the transactions tx_1 – tx_4 from Fig. 3, out of the knowledge that some of them commute, thanks to their footprint disjointness. We will denote by $\text{Owns}(\mathcal{S}, \{f_1, \dots, f_n\})$ an *ownership constraint*, meaning that the shard $\mathcal{S}$ logically owns the contract’s state components (fields or map entries) $f_1, \dots, f_n$ and, thus, *only this shard* may alter the values of those components by sequentially processing all the corresponding transactions.

Now consider two shards, $\mathcal{S}_1$ and $\mathcal{S}_2$ and the following set of ownership constraints, where bal and all denote the corresponding fields balance and allowances:

$$\left\{ \begin{array}{l} \text{Owns}(\mathcal{S}_1, \{\text{bal}[A], \text{bal}[B], \text{bal}[C], \text{all}[A][D]\}), \\ \text{Owns}(\mathcal{S}_2, \{\text{all}[A][E]\}) \end{array} \right\} \quad (1)$$

Clearly, $\mathcal{S}_1$ and $\mathcal{S}_2$ own disjoint portions of the contract’s state, thus, it will be safe to assign transactions to shards as $\mathcal{S}_1 \mapsto \{tx_1, tx_2, tx_3\}$ and $\mathcal{S}_2 \mapsto \{tx_4\}$, obtain the final result deterministically by merging their non-conflicting changes. This sharding strategy scales for more shards with designated ownership of the contract’s components and larger number of transactions with logically disjoint footprints.

Sharding Strategy 2: Commutativity of Addition. Even though transaction tx_2 modifies the component $\text{bal}[A]$, it does so in a commutative fashion and, thus, cannot affect the outcome of any other of the listed transactions (ditto for tx_1 and $\text{bal}[B]$). With this observation, we can refine the transaction footprints and the notion of ownership (Fig. 3, right), allowing for a parallel execution with three shards:

$$\left\{ \begin{array}{l} \text{Owns}(\mathcal{S}_1, \{\text{bal}[A], \text{all}[A][D]\}), \\ \text{Owns}(\mathcal{S}_2, \{\text{all}[A][E]\}), \text{Owns}(\mathcal{S}_3, \{\text{bal}[C]\}) \end{array} \right\} \quad (2)$$

In the constraints above, $\mathcal{S}_1$ no longer has to own $\text{bal}[B]$ or $\text{bal}[C]$, while shard $\mathcal{S}_3$ now needs to own $\text{bal}[C]$. In order to perform transactions allocated as $\mathcal{S}_1 \mapsto \{tx_1, tx_3\}$, $\mathcal{S}_2 \mapsto \{tx_4\}$, $\mathcal{S}_3 \mapsto \{tx_2\}$, obtaining the *same result* as in the previous case, we need to redefine the state *join* operation. Specifically, instead of *overwriting* the values in entries $\text{bal}[B]$ and $\text{bal}[A]$ upon “disjoint merging” as before, we will need to *add up the deltas* to those components resulting from token transfers in tx_1 and tx_3 , similarly to handling ordinary transfers (Sec. 1).

The Main Idea. To summarise these observations: contracts such as ERC20, whose operations only manipulate a *small part* of the state, allow for parallel conflict-free execution of their operations, if these operations commute. The *ownership constraints* state *which parts* of a contract’s state a shard must have exclusive access to in order to execute its operations without conflicts with other shards altering the same contract concurrently. The *join* defines the way to deterministically reconcile outcomes of the parallel executions.

2.3 Commutativity and State Ownership

It is common to reason about operation commutativity in terms of action traces [13]. That said, our way of thinking is inspired by the logical abstractions used for compositional verification of heap-manipulating programs [9, 10, 50].

In our setup, we are interested in parallelising executions of a family of single-contract transactions over a state-space Σ of a contract, collectively represented by a function $\mathcal{F}_x : \Sigma \rightarrow \Sigma$. Here, x denotes a vector of user inputs, *i.e.*, specifying which contract’s function to call, as well as its inputs. Two transactions identified by different user inputs x_1 and x_2 commute *iff* for any state σ , $\mathcal{F}_{x_1}(\mathcal{F}_{x_2}(\sigma)) = \mathcal{F}_{x_2}(\mathcal{F}_{x_1}(\sigma))$.

In order to enable parallelism, our goal is to identify a *commutative*, *associative*, and *partial* operation $\uplus : \Sigma \rightarrow \Sigma \rightarrow \Sigma$, such that for any σ_1, σ_2 and x , if $\mathcal{F}_x(\sigma_1)$ is defined (*i.e.*, σ_1 contains at least the *footprint* of $\mathcal{F}_x$) and $\sigma_1 \uplus \sigma_2$ is defined, then $\mathcal{F}_x(\sigma_1 \uplus \sigma_2) = \mathcal{F}_x(\sigma_1) \uplus \mathcal{F}_x(\sigma_2)$. This equality is referred to as *action locality* in the program logics literature [10], and, when it holds, it enables compositional program analyses [21] and concurrency specifications [44].² The virtue of $\mathcal{F}_x$ ’s locality for our purposes becomes apparent by observing the following chain of equalities for $\sigma = \sigma_1 \uplus \sigma_2$ when σ_1 and σ_2 are such that $\mathcal{F}_{x_1}(\sigma_1)$ and $\mathcal{F}_{x_2}(\sigma_2)$ are defined:

$$\begin{aligned} \mathcal{F}_{x_1}(\mathcal{F}_{x_2}(\sigma)) &= \mathcal{F}_{x_1}(\mathcal{F}_{x_2}(\sigma_1 \uplus \sigma_2)) = \mathcal{F}_{x_1}(\sigma_1 \uplus \mathcal{F}_{x_2}(\sigma_2)) \\ &= \mathcal{F}_{x_2}(\sigma_2) \uplus \mathcal{F}_{x_1}(\sigma_1) = \mathcal{F}_{x_2}(\sigma_2 \uplus \mathcal{F}_{x_1}(\sigma_1)) \\ &= \mathcal{F}_{x_2}(\mathcal{F}_{x_1}(\sigma_2 \uplus \sigma_1)) = \mathcal{F}_{x_2}(\mathcal{F}_{x_1}(\sigma)) \end{aligned}$$

This reasoning demonstrates the desired commutativity, and also provides a recipe for computing the final result in a *divide-and-conquer* fashion by taking it to be $\mathcal{F}_{x_1}(\sigma_1) \uplus \mathcal{F}_{x_2}(\sigma_2)$; the order does not matter, as $\uplus$ is commutative and associative. One can think of $\uplus$ as both the “logical split and join” operations, while a footprint of a transaction executing $\mathcal{F}_x$ is the minimal part σ' of the contract state, which *must be owned* by the shard executing it, so that $\mathcal{F}_x(\sigma')$ is defined.

Getting back to our motivating example of ERC20 sharding, **Strategy 1** corresponds to $\uplus$ taken as a disjoint union of the entry sets of the contract’s mapping fields (let’s call it *OwnOverwrite*). **Strategy 2** corresponds to $\uplus$ defined as a non-disjoint union with an implicit split of integer values in map entries—this way in the case of concurrent updates, the result can be obtained by summation of the per-shard portions of those values (we will call it *IntMerge*).

2.4 Pragmatic Considerations and Technical Setup

Ethereum is the most popular smart contract platform, and a number of other blockchain ecosystems also use the Ethereum Virtual Machine (EVM). While it would be desirable to implement our ideas directly in Ethereum, unfortunately, its infrastructure is currently unsuitable for our purposes:

²Readers familiar with state-of-the art program logics for concurrency can recognise that we are looking for a suitable Partial Commutative Monoid (PCM) [10, 34, 44], which would enable *framing* of contract operations $\mathcal{F}_x$.

(identifiers) $i, f, c ::= \text{alpha-numeric string}$
(types) $t ::= \text{int} \mid \text{string} \mid \text{unit} \mid \text{bool} \mid \text{map } t \ t \mid t \rightarrow t \mid \dots$
(patterns) $\text{pat} ::= _ \mid i \mid \text{constr } c \ \text{pat}$
(expressions) $e ::= \text{val } v \mid \text{var } i \mid \text{message } (\bar{i} \mapsto \bar{i}) \mid \text{constr } c \ \bar{i} \ \bar{i} \mid$
 $\text{builtin } \text{blt } \bar{i} \mid \text{let } i = e_1 \text{ in } e_2 \mid \text{fun } (i : t) \Rightarrow e \mid$
 $\text{app } i \ \bar{i}_j \mid \text{match } i \ \text{pat} \Rightarrow \bar{e} \mid \text{tfun } \alpha \Rightarrow e \mid \text{inst } i \ \bar{i}$
(statements) $s ::= i_1 \leftarrow i_2 \mid i_1 := i_2 \mid i = e \mid i_1 \ [\bar{i}_k] := i_2 \mid i_1 \leftarrow i_2 \ [\bar{i}_k] \mid$
 $i_1 \leftarrow \text{exists } i_2 \ [\bar{i}_k] \mid \text{delete } i_1 \ [\bar{i}_k] \mid i_1 \leftarrow \&i_2 \mid$
 $\text{match } i \ \text{pat} \Rightarrow \bar{s} \mid \text{accept} \mid \text{send } i \mid \text{event } i \mid \text{throw}$

Figure 4. Scilla syntax.

