

Blockchain Abstract Data Type

Emmanuelle Anceaume, Antonella Del Pozzo, Romaric Ludinard, Maria Potop-Butucaru, Sara Tucci-Piergiovanni

► **To cite this version:**

Emmanuelle Anceaume, Antonella Del Pozzo, Romaric Ludinard, Maria Potop-Butucaru, Sara Tucci-Piergiovanni. Blockchain Abstract Data Type. [Research Report] Sorbonne Université, CNRS, Laboratoire d'Informatique de Paris 6, LIP6, Paris, France. 2018, pp.1-30. <hal-01718480v2>

HAL Id: hal-01718480

<https://hal.sorbonne-universite.fr/hal-01718480v2>

Submitted on 12 May 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Blockchain Abstract Data Type

Emmanuelle Anceaume[‡], Antonella Del Pozzo^{*}, Romaric Ludinard^{**},
Maria Potop-Butucaru[†], Sara Tucci-Piergiovanni^{*}

[‡]CNRS, IRISA

^{*}CEA LIST, PC 174, Gif-sur-Yvette, 91191, France

^{**} IMT Atlantique, IRISA

[†]Sorbonne Université, CNRS, Laboratoire d’Informatique de Paris 6, LIP6, Paris, France

Abstract

The presented work continues the line of recent distributed computing community efforts dedicated to the theoretical aspects of blockchains. This paper is the first to specify blockchains as a composition of *abstract data types* all together with a hierarchy of *consistency criteria* that formally characterizes the histories admissible for distributed programs that use them. Our work is based on an original oracle-based construction that, along with new consistency definitions, captures the eventual convergence process in blockchain systems. The paper presents as well some results on implementability of the presented abstractions and a mapping of representative existing blockchains from both academia and industry in our framework.

1 Introduction

The paper proposes a new data type to formally model blockchains and their behaviors. We aim at providing consistency criteria to capture the correct behavior of current blockchain proposals in a unified framework. It is already known that some blockchain implementations solve eventual consistency of an append-only queue using Consensus [6, 5]. The question is about the consistency criterion of blockchains as Bitcoin [26] and Ethereum [31] that technically do not solve Consensus, and their relation with Consensus in general.

We advocate that the key point to capture blockchain behaviors is to define consistency criteria allowing mutable operations to create forks and restricting the values read, i.e. modeling the data structure as *an append-only tree* and not as an append-only queue. This way we can easily define a semantics equivalent to eventual consistent append-only queue but as well weaker semantics. More in detail, we define a semantic equivalent to eventual consistent append-only queue by restricting any two reads to return two chains such that one is the prefix of the other. We call this consistency property Strong Prefix (already introduced in [20]). Additionally, we define a weaker semantics restricting any two reads to return chains that have a divergent prefix for a finite interval of the history. We call this consistency property Eventual Prefix.

Another peculiarity of blockchains lies in the notion of *validity* of blocks, i.e. the blockchain must contain only blocks that satisfy a given predicate. Let us note that validity can be achieved

through proof-of-work (Dwork and Naor [15]) or other agreement mechanisms. We advocate that to abstract away implementation-specific validation mechanisms, the validation process must be encapsulated in an oracle model separated from the process of updating the data structure. Because the oracle is the only generator of valid blocks and only valid blocks can be appended, it follows that it is the oracle that grants the access to the data structure and it might also own a synchronization power to control the number of forks, in terms of branches of the tree from a given block. In this respect we define oracles models such that, depending on the model, the number of forks from a given block can be: unbounded, up to $k > 1$, and $k = 1$ (no fork) for the strongest oracle model.

The blockchain is then abstracted by an oracle-based construction in which the update and consistency of the tree data structure depends on the validation and synchronization power of the oracle.

