Cryptoeconomic Systems

The Transaction Graph for Modeling Blockchain Semantics

**Christian Cachin¹, Angelo De Caro², Pedro Moreno-Sanchez³,
Björn Tackmann⁴, Marko Vukolic²**

¹University of Bern (Work done while at IBM Research - Zurich), ²IBM Research - Zurich,

³IMDEA Software Institute (Work done while at IBM Research - Zurich),

⁴DFINITY (Work done while at IBM Research - Zurich)

Published on: Dec 02, 2020

License: [Creative Commons Attribution 4.0 International License \(CC-BY 4.0\)](https://creativecommons.org/licenses/by/4.0/)

Abstract

The advent of Bitcoin paved the way for a plethora of blockchain systems supporting diverse applications beyond cryptocurrencies. Although in-depth studies of the consensus protocols as well as the privacy of blockchain transactions are available, there is no formal model of the transaction semantics that a blockchain is supposed to guarantee.

In this work, we fill this gap, motivated by the observation that the semantics of transactions in blockchain systems can be captured by a directed acyclic graph. Such a transaction graph, or TDAG, generally consists of the states and the transactions as transitions between the states, together with conditions for the consistency and validity of transactions. We instantiate the TDAG model for three prominent blockchain systems: Bitcoin, Ethereum, and Hyperledger Fabric. We specify the states and transactions as well as the validity conditions of the TDAG for each one. This demonstrates the applicability of the model and formalizes the transaction-level semantics that these systems aim for.

1. Introduction

The success of Bitcoin [1] has sparked the development of many other blockchain systems. Whereas the first blockchains after Bitcoin (called *alt-coins*) resembled the cryptocurrency functionality offered by Bitcoin and mostly differed in the choice of certain parameters, Ethereum [2] was the pioneer of so-called *smart contract* systems that support arbitrary (deterministic) computation on the blockchain. Platforms for running smart contracts are seen to be of wide-spread interest for replacing trusted parties, whether in public blockchains where participation is open to anyone or in private blockchains inside a consortium.

Many recent blockchain platforms run generic computations, model specific asset classes, or add cryptographic privacy guarantees; prominent systems today include Hyperledger Fabric [3], R3 Corda [4], Tendermint/Cosmos [5], and Chain Core [6].

Blockchain systems have attracted attention not only from industry but also from academia. Many works have analyzed blockchains from different perspectives, for example, focusing on the underlying consensus protocols [7][8][9], their privacy guarantees [10][11][12], and many more aspects. This collection is necessarily partial; excellent surveys exist in the literature [13][14][15][16].

What is, surprisingly, missing to date is a formal model of the semantics of a blockchain, addressing the transaction-level consistency guarantees that they aim to achieve. These guarantees are intuitive and easy to grasp in the context of Bitcoin: given a proper modeling of the mining of new coins, the overall amount of bitcoins must remain invariant. For the newer, generic, and more complex blockchains, such as Ethereum or Hyperledger Fabric, a proper model of the guarantees they provide appears necessary. For instance, such a model should allow for reasoning whether the intuitively expected guarantees are indeed achieved. It should also model the operation of a blockchain at an appropriate level, such that the properties of a system appear concisely and differences across platforms become visible. In particular, it has to describe the criteria that determine whether a transaction that manipulates state is considered valid and consequently executed by the nodes.

Our contributions

We introduce a formal model, called the *transaction graph* or *TDAG* for short, a directed acyclic graph that models the *transactions* occurring on a blockchain and how they interact through *states*. In a nutshell, a TDAG is a graph consisting of transactions that link states to each other. Each transaction may consume, observe, or produce states, and occurs only with respect to an external input that triggers the transaction. The model abstracts the transaction validation into a predicate that can be evaluated locally in the graph, in the sense that validation only considers the relevant states. This corresponds to how many blockchains work, during the process of transaction validation and consensus, which must be efficient and based on local state. The TDAG is a generic model to encode properties expected from every blockchain system, such as notions of *validity* and *consistency*, and for characterizing the *invariants* that must be enforced in a blockchain.

We instantiate the TDAG model for three different prominent blockchains: Bitcoin, Ethereum, and Hyperledger Fabric. For each system, we describe the states and transactions of the TDAG, specify the notion of consistency, and define the validity of transactions. This shows the broad applicability of our model, and results in an abstract description of these real-world systems.

Related work

Atzei *et al.* provided a formal transaction model for Bitcoin [17]. While their model covers certain aspects of Bitcoin (like scripts and multi-signature) in more detail than ours, it does not allow one to model and compare with other blockchain systems. BitML [18] develops a formal language for smart contracts based on Bitcoin. Like our

work, it is more abstract than [17], but the scope is different: While BitML targets modeling and describing concrete contracts, the TDAG model aims to model the consistency of the platform as a whole. The TDAG can be seen as a refinement of the *precedence graph* (or *serialization graph*) from database concurrency theory [19], which relates transactions with conflicting data access. Additionally, the TDAG contains states as vertices, as one goal of the TDAG (besides formalizing conflicts) is to make statements about the consistency of the states.

2. Transaction graphs

This section introduces the *transaction directed acyclic graph*, abbreviated *transaction graph* or *TDAG* for representing the semantics of a blockchain. It models the *context* held by the blockchain and its evolution through *transactions* that obey *validation* rules.

We start by introducing some notation. Let $\mathcal{E} \subseteq \mathcal{X} \times \mathcal{Y}$ be a relation between sets $\mathcal{X}$ and $\mathcal{Y}$. For the predicate $(x, y) \in \mathcal{E}$, we also write $x\mathcal{E}y$. Furthermore, we denote the set $\{y : x\mathcal{E}y\}$ by $x\mathcal{E}\star$ and its size by $|x\mathcal{E}\star|$.

2.1 Definition

A *transaction graph* or TDAG is a directed acyclic graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. The vertices $\mathcal{V}$ can be partitioned into *states* $\mathcal{S}$ and *witnesses* $\mathcal{W}$, that is, $\mathcal{V} = \mathcal{S} \dot{\cup} \mathcal{W}$. At a high level, the edges $\mathcal{E}$ represent transitions between states. More precisely, an edge $e \in \mathcal{E}$ represents the relation between a state and a witness in the context of a transaction, and an edge may connect a state to a witness or vice versa. The edges can be partitioned into *consuming*, *observing*, and *producing* edges, denoted $\mathcal{E}_C$, $\mathcal{E}_O$, and $\mathcal{E}_P$, respectively, such that $\mathcal{E} = \mathcal{E}_C \dot{\cup} \mathcal{E}_O \dot{\cup} \mathcal{E}_P$. We now introduce the elements of $\mathcal{G}$ informally.

- **States** $\bigcirc$. The first type of vertex, $s \in \mathcal{S}$, denotes an atomic *state* represented by the blockchain and is depicted by a circle $\bigcirc$. It models an individual asset, a digital coin, some coins controlled by a particular cryptographic key, a variable of a smart contract at a moment in time, and so on. The complete context of the blockchain consists of all states that exist at a particular time. A state results from a transaction on the blockchain and can transition to other states through a transaction. There is a special *genesis state* $s_g \in \mathcal{S}$, which represents the initial state of the blockchain. There is a single genesis state by intention because the blockchain system is initialized exactly once.
- **Witnesses** $\square$. The second kind of vertex, $w \in \mathcal{W}$, denotes a *witness* in the context of a transaction and is depicted by a rectangle $\square$. It represents any data included in a transaction that is required for the transaction to be valid according to the validation

rules of the blockchain system. Every transaction of the blockchain system contains exactly one witness.

- **Consuming edges** $\bigcirc \longrightarrow \square$. A *consuming edge* $e \in \mathcal{E}_C$ connects a state to a witness and models that the state $\bigcirc$ is *consumed* by the transaction that involves witness $\square$, *i.e.*, the unique transaction that corresponds to $\square$. A state can be consumed exactly once, *i.e.*, it is not available for being consumed by another transaction once it has been consumed. Consuming a state means that the state is “updated” or “overwritten” by the transaction.
- **Observing edges** $\bigcirc - - \rightarrow \square$. An *observing edge* $e \in \mathcal{E}_O$ also connects a state to a witness; it models that the state enters into the transaction represented by the witness, but that it remains available for consumption by another transaction. A state can be observed by many transactions, even if it is also consumed. Intuitively a transaction that observes a state “reads” it.
- **Producing edges** $\square \longrightarrow \bigcirc$. A *producing edge* $e \in \mathcal{E}_P$ connects a witness to a state, and denotes that the state is *created* or *produced* by the transaction corresponding to the witness. Every state apart from the genesis state is produced exactly once.

With these notions, a transaction represents a transition from one state, or from some set of states, in a TDAG to another set of states according to the blockchain system. The transaction is linked to a unique witness, which makes it “valid” as described later. We say that a transaction has *input states* that are consumed or observed by the transaction and *output states* that are produced by the transaction. More formally, a transaction is also a weakly connected DAG, *i.e.*, a DAG that is connected as a graph.

Definition 2.1 (*Transaction*)

A weakly connected DAG $T = (\mathcal{V}, \mathcal{E})$ with a set of *input states* $\mathcal{S}_I$, a set of *output states* $\mathcal{S}_O$, and a witness w is called a *transaction* whenever

- Every input state in $\mathcal{S}_I$ is a source (has indegree zero);
- Every output state in $\mathcal{S}_O$ is a sink (has outdegree zero);
- $\mathcal{V} = \mathcal{S}_I \dot{\cup} \mathcal{S}_O \dot{\cup} \{w\}$;
- Every edge in $\mathcal{E}$ is either a consuming edge or an observing edge and links some input state $s_i \in \mathcal{S}_I$ to w , or it is a producing edge and links w to some output state $s_o \in \mathcal{S}_O$.

The transaction as defined above does not directly correspond to what is referred to as *transaction* in many blockchain platforms, and referred to as *witness* here:

[Definition 2.1](#) references the input states used in the execution of the transaction. In many practical systems such as Ethereum, however, the input states used in the execution depend on where the transaction is included in the blockchain, as the state may be modified by preceding transactions. [Definition 2.1](#) allows one to model the local validity conditions for transactions, which necessarily involve the input and output states.