- **Protocol-level Support.** At the time of writing, the available prototype of Ethereum 2.0 [64] does not support cross-shard transactions or smart contracts.
- **Language-level Support.** EVM bytecode, Ethereum’s low-level language, is difficult to analyse soundly, due to the lack of modularity and structured control flow [65]. One could engineer an analysis inferring ownership constraints for Solidity [25], a high-level Ethereum language, by possibly decompiling EVM contracts [29]. However, a number of Solidity’s features (e.g., inter-contract calls) and unpredictable performance of EVM decompilers make it a challenging target for an efficient sound static analysis [33].

Bringing Ethereum infrastructure to the state necessary to deploy the described ideas would be an effort going well beyond the scope of this paper. Instead, we implemented our approach as a static analysis for Scilla [57]—a strongly-typed ML-style language for smart contracts. Scilla is supported natively (via a definitional interpreter) by an industry-scale blockchain [45, 67] that (a) provides rudimentary infrastructure for sharding, (b) is available open-source, and (c) is widely used and contains dozens of contracts implemented in Scilla by users and available for evaluating our approach.

3 CoSPLIT Analysis in a Nutshell

3.1 The Language

Scilla [57] is a minimalistic memory- and type-safe functional language, similar to OCaml and Haskell for an account-based (i.e., Ethereum-style) smart contract model. Scilla provides a very small set of state-manipulating primitives for altering contract state (i.e., reading from the blockchain state and changing the values of contract fields). Its *pure* (i.e., side effect-free) fragment corresponds to System F [27, 54] without recursion (but with bounded iteration). All of the standard library as well as user-defined contract-agnostic computations are implemented in Scilla as pure functions. This design choice removes the need for inter-contract calls for the sake of code reuse and makes contract analysis scale, as pure functions need to be analysed *only once*.

Contracts in Scilla are encoded as communicating state-transition systems in the style of IO-automata [46]. That is, all interaction between contracts is done by means of *message passing*. A contract’s state changes as a result of executing

```

1 transition Transfer(to: ByStr20, amount: UInt)
2   get_from_bal ← balances[_sender];
3   match get_from_bal with
4   | Some bal ⇒
5     can_do = uint_le amount bal;
6     match can_do with
7     | True ⇒
8       new_from_bal = builtin sub bal amount;
9       balances[_sender] := new_from_bal;
10      get_to_bal ← balances[to];
11      new_to_bal = (match get_to_bal with
12      | Some bal ⇒ builtin add bal amount
13      | None ⇒ amount);
14      balances[to] := new_to_bal
15      | False ⇒ (* Report a transaction failure *)
16      msg_to_recipient = {_tag : "Accepted"; _to : to; ...};
17      msg_to_sender = {_tag : "Success"; _to : _sender; ...};
18      send msg_to_recipient; send msg_to_sender

```

Figure 5. FungibleToken Transfer transition in Scilla.

its *transitions* as reactions to received messages from the users or other contracts. While transitions in Scilla contracts are similar to functions in Solidity, they provide stronger encapsulation and atomicity guarantees, in particular, disallowing reentrancy. This model allows one to analyse each contract’s transitions *in isolation* from any other contract’s code, thus allowing for deriving their signatures *statically* without over-approximating the effects of the external calls.

Fig. 4 describes the syntax for Scilla statements and expressions, and Fig. 5 shows the code for the Transfer transition of a FungibleToken (i.e., ERC20) contract. The transition takes one implicit argument, the address `_sender` of the transaction originator, and two explicit arguments, the address of the recipient and the amount of tokens to be transferred. It reads the balance of the sender (line 2), ensures that the sender has enough tokens to transfer (line 6), and updates the sender (line 9) and recipient balances (line 14).

3.2 Inferring Transition Summaries

CoSPLIT implements a compositional inter-procedural *effect analysis* [49] inferring *transition summaries* (a set of *effects*), which describe, in a symbolic form, (1) the *state footprint* of transitions and (2) the *contributions* of the initial values of contract fields, read during a transition’s execution, to the final field values. A transition’s summary is the result of the abstract interpretation [15, 16] of the transition’s code. These summaries help answer queries such as (1) “Does the transition write to field *f*?” and (2) “Can the transition’s effect to field *f* be represented as an addition of a constant to *f*’s old value?”. The two types of queries correspond to our two sharding strategies (see the discussion in Sec. 2.2).

The *state footprint* describes the portion of a contract state affected by a transition (and, respectively, by a transaction that executes it), whereas the *contribution types* let us decide if writes that affect the same piece of state commute. We explain the analysis in two stages, corresponding to the two types of queries. The stages are intertwined in the implementation, but here we separate them to aid in presentation.

(effect)	$\varepsilon ::= \text{Read}(f) \mid \text{Write}(f, \tau) \mid \text{Condition}(\tau) \mid \text{AcceptFunds} \mid \text{SendMsg}(\tau) \mid \top$
(contrib. type)	$\tau ::= \langle \text{cs} \mapsto (\text{card}, \text{ops}), p \rangle \mid \text{EFun } i \tau \mid \top \mid \perp$
(contrib. src.)	$\text{cs} ::= \text{Field } f \mid \text{Const } c \mid \text{Formal } i \mid \top$
(cardinality)	$\text{card} ::= 0 \mid 1 \mid \omega$
(precision)	$p ::= \text{Exact} \mid \text{Inexact}$
(operation)	$\text{op} ::= \text{Builtin } \text{blt} \mid \text{Cond}$
(cardinality order)	$0 \sqsubseteq 1 \sqsubseteq \omega$
(operations order)	$\text{ops}_1 \sqsubseteq \text{ops}_2 \quad \text{iff } \text{ops}_1 \subset \text{ops}_2$
(precision order)	$\text{Exact} \sqsubseteq \text{Inexact}$
	$0 \oplus \alpha = \alpha \quad 0 \sqcup \alpha = \alpha \quad 0 \otimes \alpha = 0$ $1 \oplus 1 = \omega \quad 1 \sqcup 1 = 1 \quad 1 \otimes 1 = 1$ $\alpha \oplus \omega = \omega \quad \alpha \sqcup \omega = \omega \quad \alpha \otimes \omega = \omega$

Figure 6. Components of CoSPLIT abstract domain.

3.3 State Footprints

Fig. 6 shows the CoSPLIT abstract domain. The first stage of the analysis computes an over-approximation of the state footprints of contract transitions, expressed as a set of *effects* (denoted ε). Effects describe how the transition interacts with the blockchain state. For instance, the `AcceptFunds` effect (contributed by the **accept** statement) changes *both* the contract's and the sender's native token balances.³ Similarly, the `SendMsg` effect (contributed by **send**) might invoke transitions of other contracts or send native tokens. Finally, the `Read` and `Write` effects describe which portions of the contract's own state may be accessed by the transition. For each transition, CoSPLIT iterates over every statement in the transition's code and determines the static over-approximation of that statement's effect. In some cases, this over-approximation is the uninformative effect $\top$. Due to Scilla's design, the translation between statements and effects is direct. As an example, we show the analysis rules for map reads and writes, which can be found in the top box of Fig. 7. These rules are applied when analysing lines 2, 9, and 14 in Fig. 5. The parts shown in grey boxes, including the rule for the non-effectful `BIND` statement, appear due to the second stage of the analysis, explained below.

The `MAPGET` and `MAPUPDATE` rules extend the transition's summary Σ with the appropriate `Read` or `Write` effect, which identifies the portion f (for field) of the contract state that is operated on. For map accesses, f includes the name of the map and the symbolic names of keys $\overline{i_k}$ used to index into the map. For accesses to non-map contract fields, only the name of the field is included.

Whereas contract fields can always be described, the keys used to index into a map can be the result of a computation

³This can be expressed as a `Write` affecting the two balances, but we prefer a more informative effect since balances are treated specially by the blockchain protocol. Similar reasoning applies to `SendMsg`.