The main contribution of the paper is a formal unified framework providing blockchain consistency criteria that can be combined with oracle models in a proper *hierachy of abstract data types* [29] independent of the underlying communication and failure model. Thanks to the establishment of the formal framework the following implementability results are shown:

- The strongest oracle, guaranteeing no fork, has Consensus number ∞ in the Consensus hierarchy of concurrent objects [21] (Theorem 4.2). It must be noted that we considered Consensus defined in [11, 19, 8], in which the *Validity* property states that a valid block can be decided even if sent by a faulty process.
- The weakest oracle, which validates a potential unbounded number of blocks to be appended to a given block, has Consensus number 1 (Theorem 4.3).
- The impossibility to guarantee Strong Prefix in a message-passing system if forks are allowed (Theorem 4.8). This means that Strong Prefix needs the strongest oracle to be implemented, which is at least as strong as Consensus.
- A necessary condition (Theorem 4.7) for Eventual Prefix in a message-passing system, called *Update Agreement* stating that each update sent by a correct process must be eventually received by every correct process. The result implies that it is impossible to implement Eventual Prefix if even only one message sent by a correct process is dropped.

The proposed framework along with the above-mentioned results helps in classifying existing blockchains in terms of their consistency and implementability. We used the framework to classify several blockchain proposals. We showed that Bitcoin [26] and Ethereum [31] have a validation mechanism that maps to our weakest oracle and then they only implement Eventual prefix, while other proposals maps to our strongest oracle, falling in the class of those that guarantee Strong Prefix (e.g. Hyperledger Fabric [5], PeerCensus [12], ByzCoin [24], see Section 5 for further details). **Related Work.** Formalisation of blockchains in the lens of distributed computing has been recognized as an extremely important topic [22]. The topic is recent and to the best of our knowledge, no other attempt proposed a unified framework capturing both Consensus-based and proof-of-work blockchains, as the presented paper aims at proposing.

In [1], the authors present a study about the relationship of BFT consensus and blockchains. In order to abstract the proof-of-work mechanism the authors propose a specific oracle, in the same spirit of our oracle abstraction. While their oracle is more specific then ours, since it makes a direct reference to proof-of-work properties, it offers as well a fairness property. Note that we do

not formalize fairness properties in this paper, we only offer a generic merit parameter that can be used to define fairness. Let us note that apart from the fairness property, our oracle captures the semantics of [1]’s oracle.

In parallel and independently of the work in [3], [6] proposes a formalization of distributed ledgers modeled as an ordered list of records. The authors propose in their formalization three consistency criteria: eventual consistency, sequential consistency and linearizability. They discuss how Hyperledger Fabric implements eventual consistency and propose implementations for sequential consistency and linearizability using a total order broadcast abstraction. Interestingly, they show that a distributed ledger that provides eventual consistency can be used to solve the consensus problem. These findings confirm our results about the necessity of Consensus to solve Strong Prefix and corroborate our mapping of Hyperledger Fabric. On the other hand the proposed formalization does not propose weaker consistency semantics more suitable for proof-of-work blockchains as BitCoin. Indeed, [6] continues and it is complementary to the work on the first formalisation of Bitcoin as a distributed ledger proposed in [4] where the distributed ledger is modelled as a simple register. These works suggest different abstractions to model proof-of-work and Consensus-based blockchains, respectively. The presented paper, on the other hand, thanks to our oracle-based construction (not present in [6], [4]) generalizes both [4] and [6] to encompass both kind of blockchains in a unified framework.

Finally, [20] presents an implementation of the Monotonic Prefix Consistency (MPC) criterion and showed that no criterion stronger than MPC can be implemented in a partition-prone message-passing system. Nicely, this result and more in general solvability results for eventual consistency [13] immediately apply to our Strong Prefix criterion.

2 Preliminaries on shared object specifications based on Abstract Data Types

The basic idea underlying the use of abstract data types is to specify shared objects using two complementary facets [28]: a sequential specification that describes the semantics of the object, and a consistency criterion over concurrent histories, i.e. the set of admissible executions in a concurrent environment. In this work we are interested in consistency criteria achievable in a distributed environment in which processes are sequential and communicate through message-passing.

