






A True Concurrent Model of Smart Contracts Executions

Massimo Bartoletti¹, Letterio Galletta², and Maurizio Murgia³

¹ University of Cagliari, Cagliari, Italy

² IMT School for Advanced Studies, Lucca, Italy

³ University of Trento, Trento, Italy

Abstract. The development of blockchain technologies has enabled the trustless execution of so-called *smart contracts*, i.e. programs that regulate the exchange of assets (e.g., cryptocurrency) between users. In a decentralized blockchain, the state of smart contracts is collaboratively maintained by a peer-to-peer network of mutually untrusted nodes, which collect from users a set of *transactions* (representing the required actions on contracts), and execute them in some order. Once this sequence of transactions is appended to the blockchain, the other nodes validate it, re-executing the transactions in the same order. The serial execution of transactions does not take advantage of the multi-core architecture of modern processors, so contributing to limit the throughput. In this paper we propose a true concurrent model of smart contracts execution. Based on this, we show how static analysis of smart contracts can be exploited to parallelize the execution of transactions.

1 Introduction

Smart contracts [21] are computer programs that transfer digital assets between users without a trusted authority. Currently, smart contracts are supported by several blockchains, the first and most widespread one being Ethereum [9]. Users interact with a smart contract by sending *transactions*, which trigger state updates, and may possibly involve transfers of crypto-assets between the called contract and the users. The sequence of transactions on the blockchain determines the state of each contract, and the balance of each user.

The blockchain is maintained by a peer-to-peer network of nodes, which follow a consensus protocol to determine, at each turn, a new block of transactions to be added to the blockchain. This protocol guarantees the correct execution of contracts also in the presence of (a minority of) adversaries in the network, and ensures that all the nodes have the same view of their state. Nodes play the role of *miner* or that of *validator*. Miners gather from the network sets of transactions sent by users, and execute them *serially* to determine the new state. Once

a block is appended to the blockchain, validators re-execute all its transactions, to update their local view of the contracts state and of the users' balance. To do this, validators process the transactions exactly in the same order in which they occur in the block, since choosing a different order could potentially result in inconsistencies between the nodes (note that miners also act as validators, since they validate all the blocks received from the network).

Although executing transactions in a purely sequential fashion is quite effective to ensure the consistency of the blockchain state, in the age of multi-core processors it fails to properly exploit the computational capabilities of nodes. By enabling miners and validators to concurrently execute transactions, it would be possible to improve the efficiency and the throughput of the blockchain.

This paper exploits techniques from concurrency theory to provide a formal backbone for parallel executions of transactions. More specifically, our main contributions can be summarised as follows:

- As a first step, we formalise blockchains, giving their semantics as a function which maps each contract to its state, and each user to her balance. This semantics reflects the standard implementation of nodes, where transactions are evaluated in sequence, without any concurrency.
- We introduce two notions of *swappability* of transactions. The first is purely semantic: two adjacent transactions can be swapped if doing so preserves the semantics of the blockchain. The second notion, called *strong swappability*, is more syntactical: it checks a simple condition (inspired by Bernstein's conditions [7]) on static approximations of the variables read/written by the transactions. Theorem 2 shows that strong swappability is strictly included in the semantic relation. Further, if we transform a blockchain by repeatedly exchanging adjacent strongly swappable transactions, the resulting blockchain is observationally equivalent to the original one (Theorem 4).
- Building upon strong swappability, we devise a true concurrent model of transactions execution. To this purpose, we transform a block of transactions into an *occurrence net*, describing exactly the partial order induced by the swappability relation. We model the concurrent executions of a blockchain in terms of the *step firing sequences* (i.e. finite sequences of *sets* of transitions) of the associated occurrence net. Theorem 5 establishes that the concurrent executions and the serial one are semantically equivalent.
- We describe how miners and validators can use our results to concurrently execute transactions, exploiting the multi-core architecture available on their nodes. Remarkably, our technique is compatible with the current implementation of the Ethereum blockchain, while the other existing approaches to parallelize transactions execution would require a soft-fork.
- We apply our technique to ERC-721 tokens, one of the most common kinds of contracts in Ethereum, showing them to be suitable for parallelization.

Because of space constraints, all the proofs of our results are in [6].

2 Transactions and Blockchains

In this section we introduce a general model of transactions and blockchains, abstracting from the actual smart contracts language.

A smart contract is a finite set of functions, i.e. terms of the form $\mathbf{f}(\mathbf{x})\{S\}$, where \mathbf{f} is a function name, \mathbf{x} is the sequence of formal parameters (omitted when empty), and S is the function body. We postulate that the functions in a contract have distinct names. We abstract from the actual syntax of S , and we just assume that the semantics of function bodies is defined (see e.g. [5] for a concrete instance of syntax and semantics of function bodies).

Let **Val** be a set of *values*, ranged over by v, v', \dots , let **Const** be a set of *constant names* x, y, \dots , and let **Addr** be a set of *addresses* $\mathcal{X}, \mathcal{Y}, \dots$, partitioned into *account addresses* $\mathcal{A}, \mathcal{B}, \dots$ and *contract addresses* $\mathcal{C}, \mathcal{D}, \dots$. We assume a mapping Γ from addresses to contracts.

We assume that each contract has a key-value store, which we render as a partial function $\mathbf{Val} \rightarrow \mathbf{Val}$ from keys $k \in \mathbf{Val}$ to values. The state of the blockchain is a function $\sigma : \mathbf{Addr} \rightarrow (\mathbf{Val} \rightarrow \mathbf{Val})$ from addresses to key-value stores. We postulate that $\mathbf{balance} \in \text{dom } \sigma \mathcal{X}$ for all \mathcal{X} . A *qualified key* is a term of the form $\mathcal{X}.k$. We write $\sigma(\mathcal{X}.k)$ for $\sigma \mathcal{X}k$; when $k \notin \text{dom } \sigma \mathcal{X}$, we write $\sigma(\mathcal{X}.k) = \perp$. We use p, q, \dots to range over qualified keys, P, Q, \dots to range over sets of them, and \mathbb{P} to denote the set of all qualified keys.

To have a uniform treatment of accounts and contracts, we assume that for all account addresses \mathcal{A} , $\text{dom } \sigma \mathcal{A} = \{\mathbf{balance}\}$, and that the contract $\Gamma(\mathcal{A})$ has exactly one function, which just skips. In this way, the statement $\mathcal{A}.(n)$, which transfers n currency units to \mathcal{A} , can be rendered as a call to this function.

State updates define how values associated with qualified keys are modified.

Definition 1 (State update). *A state update $\pi : \mathbf{Addr} \rightarrow (\mathbf{Val} \rightarrow \mathbf{Val})$ is a function from qualified keys to values; we denote with $\{v/x.k\}$ the state update which maps $\mathcal{X}.k$ to v . We define $\text{keys}(\pi)$ as the set of qualified keys $\mathcal{X}.k$ such that $\mathcal{X} \in \text{dom } \pi$ and $k \in \text{dom } \pi \mathcal{X}$. We apply updates to states as follows:*

$$(\sigma\pi)\mathcal{X} = \delta_{\mathcal{X}} \quad \text{where} \quad \delta_{\mathcal{X}}k = \begin{cases} \pi \mathcal{X}k & \text{if } \mathcal{X}.k \in \text{keys}(\pi) \\ \sigma \mathcal{X}k & \text{otherwise} \end{cases} \quad \diamond$$

We denote with $\llbracket S \rrbracket_{\sigma, \rho}^{\mathcal{X}}$ the semantics of the statement S . This semantics is either a blockchain state σ' , or it is undefined (denoted by \perp). The semantics is parameterised over a state σ , an address \mathcal{X} (the contract wherein S is evaluated), and an *environment* $\rho : \mathbf{Const} \rightarrow \mathbf{Val}$, used to evaluate the formal parameters and the special names **sender** and **value**. These names represent, respectively, the caller of the function, and the amount of currency transferred along with the call. We postulate that **sender** and **value** are not used as formal parameters.

We define the auxiliary operators $+$ and $-$ on states as follows:

$$\sigma \circ (\mathcal{X} : n) = \sigma\{\sigma \mathcal{X} \mathbf{balance} \circ n / \mathcal{X} \mathbf{balance}\} \quad (\circ \in \{+, -\})$$

i.e., $\sigma + \mathcal{X} : n$ updates σ by increasing the **balance** of \mathcal{X} of n currency units.