BIND $\frac{\Gamma \vdash e : \tau \quad \Gamma' = \Gamma, i : \tau}{\Gamma; \Sigma \vdash i = e \rightsquigarrow \Gamma'; \Sigma}$		MAPUPDATE $\frac{\Gamma \vdash i_2 : \tau \quad b = \text{CanSummarise } \overline{i_k} \quad \varepsilon = \text{if } b \text{ then } \text{Write}(i_1 \overline{i_k}, \tau) \text{ else } \top}{\Gamma; \Sigma \vdash i_1 \overline{i_k} := i_2 \rightsquigarrow \Gamma; \Sigma, \varepsilon}$	
MAPGET $b = (\text{Write}(i_2 \overline{i_k}, _) \notin \Sigma) \wedge (\text{CanSummarise } \overline{i_k})$ $\frac{\tau, \varepsilon = \text{if } b \text{ then } \langle \text{Field } i_2 \overline{i_k} \mapsto (1, \emptyset), \text{Exact} \rangle, \text{Read}(i_2 \overline{i_k}) \text{ else } \top, \top}{\Gamma; \Sigma \vdash i_1 \leftarrow i_2 \overline{i_k} \rightsquigarrow \Gamma, i_1 : \tau; \Sigma, \varepsilon}$		LITERAL $\Gamma \vdash \text{val } l : \langle \text{Const } l \mapsto (1, \emptyset), \text{Exact} \rangle$	
BUILTIN $\frac{\tau' = \oplus \Gamma(\text{arg}) \quad \tau = \tau' \text{ with ops } \neq \text{blt}}{\Gamma \vdash \text{builtin } \text{blt } \text{arg} : \tau}$		CONSTR $\frac{\tau = \oplus \Gamma(i)}{\Gamma \vdash \text{constr } c \overline{i} : \tau}$	
FUN $i = \text{unique arg. identifier}$ $\frac{\Gamma' = \Gamma, i : \langle \text{Formal } i \mapsto (1, \emptyset), \text{Exact} \rangle \quad \Gamma' \vdash e : \tau_e \quad \tau = \text{EFun } i \tau_e}{\Gamma \vdash \text{fun } (i : t) \Rightarrow e : \tau}$		APP $i_j = \text{unique identifier for } j\text{-th arg.}$ $\frac{\Gamma \vdash f : \tau_f \quad \forall j, \Gamma \vdash a_j : \tau_j \quad \tau = \oplus (\tau_f[i_j] \otimes \tau_j)}{\Gamma \vdash \text{app } f \overline{a_j} : \tau}$	
LET $\frac{\Gamma' = \Gamma, i : \tau_i \quad \Gamma \vdash e_1 : \tau_i \quad \Gamma' \vdash e_2 : \tau}{\Gamma \vdash \text{let } i = e_1 \text{ in } e_2 : \tau}$		MATCH $\frac{\Gamma \vdash x : \tau_x \quad \forall i, \Gamma_i = \text{binder}(\text{pat}_i) \rightarrow \Gamma(x) \quad \forall i, \Gamma \cup \Gamma_i \vdash e_i : \tau_i \quad \tau = \text{MatchC}(x, \tau_x, \text{pat}_i, e_i, \overline{\tau_i})}{\Gamma \vdash \text{match } x \overline{\text{pat}_i} \Rightarrow \overline{e_i} : \tau}$	

Figure 7. Selected rules for statements and expressions.

and may even depend on contract state. As such, we only assign an informative effect ε to accesses when the keys used to index into the map are *transition parameters*, as they are in Fig. 5. Moreover, for nested maps, we require that the access is bottom-level, *i.e.*, it touches a primitive value rather than a map. These constraints are captured by `CanSummarise`. If the access cannot be summarised, the $\top$ effect is given. Generally, an access can be statically described whenever the keys are not dependent on the contract state, but we limit ourselves to keys that are transition parameters to simplify transaction dispatch, which is described in Sec. 4.3.

3.4 Contribution Types

The summaries produced so far are sufficient to enable the disjoint state ownership sharding strategy. But, as we have seen, we can execute in parallel even transactions with *effects over the same state*, as long as the effects commute. The second stage of the analysis annotates the effects produced by the first stage with contribution types (denoted τ), which help determining whether they commute. Concretely, for every expression e , CoSPLIT computes a *contribution type*, an over-approximation of the set of arithmetic operations and of the set of contract state components from the beginning of the transition execution that contribute to e 's result.

The type τ (cf. Fig. 6) ascribed to an expression e indicates, which *contribution sources*, *i.e.*, parts of the contract state

```

Read(balances[_sender])          Read(balances[to])
Condition(balances[_sender], amount)
Write(balances[_sender], ⟨amount&balances[_sender], 1, sub⟩)
Write(balances[to], ⟨amount&balances[to], 1, add⟩)
SendMsg(funds = zero; destination = to)
SendMsg(funds = zero; destination = _sender)

```

Figure 8. Set of effects of the Transfer transition.

(**Field** f), transition parameters/constants (**Const** c), and function parameters (**Formal** i), flow into e 's result, what operations are applied to those sources, and *how many times* each source contributes to e 's result. The precision component p in types records whether over-approximation of the set of operations has taken place due to joining control flows, *i.e.*, the analysis has lost precision. This lets us answer questions like “Can the transition’s effect to field f be represented as an *addition of a constant* to f 's old value?”. If the τ ascribed to the value written to field f has as its only *exact* contribution **Field** $f \mapsto (1, \text{Builtin add})$, and also some constants and transition parameters, the answer is *yes*.

The Importance of Cardinalities. The most important component in the type is the *cardinality* of **Field** f . If f did not show up in τ , then the written value would be a constant, and different writes would not commute, as they might have potentially different values (*e.g.*, different transition parameters). Conversely, if f 's contribution was non-linear, *i.e.*, its cardinality is ω (“many”),⁴ then even though f is modified by a commutative operation (addition), the effect would not commute. For example, the linear function $f(x) = x + 1$ commutes with $g(x) = x + 2$, but does not with $h(x) = x + x + 1$, as $f(h(a)) \neq h(f(a))$. The *linearity* (“used-once”) information attached to contribution sources lets us ensure that operations are used in ways that guarantee commutative effects. We lift the $\oplus$ operation on cardinalities to types by adding the cardinalities of matching sources and set-unioning their operations, ascribing to the result the $\sqcup$ of their precisions. For $\otimes$, which is defined only between a type and a single contribution, we multiply the cardinalities of the arguments and modify the other components analogously to the $\oplus$ lifting.

Computing Contribution Types. The values read from the mutable contract state, literals, transition and contract parameters are all contribution sources. For example, the LITERAL and MAPGET rules in Fig. 7 show how new contribution sources are introduced. A read into a binder i_1 from a location that was not overwritten, extends the typing context Γ by giving i_1 the linear contribution type τ shown in the rule. The type shows that the value of i_1 is the value of the respective “pseudo-field” $i_2 \overline{[i_k]}$ (*i.e.*, a map entry) at the beginning of the transition execution. Contributions from multiple sources are combined with their cardinalities added up *point-wise* via $\oplus$ operator (*cf.* Fig. 6). For example, the

⁴Our cardinality domain is inspired by GHC’s cardinality analysis [58].

```

(constraint)  $oc ::= \text{Owns}(f) \mid \text{UserAddr}(x) \mid \text{NoAliases}(\langle x, y \rangle) \mid$ 
SenderShard  $\mid$  ContractShard  $\mid \perp$ 
(join)  $\psi_f ::= \text{OwnOverwrite} \mid \text{IntMerge}$ 

```

Effect	Constraint
$\top \vee \text{SendMsg}(\top)$	$\perp$
AcceptFunds	SenderShard
$\text{SendMsg}(_amount.\tau \neq 0)$	ContractShard
$\text{SendMsg}(_recipient.\tau = r)$	UserAddr(r)
$\text{Read/Write}(m \ [x]), \text{Read/Write}(m \ [y])$	NoAliases($\langle x, y \rangle$)

Figure 9. Elements of sharding signatures (top). Translation between effects and non-ownership constraints (bottom).

application of a BUILTIN function adds up the contributions from its arguments and records an application of **Builtin** blt to each contribution source, recording this in the type.

Functions are given *arrow types*, **EFun** $i \tau$, with each formal parameter constituting an artificial contribution source **Formal** i . When the function is applied (rule APP), the arrow type is evaluated by substituting the formal parameter’s contribution with the actual argument’s contributions, multiplying the cardinalities. Our approach supports up to second-order Scilla functions by expressing operations on contribution types ($\oplus$, $\sqcup$, and $\otimes$) at the type level, and deferring normalisation until the arguments are known. For simplicity, we do not show this in the figures. Even though Scilla provides polymorphic functions and type applications (as in System F), they bear no significance for our analysis.

Finally, the types given to MATCH over-approximate the contributions of all clauses, whose types are combined using the MatchC operator defined as follows:

$$\begin{aligned}
\text{MatchC}(x, \tau_x, \overline{pat_i}, \overline{e_i}, \overline{\tau_i}) &\triangleq \tau_{cond} \oplus \sqcup \overline{\tau_i} \\
&\text{where} \\
\tau_{cond} &\triangleq \text{if IsKnownOp}(x, \overline{pat_i}, \overline{e_i}) \\
&\quad \text{then } \perp \text{ else } \text{AdaptC } \tau_x \\
\text{AdaptC } \langle cs \mapsto (_, _) \rangle &\triangleq \text{if SameVars}(\overline{\tau_i}) \\
&\quad \text{then } \langle cs \mapsto (0, \text{Cond}) \rangle, \text{Exact} \\
&\quad \text{else } \langle cs \mapsto (0, \text{Cond}) \rangle, \text{Inexact}
\end{aligned}$$

The additional contribution τ_{cond} accounts for whether **matching** over the scrutinee x induces non-trivial data flow (in which case its contribution is determined via AdaptC), or simply “peels off” a constructor of an option value (in which case it has no contribution). The latter is a very common special scenario, (see, *e.g.*, lines 11–13 of Fig. 5), and without this machinery the analysis would lose too much precision.