2.1 Abstract Data Type (ADT)

The model used to specify an abstract data type is a form of transducer, as Mealy’s machines, accepting an infinite but countable number of states. The values that can be taken by the data type are encoded in the abstract state, taken in a set Z . It is possible to access the object using the symbols of an input alphabet A . Unlike the methods of a class, the input symbols of the abstract data type do not have arguments. Indeed, as one authorizes a potentially infinite set of operations, the call of the same operation with different arguments is encoded by different symbols. An operation can have two types of effects. First, it can have a side-effect that changes the abstract state, the corresponding transition in the transition system being formalized by a transition function τ . Second, operations can return values taken in an output alphabet B , which depend on the state in which they are called and an output function δ . For example, the pop operation in a stack removes the element at the top of the stack (its side effect) and returns that element (its output).

The formal definition of abstract data types is as follows.

Definition 2.1. (Abstract Data Type T) An abstract data type is a 6-tuple $T = \langle A, B, Z, \xi_0, \tau, \delta \rangle$ where:

- A and B are countable sets called input alphabet and output alphabet;
- Z is a countable set of abstract states and ξ_0 is the initial abstract state;
- $\tau : Z \times A \rightarrow Z$ is the transition function;
- $\delta : Z \times A \rightarrow B$ is the output function.

Definition 2.2. (Operation) Let $T = \langle A, B, Z, \xi_0, \tau, \delta \rangle$ be an abstract data type. An *operation* of T is an element of $\Sigma = A \cup (A \times B)$. We refer to a couple $(\alpha, \beta) \in A \times B$ as α/β . We extend the transition function τ over the operations and apply τ on the operations input alphabet:

$$\tau_T : \begin{cases} Z \times \Sigma \rightarrow Z \\ (\xi, \alpha) \mapsto \tau(\xi, \alpha) \text{ if } \alpha \in A \\ (\xi, \alpha/\beta) \mapsto \tau(\xi, \alpha) \text{ if } \alpha/\beta \in A \times B \end{cases}$$

2.2 Sequential specification of an ADT

An abstract data type, by its transition system, defines the sequential specification of an object. That is, if we consider a path that traverses its system of transitions, then the word formed by the subsequent labels on the path is part of the sequential specification of the abstract data type, i.e. it is a sequential history. The language recognized by an ADT is the set of all possible words. This language defines the sequential specification of the ADT. More formally,

Definition 2.3. (Sequential specification $L(T)$) A finite or infinite sequence $\sigma = (\sigma_i)_{i \in D} \in \Sigma^\infty$, $D = \mathbb{N}$ or $D = \{0, \dots, |\sigma| - 1\}$ is a *sequential history* of an abstract data type T if there exists a sequence of the same length $(\xi_{i+1})_{i \in D} \in Z^\infty$ (ξ_0 has already been defined as the initial state) of states of T such that, for any $i \in D$,

- the output alphabet of σ_i is compatible with ξ_i : $\xi_i \in \delta_T^{-1}(\sigma_i)$;
- the execution of the operation σ_i is such that the state changed from ξ_i to ξ_{i+1} : $\tau_T(\xi_i, \sigma_i) = \xi_{i+1}$.

The *sequential specification* of T is the set of all its possible sequential histories $L(T)$.

2.3 Concurrent histories of an ADT

Concurrent histories are defined considering asymmetric event structures, i.e., partial order relations among events executed by different processes [28].

Definition 2.4. (Concurrent history H) The execution of a program that uses an abstract data type $T = \langle A, B, Z, \xi_0, \tau, \delta \rangle$ defines a concurrent history $H = \langle \Sigma, E, \Lambda, \mapsto, \prec, \nearrow \rangle$, where

- $\Sigma = A \cup (A \times B)$ is a countable set of operations;

- E is a countable set of events that contains all the ADT operations invocations and all ADT operation response events;
- $\Lambda : E \rightarrow \Sigma$ is a function which associates events to the operations in Σ ;
- \mapsto : is the process order relation over the events in E . Two events are ordered by \mapsto if they are produced by the same process;
- \prec : is the operation order, irreflexive order over the events of E . For each couple $(e, e') \in E^2$, if e is an operation invocation and e' is the response for the same operation then $e \prec e'$, if e' is the invocation of an operation occurred at time t' and e is the response of another operation occurred at time t with $t < t'$ then $e \prec e'$;
- \nearrow : is the program order, irreflexive order over E , for each couple $(e, e') \in E^2$ with $e \neq e'$ if $e \mapsto e'$ or $e \prec e'$ then $e \nearrow e'$.