As the name suggests, a transaction graph contains many transactions.

Definition 2.2 (TDAG)

A *transaction graph (TDAG)* is a directed unweighted graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \mathcal{S} \dot{\cup} \mathcal{W}$ are the vertices and $\mathcal{E} = \mathcal{E}_C \dot{\cup} \mathcal{E}_O \dot{\cup} \mathcal{E}_P$ are the edges. The set $\mathcal{S}$ denotes the states and contains a special state s_g called *genesis*. The set $\mathcal{W}$ denotes the witnesses. Edges are partitioned into three subsets, where $\mathcal{E}_C \subseteq \mathcal{S} \times \mathcal{W}$ denotes *consuming edges*, $\mathcal{E}_O \subseteq \mathcal{S} \times \mathcal{W}$ denotes *observing edges*, and $\mathcal{E}_P \subseteq \mathcal{W} \times \mathcal{S}$ denotes the *producing edges*.

It satisfies the following conditions:

1. s_g does not have any producing or observing edges and it has a single consuming edge, *i.e.*, $|\star \mathcal{E}_P s_g| = 0 \wedge |s_g \mathcal{E}_O \star| = 0 \wedge \exists! w \in \mathcal{W} : s_g \mathcal{E}_C w$.
2. Every state except for the genesis state has exactly one producing edge, *i.e.*, $\forall s \in \mathcal{S} \setminus \{s_g\} \exists! w \in \mathcal{W} : w \mathcal{E}_P s$.
3. Every state except for the genesis state may have multiple successors, but at most one among them is connected with a consuming edge, *i.e.*, $\forall s \in \mathcal{S} : |s \mathcal{E}_C \star| \leq 1$.
4. $\mathcal{G}$ is weakly connected.
5. $\mathcal{G}$ has no cycles.

The consuming and observing edges incident to a state are also called the *outgoing edges* of that state. Similarly, the consuming and observing edges incident to a witness are called *incoming edges* of that witness. The producing edges of a witness are *outgoing edges* of the witness. There is no order among the edges incident to a vertex in a TDAG. The set of all *unconsumed* states in a TDAG are the states without an incident consuming edge.

In a TDAG, every witness w corresponds to a unique transaction $t(w)$. The next definition follows naturally and is easily seen to be equivalent to [Definition 2.1](#).

Definition 2.3 (Transaction in a TDAG)

Given a TDAG $\mathcal{G} = (\mathcal{S} \dot{\cup} \mathcal{W}, \mathcal{E})$ and a witness $w \in \mathcal{W}$, the *transaction* with witness w is the unique subgraph $t = (\mathcal{S}' \dot{\cup} \{w\}, \mathcal{E}') \subseteq \mathcal{G}$, where

- $w \in \mathcal{W}$ is the *witness* of the transaction;
- $\mathcal{S}'$ is the set of states connected to w , i.e., $\mathcal{S}' = \{s \in \mathcal{S} : s\mathcal{E}_C w \vee s\mathcal{E}_O w \vee w\mathcal{E}_P s\}$; and
- $\mathcal{E}'$ are the edges with both endpoints in $\mathcal{S}' \dot{\cup} \{w\}$.

The *input states* of $t(w)$ are the states being observed or consumed by $t(w)$, and the *output states* of $t(w)$ are the states being produced by $t(w)$. With this terminology, a transaction $t \subseteq \mathcal{G}$ can have one of the following five types, which depend mostly on the number of input and output states:

- **INIT:** A unique *initialization transaction* exists in every non-empty TDAG, consisting of a consuming edge that links the genesis state to a witness w , and a set of producing edges that link w to a set of states.
- **SISO:** A *single-input, single-output transaction* consists of one consuming edge that links one input state to a witness w , and one producing edge that links w to an output state.
- **SIMO:** A *single-input, multi-output transaction* consists of one consuming edge that links an input state s to a witness w , and a set of producing edges that link w to a set of output states.
- **MISO:** A *multi-input, single-output transaction* contains a set of multiple consuming and observing edges that link distinct input states to a witness w , and one producing edge that links w to an output state.
- **MIMO:** A *multi-input, multi-output transaction* contains a set of multiple consuming and observing edges that link distinct input states to a witness w , and a set of producing edges that link w to a set of output states.

[Figure 1](#) shows the possible transaction types in a TDAG. The initialization transaction plays a special role; it represents the creation of the blockchain, which typically creates all assets represented by the states. Modeling initialization through a specific transaction is a deliberate design choice that will become clear later, in the context of transaction validation. The other types represent “ordinary” transactions that consume (and possibly observe) one or more states and produce one or more states. We note that SISO and SIMO transactions have a single input state and have no observing edges. This models that a transaction must update or overwrite at least one state for it

to make sense of being included in the blockchain, as simple read queries can be handled by inspecting the blockchain.

Figure 1: Graphical representation of transactions. States are represented by circles and witnesses are represented by boxes. Two concentric circles represent the genesis state. Observing edges are represented with a dashed arrow whereas producing edges and consuming edges are represented with solid arrows.

For the moment, it suffices to say that the initialization transaction typically creates all “assets” modeled by the blockchain or the “states” that it holds, setting them to a predefined value. This allows a subsequent transaction to be linked only with the states to which it refers and that it consumes. Otherwise, all transactions that modify any state would be linked from the genesis state (with a consuming edge), contrary to the condition that every state has at most one consuming edge. We consider this an important property of the TDAG model. A further argument for modeling only one initialization transaction goes as follows: if there were multiple INIT transactions, then it would not be easily possible to assess whether one INIT transaction is “valid” without looking also at the other ones. For instance, an INIT transaction that creates a new asset is only valid if no other INIT transaction has created the same asset beforehand.

Therefore, we purposely restrict the model so that it has a single initialization transaction for simplicity, but without loss of generality as this unique initialization transaction can create as many states as required throughout the lifetime of the blockchain.

[Figure 2](#) shows an illustrative example of a TDAG modeling a Bitcoin execution with four transactions. First, $t_0(w_0)$ represents the creation of the Bitcoin blockchain by minting all available bitcoins into a Bitcoin address containing unmined bitcoins (s_0). Here, w_0 represents the Bitcoin creation rules. Second, $t_1(w_1)$ represents a transaction that transfers some unmined bitcoins (s_0) to the Bitcoin address of a user u that successfully mined the first Bitcoin block (s_2); $t_1(w_1)$ saves the remaining unmined bitcoins (s_1) for subsequent block creations. Here, w_1 represents proof-of-work in the block mined by u . Third, $t_2(w_2)$ represents a transaction where u transfers some of her bitcoins (s_2) to another Bitcoin address (s_4). The associated transaction fee is modeled as another address (s_3). Here, w_2 represents the authorization of the transaction in the form of a digital signature by u . Finally, $t_3(w_3)$ represents a transaction that rewards a

user for creating a Bitcoin block containing $t_2(w_2)$. In that sense, $t_3(w_3)$ is similar to $t_1(w_1)$, with the difference that $t_3(w_3)$ also captures the fact that the user also receives the fees associated to $t_2(w_2)$.

Figure 2: Illustrative example of a TDAG. Here, we use the same notation as in [Figure 1](#). Graphically, each transaction $t_i(w_i)$ is the subgraph where vertices are the set composed of w_i along with the set of states sharing an edge with w_i ; and edges are the set of incoming edges and outgoing edges for w_i .

We note that this example does not contain any observing edge. This results from the fact that read-only operations are not supported in Bitcoin.

2.2 Conflicts and validity

A central goal of blockchain systems is to prevent conflicts among transactions and to ensure validity for all transactions, as a result of a consensus process executed among the participating entities. The TDAG model permits one to have a closer look at the semantics of conflicts and validity; modeling consensus is outside the scope of this work.

Intuitively, a *conflict* in a blockchain underlying a cryptocurrency such as Bitcoin occurs in an attempt to “double-spend” money. According to the example describing Bitcoin from before (and expanded in [Section 3](#)), assume that a state s in a TDAG

corresponds to bitcoins held by a particular Bitcoin address. Two transactions that double-spend such bitcoins map to two transactions that both consume s . But every state in a TDAG can be consumed at most once, hence, the TDAG model already prevents this form of conflict.

In blockchains for arbitrary smart contracts, a conflict corresponds to a situation where generic validation rules for transactions are violated. Such rules may refer to coins (such as an amount of Ether in Ethereum) or to other assets modeled in the blockchain. The TDAG model for these blockchains also imposes that every state can be consumed at most once.

When one considers an arbitrary set of transactions (not arising from the same transaction graph), such as transactions that have merely been proposed and are not executed on the blockchain yet, then conflicts among them could exist. This is the case in a cryptocurrency like Bitcoin when a miner searches for the next block, for example, and two transactions might be floating around in the network that both attempt to consume the same state s . Similarly, conflicting transactions exist in smart-contract platforms during the process of reaching consensus on a valid blockchain execution.

We now consider a set of transactions (in the form of a graph) and define what it means for them to be conflict-free.

Definition 2.4 (*Conflict-freedom*)

Consider a DAG $\mathcal{T} = (\mathcal{S}_T \dot{\cup} \mathcal{W}_T, \mathcal{E}_T)$ with states $\mathcal{S}_T$, witnesses $\mathcal{W}_T$, producing edges $\mathcal{E}_P \subseteq \mathcal{E}_T$ and consuming edges $\mathcal{E}_C \subseteq \mathcal{E}_T$ that contains a transaction for every witness $w \in \mathcal{W}_T$. We say that $\mathcal{T}$ *has no conflicts* if every state has at most one producing edge and one consuming edge, *i.e.*, $\forall s \in \mathcal{S}_T : |\star \mathcal{E}_P s| \leq 1 \wedge |s \mathcal{E}_C \star| \leq 1$.

A conflict-free set of transactions can be added to a TDAG. To ensure that its addition does not cause any conflicts with the TDAG, only simple and local conditions have to be verified.