3.5 Calculating Sharding Signatures

From the set of transition summaries of a given contract (one such summary is shown in Fig. 8), and provided with user input as to what transitions to attempt to shard and which fields can be treated *weakly* for reading (*cf.* Sec. 4.2.3), CoSPLIT derives a sharding signature, consisting of a set of constraints oc for each transition in the contract and a join

Algorithm 3.1: Derive Sharding Signature

```

input : effect summaries, selected transitions, weak reads
output: transition (ownership) constraints, field join operations

 $\bar{\Sigma} \leftarrow$  effect summaries of selected transitions
 $wrs \leftarrow$  reads the user accepted might be stale
for  $\Sigma \in \bar{\Sigma}$  do
   $cfs \leftarrow$  GetConstantFields( $\Sigma$ )
   $\Sigma \leftarrow \forall f \in cfs, \Sigma.remove(Read(f))$ 
   $\Sigma \leftarrow \Sigma.MarkConstantsInTypes(cfs)$ 
 $lcws \leftarrow$  GetTransitionCommWrites( $\bar{\Sigma}$ )
 $cws, joins \leftarrow$  TryConsolidateJoinsGlobally( $lcws$ )
 $\bar{\Sigma} \leftarrow \bar{\Sigma}.RemoveSpuriousReads(cws)$ 
if  $joins \neq \emptyset \wedge wrs = StaleReads(\bar{\Sigma}, joins)$  then
   $oc \leftarrow \{\}$ 
  for  $\Sigma \in \bar{\Sigma}$  do
     $c \leftarrow$  GenEnvironmentConstraints( $\Sigma$ )
    foreach  $Read(f) \in \Sigma$  do  $c \leftarrow c \cup Owns(f)$ 
    foreach  $Write(f) \in (\Sigma \setminus cws)$  do  $c \leftarrow c \cup Owns(f)$ 
     $oc \leftarrow oc \cup c$ 
  return ( $oc, joins$ )

```

operation $\wp_f$ for each field. The top part of Fig. 9 enumerates the constraints that summaries can impose, as well as join operations we currently support. Constraints are static symbolic representations of conditions that *must be satisfied at runtime*. They refer to mutable fields or transition parameters as symbolic values, e.g., $Owns(f)$ and $UserAddr(x)$. For instance, the $Owns(balances[_sender])$ constraint denotes that a shard executing the transition must own the $_sender$ portion of the $balances$ state component, where $_sender$ is replaced at runtime by the actual value given by the transaction. The other constraints are imposed by the blockchain environment, e.g., $SenderShard$, which must be satisfied if the contract accepts funds, or arise as preconditions for the soundness of our analysis, e.g., that keys used for map accesses do not alias. $\perp$ corresponds to an unsatisfiable constraint, meaning that the transition cannot be executed in parallel with other transactions over the same contract.

Algorithm 3.1 shows the procedure for deriving the sharding signature. The contract developer selects a set of transitions that should be executed in parallel, and the algorithm inspects their summaries to determine what constraints must be satisfied to enable parallel execution. First, it identifies which (if any) contract fields are not written to in the selected transitions, and marks their reads as non-effectful (i.e., it removes them from the summary) and their contributions as constant. Second, looking at each summary in isolation, it identifies using the types, which writes have a commutative effect, e.g. $Write(balances[to])$ in Fig. 8. Then, it determines if a join operation exists for every field, i.e., the writes across transitions are compatible, and marks reads that only flow into commutative writes, e.g. $Read(balances[to])$ (since $balances[to]$ does not appear in any other type), as non-effectful.⁵ Finally, if the developer accepts that reads from

fields that are commutatively written to might return *stale data* (cf. Sec. 4.2.3), the algorithm translates effects into constraints via the mapping in Fig. 9 and by requiring ownership of every field that is read or non-commutatively written to.

4 Enabling Parallelism with CoSPLIT

In this section, we show how the signatures inferred by CoSPLIT, as described in Sec. 3, can be used to allow for parallel transaction execution in a sharded blockchain.

4.1 The Sharding Model

We integrated CoSPLIT with Zilliqa blockchain [67]. Zilliqa is one of the first sharded chains in production. It implements the Elastico protocol for secure sharding [45] and relies on an optimised version of the Practical Byzantine Fault Tolerance (PBFT) protocol for consensus in the network [12, 61]. At the time of writing, Zilliqa mainnet has processed 9.6 million transactions and contained 28 types of *unique* smart contracts (some of them have many deployed copies). Below, we outline the relevant parts of its architecture and transaction processing logic, referring the reader to the corresponding manuscripts for details (on, e.g., security and epoch-based mining) [45, 67], which are not critical for our presentation.

Network Architecture. The Zilliqa network consists of three main components: the *lookup nodes*, the *shards*, and the *Directory Service* (aka the DS committee) (cf. Fig. 10).

Lookup nodes are the entry-point to the network. Any transaction created by a user has to be sent to the lookup nodes, which thereupon group several transactions together in a packet and dispatch them to one of the shards or the DS committee for processing. Each shard (and similarly the DS committee, which in fact is a special shard) stores the full blockchain state and runs PBFT to reach consensus on validated transactions. It then proposes a MicroBlock (MB) that contains information on the transactions that it has processed. MicroBlocks are then sent to the DS committee together with StateDelta (SD) which encodes the changes in the state of the accounts that were touched by the transactions within a MicroBlock. Once all the MicroBlocks and the corresponding StateDeltas reach the DS committee, the latter combines them all in the form of a FinalBlock (FB) and a FinalStateDelta (FSD). The FinalBlock and FinalBlock-Delta are then sent back to each shard so that all the shards have the same view of the full global state—Zilliqa shards transaction execution, but not state storage.

The Default Sharding Strategy. A client-issued transaction can be processed either by one of the shards or by the DS committee (Fig. 10). Zilliqa employs a simple deterministic transaction assignment strategy to shards to ensure that double spends are detected within a shard without complex cross-shard communication [39]. User-to-user payment transactions are deterministically assigned to shards based on the sender’s address. That is, all transactions from the

⁵The Condition effect prevents removal of reads that affect control-flow.

Figure 10. Transaction processing in a sharded architecture.

same user get handled in the same shard, so any double spend from a specific user can be detected within a single shard in the same way it gets handled in a non-sharded architecture.

For smart contracts, Zilliqa implements a non-efficient *conservative* strategy. Specifically, the network statically assigns both contracts and end users to shards. Transactions to a contract invoked by users residing in the same shard as the contract are handled within the shard, while transactions to a contract invoked by users from an outside shard are handled in the DS committee. To ensure that shards and the DS do not end up manipulating the state of the same contract concurrently, the protocol requires the DS committee to process transactions assigned to it *only after* the shards have finished processing their transactions.

Given this simple deterministic assignment, the parallelism achieved for smart contract transaction processing is quite limited. In fact, the more shards there are, the more transactions will need to be processed by the DS committee.

4.2 Revising the Account-Based Blockchain Model

In order to employ the described sharding model for COSPLIT-enabled parallelism, we need to revise a few core aspects in the design of Ethereum-style blockchains.

4.2.1 Relaxing the Nonce Mechanism. Ethereum’s account based model (adopted by Zilliqa and similar systems) uses of the *nonce* mechanism for defining a total order on all transactions emitted by a particular user. Nonces are calculated by counting the number of transactions sent from a user address and are digitally signed, addressing the following design aspects: (a) strict, gap-free, user-defined ordering of transactions, and (b) prevention of replay attacks. Thanks to the nonces, the user can send many transactions with subsequent numbers, and they are going to be processed in *this exact order*—the protocol will not process transaction with a nonce $n + 1$ before the one with nonce n .

Because of aspect (a), nonces pose a bottleneck to sharded executions. While in plain Zilliqa all transactions from a single user are guaranteed to be handled in the same shard, the

nonce mechanism prevents parallel executing of transactions with the same origin in different shards, as the total order of nonces cannot be communicated. We notice that in practice no applications rely on a specific order of user transactions *before they are committed* by the protocol.⁶ Therefore, it suffices for transactions to be processed in an increasing nonce order, without waiting for all “gaps” to be filled, treating them similarly to ballots in Paxos [41]. This relaxation requires a very small change in the protocol logic. With it, we kept the aspect (b) of the nonce mechanism, while allowing for parallel executions. For instance, this way we can execute in parallel two disjoint sets of *commuting* transactions from the same user with nonces $\{1, 3, 5\}$ and $\{2, 4\}$, respectively.

4.2.2 Parallel Gas Accounting. *Gas accounting* is a mechanism to charge users for executing transactions [65]. Such deductions must be treated sequentially to avoid overspending. We circumvent this bottleneck to parallelism by splitting a user’s balance across shards (with a larger fraction given to the shard handling money transfers from that user), so gas costs can be charged without coordinating balance changes.

4.2.3 Weak Reads. In the Transfer transition (Fig. 8), we saw that the write into `balances[to]` has a commutative effect. As a result, the processing shard does not need to own the field to execute the transition. However, allowing commutative writes means that transitions executing in a different shard might read *stale* values of `balances[_sender]`. In our example, this is fine—the sender may have more tokens than she thinks. Yet, in general, introducing commutative writes weakens the semantics of reads and a static analysis cannot determine whether this is “fine” for specific contracts. As a rule of thumb, a read can safely be marked as *weak* if the contract semantics is “monotone” in the corresponding value with respect to some lattice (as in [40]), *i.e.*, the behaviour of the contract is not affected if a higher value is read and other shards can increase the value, but not decrease it. Ideally, the programming language itself would allow contract developers to mark certain reads as *weak*, but neither Scilla nor Solidity, both designed with sequential semantics in mind, currently have this feature. For now, we require that weak reads be provided as input to [Algorithm 3.1](#).