2.4 Consistency criterion

The consistency criterion characterizes which concurrent histories are admissible for a given abstract data type. It can be viewed as a function that associates a concurrent specification to abstract data types. Specifically,

Definition 2.5. (Consistency criterion C) A consistency criterion is a function

$$C : \mathcal{T} \rightarrow \mathcal{P}(\mathcal{H})$$

where \mathcal{T} is the set of abstract data types, \mathcal{H} is a set of histories and $\mathcal{P}(\mathcal{H})$ is the sets of parts of \mathcal{H} .

Let \mathcal{C} be the set of all the consistency criteria. An algorithm A_T implementing the ADT $T \in \mathcal{T}$ is C -consistent with respect to criterion $C \in \mathcal{C}$ if all the operations terminate and all the admissible executions are C -consistent, i.e. they belong to the set of histories $C(T)$.

3 BlockTree and Token oracle ADTs

In this section we present the BlockTree and the token Oracle ADTs along with consistency criteria.

3.1 BlockTree ADT

We formalize the data structure implemented by blockchain-like systems as a *directed rooted tree* $bt = (V_{bt}, E_{bt})$ called *BlockTree*. Each vertex of the BlockTree is a *block* and any edge points backward to the root, called *genesis block*. The height of a block refers to its distance to the root. We denote by b_k a block located at height k . By convention, the root of the BlockTree is denoted by b_0 . Blocks are said valid if they satisfy a predicate P which is application dependent (for instance, in Bitcoin, a block is considered valid if it can be connected to the current blockchain and does not contain transactions that double spend a previous transaction). We represent by \mathcal{B} a countable and non empty set of blocks and by $\mathcal{B}' \subseteq \mathcal{B}$ a countable and non empty set of valid blocks, i.e., $\forall b \in \mathcal{B}'$, $P(b) = \top$. By assumption $b_0 \in \mathcal{B}'$; We also denote by \mathcal{BC} a countable non empty set of blockchains, where a blockchain is a path from a leaf of bt to b_0 . A blockchain is denoted by bc . Finally, \mathcal{F} is




Figure 1: A possible path of the transition system defined by the BT-ADT. We use the following syntax on the edges: **operation/output**.

a countable non empty set of selection functions, $f \in \mathcal{F} : \mathcal{BT} \rightarrow \mathcal{BC}$; $f(bt)$ selects a blockchain bc from the BlockTree bt (note that b_0 is not returned) and if $bt = b_0$ then $f(b_0) = b_0$. This reflects for instance the longest chain or the heaviest chain used in some blockchain implementations. The selection function f and the predicate P are parameters of the ADT which are encoded in the state and do not change over the computation.

The following notations are also deeply used: $\{b_0\} \frown f(bt)$ represents the concatenation of b_0 with the blockchain of bt ; and $\{b_0\} \frown f(bt) \frown \{b\}$ represents the concatenation of b_0 with the blockchain of bt and a block b ;

3.1.1 Sequential specification of the BlockTree

The sequential specification of the BlockTree is defined as follows.

Definition 3.1 (BlockTree ADT (*BT-ADT*)). The BlockTree Abstract Data Type is the 6-tuple $\text{BT-ADT} = \langle A = \{\text{append}(b), \text{read}(): b \in \mathcal{B}\}, B = \mathcal{BC} \cup \{\text{true}, \text{false}\}, Z = \mathcal{BT} \times \mathcal{F} \times (\mathcal{B} \rightarrow \{\text{true}, \text{false}\}) \rangle$, $\xi_0 = (bt^0, f, \tau, \delta)$, where the transition function $\tau : Z \times A \rightarrow Z$ is defined by

- $\tau((bt, f, P), \text{append}(b)) = (\{b_0\} \frown f(bt) \frown \{b\}, f, P)$ if $b \in \mathcal{B}'$; (bt, f, P) otherwise;
- $\tau((bt, f, P), \text{read}()) = (bt, f, P)$,

and the output function $\delta : Z \times A \rightarrow B$ is defined by

- $\delta((bt, f, P), \text{append}(b)) = \text{true}$ if $b \in \mathcal{B}'$; false otherwise;
- $\delta((bt, f, P), \text{read}()) = \{b_0\} \frown f(bt)$;
- $\delta((bt_0, f, P), \text{read}()) = b_0$.