A *transaction* T is a term of the form:

$$\mathcal{A} \xrightarrow{n} \mathcal{C} : \mathbf{f}(\mathbf{v})$$

Intuitively, \mathcal{A} is the address of the caller, \mathcal{C} is the address of the called contract, \mathbf{f} is the called function, n is the value transferred from \mathcal{A} to \mathcal{C} , and \mathbf{v} is the sequence of actual parameters. We denote the semantics of T in σ as $\llbracket \mathsf{T} \rrbracket_\sigma$, where the function $\llbracket \cdot \rrbracket_\sigma$ is defined in Fig. 1, which we briefly comment.

$$\frac{\begin{array}{l} \mathbf{f}(\mathbf{x})\{S\} \in \Gamma(\mathcal{C}) \\ \sigma \mathcal{A} \text{ balance} \geq n \\ \llbracket S \rrbracket_{\sigma - \mathcal{A} : n + \mathcal{C} : n, \{ \mathcal{A} / \text{sender}, n / \text{value}, \mathbf{v} / \mathbf{x} \}}^{\mathcal{C}} = \sigma' \end{array}}{\llbracket \mathcal{A} \xrightarrow{n} \mathcal{C} : \mathbf{f}(\mathbf{v}) \rrbracket_\sigma = \sigma'} \text{[Tx1]} \quad \frac{\begin{array}{l} \mathbf{f}(\mathbf{x})\{S\} \in \Gamma(\mathcal{C}) \\ \left(\sigma \mathcal{A} \text{ balance} < n \quad \text{or} \right. \\ \left. \llbracket S \rrbracket_{\sigma - \mathcal{A} : n + \mathcal{C} : n, \{ \mathcal{A} / \text{sender}, n / \text{value}, \mathbf{v} / \mathbf{x} \}}^{\mathcal{C}} = \perp \right) \end{array}}{\llbracket \mathcal{A} \xrightarrow{n} \mathcal{C} : \mathbf{f}(\mathbf{v}) \rrbracket_\sigma = \sigma} \text{[Tx2]}$$

Fig. 1. Semantics of transactions.

The semantics of a transaction $\mathsf{T} = \mathcal{A} \xrightarrow{n} \mathcal{C} : \mathbf{f}(\mathbf{v})$, in a given blockchain state σ , is a new state σ' . Rule [Tx1] handles the case where the transaction is successful: this happens when \mathcal{A} 's balance is at least n , and the function call terminates in a non-error state. Note that n units of currency are transferred to \mathcal{C} *before* starting to execute \mathbf{f} , and that the names **sender** and **value** are bound, respectively, to \mathcal{A} and n . Rule [Tx2] applies either when \mathcal{A} 's balance is not enough, or the execution of \mathbf{f} fails. In these cases, T does not alter the state.

A *blockchain* \mathbf{B} is a finite sequence of transactions; we denote with ϵ the empty blockchain. The semantics of a blockchain is obtained by folding the semantics of its transactions, starting from a given state σ :

$$\llbracket \epsilon \rrbracket_\sigma = \sigma \quad \llbracket \mathbf{TB} \rrbracket_\sigma = \llbracket \mathbf{B} \rrbracket_{\llbracket \mathbf{T} \rrbracket_\sigma}$$

Note that erroneous transactions can occur within a blockchain, but they have no effect on its semantics (as rule [Tx2] makes them identities w.r.t. the append operation). We assume that in the initial state of the blockchain, denoted by σ^* , each address \mathcal{X} has a balance $n_{\mathcal{X}}^* \geq 0$, while all the other keys are unbound.

We write $\llbracket \mathbf{B} \rrbracket$ for $\llbracket \mathbf{B} \rrbracket_{\sigma^*}$, where $\sigma^* \mathcal{X} = \{ n_{\mathcal{X}}^* / \text{balance} \}$. We say that a state σ is *reachable* if $\sigma = \llbracket \mathbf{B} \rrbracket$ for some \mathbf{B} .

Example 1. Consider the following functions of a contract at address \mathcal{C} :

$$\mathbf{f}_0()\{x:=1\} \quad \mathbf{f}_1()\{\text{if } x = 0 \text{ then } \mathcal{B}.\} \quad (1) \quad \mathbf{f}_2()\{\mathcal{B}.\} \quad (1)$$

Let σ be a state such that $\sigma \mathcal{A} \text{ balance} \geq 2$, and let $\mathbf{B} = \mathsf{T}_0 \mathsf{T}_1 \mathsf{T}_2$, where:

$$\mathsf{T}_0 = \mathcal{A} \xrightarrow{0} \mathcal{C} : \mathbf{f}_0() \quad \mathsf{T}_1 = \mathcal{A} \xrightarrow{1} \mathcal{C} : \mathbf{f}_1() \quad \mathsf{T}_2 = \mathcal{A} \xrightarrow{1} \mathcal{C} : \mathbf{f}_2()$$

By applying rule $[\text{Tx}_1]$ three times, we have that:

$$\begin{aligned} \llbracket \mathbf{T}_0 \rrbracket_\sigma &= \llbracket x := 1 \rrbracket_{\sigma, \{\mathcal{A}/\text{sender}, 0/\text{value}\}}^{\mathcal{C}} = \sigma\{1/e.x\} = \sigma' \\ \llbracket \mathbf{T}_1 \rrbracket_{\sigma'} &= \llbracket \text{if } x = 0 \text{ then } \mathcal{B}. \quad (1) \rrbracket_{\sigma' - \mathcal{A} : 1 + \mathcal{C} : 1, \{\mathcal{A}/\text{sender}, 1/\text{value}\}}^{\mathcal{C}} \\ &= \sigma' - \mathcal{A} : 1 + \mathcal{C} : 1 = \sigma'' \\ \llbracket \mathbf{T}_2 \rrbracket_{\sigma''} &= \llbracket \mathcal{B}. \quad (1) \rrbracket_{\sigma'' - \mathcal{A} : 1 + \mathcal{C} : 1, \{\mathcal{A}/\text{sender}, 1/\text{value}\}}^{\mathcal{C}} = \sigma'' - \mathcal{A} : 1 + \mathcal{B} : 1 \end{aligned}$$

Summing up, $\llbracket \mathbf{B} \rrbracket_\sigma = \sigma\{1/e.x\} - \mathcal{A} : 2 + \mathcal{B} : 1 + \mathcal{C} : 1$. \diamond

3 Swapping Transactions

We define two blockchain states to be *observationally equivalent* when they agree on the values associated to all the qualified keys. Our formalisation is parameterised on a set of qualified keys P over which we require the agreement.

Definition 2 (Observational equivalence). *For all $P \subseteq \mathbb{P}$, we define $\sigma \sim_P \sigma'$ iff $\forall p \in P : \sigma p = \sigma' p$. We say that σ and σ' are observationally equivalent, in symbols $\sigma \sim \sigma'$, when $\sigma \sim_P \sigma'$ holds for all P .* \diamond

Lemma 1. *For all $P, Q \subseteq \mathbb{P}$: (i) \sim_P is an equivalence relation; (ii) if $\sigma \sim_P \sigma'$ and $Q \subseteq P$, then $\sigma \sim_Q \sigma'$; (iii) $\sim = \sim_{\mathbb{P}}$.* \diamond

We extend the equivalence relations above to blockchains, by passing through their semantics. For all P , we define $\mathbf{B} \sim_P \mathbf{B}'$ iff $\llbracket \mathbf{B} \rrbracket_\sigma \sim_P \llbracket \mathbf{B}' \rrbracket_\sigma$ holds for all reachable σ (note that all the definitions and results in this paper apply to reachable states, since the unreachable ones do not represent actual contract executions). We write $\mathbf{B} \sim \mathbf{B}'$ when $\mathbf{B} \sim_P \mathbf{B}'$ holds for all P . The relation \sim is a *congruence* with respect to the append operation, i.e. if $\mathbf{B} \sim \mathbf{B}'$ then we can replace \mathbf{B} with \mathbf{B}' in a larger blockchain, preserving its semantics.