4.3 CoSPLIT in Action

Using CoSPLIT assumes two modes: offline and online. In the former, the user who is about to deploy her contract to the blockchain, provides hints to the tool in order to choose the most suitable sharding signature (Fig. 11). In the latter, CoSPLIT is run automatically by the miners as a part of the validation pipeline for contracts proposed to be deployed.

⁶Furthermore, the UTxO blockchain model adopted by, e.g., Bitcoin [48] promotes this kind of weak notion of consistency, in which the user *cannot* predict the order in which her transactions are committed.

Figure 11. CoSPLIT for a contract developer (offline mode).

Offline Mode. Prior to the deployment of her contract C (written in Scilla), the developer runs the CoSPLIT analysis on C 's code, producing a *transition summary*. Such a summary (cf. Sec. 3.2) itself is not yet useful for sharding. The developer next sends a query (a (sub)set $\{\mathcal{T}_1, \mathcal{T}_2, \dots\}$ of C 's transitions and a subset $\{f_1, f_2, \dots\}$ of C 's fields that can be treated *weakly* for reading) to the sharding solver, obtaining as a result a sharding signature: a set of ownership constraints $\overline{oc}$ and a dictionary of *per-field* join functions $\overline{\Psi}_f$ suitable for sharding those transitions. CoSPLIT returns a pair $(\overline{oc}, \overline{\Psi}_f)$ of encodings, which the developer submits along with C 's code in a contract-deploying transaction.

Validating Sharding Signatures. Each node in the shard, upon receiving a transaction with a contract C and its signature $(\overline{oc}, \overline{\Psi}_f)$, runs CoSPLIT ensuring that the signature is valid (the set of sharding transitions can be obtained from the constraints $\overline{oc}$). The signature is then broadcast, along with the contract code and metadata, to all nodes in the system.

The signatures of different contracts are independent, so the deployment of a new contract does not affect the validity of previously computed sharding signatures.

Assigning Transactions to Shards. A lookup node, upon receiving a transaction $tx = \langle C, \mathcal{T}, x \rangle$ invoking a transition $\mathcal{T}$ of a contract C with inputs x , first conservatively checks whether the transaction is indeed single-contract (i.e., no more than one transition is invoked). If it cannot validate it as such, the transaction is routed to the DS committee. Otherwise, the lookup node exercises the *sharding strategy* induced by C 's signature $(\overline{oc}, \overline{\Psi}_f)$, if such signature is available.

To do so, the lookup node invokes the $dispatch_{\overline{oc}}(\mathcal{T}, x)$ procedure, which identifies a shard S that satisfies $\mathcal{T}$'s constraints, and routes tx to that shard. If no satisfying shard exists, the transaction is routed to the DS committee.

Executing Sharded Transactions. Transactions are applied to a contract in *epochs*, so all nodes in all shards share *the same* contract state at each epoch's beginning and end. Upon receiving a transaction that invokes a transition $\mathcal{T}$ of C with inputs x , a node in a shard S executes it by applying the transition's logic $\mathcal{F}_{\mathcal{T},x}$ to the contract state σ at hand.

To understand why this is correct and leads to deterministic parallelism, let us think of a “logical split” of the state σ using the per-field join operations $\Psi = \overline{\Psi}_f$ from the contract's sharding signature. Here, we consider the contract state σ as a join of the two states σ_{self}^S and σ_{other}^S [44], where σ_{self}^S is

owned by the node's shard S , while σ_{other}^S is not owned, and, hence will not be affected by executing $\mathcal{F}_{\mathcal{T},x}$. Therefore, by running the contract's transition, i.e., executing $\mathcal{F}_{\mathcal{T},x}(\sigma)$, the node obtains the state $\sigma'^S = \sigma_{\text{self}}'^S \Psi \sigma_{\text{other}}^S$, where only the owned part has been changed from σ_{self}^S to $\sigma_{\text{self}}'^S$. Due to the *locality* of $\mathcal{F}_{\mathcal{T},x}$ with respect to the state owned by the node's shard (which is ensured by dispatching the transaction via the ownership constraints), all changes done by the nodes of the shard S to the state of the contract can be accumulated into the shard-specific state delta σ_{Δ}^S , which only includes affected parts of the contract's state that are “logically disjoint” from any changes other shards have made concurrently.

The final state delta σ_{Δ}^S is sent to the DS committee that combines the initial contract state σ with all the state deltas by executing a *three-way merge* [35, 43] using the per-field operations $\overline{\Psi}_f$ from the signature. This deterministic merge *always succeeds*, thanks to logical disjointness of the state deltas from different shards. Finally, the DS committee executes all the remaining potentially conflicting transactions, obtaining the final contract state, which is broadcast system-wide to all other shards at the end of the epoch.

5 Evaluation

We implemented CoSPLIT in OCaml as a pluggable checker, adding an optional phase to the existing Scilla type checking pipeline. Put together, the analysis, query solver, transaction dispatcher, and state delta merger measure 2900 lines of OCaml code. All interaction between CoSPLIT and the nodes of Zilliqa network happens via JSON-RPC; that is, the approach can be reused by any other system that provides a way to serialise/deserialise the state of Scilla contracts.

In our evaluation of CoSPLIT, we focus on two aspects of the tool: the quality of the analysis for sharding signatures (Sec. 5.1) and the impact of using the signatures to the system throughput when executing transactions to popular contracts (Sec. 5.2) from the Zilliqa blockchain.⁷

5.1 Evaluating the Analysis

We are interested in quantitative answers to the following questions with regards to the CoSPLIT analysis:

- Is the analysis fast enough to be used for validating contract-deploying transactions (Sec. 5.1.1)?
- Is the analysis capable of inferring non-trivial sharding signatures for real-world contracts (Sec. 5.1.2)?

To answer these questions, we have run the analysis on all 49 unique contracts from Zilliqa mainnet and testnet.

5.1.1 Analysis Performance. All Scilla contracts, upon deployment to the blockchain, are validated by the miners that are forced to parse their code and run the type-checker. We ran the contract deployment pipeline (parsing,

⁷The artefact, including CoSPLIT, its integration with Zilliqa, and the benchmark suite used for evaluation, is archived and made available [53].

Figure 12. Parsing, type checking, and analysis times (μ s).

type-checking, and sharding analysis) for the contracts in our sample and measured how much time is spent in each deployment stage. Fig. 12 shows the performance breakdown of those three components. The numbers were obtained on a PC with an Intel Core i7-8665U CPU by executing the pipeline for each contract a thousand times and averaging the resulting time for each stage. The analysis adds a significant but acceptable overhead of around 46% to the total contract deployment time. Since contract deployments form a very small fraction of all transactions and happen only once during the lifetime of a contract, the analysis can be run by the miners without affecting overall throughput.

5.1.2 Analysis Efficacy.

The bar chart on the right summarises the number of transitions (from 1 to 18) for our 49 contracts. While it may be more likely that a contract with a large number of transitions can be sharded efficiently, having many transitions might also indicate having complex logic, making it difficult to infer a useful signature. To quantify the efficacy of the analysis, we introduce some new terminology.

Definition 5.1 (Hogged fields). A contract’s transition $\mathcal{T}$ hogs a field f in a sharding signature sg iff sg ’s ownership constraints require a shard to *fully* own f to execute $\mathcal{T}$.

Definition 5.2. A sharding signature sg is *good enough* (GE) for its selection of k contract transitions, iff either

- $k = 1$ and the selected transition does not hog fields, or
- $k > 1$ and any field is hogged by at most one transition.

(a) Largest good enough signatures (b) Maximal GE signatures

Figure 13. Statistics for Good Enough signatures.

Intuitively, a sharding signature is good enough if it allows for the existence of a contract state, in which some k of the transitions *can be* run in parallel by different shards. Fig. 13a shows the sizes of the largest good enough signatures for our contract selection. It is worth noting that a larger GE signature (in terms of a number k) might perform worse under real-world load than one with a smaller k , which shards different but *more frequently used* transitions. The following definition outlines the signatures worth comparing.

Definition 5.3 (Maximal GE signature). A GE sharding signature is *maximal* if its selection of transitions is not a proper subset of some other GE signature’s selection.

A contract might have a number of maximal signatures of various sizes. The plot in Fig. 13b depicts those numbers for our contracts. Computing the maximal signatures at the mining time is impractical, as it requires making $\sum_{k=1}^n \binom{n}{k}$ queries to the sharding solver (Fig. 11). Luckily, this computation can be done offline by a contract implementer, who decides prior to the deployment which of the signatures to propose. The miners need to validate only that one signature.

These findings suggest that CoSPLIT indeed uncovers many opportunities for parallel execution of smart contracts.

5.2 Evaluating Sharded Executions

We evaluate the integration of CoSPLIT with Zilliqa by measuring the impact on throughput for five representative Scilla contracts. The chosen contracts are: (1) the most popular contracts on Zilliqa (e.g., UD), (2) equivalent to their popular counterparts on Ethereum currently (e.g., FungibleToken and NonfungibleToken) and in the past (e.g., Crowdfunding).