Definition 2.5 (*Adding transactions to a TDAG*)

Consider a TDAG $\mathcal{G} = (\mathcal{S} \dot{\cup} \mathcal{W}, \mathcal{E})$ and a DAG $\mathcal{T} = (\mathcal{S}_T \dot{\cup} \mathcal{W}_T, \mathcal{E}_T)$ containing a conflict-free set of transactions such that

1. No witness of $\mathcal{T}$ is in $\mathcal{G}$, *i.e.*, $\mathcal{W} \cap \mathcal{W}_T = \emptyset$;
2. Every input state of $\mathcal{T}$ is an unconsumed output state of $\mathcal{G}$, *i.e.*, $\{s \in \mathcal{S}_T : |\star \mathcal{E}_P s| = 0\} \subseteq \{s \in \mathcal{S} : |s \mathcal{E}_C \star| = 0\}$;

3. The output states of $\mathcal{T}$ do not exist in $\mathcal{G}$, i.e., $\{s \in \mathcal{S}_T : |s\mathcal{E}_C \star| = 0\} \cap \mathcal{S} = \emptyset$.

Then the result of *adding $\mathcal{T}$ to $\mathcal{G}$* is the DAG $\bar{\mathcal{G}} = (\bar{\mathcal{S}} \dot{\cup} \bar{\mathcal{W}}, \bar{\mathcal{E}})$, with $\bar{\mathcal{S}} = \mathcal{S} \cup \mathcal{S}_T$, $\bar{\mathcal{W}} = \mathcal{W} \dot{\cup} \mathcal{W}_T$, and $\bar{\mathcal{E}} = \mathcal{E} \dot{\cup} \mathcal{E}_T$.

Theorem 2.6

When a conflict-free set of transactions $\mathcal{T} = (\mathcal{S}_T \dot{\cup} \mathcal{W}_T, \mathcal{E}_T)$ is added to a TDAG $\mathcal{G} = (\mathcal{S} \dot{\cup} \mathcal{W}, \mathcal{E})$, then the resulting graph $\bar{\mathcal{G}} = (\bar{\mathcal{S}} \dot{\cup} \bar{\mathcal{W}}, \bar{\mathcal{E}})$ is also a TDAG.

PROOF: Here we show that $\bar{\mathcal{G}}$ satisfies the conditions to be a TDAG.

1. The genesis state must not have producing or observing edges, and it must have a single consuming edge. This condition is fulfilled since $\mathcal{G}$ is a TDAG and $\mathcal{T}$ does not contain the genesis state if it is already consumed in $\mathcal{G}$.
2. Every state, other than genesis, must have a single producing edge. This condition is fulfilled in $\mathcal{G}$ and in $\mathcal{T}$ by definition. Now, the addition of t to $\mathcal{G}$ does not create new edges. Therefore, this condition holds also in $\bar{\mathcal{G}}$.
3. Every state, other than the genesis, can have multiple successors, but at most one among them is connected with a consuming edge. It is easy to see that $\bar{\mathcal{G}}$ fulfills this condition following an argument similar as before.
4. The graph must be weakly connected. Note that by the definition of TDAG, each vertex $v \in \mathcal{S} \dot{\cup} \mathcal{W}$ is weakly connected to every unconsumed state in $\mathcal{G}$. Moreover, every vertex $v' \in \mathcal{S}_T \dot{\cup} \mathcal{W}_T$ is weakly connected to at least one input state of $\mathcal{T}$. Now, as the set of input states in $\mathcal{T}$ is a subset of the unconsumed states in $\mathcal{G}$, it follows that $\bar{\mathcal{G}}$ is weakly connected.
5. The graph must not have cycles. According to the assumptions on $\mathcal{T}$ and because $\mathcal{G}$ is a DAG, and through the way in which $\bar{\mathcal{G}}$ is constructed, it is easy to see that $\bar{\mathcal{G}}$ has no cycles.

We now introduce the notion of *validity* for transactions in a TDAG, which models the fact that on a blockchain only “valid” transactions are executed. As an important design choice of the model, the validity of a transaction in a TDAG must be decidable locally—that is, from the transaction alone—considering only its input states, the witness, and the output states. To capture this notion, we assume that the blockchain context defines a boolean *validation predicate* $\mathbb{P}(\cdot)$ on the space of all transactions included in the blockchain.

Definition 2.7 (Validity)

Let t be a transaction in a TDAG $\mathcal{G}$. Then t is *valid* whenever $\mathbb{P}(t) = \text{True}$. Furthermore, $\mathcal{G}$ is a *valid transaction graph* if all transactions in $\mathcal{G}$ are valid.

Combined with the locally checkable conditions for adding transactions to a TDAG, the fact that the validity of a transaction is locally decidable defines, in an influential way, how many blockchain systems work during consensus, validation, and the execution of new transactions. The only steps needed for validation are to ensure the validity predicate of a candidate transaction plus the checks according to [Definition 2.5](#) involving the states to which the transaction refers.

Transaction validation also relies on the property that all states in the TDAG are distinct. In a typical blockchain, the validation function relies on a cryptographic hash of the states to which it refers; this directly ensures uniqueness. For example, consider an execution of a smart contract that holds state on the blockchain in the form of a local variable *var*. The contract may update *var* multiple times, and it may write the same value to *var* more than once. To make the resulting states in the TDAG different, the model will usually include a version number in the state that makes each assignment unique.

At this point, let us review our design choice of a single INIT transaction. Using a single transaction to create all assets represented by the states enables one to locally check the validity of the initialization of the blockchain as well as preserve the locally checkable conditions for further transactions consuming those states.

2.3 Composition of transaction graphs

In Bitcoin (and many other cryptocurrencies), all the miners participate in the consensus protocol to decide about the validity of every single transaction. The permissionless nature of this consensus mechanism heavily limits the transaction throughput. One alternative to overcome this scalability issue is called sharding and consists of organizing disjoint sets of miners, letting each of these sets reach consensus about a subset of the transactions. The composition of those subsets of transactions is required, then, to shape the blockchain.

In the following, we describe the composition of transaction graphs, which states the conditions under which two TDAGs can be merged into a single one. One may then reason about their consistency and validity in a unified manner. Composition of transaction graphs can be used to model the goal of protocols for cross-chain

transactions, namely that the combined state of both chains achieves the expected consistency properties.

Definition 2.8 (TDAG composition).

Consider two TDAGs, $\mathcal{G} := (\mathcal{S} \dot{\cup} \mathcal{W}, \mathcal{E})$ and $\mathcal{G}' := (\mathcal{S}' \dot{\cup} \mathcal{W}', \mathcal{E}')$. Assume that $t(w)$ denotes the INIT transaction in $\mathcal{G}$ and $t'(w')$ denotes the INIT transaction in $\mathcal{G}'$. Further assume that $\hat{t}(\hat{w})$ denotes a INIT transaction where $\hat{w} = (w, w')$ and the output states are the union of output states from $t(w)$ and $t'(w')$. Then, the *composition of $\mathcal{G}$ and $\mathcal{G}'$* is the TDAG $\hat{\mathcal{G}} = \mathcal{T}_{\mathcal{G}} \setminus \{t(w)\} \cup \mathcal{T}_{\mathcal{G}'} \setminus \{t'(w')\} \cup \hat{t}(\hat{w})$.

Theorem 2.9 (Composition of two TDAGs into one TDAG).

The composition of two TDAGs $\mathcal{G}$ and $\mathcal{G}'$ results in a graph $\hat{\mathcal{G}}$, which is also a TDAG.

PROOF: Here we show that $\hat{\mathcal{G}}$ satisfies the conditions to be a TDAG.

1. The genesis state must not have producing or observing edges, and it must have a single consuming edge. This condition is fulfilled by our definition of the INIT transaction $\hat{t}(\hat{w})$.
2. Every state, other than genesis, must have a single producing edge. As $\mathcal{G}$ and $\mathcal{G}'$ are two TDAGs, it is easy to see that each state in $\mathcal{T}_{\mathcal{G}} \setminus \{t(w)\}$ and $\mathcal{T}_{\mathcal{G}'} \setminus \{t'(w')\}$ has a single producing edge. Moreover, by definition of INIT transaction, each output state in $\hat{t}(\hat{w})$ has a single producing edge.
3. Every state, other than the genesis, can have multiple successors, but at most one among them is connected with a consuming edge. It is easy to see that $\hat{\mathcal{G}}$ fulfills this condition along the lines of previous argument.
4. The graph must be weakly connected. $\mathcal{T}_{\mathcal{G}} \setminus \{t(w)\}$ and $\mathcal{T}_{\mathcal{G}'} \setminus \{t'(w')\}$ are connected by definition, as $\mathcal{G}$ and $\mathcal{G}'$ are two TDAGs. Moreover, the definition of the INIT transaction $\hat{t}(\hat{w})$ ensures that any vertex in $\mathcal{T}_{\mathcal{G}} \setminus \{t(w)\}$ is connected to any vertex in $\mathcal{T}_{\mathcal{G}'} \setminus \{t'(w')\}$ through $\hat{w}$.
5. The graph must not have cycles. $\mathcal{T}_{\mathcal{G}} \setminus \{t(w)\}$ and $\mathcal{T}_{\mathcal{G}'} \setminus \{t'(w')\}$ are acyclic by definition, as $\mathcal{G}$ and $\mathcal{G}'$ are two TDAGs. Moreover, the addition of $\hat{t}(\hat{w})$ clearly does not introduce any cycle.

3. Applications

In this section, we describe how executions of different blockchain systems are modeled by transaction graphs. We cover three prominent blockchains: Bitcoin, Ethereum, and Hyperledger Fabric. They differ in how they store assets in their state. Bitcoin, for example, does not have state “variables” but maintains an asset only in the

context of the transaction that created it. Ethereum, on the other hand, uses variables and accounts for its state. The data model in Fabric is a key-value store (KVS), which can be mapped to a local database on each node. Due to lack of space, this section only gives a short overview, and more details appear in the full version [20].

Throughout this section, we denote by $y \leftarrow H(x)$ a cryptographic, collision-free hash function that takes as input a bit-string $x \in \{0, 1\}^*$ of arbitrary length and returns a fixed-length string $y \in \{0, 1\}^l$.