Lemma 2. $\mathbf{B} \sim \mathbf{B}' \implies \forall \mathbf{B}_0, \mathbf{B}_1 : \mathbf{B}_0 \mathbf{B} \mathbf{B}_1 \sim \mathbf{B}_0 \mathbf{B}' \mathbf{B}_1$. \diamond

Two transactions are *swappable* when exchanging their order preserves observational equivalence.

Definition 3 (Swappability). *Two transactions $\mathbf{T} \neq \mathbf{T}'$ are swappable, in symbols $\mathbf{T} \rightleftharpoons \mathbf{T}'$, when $\mathbf{T} \mathbf{T}' \sim \mathbf{T}' \mathbf{T}$.* \diamond

Example 2. Recall the transactions in Example 1. We have that $\mathbf{T}_0 \rightleftharpoons \mathbf{T}_2$ and $\mathbf{T}_1 \rightleftharpoons \mathbf{T}_2$, but $\mathbf{T}_0 \not\rightleftharpoons \mathbf{T}_1$ (see Fig. 5 in Appendix A of [6]). \diamond

We shall use the theory of trace languages originated from Mazurkiewicz's works [17] to study observational equivalence under various swapping relations. Below, we fix the alphabet of trace languages as the set \mathbf{Tx} of all transactions.

Definition 4 (Mazurkiewicz equivalence). *Let I be a symmetric and irreflexive relation on \mathbf{Tx} . The Mazurkiewicz equivalence \simeq_I is the least congruence in the free monoid \mathbf{Tx}^* such that: $\forall \mathbf{T}, \mathbf{T}' \in \mathbf{Tx} : \mathbf{T} I \mathbf{T}' \implies \mathbf{T} \mathbf{T}' \simeq_I \mathbf{T}' \mathbf{T}$.*

Theorem 1 below states that the Mazurkiewicz equivalence constructed on the swappability relation \rightleftharpoons is an observational equivalence. Therefore, we can transform a blockchain into an observationally equivalent one by a finite number of exchanges of adjacent swappable transactions.

Theorem 1. $\simeq_{\rightleftharpoons} \subseteq \sim$. ◇

Example 3. We can rearrange the transactions in Example 1 as $T_0T_1T_2 \sim T_0T_2T_1 \sim T_2T_0T_1$. Instead, $T_1T_0T_2 \not\sim T_2T_0T_1$ (e.g., starting from a state σ such that $\sigma \mathcal{A}\text{balance} = 2$ and $\sigma \mathcal{C}x = 0$, see Fig. 6 in Appendix A of [6]). ◇

Note that the converse of Theorem 1 does not hold: indeed, $\mathbf{B} \simeq_{\rightleftharpoons} \mathbf{B}'$ requires that \mathbf{B} and \mathbf{B}' have the same length, while $\mathbf{B} \sim \mathbf{B}'$ may also hold for blockchains of different length (e.g., $\mathbf{B}' = \mathbf{B}\mathbf{T}$, where \mathbf{T} does not alter the state).

Safe Approximations of Read/Written Keys. Note that the relation \rightleftharpoons is undecidable whenever the contract language is Turing-equivalent. So, to detect swappable transactions we follow a static approach, consisting of two steps. First, we over-approximate the set of keys read and written by transactions, by statically analysing the code of the called functions. We then check a simple condition on these approximations (Definition 7), to detect if two transactions can be swapped. Since static analyses to over-approximate read and written variables are quite standard [18], here we just rely on such approximations, by only assuming their safety. In Definition 5 we state that a set P safely approximates the keys *written* by \mathbf{T} , when \mathbf{T} does not alter the state of the keys not in P . Defining set of *read* keys is a bit trickier: intuitively, we require that if we execute the transaction starting from two states that agree on the values of the keys in the read set, then these executions should be equivalent, in the sense that they do not introduce new differences between the resulting states (with respect to the difference already existing before).

Definition 5 (Safe approximation of read/written keys). *Given a set of qualified keys P and a transaction \mathbf{T} , we define:*

$$\begin{aligned}
 P \models^w \mathbf{T} & \text{ iff } \forall Q : Q \cap P = \emptyset \implies \mathbf{T} \sim_Q \epsilon \\
 P \models^r \mathbf{T} & \text{ iff } \forall \mathbf{B}, \mathbf{B}', Q : \mathbf{B} \sim_P \mathbf{B}' \wedge \mathbf{B} \sim_Q \mathbf{B}' \implies \mathbf{B}\mathbf{T} \sim_Q \mathbf{B}'\mathbf{T} \quad \diamond
 \end{aligned}$$

Example 4. Let $\mathbf{T} = \mathcal{A} \xrightarrow{1} \mathcal{C} : \mathbf{f}()$, where $\mathbf{f}()$ { \mathcal{B} . (1)} is a function of \mathcal{C} . The execution of \mathbf{T} affects the **balance** of \mathcal{A} , \mathcal{B} and \mathcal{C} ; however, $\mathcal{C}.\text{balance}$ is first incremented and then decremented, and so its value remains unchanged. Then, $\{\mathcal{A}.\text{balance}, \mathcal{B}.\text{balance}\} \models^w \mathbf{T}$, and it is the smallest safe approximation of the keys written by \mathbf{T} . To prove that $P = \{\mathcal{A}.\text{balance}\} \models^r \mathbf{T}$, assume two blockchains \mathbf{B} and \mathbf{B}' and a set of keys Q such that $\mathbf{B} \sim_P \mathbf{B}'$ and $\mathbf{B} \sim_Q \mathbf{B}'$. If $\llbracket \mathbf{B} \rrbracket \mathcal{A}\text{balance} < 1$, then by [Tx2] we have $\llbracket \mathbf{B}\mathbf{T} \rrbracket = \llbracket \mathbf{B} \rrbracket$. Since $\mathbf{B} \sim_P \mathbf{B}'$, then also $\llbracket \mathbf{B}' \rrbracket \mathcal{A}\text{balance} < 1$, and so by [Tx2] we have $\llbracket \mathbf{B}'\mathbf{T} \rrbracket = \llbracket \mathbf{B}' \rrbracket$. Then, $\mathbf{B}\mathbf{T} \sim_Q \mathbf{B}'\mathbf{T}$. Otherwise, if $\llbracket \mathbf{B} \rrbracket \mathcal{A}\text{balance} = n \geq 1$, then by [Tx1] the execution of \mathbf{T} transfers one unit of currency from \mathcal{A} to \mathcal{B} , so the execution of \mathbf{T} affects

exactly $\mathcal{A}.\text{balance}$ and $\mathcal{B}.\text{balance}$. So, it is enough to show that $\mathbf{B} \sim_{\{q\}} \mathbf{B}'$ implies $\mathbf{BT} \sim_{\{q\}} \mathbf{B}'\mathbf{T}$ for $q \in \{\mathcal{A}.\text{balance}, \mathcal{B}.\text{balance}\}$. For $q = \mathcal{A}.\text{balance}$, we have that $\llbracket \mathbf{B}'\mathbf{T} \rrbracket \mathcal{A}.\text{balance} = n - 1 = \llbracket \mathbf{BT} \rrbracket \mathcal{A}.\text{balance}$. For $q = \mathcal{B}.\text{balance}$, we have that $\llbracket \mathbf{B}'\mathbf{T} \rrbracket \mathcal{B}.\text{balance} = \llbracket \mathbf{B}' \rrbracket \mathcal{B}.\text{balance} + 1 = \llbracket \mathbf{B} \rrbracket \mathcal{B}.\text{balance} + 1 = \llbracket \mathbf{BT} \rrbracket \mathcal{B}.\text{balance}$. Therefore, we conclude that $P \models^r \mathbf{T}$. \diamond

Widening a safe approximation (either of read or written keys) preserves its safety; further, the intersection of two safe write approximations is still safe (see Lemma 6 in Appendix A of [6]). From this, it follows that there exists a *least* safe approximation of the keys written by a transaction.

Strong Swappability. We use safe approximations of the read/written keys to detect when two transactions are swappable. To achieve that, we check whether two transactions \mathbf{T} and \mathbf{T}' operate on disjoint portions of the blockchain state. More specifically, we recast in our setting Bernstein's conditions [7] for the parallel execution of processes: it suffices to check that the set of keys written by \mathbf{T} is disjoint from those written or read by \mathbf{T}' , and vice versa. When this happens we say that the two transactions are *strongly swappable*.