Contract	LOC	#Trans	Larg.GES	#Max.GES
FungibleToken	439	10	6	2
Crowdfunding	186	3	2	1
NonfungibleToken	288	5	3	2
ProofPFS	289	10	8	2
UD Registry	500	11	6	2

The table above summarises contract sizes, number of transitions, largest GE strategy, and the number of maximal GE signatures. We set out to answer two main questions:

- What throughput improvement, in terms of transactions per second, can CoSPLIT help to achieve (Sec. 5.2.1)?
- What is the impact of the overheads imposed by CoSPLIT-enabled sharding (Sec. 5.2.2)?

Experimental Setup. To obtain the throughput figures, we deployed small-scale testnets in various configurations on Amazon EC2 containers. Each node runs on a *t2.medium* machine with 2 logical CPUs and 4GB of RAM, running Ubuntu 16.04. These specifications reflect the minimum requirements needed to run a Zilliqa node. For our benchmarks, we fix the shard size to be 5 nodes per shard and measure the effect on throughput of increasing the number of shards. We use the same shard and DS gas limits as the Zilliqa mainnet.

Selection of Sharding Signatures. We deploy each of the five contracts in two configurations, one with no sharding information (baseline), and one with a “reasonable” sharding signature, informed by expected usage of the contract. For our experiments, we make the choices as follows:

- For FungibleToken (FT, Zilliqa’s ERC20), we shard the Mint, Transfer and TransferFrom, but not IncreaseAllowance, Burn, or other administrative transitions.
- For NonfungibleToken (NFT, Zilliqa’s ERC-721 [59]), we shard Mint and Transfer (which includes transfer-from functionality), but not Burn and Approve.
- For ProofIPFS, we shard the transition that notarises a hash, but not the one that removes it from the contract.
- For the Unstoppable Domains (UD) registry, we shard granting a new domain name and updating the record associated with a name, but not transfers of ownership.

For the Crowdfunding contract, there is only one possible choice, which is to shard the Donate and Claimback (if the goal was not reached) transitions. We argue that our choices reflect what a reasonable contract deployer would select and, as such, the measured throughput reflects probable scenarios.

We remark that we had to slightly rewrite the NFT and UD contracts, compared to their mainnet definitions, to make them shardable. These rewrites did not affect the semantics of the contracts. We discuss the details of the changes, as well as the potential to automate such modifications, in Sec. 6.

5.2.1 Measuring Throughput. After deployment, we subject the contracts, in sequence, to different workloads sustained over 10 epochs (roughly 8.5 minutes) and measure the resulting throughput. As Fig. 14 shows, for most of the workloads in the benchmark, we obtain a roughly linear TPS increase as the number of shards goes up. The two exceptions are the “FT fund” and the “ProofIPFS register” workloads. The former transfers fungible tokens from a single source to multiple destinations (all transactions go to the source’s shard). The latter notarises a hash, but also keeps a list of notarised items for each user, and thus accesses *two* separate fields, which typically will be *owned by different shards*, hence many transactions need to be processed by the DS Committee. We note that for workloads that do not shard well, performance does not degrade as compared to the baseline, and in some cases (e.g. ProofIPFS) marginally improves.

Figure 14. Average TPS for different contract transitions as a function of number of shards, over 10 epochs.

The “FT transfer” workload sends tokens from random sources to random destinations. In the baseline configuration, the throughput is the same as for the single-source workload. The CoSPLIT-empowered sharding strategy, on the other hand, fully utilises the shards’ processing capacity and we get an almost linear increase in throughput as the number of shards increases. A similar effect arises for crowdfunding donations. Interestingly, the “NFT mint” workload (which creates new tokens) is also single-source, just like “FT fund”. However, the relevant transition does not affect state depending on the identity of the transaction sender, but only on the identifier of the created token. As such, we can obtain linear scaling even for a single-source workload. This is only possible because of the changes to the account-based model that we detailed in Sec. 4.2. Finally, Unstoppable Domains is the most popular smart contract on the Zilliqa mainnet, accounting for over half of the smart contract executions. We manage to shard the most popular transitions on this contract, which account for 90% of usage, and show linear increases in throughput for them as well.

5.2.2 Introduced Overheads. Integrating CoSPLIT adds overheads to transaction dispatch and state delta merging. Concretely, we see transaction dispatch time increase from an average 8 μ s to an average 475 μ s, and the state delta merging increasing from 0.8 μ s to 48.65 μ s per changed state field. This amounts to a 60x slowdown of these operations, most of it as a result of serialisation and deserialisation costs, but this is fully justified by the resulting increase in *overall system throughput*. The overall performance gain comes from the fact that applying a delta is much faster than executing all the transactions that resulted in it. For instance, for FungibleToken, the effects of 50 seconds of transaction execution time can be merged in roughly 0.5 seconds.

5.2.3 Ownership versus Commutativity. In terms of the contribution of the two sharding strategies to the throughput improvement, we observe that contracts that manipulate non-fungible quantities (e.g., NFTs, domains, notary) benefit from the disjoint state ownership analysis and contracts that manipulate fungible quantities (e.g., FT) benefit from the commutativity analysis. Contracts that have a mixture of non-fungible and fungible quantities, e.g., a voting contract

that keeps track of who voted and of the total number of votes, can benefit from both strategies.

6 Discussion and Future Work

Handling Integer Overflows. CoSPLIT's signature inference does not take possible integer overflows into account. Overflows (and underflows) may cause a problem in the case when `IntMerge` is used to join state deltas from different shards that individually do not cause an overflow, but do so when joined. At the moment, our implementation ignores this issue. A working solution would be to modify the Scilla interpreter, providing it with information about the number N of shards. Specifically, it should perform additional post-hoc validation of a transaction, which will fail if the difference between the initial (per epoch) value v of any affected integer-valued component and its value after the transaction is executed exceeds $\lfloor \frac{\text{MAX_INT} - v}{N} \rfloor$. The information about such components is already available in the sharding signatures. Furthermore, a user might be given an option to pay a higher gas fee, in order to reduce a risk for her transaction to be rejected due to this conservative check. Such transactions will be routed directly to the DS committee, and, thus, processed sequentially.

Automated Contract Repair. As mentioned in Sec. 3.3, the analysis can describe map accesses only if the keys used to index into the map are transition parameters. During the evaluation, we discovered a small number of contracts that CoSPLIT cannot shard due to this limitation, but which can be made shardable by a simple refactoring. For instance, an NFT contract's `Transfer` transition ensures that the transaction's sender is authorized by the token owner to initiate the transfer by checking for inclusion in `approvals[tokenOwner]`, where `tokenOwner` is read from the contract's state. This cannot be sharded. However, if we make `tokenOwner` a parameter and rewrite the transition to check that the supplied value matches the value in the contract's state before attempting the transfer (akin to `compare-and-swap`), the transition becomes shardable. In future work, we plan to address this using program repair techniques, suggesting the shardable contract version to the developer before deployment.

CoSPLIT and Other Blockchains. As demonstrated in Sec. 3, the core analysis of CoSPLIT does not rely on any specific features of Scilla but is easy to implement for it due to the language's minimalism and restrictions (e.g., very limited set of side effects). Should other *account-based* blockchains, e.g., Tezos [28] and Ethereum [65], provide a sharded architecture in the future, we believe a similar analysis could be implemented for them as well. The key challenge of developing CoSPLIT for Michelson, the language of Tezos [63], is in reasoning about its stack-based executions, tracking the provenance and cardinality of pushed and popped values. Ethereum's EVM could be supported through decompilation into a high-level language [29]. The approach will most likely have to be restricted to contracts with no external calls.

7 Related Work

Sharding Contracts in Blockchains. Several industry proposals outline approaches for smart contract sharding, yet none of them provide an efficient solution for sharding same-contract transactions. For example, the Elrond protocol moves a smart contract to the same shard where its static dependencies lie [23], which takes at least 5 rounds of consensus for a transaction to be finalised. Harmony [30] allows one to deploy contracts in individual shards, with no cross-shard communication allowed. Ethereum 2.0 proposes a cross-shard yanking scheme where the contract code and data is moved into a shard at runtime [64]. The shard then locks the contract to block any parallel execution of other transactions affecting it. At the time of writing, none of these solutions appear to be fully implemented.

The Chainspace protocol allows for sharded execution of smart contracts (as well as state sharding) by representing state evolution as a directed acyclic graph [1], similar to Bitcoin's UTXO transaction model. In Chainspace, contracts have to be written in a specific fashion, so their execution happens *off the chain*. For an unspecified kind of transaction, the authors of Chainspace report linearly increasing throughput of approximately 75 TPS per each two shards of four nodes in each [1, Fig. 6]. Importantly, Chainspace does not address the scalability problem with same-contract transactions (which we do): under high contention for the same contract the rate of aborted transactions rises. This is because its protocol, *S-BAC* (a combination of PBFT [12] and Two-Phase Commit for inter-shard communication), implements a variant of optimistic concurrency control, whereas CoSPLIT allows for pessimistic (race-free) concurrency.