3.1 Bitcoin

Since Bitcoin (bitcoin.org) is the prototype of all blockchain systems, there are many publicly available descriptions [1][21] and we keep the background short. Likewise, the discussion here applies to all *alt-coins* patterned after Bitcoin.

Bitcoin combines transaction validation, coin mining, and agreement on the ledger with the “Nakamoto protocol” that uses proof-of-work and ensures consensus. A *block* in Bitcoin can hold two types of *transactions*:

- A *coinbase transaction* that transfers as-of-yet unmined bitcoins to a Bitcoin address as chosen by the miner of the corresponding block, as a reward for creating the block. This transaction is valid if (i) it transfers a number of bitcoins according to the height of the block to a Bitcoin address; and (ii) it contains the solution to the proof-of-work puzzle for successful mining of the block.
- A *regular transaction* transfers bitcoins from a set of Bitcoin (input) addresses to another set of Bitcoin (output) addresses. It also incurs a fee, defined as the difference between the bitcoin amounts in the input and output, which is assigned to the miner of the block in which the transaction appears. A regular transaction is valid if it includes a confirmation for each input for the amount and output and if it does not create new bitcoins.

Bitcoin value exists in the blockchain in the form of *unspent transaction output*, often abbreviated *UTXO*, which has been assigned to an address, representing a digital-signature public key. This value is controlled by the holder of the corresponding private key. It can be spent and transferred to another address by signing a transaction with the private key.

In the TDAG modeling Bitcoin, we let every state be a tuple of the form **(addr, val, hash, height)** where *addr* denotes an address, *val* denotes the amount of bitcoins held in this state, *txhash* is the cryptographic hash of the Bitcoin transaction generating

this state, and *index* is the sequential index of the output among all outputs generated in that transaction.

In contrast to the Bitcoin code, we model transaction fees and unmined bitcoins as held by or associated to an (imaginary) address. This allows a coherent model for the TDAG. Thus, the state resulting from the special INIT transaction is fixed to $(addr_0, 21M, H(\emptyset), 0)$, holding all 21 million bitcoins that ever exist.

The form of a witness depends on the transaction type: the witness for a coinbase transaction is the solution for the proof-of-work to assign the bitcoins to the address designated by the miner. For a regular transaction, the witness consists of a set of confirmations for the transfer of bitcoin, in the form of a digital signature for each UTXO, over the input and output addresses of the transfer. Finally, the INIT transaction does not require any witness.

The TDAG for Bitcoin contains producing and consuming edges but no observing edges. For a coinbase transaction, the input states are the unconsumed state of unmined bitcoins and the fee states for the transactions included in the mined block. One producing edge leads to a state for collecting the fees and the mining reward, another one to a state containing the remaining unmined bitcoins. Its witness is the mining proof. For a regular transaction, the input states are the unconsumed states representing the transaction inputs and the produced output states correspond to the transactions output addresses. The witness holds a set of confirmations (digital signatures), confirming for each input state the transfer of some bitcoins to the corresponding output addresses.

The transaction predicate incorporates the validation rules of Bitcoin, as expressed in the states, witnesses, and transactions of the TDAG.

With these definitions, one can then show the intuitive result that, except with negligible probability, every (legal) execution of Bitcoin, considering only transactions that are “deep enough” in the blockchain (*e.g.*, six blocks deep) [7] gives rise to a TDAG constructed like this. The formal analysis of this result exploits the fact that the DAG formed by the hash-function applications among states has no cycles, and therefore satisfies the properties of a TDAG.

3.2 Ethereum

Ethereum [2] is the most prominent public blockchain and cryptocurrency supporting generic smart contracts today. In Ethereum, there exist two types of accounts, called

externally owned accounts and *contract accounts*. Externally owned accounts, by and large, resemble the accounts of other cryptocurrencies such as Bitcoin. In these accounts, users maintain their currency balance in Ether, owned by them. But the main innovation of Ethereum lies in contract accounts, which represent a smart contract (an arbitrary piece of code in the platform-specific language) that executes a set of instructions upon receiving suitable input. A contract account also holds and controls its own Ether balance.

Ethereum supports several types of transactions. First, a transaction in Ethereum can be used to transfer Ether between two externally owned accounts. This type of transaction is like the exchange of coins in other cryptocurrencies. Second, a transaction can be used to create a contract with the code of the contract and an externally owned account as inputs. It outputs a contract account with the information required to initialize the implemented code (*e.g.*, the inputs for the init function). Finally, a transaction can be used to invoke an existing contract on the blockchain.

An Ethereum transaction includes as input the sender's address (an externally owned account), a recipient address (another account), a transaction value to be transferred from the sender's address to the recipient, and some arguments with parameters for the contract. As the transaction sender has to pay *gas* for the execution of the contract, the transaction also specifies a *gas price* that determines the price the sender is willing to pay for each executed instruction, and a *gas limit*, specifying a maximum overall price for the execution. A contract may also call functions of other contracts; however, this will not give rise to new transactions, as these calls take place in the context of the original transaction.

To model an Ethereum execution as a TDAG, we let each state consist of a tuple **(addr, account-type, code, local-state, gas-price, val)**. Here, *addr* denotes the account address that produced the state, *account-type* determines whether this is a state of an external account or a contract account, *code* is a hash of the smart contract's code, *local-state* denotes, collectively, all variables held by the contract, and *val* is the Ether balance held by the account after the execution that produced the state. If *account-type* specifies an externally owned account, then the smart contract is the fixed logic to validate payments from such accounts.

There is also a genesis state that models the creation of an Ethereum blockchain. In contrast to Bitcoin, there is currently no bound on the amount of Ether that will exist in the public Ethereum blockchain; the creation of new Ether is therefore subsumed into the mining operation and its validation.

A transaction in the TDAG is determined by the witness. It corresponds to an invocation of a smart contract and contains a gas limit and regular input arguments that validate the transaction. For instance, these arguments must contain a digital signature valid under the public key associated to the invoking external account that runs the transaction.

The transaction contains the state of the invoking account and the state of the contract as input states, with consuming edges to the witness. It also produces two states, an updated state of the invoking account and an updated state of the contract, as resulting from running the contract with the given gas limit and input arguments. If the contract calls functions of other contracts and they modify their state, then the states representing these contracts are also part of the transaction in the TDAG (as input states and output states). The validation predicate simply executes the code.

For mining new Ether, running transactions, and collecting the corresponding fees, similar states and validation logic as in the TDAG model of Bitcoin are added. Given these notions, one can show that every (legal) execution of Ethereum, considering as in Bitcoin only those transactions that are deep enough in the blockchain, produces a valid TDAG.

3.3 Hyperledger Fabric

[Hyperledger Fabric](#), or *Fabric* for short, is a permissioned blockchain framework, designed to support modular implementations of different components, including its consensus protocol, membership provider, and cryptography library [3]. The nodes executing the transactions in Fabric are called *peers*.

An instance of Fabric may contain multiple channels that may run on different sets of peers, where each channel operates like a blockchain system independent of the others, apart from using some of the same code infrastructure, ordering protocol, and other components. We therefore consider only one channel here, modeling one blockchain.

On a channel, a *configuration* transaction (*configtx*) sets the initial values used for transaction processing, such as the credentials of the peers or organizations controlling the channel, the implementation of its ordering service, and so on. Once a channel has been prepared like this, it is ready to execute operations on its peers. Transactions in Fabric are executed by smart contracts called *chaincode*.

Chaincode is first *installed* on the peer and may later be *upgraded*; it must be *instantiated* for a specific channel before it can process transactions. Once instantiated on the channel, a chaincode supports two types of transactions: *init* and *invoke*. An *init* transaction is executed once after the chaincode has been installed or upgraded; it specifies an *endorsement policy* that determines how any subsequent transaction of this chaincode should be authorized. A chaincode determines which peers it executes on through the endorsement policy: whether all peers in the channel execute it, or only some, and which peers or which set of peers are sufficient to authorize the execution of the transaction.

An *invoke* transaction is used to execute a computation that may read and modify the state of the chaincode, which is a set of key-value pairs. The operations to access the state are $GetState(k) \rightarrow v$ (given a key k , return the last value v written to it) and $PutState(k, v)$ (write the value v to storage under the key k).

The processing of a transaction on Fabric proceeds like this [22]:

1. A client creates and signs a transaction for a particular chaincode and sends it to the respective endorsing peers.
2. The endorsing peers simulate the transaction on their current copy of the key-value store (KVS), verifying that the client is authorized to execute it. If successful, each endorsing peer returns the result of the execution to the client. This is also called an *endorsement*. It comes in the form of a signed readset and writeset (with the key-value pairs accessed during simulation, including a *version* for every value in the readset, determined by the logical time when this value was written). The endorsement serves as a static representation of the chaincode execution.
3. When the client has assembled enough endorsements that produce the same KVS changes and that satisfy the endorsement policy, it combines them to a transaction proposal. Then the client broadcasts this transaction proposal to the ordering service, which simply orders transactions without considering their semantics. Currently an ordering service based on Apache Kafka (kafka.apache.org) running in a cluster is supported, and an ordering service using BFT consensus is under development [23][9].
4. The ordering service disseminates an ordered stream of transactions (grouped into blocks) to the peers on the channel. Each peer on its own then validates each transaction, by verifying that the endorsement policy is satisfied and that there were no changes to the key-value pairs contained in the readset (since transaction simulation).

5. If successful, the peer appends the block to the blockchain (of the channel) and performs the updates from the writeset to its local copy of the KVS. This assigns a version to the modified key-value pairs. Since the validation is deterministic, the states and versions are the same for all correct peers.

In the TDAG for Fabric, the states correspond to the entries in the KVS. Every state is a tuple containing at least **(key, version)**. It is assumed that an *init* transaction implicitly initializes every *key* used by the chaincode later with a default value (—). The *init* transaction is always valid.

Furthermore, every *invoke* transaction that reads or writes a set of keys $\mathcal{K}$ contains an observing edge for every $k \in \mathcal{K}$ accessed by an operation $GetState(k)$ but not by an operation $PutState(k, \star)$, and a consuming edge for every k that is written using an operation $PutState(k, \star)$. In other words, every key is implicitly read before it is written and, thus, a transaction in the TDAG modeling a Fabric execution has the same number of consuming edges as the number of producing edges.