Definition 6 (Strong swappability). *We say that two transactions $\mathbf{T} \neq \mathbf{T}'$ are strongly swappable, in symbols $\mathbf{T} \# \mathbf{T}'$, when there exist $W, W', R, R' \subseteq \mathbb{P}$ such that $W \models^w \mathbf{T}$, $W' \models^w \mathbf{T}'$, $R \models^r \mathbf{T}$, $R' \models^r \mathbf{T}'$, and:*

$$(R \cup W) \cap W' = \emptyset = (R' \cup W') \cap W \quad \diamond$$

Example 5. Let $\mathbf{f}_1() \{\text{skip}\}$ and $\mathbf{f}_2(x) \{x. \quad (\text{value})\}$ be functions of the contracts \mathcal{C}_1 and \mathcal{C}_2 , respectively, and consider the following transactions:

$$\mathbf{T}_1 = \mathcal{A} \xrightarrow{1} \mathcal{C}_1 : \mathbf{f}_1() \quad \mathbf{T}_2 = \mathcal{B} \xrightarrow{1} \mathcal{C}_2 : \mathbf{f}_2(\mathcal{F})$$

where \mathcal{A} , \mathcal{B} , and \mathcal{F} are account addresses. To prove that $\mathbf{T}_1 \# \mathbf{T}_2$, consider the following safe approximations of the written/read keys of \mathbf{T}_1 and \mathbf{T}_2 , respectively:

$$\begin{aligned} W_1 &= \{\mathcal{A}.\text{balance}, \mathcal{C}_1.\text{balance}\} \models^w \mathbf{T}_1 & R_1 &= \{\mathcal{A}.\text{balance}\} \models^r \mathbf{T}_1 \\ W_2 &= \{\mathcal{B}.\text{balance}, \mathcal{F}.\text{balance}\} \models^w \mathbf{T}_2 & R_2 &= \{\mathcal{B}.\text{balance}\} \models^r \mathbf{T}_2 \end{aligned}$$

Since $(W_1 \cup R_1) \cap W_2 = \emptyset = (W_2 \cup R_2) \cap W_1$, the two transactions are strongly swappable. Now, let:

$$\mathbf{T}_3 = \mathcal{B} \xrightarrow{1} \mathcal{C}_2 : \mathbf{f}_2(\mathcal{A})$$

and consider the following safe approximations W_3 and R_3 :

$$W_3 = \{\mathcal{B}.\text{balance}, \mathcal{A}.\text{balance}\} \models^w \mathbf{T}_3 \quad R_3 = \{\mathcal{B}.\text{balance}\} \models^r \mathbf{T}_3$$

Since $W_1 \cap W_3 \neq \emptyset \neq W_2 \cap W_3$, then $\neg(\mathbf{T}_1 \# \mathbf{T}_3)$ and $\neg(\mathbf{T}_2 \# \mathbf{T}_3)$. \diamond

The following theorem ensures the soundness of our approximation, i.e. that if two transactions are strongly swappable, then they are also swappable. The converse implication does not hold, as witnessed by Example 6.

Theorem 2. $\top \# \top' \implies \top \rightleftharpoons \top'$. ◇

Example 6 (Swappable transactions, not strongly). Consider the following functions and transactions of a contract at address \mathcal{C} :

$$\begin{aligned} \mathbf{f}_1() \{ \text{if sender} = \mathcal{A} \quad k_1 = 0 \text{ then } k_1 := 1 \text{ else throw} \} \quad \top_1 = \mathcal{A} \xrightarrow{1} \mathcal{C} : \mathbf{f}_1() \\ \mathbf{f}_2() \{ \text{if sender} = \mathcal{B} \quad k_2 = 0 \text{ then } k_2 := 1 \text{ else throw} \} \quad \top_2 = \mathcal{B} \xrightarrow{1} \mathcal{C} : \mathbf{f}_2() \end{aligned}$$

We prove that $\top_1 \rightleftharpoons \top_2$. First, consider a state σ such that $\sigma \mathcal{A} \text{balance} > 1$, $\sigma \mathcal{B} \text{balance} > 1$, $\sigma \mathcal{C} \text{balance} = n$, $\sigma \mathcal{C} k_1 = 0$ and $\sigma \mathcal{C} k_2 = 0$. We have that:

$$\llbracket \top_1 \top_2 \rrbracket_\sigma = \sigma \{ 1/\mathcal{C}.k_1, 1/\mathcal{C}.k_2, n+2/\mathcal{C}.balance \} = \llbracket \top_2 \top_1 \rrbracket_\sigma$$

In the second case, let σ be such that $\sigma \mathcal{A} \text{balance} < 1$, or $\sigma \mathcal{B} \text{balance} < 1$, or $\sigma \mathcal{C} k_1 \neq 0$, or $\sigma \mathcal{C} k_2 \neq 0$. It is not possible that the guards in \mathbf{f}_1 and \mathbf{f}_2 are both true, so \top_1 or \top_2 raise an exception, leaving the state unaffected. Then, also in this case we have that $\llbracket \top_1 \top_2 \rrbracket_\sigma = \llbracket \top_2 \top_1 \rrbracket_\sigma$, and so \top_1 and \top_2 are swappable. However, they are *not* strongly swappable if there exist reachable states σ, σ' such that $\sigma \mathcal{C} k_1 = 0 = \sigma' \mathcal{C} k_2$. To see why, let $W_1 = \{ \mathcal{A}.balance, \mathcal{C}.balance, \mathcal{C}.k_1 \}$. From the code of \mathbf{f}_0 we see that W_1 is the least safe over-approximation of the written keys of \top_1 ($W_1 \models^w \top_1$). This means that every safe approximation of \top_1 must include the keys of W_1 . Similarly, $W_2 = \{ \mathcal{B}.balance, \mathcal{C}.balance, \mathcal{C}.k_2 \}$ is the least safe over-approximation of the written keys of \top_2 ($W_2 \models^w \top_2$). Since the least safe approximations of the keys written by \top_1 and \top_2 are not disjoint, $\top_1 \# \top_2$ does not hold. ◇

Theorem 3 states that the Mazurkiewicz equivalence $\simeq_{\#}$ is stricter than $\simeq_{\rightleftharpoons}$. Together with Theorem 1, if \mathbf{B} is transformed into \mathbf{B}' by exchanging adjacent strongly swappable transactions, then \mathbf{B} and \mathbf{B}' are observationally equivalent.

Theorem 3. $\simeq_{\#} \subseteq \simeq_{\rightleftharpoons}$. ◇

Note that if the contract language is Turing-equivalent, then finding approximations which satisfy the disjointness condition in Definition 6 is not computable, and so the relation $\#$ is undecidable.

Parameterised Strong Swappability. Strongly swappability abstracts from the actual static analysis used to obtain the safe approximations: it is sufficient that such an analysis exists. Definition 7 below parameterises strong swappability over a static analysis, which we represent as a function from transactions to sets of qualified keys, just requiring it to be a safe approximation. Formally, we say that W is a *static analysis of written keys* when $W(\top) \models^w \top$, for all \top ; similarly, R is a *static analysis of read keys* when $R(\top) \models^r \top$, for all \top .

Definition 7 (Parameterised strong swappability). *Let W and R be static analyses of written/read keys. We say that \top, \top' are strongly swappable w.r.t. W and R , in symbols $\top \#_{\#}^W \top'$, if:*

$$(R(\top) \cup W(\top)) \cap W(\top') = \emptyset = (R(\top') \cup W(\top')) \cap W(\top) \quad \diamond$$

Note that an effective procedure for computing W and R gives an effective procedure to determine whether two transactions are (strongly) swappable.

Lemma 3. *For all static analyses W and R : (i) $\#_R^W \subseteq \#$; (ii) if W and R are computable, then $\#_R^W$ is decidable. \diamond*

From the inclusion in item (i) of Lemma 3 and from Theorem 3 we obtain:

Theorem 4. $\simeq_{\#_R^W} \subseteq \simeq_{\#} \subseteq \simeq_{\Rightarrow}$. \diamond

4 True Concurrency for Blockchains

Given a swappability relation \mathcal{R} , we transform a sequence of transactions \mathbf{B} into an *occurrence net* $N_{\mathcal{R}}(\mathbf{B})$, which describes the partial order induced by \mathcal{R} . Any concurrent execution of the transactions in \mathbf{B} which respects this partial order is equivalent to the serial execution of \mathbf{B} (Theorem 5).