Smart Contracts and Concurrency. Dickerson *et al.* [18] describe a commutativity-based approach where miners execute smart contracts in parallel *locally*, using software transactional memory techniques [31, 32]. Unlike our approach that detects possible transaction conflicts pessimistically by means of a static analysis prior to contract deployment, the work by Dickerson *et al.* exercises the optimistic approach, where conflicts are identified on-the-fly by the miners, with the corresponding executions aborted. It is not clear how to apply these ideas from optimistic concurrency for efficient *sharding* of general smart contracts in a Byzantine setting, where an adversary can craft cheap (in terms of gas costs) conflict-producing transactions that force the shards to re-execute expensive transactions from the same batch. A more fine-grained identification of conflicts would most probably require an analysis similar to ours, as conjectured in [56]. At the same time, our solution is complementary to Dickerson *et al.*'s work and other similar approaches. For instance, one can add single-node parallelism on top of CoSPLIT-enabled sharding, and the analysis can help identify which transactions are guaranteed not to require rollback and re-execution, even within a single shard.

Recent work by Bartoletti *et al.* makes an observation similar to ours that commutativity (which they dub *swappability*) of transactions manipulating Ethereum contracts enables their parallel execution [6]. Their syntactic criterion for inferring swappability is, however, more coarse-grained than our analysis, and is based on determining disjointness of transaction footprints, without taking into account commutativity of operations such as addition. As such, their approach would not allow to shard individual non-conflicting updates in an ERC20 contract as we do. Bartoletti *et al.*'s approach has not been implemented in practice.

Inferring Commutativity. Reasoning about commutativity between program parts is an important problem with applications including parallelising compilers [38, 55], speculative execution [31], and race detection [19]. Most of existing techniques for inferring commutativity are based on analysing dynamic executions [2, 19, 26] or by solving SMT constraints [5] and, thus, cannot be used efficiently as a part of transaction validation. Our analysis is close in spirit to the work (in [55]) on static analysis to determine operation commutativity for compile-time parallelisation. Our analysis uses simpler abstractions, allowing it to be implemented in a compositional fashion and have linear execution cost.

8 Conclusion

We presented a new approach to shard the execution of smart contracts in an account-based blockchain model, based on inferring ownership constraints and commutativity for state-manipulating contract operations. In our approach, smart contracts are first processed by a static analysis tool that produces their sharding signatures, which are then used for shard allocation that maximises parallelism. We have demonstrated that our approach, when integrated into the sharded production-scale blockchain, allows for linear scaling of the transaction throughput for a selection of common smart contracts that were considered unavoidable execution bottlenecks in existing blockchain systems.

Acknowledgments

We thank the OSDI'20 and PLDI'21 PC and AEC reviewers, and our shepherd, Gustavo Petri, for their valuable feedback on earlier drafts of this paper and the artefact. We thank Aquinas Hobor, Yaoqi Jia, and Prateek Saxena who participated in the discussions on smart contract sharding at earlier stages of this work. We are grateful to Haichuan Liu, Antonio Nicolas Nunez, and Jun Hao Tan, for their help in integrating CoSPLIT with the Zilliqa protocol, and to Bryan Tan for conducting preliminary experiments. We benefited a lot from discussions and feedback on preliminary versions of this work by Andreea Costea, Kiran Gopinathan, Jacob Johannsen, Vaivaswatha Nagaraj, and Anton Trunov. Ilya Sergey's work has been supported by the grant of Singapore NRF National Satellite of Excellence in Trustworthy Software Systems (NSoE-TSS) and by the NUS Crystal Centre.

References

- [1] Mustafa Al-Bassam, Alberto Sonnino, Shehar Bano, Dave Hrycyszyn, and George Danezis. 2018. Chainspace: A Sharded Smart Contracts Platform. In *NDSS*. <https://doi.org/10.14722/ndss.2018.23241>
- [2] Farhana Aleen and Nathan Clark. 2009. Commutativity analysis for software parallelization: letting program transformations see the big picture. In *ASPLOS*. ACM. <https://doi.org/10.1145/1508244.1508273>
- [3] Hagit Attiya, Rachid Guerraoui, Danny Hendler, Petr Kuznetsov, Maged M. Michael, and Martin T. Vechev. 2011. Laws of order: expensive synchronization in concurrent algorithms cannot be eliminated. In *POPL*. ACM, 487–498. <https://doi.org/10.1145/1925844.1926442>
- [4] Shehar Bano, Alberto Sonnino, Mustafa Al-Bassam, Sarah Azouvi, Patrick McCorry, Sarah Meiklejohn, and George Danezis. 2019. SoK: Consensus in the Age of Blockchains. In *Advances in Financial Technologies*. ACM, 183–198. <https://doi.org/10.1145/3318041.3355458>
- [5] Kshitij Bansal, Eric Koskinen, and Omer Tripp. 2018. Automatic Generation of Precise and Useful Commutativity Conditions. In *TACAS (LNCS, Vol. 10805)*. Springer, 115–132. https://doi.org/10.1007/978-3-319-89960-2_7
- [6] Massimo Bartoletti, Letterio Galletta, and Maurizio Murgia. 2020. A True Concurrent Model of Smart Contracts Executions. In *COORDINATION (LNCS, Vol. 12134)*. Springer, 243–260. https://doi.org/10.1007/978-3-030-50029-0_16
- [7] Sebastian Burckhardt, Alexandro Baldassin, and Daan Leijen. 2010. Concurrent programming with revisions and isolation types. In *OOPSLA*. ACM, 691–707. <https://doi.org/10.1145/1869459.1869515>
- [8] Sebastian Burckhardt and Daan Leijen. 2011. Semantics of Concurrent Revisions. In *ESOP (LNCS, Vol. 6602)*. Springer, 116–135. https://doi.org/10.1007/978-3-642-19718-5_7
- [9] Cristiano Calcagno, Dino Distefano, Peter W. O'Hearn, and Hongseok Yang. 2011. Compositional Shape Analysis by Means of Bi-Abduction. *J. ACM* 58, 6 (2011), 1–66. <https://doi.org/10.1145/2049697.2049700>
- [10] Cristiano Calcagno, Peter W. O'Hearn, and Hongseok Yang. 2007. Local Action and Abstract Separation Logic. In *LICS*. IEEE, 366–378. <https://doi.org/10.1109/lics.2007.30>
- [11] Lásaro J. Camargos, Rodrigo Schmidt, and Fernando Pedone. 2007. Multicoordinated Paxos. In *PODC*. ACM, 316–317. <https://doi.org/10.1145/1281100.1281150>
- [12] Miguel Castro and Barbara Liskov. 1999. Practical Byzantine Fault Tolerance. In *OSDI*. USENIX Association, 173–186.
- [13] Austin T. Clements, M. Frans Kaashoek, Nickolai Zeldovich, Robert Tappan Morris, and Eddie Kohler. 2013. The scalable commutativity rule: designing scalable software for multicore processors. In *SOSP*. ACM, 1–17. <https://doi.org/10.1145/2517349.2522712>
- [14] ConsenSys. 2018. The Inside Story of the CryptoKitties Congestion Crisis.
- [15] Patrick Cousot and Radhia Cousot. 1977. Abstract Interpretation: A Unified Lattice Model for Static Analysis of Programs by Construction or Approximation of Fixpoints. In *POPL*. ACM, 238–252. <https://doi.org/10.1145/512950.512973>
- [16] Patrick Cousot and Radhia Cousot. 1979. Systematic Design of Program Analysis Frameworks. In *POPL*. ACM Press, 269–282. <https://doi.org/10.1145/567752.567778>
- [17] Hung Dang, Tien Tuan Anh Dinh, Dumitrel Loghin, Ee-Chien Chang, Qian Lin, and Beng Chin Ooi. 2019. Towards Scaling Blockchain Systems via Sharding. In *SIGMOD*. ACM, 123–140. <https://doi.org/10.1145/3299869.3319889>
- [18] Thomas D. Dickerson, Paul Gazzillo, Maurice Herlihy, and Eric Koskinen. 2017. Adding Concurrency to Smart Contracts. In *PODC*. ACM, 303–312. <https://doi.org/10.1145/3087801.3087835>
- [19] Dimitar Dimitrov, Veselin Raychev, Martin T. Vechev, and Eric Koskinen. 2014. Commutativity race detection. In *PLDI*. ACM, 305–315. <https://doi.org/10.1145/2594291.2594322>