A witness in the TDAG corresponds to a valid endorsement, in the form of signatures from the endorsers issued on the same readset/writeset pair from the transaction proposal. The validation predicate $\mathbb{P}(\cdot)$ contains the steps that each peer takes to validate a transaction coming from the ordering service, with respect to its local KVS. Notice that this validation only accesses the versions in the readset, but no other state entry in the KVS. Since these states are also contained in the transaction in the TDAG, the evaluation of $\mathbb{P}(\cdot)$ in the graph is local.

Given that the ordering service of Fabric outputs the same stream of blocks with transactions to every connected peer, it is easy to verify that the graph resulting from any execution of Fabric is a TDAG.

4. Conclusion

Blockchains and distributed ledger platforms are of great interest for the financial industry today, due to their role as trustless intermediaries gained from their resilience to attacks and subversion. For gaining confidence in a new technology, it is paramount to study its security with formal models.

This work has proposed transaction graphs or TDAGs as a discrete model for the semantics of the interactions in a blockchain system. In contrast to existing event-based models for generic distributed and concurrent systems, it explicitly takes into account the validation of transactions, which is an important aspect of blockchains. For

instance, the TDAG model allows one to model assets and their transfer among different entities. It also facilitates comparisons among different technologies available today.

We envision that richer semantics can be expressed by refining the TDAG model. For instance, one may argue about further invariants of the blockchain system as properties of the TDAG, similar to modeling Bitcoin’s fixed coin supply. One might also use a TDAG to formally model the provenance for generic assets that are handled by smart contracts, building on the paths through which the asset was transferred in the TDAG. One could also leverage a TDAG to formally describe the guarantees provided by a blockchain equipped with a pruning mechanism, reasoning about the remaining states in the TDAG after pruning. Finally, we additionally foresee that the TDAG can be extended to model invariants required for payment channels—for instance, payment channel transactions should be free of conflicts with those included in the TDAG.

Appendix

A. Transaction graph for Bitcoin

We start with the description of an execution of the Bitcoin system as represented by the corresponding blockchain. A Bitcoin blockchain is composed of *blocks*, where each block is created as a result of successfully executing the Bitcoin mining process [1]. The miner of such block (*i.e.*, user showing a valid proof of successful mining) chooses a set of *regular* transactions to be added in the block along with a single *coinbase* transaction. There exists a special block, denoted as *genesis* block, that represents the initialization of the blockchain.

A coinbase transaction transfers unmined bitcoins to a (set of) Bitcoin address, chosen by the corresponding miner, as a *reward* for creating the block. A coinbase transaction is valid if it transfers only the number of bitcoins set as reward according to the height of the mined block. A regular transaction transfers bitcoins from a set of Bitcoin *addresses* (*i.e.*, input addresses) to another set of Bitcoin addresses (*i.e.*, output addresses). A regular transaction is valid if: (i) it includes a *confirmation* for each input address; and (ii) it does not create new bitcoins. Finally, a regular transaction has an associated fee (*i.e.*, between the bitcoins held at input and output addresses).

Definition A.1 (Bitcoin execution)

A Bitcoin execution $\mathcal{L}_{\text{BTC}}$ is a set of blocks $\mathcal{B} := \{B_g, B_1, \dots, B_n\}$, where B_g denotes the genesis block and contains a single initialization transaction. Each other block

$B_i := (MP), CBTX, RTX_1, \dots, RTX_n$) is a tuple composed of a proof of successful mining MP , and a set of transactions containing a coinbase transaction $CBTX$ and regular transactions RTX_i . A $CBTX$ contains a Bitcoin address $ADDR$. A RTX is a tuple $(ADDR_{in}, \mathcal{F}, ADDR_{out})$, where $ADDR_{in}$ and $ADDR_{out}$ are two sets of Bitcoin addresses and $\mathcal{F}$ is a set of confirmations CF .

We now describe our modeling of a given execution of Bitcoin as a TDAG. A state represents a Bitcoin address that holds a group of bitcoins, a transaction fee, or the yet unmined bitcoins. We note that fees and unmined Bitcoins are not associated to an address in the real Bitcoin, but we model them as held by an address to have a coherent transaction graph model. The genesis state represents a Bitcoin address holding the only 21 million bitcoins that will ever exist in the Bitcoin system. Each witness represents either a proof of successful mining for a block or the (set of) confirmations required in a regular transaction. Finally, we consider two types of edges: producing and consuming edges. A producing edge links unconsumed addresses for unmined bitcoins and transaction fees to the mining proof for the corresponding coinbase transaction; or an input address to the corresponding confirmation in a regular transaction. A consuming edge links a mining proof to the Bitcoin addresses getting the reward, or a set of confirmations to the corresponding output addresses receiving (part of) the transferred bitcoins.

Definition A.2 (Transaction graph for Bitcoin)

We model an execution of Bitcoin system as a graph $\mathcal{G}_{BTC} := (\mathcal{S}_{BTC} \dot{\cup} \mathcal{W}_{BTC}, \mathcal{E}_{BTC})$, defined as follows:

- **State:** Each state $s \in \mathcal{S}_{BTC}$ is defined as a tuple $(addr, val, txhash, index)$, where $addr$ denotes a Bitcoin address, val denotes the amount of Bitcoins held at $addr$, $txhash$ denotes the hash of the transaction generating s , and $index$ the index of the output described by s . The genesis state s_g is defined as the fixed tuple $(ADDR_0, 21M, H(\emptyset), 0)$.
- **Witness:** Each witness $w \in \mathcal{W}_{BTC}$ is defined by a tuple $(txtype, \mathcal{F})$, where $txtype$ denotes the type of the transaction and determines the content of $\mathcal{F}$. In particular, $(TINITX, \emptyset)$ is the witness for the initialization transaction, $(TCBTX, MP)$ denotes a witness for a coinbase transaction, and $(TRTX, \{cf_i\})$ denotes the witness for a regular transaction.
- **Edge:** Each edge $e \in \mathcal{E}$ is defined either as consuming edge or producing edge.

The transaction graph presented here determines the modeling of the possible transactions in a Bitcoin execution. The next definition maps transactions in a Bitcoin execution to transaction types supported in a TDAG.

Definition A.3 (*Transaction types*)

A coinbase transaction is modeled as a SIMO transaction. A regular transaction is modeled as a SISO, SIMO, MISO or MIMO transaction depending on $|ADDR_{in}|$ and $|ADDR_{out}|$. For instance, SISO models a regular transaction where $|ADDR_{in}| = 1 \wedge |ADDR_{out}| = 1$. The rest are derived accordingly. Finally, we define the initialization transaction included in the genesis block as an INIT transaction of the form $t := (\{s_g, s, w\}, \{(s_g, w), (w, s)\})$, where $(s_g, w) \in \mathcal{E}_C$, $(w, s) \in \mathcal{E}_P$ and $s := (ADDR_m, 21M, H(\{s_g, w\}), 0)$, where $ADDR_m$ denotes a Bitcoin address that contains unmined bitcoins. s_g and w are as defined in [Definition A.2](#).

Finally, we complete our description of the Bitcoin context with the corresponding transaction predicate $\mathbb{P}$. For that, we use $\text{VerifyContract}(ADDR, CF)$ as a function that, upon the input of a Bitcoin address $ADDR$ and a confirmation CF , returns $TRUE$ if CF encodes a valid confirmation to spend the bitcoins held at $ADDR$. Otherwise, it returns $FALSE$. Additionally, we use $\text{VerifyWork}(MP)$ as a function that, upon the input of a mining proof MP , returns $TRUE$ if MP is a valid proof-of-work for the corresponding block, or $FALSE$ otherwise. We thereby abstract away the implementation details for validation of Bitcoin scripts and mining proofs.

Definition A.4 (*Transaction predicate in Bitcoin*).

Consider a transaction $t := (\mathcal{S} \dot{\cup} \{w\}, \mathcal{E})$. Then, $\mathbb{P}(t)$ returns $TRUE$ if the following conditions hold and $FALSE$ otherwise.

1. If t is a regular transaction ($w.txtype = TRTX$), the witness holds a valid confirmation for each input state *i.e.*, $\forall s \in \star \mathcal{E} w, \exists CF \in w.\mathcal{F} : \text{VerifyContract}(s.addr, CF)$.
2. If t is a coinbase transaction ($w.txtype = TCBTX$), the witness contains a valid mining proof, *i.e.*, $w.\mathcal{F} := \{MP\} \wedge \text{VerifyWork}(MP)$.
3. Each output state represents a positive number of bitcoins, *i.e.*, $\forall s \in w\mathcal{E}\star : s.val > 0$.
4. The sum of bitcoins held at the input states must be equal to the sum of bitcoins held at the output states, *i.e.*, $\sum_{s \in \star \mathcal{E} w} s.val = \sum_{s' \in w\mathcal{E}\star} s'.val$.

Observe that it is unnecessary to verify the *txhash* and *index* components of the states. Their sole purpose is to separate different states: as long as no collision occurs between different transaction hashes on the Bitcoin blockchain, all states in our

definition will be distinct. This property is obviously not achieved by *ADDR* and *val* alone.

A.1 Model analysis

We start this section by analyzing the definition of transaction graph presented in the previous section. We start by showing that it is a TDAG. Here, we consider *legal* a Bitcoin execution that contains only transactions that are “deep enough” in the blockchain (*e.g.*, six blocks deep). We thereby enable the study of any Bitcoin execution in terms of the properties of a TDAG such as conflict-freedom or validity.

Theorem A.5

Assume H is a collision-resistant hash function [24] and assume that $\mathcal{L}_{\text{BTC}}$ is a legal Bitcoin execution. Then, the graph $\mathcal{G}_{\text{BTC}}$ resulting from modeling $\mathcal{L}_{\text{BTC}}$ is a TDAG.