From Blockchains to Occurrence Nets. We start by recapping the notion of Petri net [19]. A *Petri net* is a tuple $N = (P, \text{Tr}, F, m_0)$, where P is a set of *places*, Tr is a set of *transitions* (with $P \cap \text{Tr} = \emptyset$), and $F : (P \times \text{Tr}) \cup (\text{Tr} \times P) \rightarrow \mathbb{N}$ is a *weight function*. The state of a net is a *marking*, i.e. a multiset $m : P \rightarrow \mathbb{N}$ defining how many *tokens* are contained in each place; we denote with m_0 the initial marking. The behaviour of a Petri net is specified as a transition relation between markings: intuitively, a transition t is enabled at m when each place p has at least $F(p, t)$ tokens in m . When an enabled transition t is fired, it consumes $F(p, t)$ tokens from each p , and produces $F(t, p')$ tokens in each p' . Formally, given $x \in P \cup \text{Tr}$, we define the *preset* $\bullet x$ and the *postset* $x \bullet$ as multisets: $\bullet x(y) = F(y, x)$, and $x \bullet(y) = F(x, y)$. A transition t is *enabled* at m when $\bullet t \subseteq m$. The transition relation between markings is defined as $m \xrightarrow{t} m - \bullet t + t \bullet$, where t is enabled. We say that $t_1 \cdots t_n$ is a *firing sequence from m to m'* when $m \xrightarrow{t_1} \cdots \xrightarrow{t_n} m'$, and in this case we say that m' is *reachable from m* . We say that m' is *reachable* when it is reachable from m_0 .

An *occurrence net* [8] is a Petri net such that: (i) $|\bullet p| \leq 1$ for all p ; (ii) $|\bullet p| = 1$ if $p \notin m_0$, and $|\bullet p| = 0$ if $p \in m_0$; (iii) F is a relation, i.e. $F(x, y) \leq 1$ for all x, y ; (iv) F^* is a acyclic, i.e. $\forall x, y \in P \cup \text{Tr} : (x, y) \in F^* \wedge (y, x) \in F^* \implies x = y$ (where F^* is the reflexive and transitive closure of F).

In Fig. 2 we transform a blockchain $\mathbf{B} = T_1 \cdots T_n$ into a Petri net $N_{\mathcal{R}}(\mathbf{B})$, where \mathcal{R} is an arbitrary relation between transactions. Although any relation \mathcal{R} ensures that $N_{\mathcal{R}}(\mathbf{B})$ is an occurrence net (Lemma 4 below), our main results hold when \mathcal{R} is a strong swappability relation. The transformation works as follows: the i -th transaction in \mathbf{B} is rendered as a transition (T_i, i) in $N_{\mathcal{R}}(\mathbf{B})$, and transactions related by \mathcal{R} are transformed into concurrent transitions. Technically, this concurrency is specified as a relation $<$ between transitions, such that $(T_i, i) < (T_j, j)$ whenever $i < j$, but T_i and T_j are not related by \mathcal{R} . The places, the weight function, and the initial marking of $N_{\mathcal{R}}(\mathbf{B})$ are chosen to ensure that the firing of transitions respects the relation $<$.

$$\begin{aligned}
 \text{Tr} &= \{(\mathsf{T}_i, i) \mid 1 \leq i \leq n\} \\
 \text{P} &= \{(*, \mathsf{t}) \mid \mathsf{t} \in \text{Tr}\} \cup \{(\mathsf{t}, *) \mid \mathsf{t} \in \text{Tr}\} \cup \{(\mathsf{t}, \mathsf{t}') \mid \mathsf{t} < \mathsf{t}'\} \\
 &\quad \text{where } (\mathsf{T}, i) < (\mathsf{T}', j) \triangleq (i < j) \wedge \neg(\mathsf{T} \mathcal{R} \mathsf{T}') \\
 F(x, y) &= \begin{cases} 1 & \text{if } y = \mathsf{t} \text{ and } (x = (*, \mathsf{t}) \text{ or } x = (\mathsf{t}', \mathsf{t})) \\ 1 & \text{if } x = \mathsf{t} \text{ and } (y = (\mathsf{t}, *) \text{ or } y = (\mathsf{t}, \mathsf{t}')) \\ 0 & \text{otherwise} \end{cases} \quad m_0(\mathsf{p}) = \begin{cases} 1 & \text{if } \mathsf{p} = (*, \mathsf{t}) \\ 0 & \text{otherwise} \end{cases}
 \end{aligned}$$

Fig. 2. Construction of a Petri net from a blockchain $\mathbf{B} = \mathsf{T}_1 \cdots \mathsf{T}_n$.

Example 7. Consider the following transactions and functions of a contract \mathcal{C} :

$$\begin{aligned}
 \mathsf{T}_f &= \mathcal{A} \xrightarrow{0} \mathcal{C} : \mathsf{f}() && \mathsf{f}() \{\text{if } x = 0 \text{ then } y := 1 \text{ else throw}\} \\
 \mathsf{T}_g &= \mathcal{A} \xrightarrow{0} \mathcal{C} : \mathsf{g}() && \mathsf{g}() \{\text{if } y = 0 \text{ then } x := 1 \text{ else throw}\} \\
 \mathsf{T}_h &= \mathcal{A} \xrightarrow{0} \mathcal{C} : \mathsf{h}() && \mathsf{h}() \{z := 1\}
 \end{aligned}$$

Let $P_f^w = P_g^r = \{\mathcal{C}.y\}$, $P_f^r = P_g^w = \{\mathcal{C}.x\}$, $P_h^w = \{\mathcal{C}.z\}$, $P_h^r = \emptyset$. It is easy to check that these sets are safe approximations of their transactions (e.g., P_f^w safely approximates the keys written by T_f). By Definition 6 we have that $\mathsf{T}_f \# \mathsf{T}_h$, $\mathsf{T}_g \# \mathsf{T}_h$, but $\neg(\mathsf{T}_f \# \mathsf{T}_g)$. We display $N_{\#}(\mathsf{T}_f \mathsf{T}_h \mathsf{T}_g)$ in Fig. 3, where $\mathsf{t}_f = (\mathsf{T}_f, 1)$, $\mathsf{t}_h = (\mathsf{T}_h, 2)$, and $\mathsf{t}_g = (\mathsf{T}_g, 3)$. Note that t_g can only be fired after t_f , while t_h can be fired independently from t_f and t_g . This is coherent with the fact that T_h is swappable with both T_f and T_g , while T_f and T_g are not swappable. \diamond




Fig. 3. Occurrence net for Example 7.

Lemma 4. $N_{\mathcal{R}}(\mathbf{B})$ is an occurrence net, for all \mathcal{R} and \mathbf{B} .

Step Firing Sequences. Theorem 5 below establishes a correspondence between concurrent and serial execution of transactions. Since the semantics of serial executions is given in terms of blockchain states σ , to formalise this correspondence we use the same semantics domain also for concurrent executions. This is obtained in two steps. First, we define concurrent executions of \mathbf{B} as the *step firing sequences* (i.e. finite sequences of *sets* of transitions) of $N_{\#}(\mathbf{B})$. Then, we give a semantics to step firing sequences, in terms of blockchain states.

We denote finite sets of transitions, called *steps*, as U, U', \dots . Their preset and postset are defined as $\bullet U = \sum_{p \in U} \bullet p$ and $U \bullet = \sum_{p \in U} p \bullet$, respectively. We say that U is *enabled at* m when $\bullet U \leq m$, and in this case firing U results in the move $m \xrightarrow{U} m - \bullet U + U \bullet$. Let $\mathbf{U} = U_1 \cdots U_n$ be a finite sequence of steps. We say that \mathbf{U} is a *step firing sequence from* m *to* m' if $m \xrightarrow{U_1} \dots \xrightarrow{U_n} m'$, and in this case we write $m \xrightarrow{\mathbf{U}} m'$.

Concurrent Execution of Transactions. We now define how to execute transactions in parallel. The idea is to execute transactions in *isolation*, and then merge their changes, whenever they are mutually disjoint. The state updated resulting from the execution of a transaction are formalised as in Definition 1.