- [20] Thomas Dinsdale-Young, Lars Birkedal, Philippa Gardner, Matthew J. Parkinson, and Hongseok Yang. 2013. Views: compositional reasoning for concurrent programs. In *POPL*. ACM, 287–300. <https://doi.org/10.1145/2429069.2429104>
- [21] Dino Distefano, Manuel Fähndrich, Francesco Logozzo, and Peter W. O'Hearn. 2019. Scaling static analyses at Facebook. *Commun. ACM* 62, 8 (2019), 62–70. <https://doi.org/10.1145/3338112>
- [22] Tayfun Elmas, Shaz Qadeer, and Serdar Tasiran. 2009. A calculus of atomic actions. In *POPL*. ACM, 2–15. <https://doi.org/10.1145/1480881.1480885>
- [23] The Elrond Team. 2019. Elrond: A Highly Scalable Public Blockchain via Adaptive State Sharding and Secure Proof of Stake.
- [24] Ethereum Foundation. 2018. ERC20 Token Standard.
- [25] Ethereum Foundation. 2018. Solidity Documentation.
- [26] Timon Gehr, Dimitar Dimitrov, and Martin T. Vechev. 2015. Learning Commutativity Specifications. In *CAV (LNCS, Vol. 9206)*. Springer. https://doi.org/10.1007/978-3-319-21690-4_18
- [27] Jean-Yves Girard. 1972. *Interprétation fonctionnelle et élimination des coupures de l'arithmétique d'ordre supérieur*. Thèse d'État. Université de Paris VII, Paris, France.
- [28] L.M. Goodman. 2014. Tezos—a Self-Amending Crypto-Ledger.
- [29] Neville Grech, Lexi Brent, Bernhard Scholz, and Yannis Smaragdakis. 2019. Gigahorse: thorough, declarative decompilation of smart contracts. In *ICSE*. IEEE / ACM, 1176–1186. <https://doi.org/10.1109/icse.2019.00120>
- [30] Harmony Team. 2018. Harmony: Technical Whitepaper.
- [31] Maurice Herlihy and Eric Koskinen. 2008. Transactional boosting: a methodology for highly-concurrent transactional objects. In *PPoPP*. ACM, 207–216. <https://doi.org/10.1145/1345206.1345237>
- [32] Maurice Herlihy, Victor Luchangco, Mark Moir, and William N. Scherer III. 2003. Software transactional memory for dynamic-sized data structures. In *PODC*. ACM, 92–101. <https://doi.org/10.1145/872035.872048>
- [33] Sungjae Hwang and Sukyoung Ryu. 2020. Gap between Theory and Practice: An Empirical Study of Security Patches in Solidity. In *ICSE*. ACM, 542–553. <https://doi.org/10.1145/3377811.3380424>
- [34] Ralf Jung, David Swasey, Filip Sieczkowski, Kasper Svendsen, Aaron Turon, Lars Birkedal, and Derek Dreyer. 2015. Iris: Monoids and Invariants as an Orthogonal Basis for Concurrent Reasoning. In *POPL*. ACM, 637–650. <https://doi.org/10.1145/2676726.2676980>
- [35] Gowtham Kaki, Swarn Priya, K. C. Sivaramakrishnan, and Suresh Jaganathan. 2019. Mergeable replicated data types. *Proc. ACM Program. Lang.* 3, OOPSLA (2019), 154:1–154:29. <https://doi.org/10.1145/3360580>
- [36] Eleftherios Kokoris-Kogias, Philipp Jovanovic, Linus Gasser, Nicolas Gailly, Ewa Syta, and Bryan Ford. 2018. OmniLedger: A Secure, Scale-Out, Decentralized Ledger via Sharding. In *IEEE Symposium on Security and Privacy*. IEEE, 583–598. <https://doi.org/10.1109/sp.2018.000-5>
- [37] Tim Kraska, Gene Pang, Michael J. Franklin, Samuel Madden, and Alan Fekete. 2013. MDCC: multi-data center consistency. In *EuroSys*. ACM, 113–126. <https://doi.org/10.1145/2465351.2465363>
- [38] Milind Kulkarni, Keshav Pingali, Bruce Walter, Ganesh Ramnarayanan, Kavita Bala, and L. Paul Chew. 2007. Optimistic parallelism requires abstractions. In *PLDI*. ACM, 211–222. <https://doi.org/10.1145/1250734.1250759>
- [39] Amrit Kumar. 2018. Provisioning Sharding for Smart Contracts: A Design for Zilliqa.
- [40] Lindsey Kuper, Aaron Turon, Neelakantan R. Krishnaswami, and Ryan R. Newton. 2014. Freeze after writing: quasi-deterministic parallel programming with LVars. In *POPL*. ACM, 257–270. <https://doi.org/10.1145/2535838.2535842>
- [41] Leslie Lamport. 1998. The Part-Time Parliament. *ACM Trans. Comput. Syst.* 16, 2 (1998), 133–169. <https://doi.org/10.1145/279227.279229>
- [42] Leslie Lamport. 2005. *Generalized Consensus and Paxos*. Technical Report 33. Microsoft Research.
- [43] Daan Leijen, Manuel Fähndrich, and Sebastian Burckhardt. 2011. Prettier concurrency: purely functional concurrent revisions. In *Haskell Symposium*. ACM, 83–94. <https://doi.org/10.1145/2034675.2034686>
- [44] Ruy Ley-Wild and Aleksandar Nanevski. 2013. Subjective auxiliary state for coarse-grained concurrency. In *POPL*. ACM, 561–574. <https://doi.org/10.1145/2429069.2429134>
- [45] Loi Luu, Viswesh Narayanan, Chaodong Zheng, Kunal Baweja, Seth Gilbert, and Prateek Saxena. 2016. A Secure Sharding Protocol For Open Blockchains. In *CCS*. ACM, 17–30. <https://doi.org/10.1145/2976749.2978389>
- [46] Nancy A. Lynch and Mark R. Tuttle. 1989. An Introduction to Input/Output Automata. *CWI Quarterly* 2 (1989), 219–246.
- [47] Iulian Moraru, David G. Andersen, and Michael Kaminsky. 2013. There is more consensus in Egalitarian parliaments. In *SOSP*. ACM, 358–372. <https://doi.org/10.1145/2517349.2517350>
- [48] Satoshi Nakamoto. 2008. Bitcoin: A peer-to-peer electronic cash system.
- [49] Flemming Nielson, Hanne Riis Nielson, and Chris Hankin. 1999. *Principles of Program Analysis*. Springer.
- [50] Peter W. O'Hearn. 2007. Resources, concurrency, and local reasoning. *Theor. Comput. Sci.* 375, 1–3 (2007), 271–307. <https://doi.org/10.1016/j.tcs.2006.12.035>
- [51] Ontology Research Team. 2019. Ontology Sharding, Draft v0.2.
- [52] Prakash Prabhu, Soumyadeep Ghosh, Yun Zhang, Nick P. Johnson, and David I. August. 2011. Commutative set: a language extension for implicit parallel programming. In *PLDI*. ACM, 1–11. <https://doi.org/10.1145/1993498.1993500>
- [53] George Pirlea, Amrit Kumar, and Ilya Sergey. 2021. *CoSplit (PLDI 2021 Artefact)*. <https://doi.org/10.5281/zenodo.4674301>
- [54] John C. Reynolds. 1974. Towards a theory of type structure. In *Programming Symposium (LNCS, Vol. 19)*. Springer, 408–425. https://doi.org/10.1007/3-540-06859-7_148
- [55] Martin C. Rinard and Pedro C. Diniz. 1997. Commutativity Analysis: A New Analysis Technique for Parallelizing Compilers. *ACM Trans. Program. Lang. Syst.* 19, 6 (1997), 942–991. <https://doi.org/10.1145/267959.269969>
- [56] Vikram Saraph and Maurice Herlihy. 2019. An Empirical Study of Speculative Concurrency in Ethereum Smart Contracts. In *Tokenomics (OASIS, Vol. 71)*. Schloss Dagstuhl, 4:1–4:15.
- [57] Ilya Sergey, Vaivaswatha Nagaraj, Jacob Johannsen, Amrit Kumar, Anton Trunov, and Ken Chan Guan Hao. 2019. Safer smart contract programming with Scilla. *PACMPL* 3, OOPSLA (2019), 185:1–185:30. <https://doi.org/10.1145/3360611>
- [58] Ilya Sergey, Dimitrios Vytiniotis, and Simon L. Peyton Jones. 2014. Modular, higher-order cardinality analysis in theory and practice. In *POPL*. ACM, 335–348. <https://doi.org/10.1145/2535838.2535861>
- [59] Dieter Shirley. 2018. ERC-721 contract specification. <http://erc721.org/>.
- [60] Alex Skidanov and Illia Polosukhin. 2019. Nightshade: Near Protocol Sharding Design. <https://near.org/downloads/Nightshade.pdf>.
- [61] Ewa Syta, Iulia Tamas, Dylan Visser, David Isaac Wolinsky, Philipp Jovanovic, Linus Gasser, Nicolas Gailly, Ismail Khoffi, and Bryan Ford. 2016. Keeping Authorities “Honest or Bust” with Decentralized Witness Cosigning. In *IEEE Symposium on Security and Privacy*. IEEE Computer Society, 526–545. <https://doi.org/10.1109/SP.2016.38>
- [62] Nick Szabo. 1994. Smart Contracts. Online manuscript.
- [63] Tezos. 2018. Michelson: the language of Smart Contracts in Tezos.
- [64] Ethereum Wiki. 2020. On Sharding Blockchains.
- [65] Gavin Wood. 2014. Ethereum: A Secure Decentralized Generalised Transaction Ledger.
- [66] Mahdi Zamani, Mahnush Movahedi, and Mariana Raykova. 2018. RapidChain: Scaling Blockchain via Full Sharding. In *CCS*. ACM, 931–948. <https://doi.org/10.1145/3243734.3243853>
- [67] Zilliqa Team. 2017. The Zilliqa Technical Whitepaper.