PROOF: Here, we show that $\mathcal{G}_{\text{BTC}} = (\mathcal{S}_{\text{BTC}} \dot{\cup} \mathcal{W}_{\text{BTC}}, \mathcal{E}_{\text{BTC}})$ fulfills the conditions to be a TDAG.

1. The genesis state must not have producing or observing edges and it must have a single producing edge. Our designed INIT transaction ensures this.
2. Every state, other than the genesis, must have a single producing edge. Assume by contradiction that it is not fulfilled. Then, there is a state $s \in \mathcal{S}_{\text{BTC}}$ with at least two producing edges and that implies that there exists two different sets $\mathcal{V} := \mathcal{S} \dot{\cup} \{w\}$ and $\mathcal{V}' := \mathcal{S}' \dot{\cup} \{w'\}$ such that $H(\mathcal{V}) = H(\mathcal{V}')$. However, $\mathcal{V}$ and $\mathcal{V}'$ contradict the assumption that H is collision resistant.
3. Every state other than the genesis can have multiple successors, but at most one among them is connected with a consuming edge. Each Bitcoin address is consumed only once in a legal Bitcoin execution. Therefore, this condition is fulfilled.
4. The graph must be weakly connected. Each new transaction consumes a previously unconsumed state in the graph, *i.e.*, either a unspent Bitcoin address, or mines as-of-yet unmined bitcoins and consumes unclaimed fees. Therefore, the overall graph is weakly connected.
5. The graph must not have cycles. Assume by contradiction that there is a cycle in $\mathcal{G}_{\text{BTC}}$. This, however, implies that there are two different transactions t and t' that produce the same state. However, as we have seen before, this only occurs in case two different transactions have the same hash, which contradicts the fact that the underlying hash function is collision resistant.

Remember from [Definition 2.7](#) that a TDAG is *valid* if each transaction individually is valid according to a transaction predicate $\mathbb{P}$. Next, we show that validating Bitcoin transactions individually in our model, suffices to safely consider that unconsumed states represent all bitcoins in the system.

Definition A.6 (Unspent bitcoins)

Consider $\mathcal{G}_{\text{BTC}}$ a TDAG modeling a Bitcoin execution. Then, the *unspent bitcoins in $\mathcal{G}_{\text{BTC}}$* are the sum of bitcoins held at unconsumed states of $\mathcal{G}_{\text{BTC}}$.

Theorem A.7 (Unspent bitcoins are all bitcoins in the system)

Consider $\mathcal{G}_{\text{BTC}}$ a valid TDAG that models a Bitcoin execution. Then, the amount of unspent bitcoins in $\mathcal{G}_{\text{BTC}}$ is equal to all bitcoins ever existing in the system. More formally, let $\mathcal{S}'$ be the set of unconsumed states in $\mathcal{G}_{\text{BTC}}$, then $\sum_{s \in \mathcal{S}} s.val = \sum_{s' \in \mathcal{S}'} s'.val$.

PROOF: Assume by contradiction that [Theorem A.7](#) does not hold. Then, there must exist a transaction $t := (\mathcal{S} \dot{\cup} \{w\}, \mathcal{E})$ in $\mathcal{T}_{\mathcal{G}_{\text{BTC}}}$ such that $\sum_{s \in \star \mathcal{E}_w} s.val \neq \sum_{s' \in w \mathcal{E}_\star} s'.val$. This, however, clearly implies that $\mathbb{P}(t)$ returns *FALSE*, which contradicts the assumption that $\mathcal{G}_{\text{BTC}}$ is a valid TDAG.

A.2 Modeling an example of bitcoin execution

Here, we describe our modeling for an illustrative example of Bitcoin execution. We assume for simplicity that the block reward is fixed to a value of 50 bitcoins, as this was the first reward set in the Bitcoin system. Additionally, we assume that the transaction fee is fixed to 1 bitcoin. We stress, however, that the TDAG model is expressive enough to relax these assumptions.

We focus on the illustrative example depicted in [Figure 3](#). In particular, [Figure 3a](#) shows a possible Bitcoin execution $\mathcal{L}_{\text{BTC}} := \{B_g, B_1, B_2\}$, where $b_g := (\emptyset, \{t_0\})$, $b_1 := (MP, \{t_1\})$ and $b_2 := (MP', \{t_2, t_3, t_4\})$. We note that this example is similar to that in [Figure 2](#) and due to lack of space we do not describe it here again. However, we remark that it is expanded here with an extra MIMO transaction (*i.e.*, $t_3(w_3)$) to show how we model transactions that involve multiple payers and multiple payees. Instead, we focus on the description of $\mathcal{G}_{\text{BTC}} := (\mathcal{S} \dot{\cup} \mathcal{W}, \mathcal{E}_P \dot{\cup} \mathcal{E}_C)$, a transaction graph modeling the aforementioned Bitcoin execution as depicted in [Figure 3b](#).

Figure 3(a): Example of Bitcoin execution $\mathcal{L}_{BTC} := \{B_g, B_1, B_2\}$. Only involved blocks and graphical description of transactions $t_1 - t_4$ have been shown. Here, t_0 represents the initialization transaction.

Figure 3(b): Example of a $\mathcal{G}_{\text{BTC}}$ instance.

- $t_0 := (\{s_g, s_0, w\}, \{(s_g, w), (w, s_0)\})$, where $(s_g, w) \in \mathcal{E}_C$ and $(w, s_0) \in \mathcal{E}_P$. This represents the initialization transaction where $s_g := (\text{ADDR}_0, 21M, \text{H}(\emptyset), 0)$, $w_0 := (\text{TINITX}, \emptyset)$ and $s_0 := (\text{ADDR}_m, 21M, \text{txhash}_0, 0)$.
- $t_1 := (\{s_0, s_1, s_2, w_1\}, \{(s_0, w_1), (w_1, s_1), (w_1, s_2)\})$, where $(s_0, w_1) \in \mathcal{E}_C$ and $\{(w_1, s_1), (w_1, s_2)\} \subseteq \mathcal{E}_P$. A SIMO transaction that issues bitcoins to Alice after she has successfully mined a block. In a bit more detail, $w_1 := (\text{TCBTX}, MP)$, $s_1 := (\text{ADDR}_m, (21M - 50), \text{txhash}_1, 0)$ and $s_2 := (\text{ADDR}_{\text{Alice}}, 50, \text{txhash}_1, 1)$, where $\text{ADDR}_{\text{Alice}}$ denotes a Bitcoin address owned by Alice. We follow this notion in the rest of the example for the addresses owned by the example users. s_0 is defined as in t_0 .
- $t_2 := (\{s_2, s_3, s_4, s_5, w_2\}, \{(s_2, w_2), (w_2, s_3), (w_2, s_4), (w_2, s_5)\})$, where $(s_2, w_2) \in \mathcal{E}_C$ and $\{(w_2, s_3), (w_2, s_4), (w_2, s_5)\} \subseteq \mathcal{E}_P$. A SIMO transaction that pays 2 bitcoins to Bob, and the remaining bitcoins are sent back to Alice. In a bit more detail, $w_2 := (\text{TRTX}, \{\text{CF}_{\text{Alice}}\})$, $s_3 := (\text{ADDR}'_m, 1, \text{txhash}_2, 0)$, $s_4 := (\text{ADDR}_{\text{Bob}}, 2, \text{txhash}_2, 1)$

and $\{(w_2, s_3), (w_2, s_4), (w_2, s_5)\} \subseteq \mathcal{E}_P$. A SIMO transaction that pays 2 bitcoins to Bob, and the remaining bitcoins are sent back to Alice. In a bit more detail,

$w_2 := (TRTX, \{CF_{Alice}\})$, $s_3 := (ADDR'_m, 1, txhash_2, 0)$, $s_4 := (ADDR_{Bob}, 2, txhash_2, 1)$ and $s_5 := (ADDR'_{Alice}, 47, txhash_2, 2)$. The rest of states are defined as in previous transactions.

- $t_3 := (\{s_4, s_5, s_6, s_7, s_8, s_9, w_3\}, \{(s_4, w_3), (s_5, w_3), (w_3, s_6), (w_3, s_7), (w_3, s_8), (w_3, s_9)\})$, where $\{(s_4, w_3), (s_5, w_3)\} \subseteq \mathcal{E}_C$ and $\{(w_3, s_6), (w_3, s_7), (w_3, s_8), (w_3, s_9)\} \subseteq \mathcal{E}_P$. A MIMO transaction that pays 3 bitcoins to Charles, jointly by Alice and Bob. In a bit more detail, $w_3 := (TRTX, \{CF'_{Alice}, CF_{Bob}\})$, $s_6 := (ADDR''_m, 1, txhash_3, 0)$, $s_7 := (ADDR'_{Bob}, 1, txhash_3, 1)$, $s_8 := (ADDR''_{Alice}, 44, txhash_3, 2)$, and $s_9 := (ADDR_{Charles}, 3, txhash_3, 3)$. The rest of states are defined as previous transactions.
- $t_4 := (\{s_1, s_3, s_6, s_{10}, s_{11}, w_4\}, \{(s_1, w_4), (s_3, w_4), (s_6, w_4), (w_4, s_{10}), (w_4, s_{11})\})$, where $\{(s_1, w_4), (s_3, w_4), (s_6, w_4)\} \in \mathcal{E}_C$ and $\{(w_4, s_{10}), (w_4, s_{11})\} \subseteq \mathcal{E}_P$. A MIMO transaction that issues bitcoins to Diana after she has successfully mined a block. Additionally, Diana claims the transaction fees for transactions t_2 and t_3 . In a bit more detail, $w_4 := (TCBTX, MP)$, $s_{10} := (ADDR'''_m, (21M - 100), txhash_4, 0)$, and $s_{11} := (ADDR_{Diana}, 52, txhash_4, 1)$.

B. Transaction graph for Hyperledger Fabric

In this section, we study the Hyperledger Fabric [3] (Fabric) blockchain-based system. We start by the description of an execution of Fabric. An execution of Fabric is represented as a set of blockchains, one per channel. However, as each single blockchain evolves independently from each other, we restrict our description here to a single blockchain. This description, however, can be easily extended to model a Fabric execution with multiple channels.