An *update collector* is a function Π that, given a state σ and a transaction \mathbf{T} , gives an update $\pi = \Pi(\sigma, \mathbf{T})$ which maps (at least) the updated qualified keys to their new values. In practice, update collectors can be obtained by instrumenting the run-time environment of smart contracts, so to record the state changes resulting from the execution of transactions. We formalise update collectors abstracting from the implementation details of such an instrumentation:

Definition 8 (Update collector). *We say that a function Π is an update collector when $\llbracket \mathbf{T} \rrbracket_\sigma = \sigma(\Pi(\sigma, \mathbf{T}))$, for all σ and \mathbf{T} .* \diamond

There exists a natural ordering of collectors, which extends the ordering between state updates (i.e., set inclusion, when interpreting them as sets of substitutions): namely, $\Pi \sqsubseteq \Pi'$ holds when $\forall \sigma, \mathbf{T} : \Pi(\sigma, \mathbf{T}) \subseteq \Pi'(\sigma, \mathbf{T})$. The following lemma characterizes the least update collector w.r.t. this ordering.

Lemma 5 (Least update collector). *Let $\Pi^*(\sigma, \mathbf{T}) = \llbracket \mathbf{T} \rrbracket_\sigma - \sigma$, where we define $\sigma' - \sigma$ as $\bigcup_{\sigma' p \neq \sigma p} \{\sigma' p / p\}$. Then, Π^* is the least update collector.* \diamond

The merge of two state updates is the union of the corresponding substitutions; to avoid collisions, we make the merge operator undefined when the domains of the two updates overlap.

Definition 9 (Merge of state updates). *Let π_0, π_1 be state updates. When $\text{keys}(\pi_0) \cap \text{keys}(\pi_1) = \emptyset$, we define $\pi_0 \oplus \pi_1$ as follows:*

$$(\pi_0 \oplus \pi_1)p = \begin{cases} \pi_0 p & \text{if } p \in \text{keys}(\pi_0) \\ \pi_1 p & \text{if } p \in \text{keys}(\pi_1) \\ \perp & \text{otherwise} \end{cases} \quad \diamond$$

The merge operator enjoys the commutative monoidal laws, and can therefore be extended to (finite) sets of state updates.

We now associate step firing sequences with state updates. The semantics of a step $U = \{(T_1, 1), \dots, (T_n, n)\}$ in σ is obtained by applying to σ the merge of the updates $\Pi(\sigma, T_i)$, for all $i \in 1..n$ —whenever the merge is defined. The semantics of a step firing sequence is then obtained by folding that of its steps.

Definition 10 (Semantics of step firing sequences). We define the semantics of step firing sequences, given Π and σ , as:

$$\llbracket \epsilon \rrbracket_{\sigma}^{\Pi} = \sigma \quad \llbracket \mathbf{U} \mathbf{U} \rrbracket_{\sigma}^{\Pi} = \llbracket \mathbf{U} \rrbracket_{\sigma'}^{\Pi} \quad \text{where } \sigma' = \llbracket \mathbf{U} \rrbracket_{\sigma}^{\Pi} = \sigma \bigoplus_{(\mathbf{T}, i) \in \mathbf{U}} \Pi(\sigma, \mathbf{T}) \quad \diamond$$

Example 8. Let $\mathbf{t}_f, \mathbf{t}_g$, and \mathbf{t}_h be as in Example 7, and let $\sigma \mathcal{C}x = \sigma \mathcal{C}y = 0$. Since $\Pi^*(\sigma, \mathbf{T}_f) = \{1/e.y\}$, $\Pi^*(\sigma, \mathbf{T}_g) = \{1/e.x\}$, and $\Pi^*(\sigma, \mathbf{T}_h) = \{1/e.z\}$, we have:

$$\begin{aligned} \llbracket \{\mathbf{t}_f, \mathbf{t}_h\} \rrbracket_{\sigma}^{\Pi^*} &= \sigma(\{1/e.y\} \oplus \{1/e.z\}) = \sigma\{1/e.y, 1/e.z\} \\ \llbracket \{\mathbf{t}_g, \mathbf{t}_h\} \rrbracket_{\sigma}^{\Pi^*} &= \sigma(\{1/e.x\} \oplus \{1/e.z\}) = \sigma\{1/e.x, 1/e.z\} \\ \llbracket \{\mathbf{t}_f, \mathbf{t}_g\} \rrbracket_{\sigma}^{\Pi^*} &= (\sigma\{1/e.y\} \oplus \{1/e.x\}) = \sigma\{1/e.y, 1/e.x\} \end{aligned}$$

Note that, for all σ :

$$\begin{aligned} \llbracket \mathbf{T}_f \mathbf{T}_h \rrbracket_{\sigma} &= \llbracket \mathbf{T}_h \mathbf{T}_f \rrbracket_{\sigma} = \sigma\{1/e.y, 1/e.z\} = \llbracket \{\mathbf{t}_f, \mathbf{t}_h\} \rrbracket_{\sigma}^{\Pi^*} \\ \llbracket \mathbf{T}_g \mathbf{T}_h \rrbracket_{\sigma} &= \llbracket \mathbf{T}_h \mathbf{T}_g \rrbracket_{\sigma} = \sigma\{1/e.x, 1/e.z\} = \llbracket \{\mathbf{t}_g, \mathbf{t}_h\} \rrbracket_{\sigma}^{\Pi^*} \end{aligned}$$

So, the serial execution of \mathbf{T}_f and \mathbf{T}_h (in both orders) is equal to their concurrent execution (similarly for \mathbf{T}_g and \mathbf{T}_h). Instead, for all σ such that $\sigma \mathcal{C}x = \sigma \mathcal{C}y = 0$:

$$\llbracket \mathbf{T}_f \mathbf{T}_g \rrbracket_{\sigma} = \sigma\{1/e.y\} \quad \llbracket \mathbf{T}_g \mathbf{T}_f \rrbracket_{\sigma} = \sigma\{1/e.x\} \quad \llbracket \{\mathbf{t}_f, \mathbf{t}_g\} \rrbracket_{\sigma}^{\Pi^*} = \sigma\{1/e.y, 1/e.x\}$$

So, concurrent executions of \mathbf{T}_f and \mathbf{T}_g may differ from serial ones. This is coherent with the fact that, in Fig. 3, \mathbf{t}_f and \mathbf{t}_g are *not* concurrent. \diamond

Concurrent Execution of Blockchains. Theorem 5 relates serial executions of transactions to concurrent ones (which are rendered as step firing sequences). Item (a) establishes a confluence property: if two step firing sequences lead to the same marking, then they also lead to the same blockchain state. Item (b) ensures that the blockchain, interpreted as a sequence of transitions, is a step firing sequence, and it is *maximal* (i.e., there is a bijection between the transactions in the blockchain and the transitions of the corresponding net). Finally, item (c) ensures that executing maximal step firing sequences is equivalent to executing serially the blockchain.

Theorem 5. Let $\mathbf{B} = \mathbf{T}_1 \cdots \mathbf{T}_n$. Then, in $\mathbf{N}_{\#}(\mathbf{B})$:

- (a) if $m_0 \xrightarrow{\mathbf{U}} m$ and $m_0 \xrightarrow{\mathbf{U}'} m$, then $\llbracket \mathbf{U} \rrbracket_{\sigma}^{\Pi^*} = \llbracket \mathbf{U}' \rrbracket_{\sigma}^{\Pi^*}$, for all reachable σ ;
- (b) $\{(\mathbf{T}_1, 1)\} \cdots \{(\mathbf{T}_n, n)\}$ is a maximal step firing sequence;
- (c) for all maximal step firing sequences \mathbf{U} , for all reachable σ , $\llbracket \mathbf{U} \rrbracket_{\sigma}^{\Pi^*} = \llbracket \mathbf{B} \rrbracket_{\sigma}$.

Remarkably, the implications of Theorem 5 also apply to $\mathbf{N}_{\#R}^W(\mathbf{B})$.