A blockchain is composed of blocks. We denote the first block as *genesis* block and each subsequent block is created by the *ordering service*. Such ordering service chooses the sorted set of transactions to be included in each block. Fabric supports two types of transactions: *Init* and *Invoke*. An *init* transaction is included in the genesis block and it is used to initialize every key used in the blockchain to a default value – and includes an *endorsement policy*, that determines how any subsequent transaction should be authorized. We consider that an initialization transaction is always valid.

An *invoke* transaction is used to carry out updates in a set of key-value pairs for the local current key-value store (KVS) through two operations: (i) $GetState(k) \rightarrow v$, that given a key k provides the most current value v associated to it; and (ii) $PutState(k, v)$,

that updates value associated to a given key k to the newly provided value v . An *invoke* transaction is valid if it contains enough *endorsements* from the set of *endorsers* specified in the *endorsement policy*.

Definition B.1 (HLF execution)

A Fabric execution $\mathcal{L}_{\text{HLF}}$ is a set of blocks $\mathcal{B} := \{b_g, b_1, \dots, b_n\}$, where b_g denotes the genesis block that contains a single init transaction, denoted by $INITX$. Each block $B_i := \{INVTX_1, \dots, INVTX_n\}$ is a set of invoke transactions $INVTX_i$. An $INITX$ contains a single endorsement policy EP . Each transaction $INVTX_i$ is defined as a tuple $(\mathcal{F}, \mathcal{U})$, where $\mathcal{F}$ denotes the set of endorsements $(\{END_1, \dots, END_n\})$, and $\mathcal{U}$ denotes the set of $\{GetState(\star), PutState(\star, \star)\}$ operations to update key-value pairs.

We continue by describing the modeling of a Fabric execution. Informally, each state in our model represents a key-value pair. Each witness represents the set of endorsements required for a transaction to be valid. Finally, here we consider three type of edges: observing, consuming and producing edges. An observing edge links a key k to the endorsement specified in a transaction that reads k but does not modify it (*e.g.*, an invoke transaction that contains only a $GetState(k)$ operation). If the key k is modified (*e.g.*, an invoke transaction that contains $PutState(k, \star)$ operation), a consuming edge links then the key k with the endorsements for such transaction. Finally, a producing edge links the endorsements to a key k a transaction has modified it (*e.g.*, by means of a $PutState(k, \star)$ operation).

Definition B.2 (Model for HLF execution)

We model a Fabric execution $\mathcal{L}_{\text{HLF}}$ as a graph $\mathcal{G}_{\text{HLF}} := (\mathcal{S}_{\text{HLF}} \dot{\cup} \mathcal{W}_{\text{HLF}}, \mathcal{E}_{\text{HLF}})$ defined as follows:

- **States:** Each state $s \in \mathcal{S}_{\text{HLF}}$ is defined as a tuple $(key, version)$, where *key* denotes the key part of a key-value pair and *version* denotes the current version number of the key-value pair. The genesis state is defined as $s_g := (params, 0)$ and denotes a special key-value pair that holds the configuration parameters for a channel as indicated in channel initialization.
- **Witness:** Each witness $w \in \mathcal{W}_{\text{HLF}}$ is defined as a tuple $(txtype, \mathcal{F})$, where *txtype* set to $TINITX$ indicates an init transaction, and *txtype* set to $TINVTX$ indicates an invoke transaction. $\mathcal{F}$ denotes an endorsement policy EP if *txtype* = $TINITX$ or a set of endorsements $\{END_i\}$ if *txtype* = $TINVTX$. For simplicity, we assume that an endorsement END also contains the corresponding set of operations $GetState(\star)$ and $PutState(\star, \star)$.

- **Edges:** Each edge $e \in \mathcal{E}_{\text{HLF}}$ is defined as either an observing, consuming, or producing edge.

Definition B.3 (Transaction types)

An invoke transaction is modeled as a SISO, MISO or MIMO transaction depending on the set of operations $GetState(\star)$ and $PutState(\star, \star)$ that it uses. For instance, a SISO transaction models a transaction that uses a single $PutState(k, \star)$ operation for a key k . A MISO transaction models a transaction that updates a single key k and reads at least one additional key k' (e.g., $\{GetState(k), PutState(k', v)\}$). Finally, a MIMO transaction models a transaction that updates several keys and possibly reads other additional keys (e.g., $\{GetState(k), PutState(k', v), Put(k'', v')\}$). An init transaction is of type INIT and is defined as $t := (\{s_g, w\} \dot{\cup} \{s_i\}, \{(s_g, w)\} \dot{\cup} \{(w, s_1), \dots, (w, s_n)\})$, where $(s_g, w) \in \mathcal{E}_C$, $\{(w, s_1), \dots, (w, s_n)\} \subseteq \mathcal{E}_P$, $w := (TINITX, EP)$, and each $s_i := (k_i, -)$. The genesis state s_g is defined in [Definition B.2](#).

We make two observations in the definition of the transaction types. First, MISO and MIMO types are restricted in the sense that they must have the same number of consuming and producing edges. This is due to the fact that we model each $PutState(\star, \star)$ operation as a consuming edge from the state of the key being updated and a producing edge to the state corresponding to the updated key-value pair. We note, however, that this is a characteristic inherent to all systems based on key-value stores and not a particular limitation of Fabric.

Second, as any system based in a key-value store, each key must exist only once. For that, we model our initialization transaction such that all the keys used in the given Fabric's execution are created and initialized to a fixed initial value $(-)$.

Now, we finalize the description of our model by defining the transaction predicate for Fabric. Here, we denote, by $VerifyEndorsement(\{END_i\})$, a boolean function that takes a set of endorsements $\{END_i\}$ and returns $TRUE$ if $\{END_i\}$ represents a valid set of endorsements according to the endorsement policy EP , and $FALSE$ otherwise. Here, we assume that EP is obtained from the initialization transaction included in the corresponding Fabric execution.

Definition B.4 (Transaction predicate in HLF)

Consider a transaction $t := (\mathcal{S} \dot{\cup} \{w\}, \mathcal{E}_O \dot{\cup} \mathcal{E}_P \dot{\cup} \mathcal{E}_C)$. Then, $\mathbb{P}(t)$ returns $TRUE$ if the following conditions hold and $FALSE$ otherwise.

1. If t is an invoke transaction, the witness must contain a set of valid endorsements, *i.e.*, $w.txtype = TINVTX \Rightarrow \text{VerifyEndorsement}(w.\mathcal{F})$.
2. If t is an invoke transaction, each output state must represent an update of a key included in a input state. Moreover, the version number for the output state must be bigger than the version number for the input state representing the same key, *i.e.*, $w.txtype = TINVTX \Rightarrow \forall s' \in w\mathcal{E}_P\star, \exists s \in \star\mathcal{E}_Cw : s'.key = s.key \wedge s'.version > s.version$.

B.1. Model Analysis

In this section, we analyze our model for the execution of the Fabric system. We start by showing that any legal Fabric execution modeled in the aforementioned manner results in a TDAG. Here, we consider as legal a Fabric execution that contains only blocks included in the blockchain that have been produced by the ordering service.

Theorem B.5.

Assume that $\mathcal{L}_{\text{HLF}}$ is a legal Fabric execution. Then, the $\mathcal{G}_{\text{HLF}}$ instance modeling $\mathcal{L}_{\text{HLF}}$ is a TDAG.

PROOF: Here, we show that $\mathcal{G}_{\text{HLF}}$ fulfills all the conditions required in [Definition 2.2](#).

1. The genesis state must not have any producing or observing edges, and it must have a single producing edge. This condition is ensured by our definition of initialization transaction.
2. Every state, other than the genesis, must have a single producing edge. Assume by contradiction that $\exists s \in \mathcal{S}_{\text{HLF}} \setminus \{s_g\} : |\star\mathcal{E}_P s| > 1$. We rule out the case $|\star\mathcal{E}_P s| = 0$ because a state only exists in $\mathcal{G}_{\text{HLF}}$ if it has been produced by a transaction. This implies that there are at least two transactions, t and t' , in $\mathcal{G}_{\text{HLF}}$ that update the same key-value pair simultaneously. This, however, contradicts the assumption that a valid execution contains only transactions *sorted* by an ordering service.
3. Every state other than the genesis can have multiple successors, but at most one among them is connected with a consuming edge. The proof for this condition holds along the same lines as for the previous condition.
4. The graph must be weakly connected. Each new transaction reads or updates a key represented by an unconsumed state in the graph. Therefore, the overall graph is weakly connected.
5. The graph must not have cycles. Assume by contradiction that there is a cycle in $\mathcal{G}_{\text{HLF}}$. This necessarily implies that there are two transactions that produce the same state. However, as we argued before, this contradicts the fact that the ordering service establishes a total order among the transactions.

As we did with the Bitcoin model, here we show that validating Fabric transactions individually suffices to reason about properties of the complete Fabric execution. In particular, we show that if $\mathcal{G}_{\text{HLF}}$ is a valid TDAG, then the highest version number (*i.e.*, most recent) for any given key is represented in a unconsumed state of $\mathcal{G}_{\text{HLF}}$.

Definition B.6 (Most recent key-value pairs)

Consider that $\mathcal{G}_{\text{HLF}}$ is a TDAG modeling a Fabric execution. Then, we define *the states representing the most recent key-value pairs* as the set of states with the highest version number for each key, *i.e.*,

$$\{s \in \mathcal{S}_{\text{HLF}} : s' \in \mathcal{S}_{\text{HLF}} \wedge s.\text{key} = s'.\text{key} \Rightarrow s.\text{version} > s'.\text{version}\}.$$

Theorem B.7 (Unconsumed states represent most recent key-value pairs)

Assume that $\mathcal{G}_{\text{HLF}}$ is a valid TDAG and models a legal Fabric execution $\mathcal{L}_{\text{HLF}}$. Then, the unconsumed states of $\mathcal{G}_{\text{HLF}}$ represent the most recent key-value pairs.