Example 9. Recall $\mathbf{B} = \mathbf{T}_f \mathbf{T}_h \mathbf{T}_g$ and $\mathbf{N}_{\#}(\mathbf{B})$ from Example 7, let $\mathbf{U} = \{\mathbf{t}_f, \mathbf{t}_h\} \{\mathbf{t}_g\}$, and let σ be such that $\sigma \mathcal{C}x = \sigma \mathcal{C}y = 0$. As predicted by item (c) of Theorem 5:

$$\llbracket \mathbf{B} \rrbracket_{\sigma} = \sigma\{1/e.y\} \{1/e.z\} = \llbracket \mathbf{U} \rrbracket_{\sigma}^{\Pi^*}$$

Let $\mathbf{U}' = \{t_f\}\{t_g, t_h\}$. We have that \mathbf{U} and \mathbf{U}' lead to the same marking, where the places $(t_f, *)$, $(t_g, *)$ and $(t_h, *)$ contain one token each, while the other places have no tokens. By item (a) of Theorem 5 we conclude that $\llbracket \mathbf{U} \rrbracket_\sigma^{II^*} = \llbracket \mathbf{U}' \rrbracket_\sigma^{II^*}$. Now, let $\mathbf{U}'' = \{t_h\}\{t_f, t_g\}$. Note that, although \mathbf{U}'' is maximal, it is not a step firing sequence, since the second step is not enabled (actually, t_f and t_g are not concurrent, as pointed out in Example 8). Therefore, the items of Theorem 5 do not apply to \mathbf{U}'' , coherently with the fact that \mathbf{U}'' does not represent any sequential execution of \mathbf{B} . \diamond

5 Case Study: ERC-721 Token

We now apply our theory to an archetypal Ethereum smart contract, which implements a “non-fungible token” following the standard ERC-721 interface [14, 15]. This contract defines the functions to transfer tokens between users, and to delegate their trade to other users. Currently, token transfers involve ~50% of the transactions on the Ethereum blockchain [1], with larger peaks due to popular contracts like Cryptokitties [22].

We sketch below the implementation of the contract, using Solidity, the main high-level smart contract language in Ethereum (see Appendix B of [6] for the full implementation).

The contract state is defined by the following mappings:

Each token is uniquely identified by an integer value (of type `uint256`), while users are identified by an `address`. The mapping `tokenToOwner` maps tokens to their owners’ addresses (the zero address is used to denote a dummy owner). The mapping `tokenCreated` tells whether a token has been created or not, while `tokenBalance` gives the number of tokens owned by each user. The mapping `delegate` allows a user to delegate the transfer of all her tokens to third parties.

The function `transfer` transfers a token from the owner to another user. The `transfer` assertion rules out some undesirable cases, e.g., if the token does not exist, or it is not owned by the `from` user, or the user attempts to transfer the token to himself. Once all these checks are passed, the transfer succeeds if the `to` of the transaction owns the token, or if he has been delegated by the owner. The mappings `tokenToOwner` and `tokenBalance` are updated as expected.

The function `delegates` delegates the transfers of all the tokens of the `to` to the `when` when the boolean `is true`, otherwise it revokes the delegation.

Assume that user \mathcal{A} owns two tokens, identified by the integers 1 and 2, and consider the following transactions:

$$\begin{aligned}
 T_1 &= \mathcal{A} \xrightarrow{0} \text{Token} : \text{transferFrom}(\mathcal{A}, \mathcal{P}, 1) \\
 T_2 &= \mathcal{A} \xrightarrow{0} \text{Token} : \text{setApprovalForAll}(\mathcal{B}, \text{true}) \\
 T_3 &= \mathcal{B} \xrightarrow{0} \text{Token} : \text{transferFrom}(\mathcal{A}, \mathcal{Q}, 2) \\
 T_4 &= \mathcal{P} \xrightarrow{0} \text{Token} : \text{transferFrom}(\mathcal{P}, \mathcal{B}, 1)
 \end{aligned}$$

We have that $T_1 \# T_2$, $T_2 \# T_4$, and $T_3 \# T_4$ (this can be proved e.g. by using the static approximations in Appendix B of [6]), while the other combinations are not swappable. Let $\mathbf{B} = T_1 T_2 T_3 T_4$. The resulting occurrence net is displayed in Fig. 4. For instance, let $\mathbf{U} = \{T_1, T_2\} \{T_3, T_4\}$, i.e. T_1 and T_2 are executed concurrently, as well as T_3 and T_4 . From item (c) of Theorem 5 we have that this concurrent execution is equivalent to the serial one.

Although this example deals with the marginal case where the sender and the receiver of tokens overlap, in practice the large majority of transactions in a block either involves distinct users, or invokes distinct ERC-721 interfaces, making it possible to increase the degree of concurrency of transactions.




Fig. 4. Occurrence net for the blockchain $\mathbf{B} = T_1 T_2 T_3 T_4$ of the ERC-721 token.

6 Related Work and Conclusions

We have proposed a static approach to improve the performance of blockchains by concurrently executing transactions. We have started by introducing a model of transactions and blockchains. We have defined two transactions to be *swappable* when inverting their order does not affect the blockchain state. We have then introduced a static approximation of swappability, based on a static analysis of the sets of keys read/written by transactions. We have rendered concurrent executions of a sequence of transactions as *step firing sequences* in the associated occurrence net. Our main technical result, Theorem 5, shows that these concurrent executions are semantically equivalent to the sequential one.

We can exploit our results in practice to improve the performances of miners and validators. Miners should perform the following steps to mine a block:

1. gather from the network a set of transactions, and put them in an arbitrary linear order \mathbf{B} , which is the mined block;
2. compute the relation $\#_R^W$ on \mathbf{B} , using a static analysis of read/written keys;
3. construct the occurrence net $N_{\#_R^W}(\mathbf{B})$;
4. execute transactions concurrently according to the occurrence net, exploiting the available parallelism.

The behaviour of validators is almost identical to that of miners, except that in step (1), rather than choosing the order of transactions, they should adhere to the ordering of the mined block \mathbf{B} . Note that in the last step, validators can execute any maximal step firing sequence which is coherent with their degree of parallelism: item (c) of Theorem 5 ensures that the resulting state is equal to the state obtained by the miner. The experiments in [12] suggest that parallelization may lead to a significant improvement of the performance of nodes: the benchmarks on a selection of representative contracts show an overall speedup of $1.33\times$ for miners and $1.69\times$ for validators, using only three cores.

Note that malevolent users could attempt a denial-of-service attack by publishing contracts which are hard to statically analyse, and therefore are not suitable for parallelization. This kind of attacks can be mitigated by adopting a mining strategy that gives higher priority to parallelizable transactions.

Applying Our Approach to Ethereum. Applying our theory to Ethereum would require a static analysis of read/written keys at the level of EVM bytecode. As far as we know, the only tool implementing such an analysis is ES-ETH [16]. However, the current version of the tool has several limitations, like e.g. the compile-time approximation of dictionary keys and of values shorter than 32 bytes, which make ES-ETH not directly usable to the purposes of our work. In general, precise static analyses at the level of the Ethereum bytecode are difficult to achieve, since the language has features like dynamic dispatching and pointer aliasing which are notoriously a source of imprecision for static analysis. However, coarser approximations of read/written keys may be enough to speed-up the execution of transactions. For instance, in Ethereum, blocks typically contain many transactions which transfer tokens between participants, and many

of them involve distinct senders and receivers. A relatively simple analysis of the code of token contracts (which is usually similar to that in Sect. 5) may be enough to detect that these transactions are swappable.

Aiming at minimality, our model does not include the *gas mechanism*, which is used in Ethereum to pay miners for executing contracts. The sender of a transaction deposits into it some crypto-currency, to be paid to the miner which appends the transaction to the blockchain. Each instruction executed by the miner consumes part of this deposit; when the deposit reaches zero, the miner stops executing the transaction. At this point, all the effects of the transaction (except the payment to the miner) are rolled back. Our transaction model could be easily extended with a gas mechanism, by associating a cost to statements and recording the gas consumption in the environment. Remarkably, adding gas does not invalidate approximations of read/written keys which are correct while neglecting gas. However, a gas-aware analysis may be more precise of a gas-oblivious one: for instance, in the statement `if k then $f_{long}()$; $x:=1$ else $y:=1$` (where f_{long} is a function which exceeds the available gas) a gas-aware analysis would be able to detect that x is not written.

Related Work. A few works study how to optimize the execution of smart contracts on Ethereum, using dynamic techniques adopted from software transactional memory [4, 12, 13]. These works are focussed on empirical aspects (e.g., measuring the speedup obtained on a given benchmark), while we focus on the theoretical counterpart. In [12, 13], miners execute a set of transactions speculatively in parallel, using abstract locks and inverse logs to dynamically discover conflicts and to recover from inconsistent states. The obtained execution is guaranteed to be equivalent to a serial execution of the same set of transactions. The work [4] proposes a conceptually similar technique, but based on optimistic software transactional memory. Since speculative execution is non-deterministic, in both approaches miners need to communicate the chosen schedule of transactions to validators, to allow them to correctly validate the block. This schedule must be embedded in the mined block: since Ethereum does not support this kind of block metadata, these approaches would require a “soft-fork” of the blockchain to be implemented in practice. Instead, our approach is compatible with the current Ethereum, since miners only need to append transactions to the blockchain. Compared to [4, 12], where conflicts are detected dynamically, our approach relies on a static analysis to detect potential conflicts. Since software transactional memory introduces a run-time overhead, in principle a static technique could allow for faster executions, at the price of a preprocessing phase. Saraph and Herlihy [20] study the effectiveness of speculatively executing smart contracts in Ethereum. They sample past blocks of transactions (from July 2016 to December 2017), replay them by using a speculative execution engine, and measure the speedup obtained by parallel execution. Their results show that simple speculative strategies yield non-trivial speed-ups. Further, they note that many of the data conflicts (i.e. concurrent read/write accesses to the same state location) arise in periods of high traffic, and they are caused by a small number of popular contracts, like e.g. tokens.

In the permissioned setting, Hyperledger Fabric [3] follows the “execute first and then order” paradigm: transactions are executed speculatively, and then their ordering is checked for correctness [2]. In this paradigm, appending a transaction requires a few steps. First, a client proposes a transaction to a set of “endorsing” peers, which simulate the transaction without updating the blockchain. The output of the simulation includes the state updates of the transaction execution, and the sets of read/written keys. These sets are then signed by the endorsing peers, and returned to the client, which submits them to the “ordering” peers. These nodes order transactions in blocks, and send them to the “committing” peers, which validate them. A block $T_1 \cdots T_n$ is valid when, if a key k is read by transaction T_i , then k has not been written by a transaction T_j with $j < i$. Finally, validated blocks are appended to the blockchain. Our model is coherent with Ethereum, which does not support speculative execution.

Future Works. A relevant line of research is the design of domain-specific languages for smart contracts that are directly amenable to techniques that, like ours, increase the degree of concurrency of executions. For this purpose, the language should support static analyses of read/written keys, like the one we use to define the strong swappability relation. Although the literature describes various static analyses of smart contracts, most of them are focussed on finding security vulnerabilities, rather than enhancing concurrency.

Outside the realm of smart contracts, a few papers propose static analyses of read/written variables. The paper [11] describes an analysis based on separation logic, and applies it to resolve conflicts in the setting of *snapshot isolation* for transactional memory in Java. When a conflict is detected, the read/write sets are used to determine how the code can be modified to resolve it. The paper [10] presents a static analysis to infer read and write locations in a C-like language with atomic sections. The analysis is used to translate atomic sections into standard lock operations. The design of new smart contract languages could take advantage of these analyses.

Acknowledgements. Massimo Bartoletti is partially supported by Aut. Reg. Sardinia projects “*Smart collaborative engineering*” and “*Sardcoin*”. Letterio Galletta is partially supported by IMT Lucca project “*PAI VeriOSS*” and by MIUR project PRIN 2017FTXR7S “*Methods and Tools for Trustworthy Smart Systems*”. Maurizio Murgia is partially supported by MIUR PON “*Distributed Ledgers for Secure Open Communities*” and by Aut. Reg. Sardinia project “*Smart collaborative engineering*”.

References

1. Ethereum token dynamics. <https://stat.bloxy.info/superset/dashboard/tokens>
2. Hyperledger fabric: read-write set semantics. <https://hyperledger-fabric.readthedocs.io/en/release-1.4/readwrite.html>
3. Androulaki, E., et al.: Hyperledger fabric: a distributed operating system for permissioned blockchains. In: EuroSys, pp. 30:1–30:15 (2018). <https://doi.org/10.1145/3190508.3190538>

4. Anjana, P.S., Kumari, S., Peri, S., Rathor, S., Somani, A.: An efficient framework for optimistic concurrent execution of smart contracts. In: PDP, pp. 83–92 (2019). <https://doi.org/10.1109/EMPDP.2019.8671637>
5. Bartoletti, M., Galletta, L., Murgia, M.: A minimal core calculus for solidity contracts. In: Pérez-Solà, C., Navarro-Arribas, G., Biryukov, A., Garcia-Alfaro, J. (eds.) DPM/CBT -2019. LNCS, vol. 11737, pp. 233–243. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-31500-9_15
6. Bartoletti, M., Galletta, L., Murgia, M.: A true concurrent model of smart contracts executions. CoRR abs/1905.04366 (2020). <http://arxiv.org/abs/1905.04366>
7. Bernstein, A.J.: Analysis of programs for parallel processing. IEEE Trans. Electron. Comput. **EC-15**(5), 757–763 (1966). <https://doi.org/10.1109/PGEC.1966.264565>
8. Best, E., Devillers, R.R.: Sequential and concurrent behaviour in Petri net theory. Theor. Comput. Sci. **55**(1), 87–136 (1987). [https://doi.org/10.1016/0304-3975\(87\)90090-9](https://doi.org/10.1016/0304-3975(87)90090-9)
9. Buterin, V.: Ethereum: a next generation smart contract and decentralized application platform (2013). <https://github.com/ethereum/wiki/wiki/White-Paper>
10. Cherem, S., Chilimbi, T.M., Gulwani, S.: Inferring locks for atomic sections. In: ACM SIGPLAN PLDI, pp. 304–315 (2008). <https://doi.org/10.1145/1375581.1375619>
11. Dias, R.J., Lourenço, J.M., Preguiça, N.M.: Efficient and correct transactional memory programs combining snapshot isolation and static analysis. In: USENIX Conference on Hot Topics in Parallelism (HotPar) (2011)
12. Dickerson, T.D., Gazzillo, P., Herlihy, M., Koskinen, E.: Adding concurrency to smart contracts. In: ACM PODC, pp. 303–312 (2017). <https://doi.org/10.1145/3087801.3087835>
13. Dickerson, T., Gazzillo, P., Herlihy, M., Koskinen, E.: Adding concurrency to smart contracts. Distrib. Comput. **33**, 209–225 (2020). <https://doi.org/10.1007/s00446-019-00357-z>
14. Entriken, W., Shirley, D., Evans, J., Sachs, N.: EIP 721: ERC-721 non-fungible token standard. <https://eips.ethereum.org/EIPS/eip-721>
15. Fröwis, M., Fuchs, A., Böhme, R.: Detecting token systems on ethereum. In: Goldberg, I., Moore, T. (eds.) FC 2019. LNCS, vol. 11598, pp. 93–112. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32101-7_7
16. Marcia, D.: ES-ETH: ethereum state change examiner (2019). <https://github.com/DiegoMarcia/ES-ETH>
17. Mazurkiewicz, A.: Basic notions of trace theory. In: de Bakker, J.W., de Roever, W.-P., Rozenberg, G. (eds.) REX 1988. LNCS, vol. 354, pp. 285–363. Springer, Heidelberg (1989). <https://doi.org/10.1007/BFb0013025>
18. Nielson, F., Nielson, H.R., Hankin, C.: Principles of Program Analysis (1999). <https://doi.org/10.1007/978-3-662-03811-6>
19. Reisig, W.: Petri Nets: An Introduction. Monographs in Theoretical Computer Science: An EATCS Series, vol. 4. Springer, Heidelberg (1985). <https://doi.org/10.1007/978-3-642-69968-9>
20. Saraph, V., Herlihy, M.: An empirical study of speculative concurrency in Ethereum smart contracts. CoRR abs/1901.01376 (2019). <http://arxiv.org/abs/1901.01376>
21. Szabo, N.: Formalizing and securing relationships on public networks. First Monday **2**(9) (1997). <http://firstmonday.org/htbin/cgiwrap/bin/ojs/index.php/fm/article/view/548>
22. Young, J.: CryptoKitties sales hit \$12 million, could be Ethereum’s killer app after all (2017). <https://cointelegraph.com/news/cryptokitties-sales-hit-12-million-could-be-ethereums-killer-app-after-all>