PROOF: Assume by contradiction that [Theorem B.7](#) does not hold. Then, there must exist at least one transaction where the *version* field in the output state for a key is smaller than the *version* field in the input state for the same key, *i.e.*,

$$\exists t := (\mathcal{S} \dot{\cup} \{w\}, \mathcal{E}) \in \mathcal{T}_{\mathcal{G}_{\text{HLF}}}, \exists s, s' \in \mathcal{S} : (s, w) \in \mathcal{E} \wedge (w, s') \in \mathcal{E} \wedge s.\text{key} = s'.\text{key} \wedge s.\text{version} < s'.\text{version}.$$

However, $\mathbb{P}(t)$ would return *FALSE*, which contradicts the fact that $\mathcal{G}_{\text{HLF}}$ is a valid TDAG.

B.2 Modeling an Example of Execution for Fabric

Here we describe how we model an illustrative example of Fabric execution. We assume for simplicity that the endorsement policy requires a single endorsement for each transaction.

Throughout our description, we focus in the illustrative example depicted in [Figure 4](#). In particular, [Figure 4b](#) shows the Fabric execution $\mathcal{L}_{\text{HLF}} := \{b_g, b_1, b_2\}$, where $b_g := \{t_0 := EP\}$, $b_1 := \{t_1 := (END, f_1)\}$ and $b_2 := \{t_2 := (END', f_2)\}$. In a bit more detail, t_0 represents the initialization transaction that initializes the key-value pairs used later in the execution, and it is included in the genesis block. Moreover, t_1 represents an invoke transaction that calls the function f_1 , and t_2 represents another invoke transaction that calls f_2 in this case.

Now, we describe how we model such Fabric execution as an instance of $\mathcal{G}_{\text{HLF}}$ as shown in [Figure 4c](#):

- $t_0 := (\{s_g, s_0, s_1, s_2, w_0\}, \{(s_g, w_0), (w_0, s_0), (w_0, s_1), (w_0, s_2)\})$, where $(s_g, w_0) \in \mathcal{E}_C$ and $\{(w_0, s_0), (w_0, s_1), (w_0, s_2)\} \subseteq \mathcal{E}_P$. It represents the initialization transaction as described above. In more detail, $s_g := (params, 0)$, $w_0 := (TINITX, EP)$, $s_0 := (a, 1)$, $s_1 := (b, 1)$ and $s_2 := (c, 1)$.
- $t_1 := (\{s_0, s_1, s_2, s_3, s_4, s_5, w_1\}, \{(s_0, w_1), (s_1, w_1), (s_2, w_1), (w_1, s_3), (w_1, s_4), (w_1, s_5)\})$, where $\{(s_0, w_1), (s_1, w_1), (s_2, w_1)\} \subseteq \mathcal{E}_C$ and $\{(w_1, s_3), (w_1, s_4), (w_1, s_5)\} \subseteq \mathcal{E}_P$. A MIMO transaction that updates the values associates to keys a, b, c. In a bit more detail, $w_1 := (TINVTX, END)$, $s_3 := (a, 2)$, $s_4 := (b, 2)$ and $s_5 := (c, 2)$. The rest of states are defined as described for t_0 .
- $t_2 := (\{s_3, s_4, s_5, s_6, w_2\}, \{(s_3, w_2), (s_4, w_2), (s_5, w_2), (w_2, s_6)\})$, where $(s_3, w_2) \in \mathcal{E}_C$, $\{(s_4, w_2), (s_5, w_2)\} \subseteq \mathcal{E}_O$ and $(w_2, s_6) \in \mathcal{E}_P$. A MISO transaction that reads the values associates to keys a, b, c, and updates the value associated to key a. In a bit more detail, $w_2 := (TINVTX, END')$ and $s_6 := (a, 3)$. The rest of states are defines as described for t_1 .

ALGORITHM 1: Function f_1

```

PUTSTATE ("a", 0);
PUTSTATE ("b", 5);
PUTSTATE ("c", 3);

```

ALGORITHM 2: Function f_2

```

a ← GETSTATE ("a");
b ← GETSTATE ("b");
c ← GETSTATE ("c");
PUTSTATE ("a", a + b + c);

```

(a) Set of GETSTATE and PUTSTATE operations for each function defined in this example.

(b) Example HLF execution $\mathcal{L}_{HLF} := \{B_g, B_1, B_2\}$. Here, t_1 invokes f_1 and t_2 invokes f_2 .

(c) Example of $\mathcal{G}_{HLF}$ instance.

Figure 4: Illustrative example of the modeling of an execution of HLF. We model an execution that contains the setup transaction (t_0), followed by an invocation to f_1 (modeled in t_1) and finally an invocation to f_2 (modeled in t_2).

Citations

1. Nakamoto, S. (2008). *Bitcoin: A Peer-to-Peer Electronic Cash System*. Retrieved from <http://www.bitcoin.org/bitcoin.pdf> [↗](#)

2. *Ethereum, A Next-Generation Smart Contract and Decentralized Application Platform*. (2017). White Paper. [↵](#)
3. Cachin, C. (2016). *Architecture of the Hyperledger Blockchain Fabric*. [↵](#)
4. Hearn, M., & Brown, R. G. (2019). *Corda: A Distributed Ledger*. Technical White Paper. [↵](#)
5. Kwon, J., & Buchman, E. (2017). *Cosmos: A Network of Distributed Ledgers*. White Paper. [↵](#)
6. *Chain Protocol Whitepaper*. (2017). White Paper. [↵](#)
7. Garay, J. A., Kiayias, A., & Leonardos, N. (2015). The Bitcoin Backbone Protocol: Analysis and Applications. In *Advances in Cryptology: Eurocrypt 2015* (Vol. 9057, pp. 281–310). Berlin, Heidelberg: Springer. [↵](#)
8. Eyal, I., Gencer, A. E., Sirer, E. G., & Renesse, R. V. (2016). Bitcoin-ng: A Scalable Blockchain Protocol. In *USENIX symposium on networked systems design and implementation (NSDI)* (pp. 45–59). Berkeley, CA: USENIX. [↵](#)
9. Cachin, C., & Vukolić, M. (2017). *Blockchain Consensus Protocols in the Wild*. arXiv:1707.01873v2. [↵](#)
10. Meiklejohn, S., Pomarole, M., Jordan, G., Levchenko, K., McCoy, D., Voelker, G. M., & Savage, S. (2013). A Fistful of Bitcoins: Characterizing Payments Among Men with No Names. In *ACM internet measurement conference (imc)* (pp. 127–140). New York, NY: ACM Press. [↵](#)
11. Ben-Sasson, E., Chiesa, A., Garman, C., Green, M., Miers, I., Tromer, E., & Virza, M. (2014). Zerocash: Decentralized Anonymous Payments from Bitcoin. In *IEEE symposium on security and privacy* (pp. 459–474). New York: IEEE Press. [↵](#)
12. Ruffing, T., & Moreno-Sanchez, P. (2017). Mixing Confidential Transactions: Comprehensive Transaction Privacy for Bitcoin. In M. Brenner, K. Rohloff, J. Bonneau, J. Miller, P. Y. A. Ryan, V. Teague, ... M. Jakobsson (Eds.), *Bitcoin and Blockchain Research (bitcoin)* (Vol. 10323, pp. 133–154). Berlin, Heidelberg: Springer. [↵](#)
13. Bonneau, J., Miller, A., Clark, J., Narayanan, A., Kroll, J. A., & Felten, E. W. (2015). SoK: Research Perspectives and Challenges for Bitcoin and Cryptocurrencies. In

IEEE Symposium on Security and Privacy (pp. 104–121). New York: IEEE Press.

<https://doi.org/10.1109/SP.2015.14> [↵](#)

14. Tschorsch, F., & Scheuermann, B. (2016). Bitcoin and Beyond: A Technical Survey on Decentralized Digital Currencies. *IEEE Communications Surveys & Tutorials*, 18(3), 2084–2123. [↵](#)

15. Armknecht, F., Karame, G. O., Mandal, A., Youssef, F., & Zenner, E. (2015). Ripple: Overview and Outlook. In *Proc. Trust and Trustworthy Computing (trust)* (pp. 163–180). Berlin, Heidelberg: Springer. [↵](#)

16. Narayanan, A., Bonneau, J., Felten, E. W., Miller, A., & Goldfeder, S. (2016). *Bitcoin and Cryptocurrency Technologies – A Comprehensive Introduction*. Princeton, NJ: Princeton University Press. [↵](#)

17. Atzei, N., Bartoletti, M., Lande, S., & Zunino, R. (2018). A Formal Model of Bitcoin Transactions. In S. Meiklejohn & K. Sako (Eds.), *Financial Cryptography and Data Security* (Vol. 10957, pp. 541–560). Berlin, Heidelberg: Springer. [↵](#)

18. Bartoletti, M., & Zunino, R. (2018). BitML: A Calculus for Bitcoin Smart Contracts. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security* (pp. 83–100). [↵](#)

19. Elmasri, R., & Navathe, S. B. (2011). *Fundamentals of Database Systems* (6th ed.). Upper Saddle River, NJ: Addison-Wesley. [↵](#)

20. Cachin, C., Caro, A. D., Moreno-Sanchez, P., Tackmann, B., & Vukolić, M. (2017). *The Transaction Graph for Modeling Blockchain Semantics*. Cryptology ePrint Archive, Report 2017/1070. [↵](#)

21. Antonopoulos, A. (2014). *Mastering Bitcoin*. Sebastopol, CA: O’Reilly. [↵](#)

22. Androulaki, E., Cachin, C., Christidis, K., Murthy, C., Nguyen, B., & Vukolić, M. (2016). *Next Consensus Architecture Proposal*. Hyperledger Wiki, Fabric Design Documents. Retrieved from <https://github.com/hyperledger/fabric/blob/master/proposals/r1/Next-Consensus-Architecture-Proposal.md> [↵](#)

23. Vukolić, M. (2017). Rethinking Permissioned Blockchains. In *Proc. ACM Workshop on Blockchain, Cryptocurrencies and Contracts (bcc ’17)* (pp. 3–7). New York, NY: ACM Press. [↵](#)

24. Goldwasser, S., & Bellare, M. (n.d.). *Lecture Notes on Cryptography*. [↵](#)