The semantic of the `read` and the `append` operations directly depend on the selection function $f \in \mathcal{F}$. In this work we let this function generic to suit the different blockchain implementations. In the same way, predicate P is let unspecified. The predicate P mainly abstracts the creation process of a block, which may fail or successfully terminate. This process will be further specified in Section 3.2.

3.1.2 Concurrent specification of a BT-ADT and consistency criteria

The concurrent specification of the BT-ADT is the set of concurrent histories. A *BT-ADT* consistency criterion is a function that returns the set of concurrent histories admissible for a BlockTree abstract data type. We define two *BT* consistency criteria: *BT Strong consistency* and *BT Eventual consistency*. For ease of readability, we employ the following notations:

- $E(a^*, r^*)$ is an infinite set containing an infinite number of `append()` and `read()` invocation and response events;
- $E(a, r^*)$ is an infinite set containing (i) a finite number of `append()` invocation and response events and (ii) an infinite number of `read()` invocation and response events;
- $e_{inv}(o)$ and $e_{rsp}(o)$ indicate respectively the invocation and response event of an operation o ; and $e_{rsp}(r) : bc$ denotes the returned blockchain bc associated with the response event $e_{rsp}(r)$;
- $\text{score} : \mathcal{BC} \rightarrow \mathbb{N}$ denotes a monotonic increasing deterministic function that takes as input a blockchain bc and returns a natural number s as score of bc , which can be the height, the weight, etc. Informally we refer to such value as the score of a blockchain; by convention we refer to the score of the blockchain uniquely composed by the genesis block as s_0 , i.e. $\text{score}(\{b_0\}) = s_0$. Increasing monotonicity means that $\text{score}(bc \frown \{b\}) > \text{score}(bc)$;
- $\text{mcps} : \mathcal{BC} \times \mathcal{BC} \rightarrow \mathbb{N}$ is a function that given two blockchains bc and bc' returns the score of the maximal common prefix between bc and bc' ;
- $bc \sqsubseteq bc'$ iff bc prefixes bc' .

BT Strong consistency. The BT Strong Consistency criterion is the conjunction of the following four properties. The block validity property imposes that each block in a blockchain returned by a `read()` operation is *valid* (i.e., satisfies predicate P) and has been inserted in the BlockTree with the `append()` operation. The Local monotonic read states that, given the sequence of `read()` operations at the same process, the score of the returned blockchain never decreases. The Strong prefix property states that for each couple of read operations, one of the returned blockchains is a prefix of the other returned one (i.e., the prefix never diverges). Finally, the Ever growing tree states that scores of returned blockchains eventually grow. More precisely, let s be the score of the blockchain returned by a read response event r in $E(a^*, r^*)$, then for each `read()` operation r , the set of `read()` operations r' such that $e_{rsp}(r) \not\prec e_{inv}(r')$ that do not return blockchains with a score greater than s is finite. More formally, the BT Strong consistency criterion is defined as follows:

Definition 3.2 (BT Strong Consistency criterion (*SC*)). A concurrent history $H = \langle \Sigma, E, \Lambda, \mapsto, \prec, \succ, \nearrow \rangle$ of the system that uses a BT-ADT verifies the BT Strong Consistency criterion if the following properties hold:

- **Block validity:** $\forall e_{rsp}(r) \in E, \forall b \in e_{rsp}(r) : bc, b \in \mathcal{B}' \wedge \exists e_{inv}(\text{append}(b)) \in E, e_{inv}(\text{append}(b)) \nearrow e_{rsp}(r)$.
- **Local monotonic read:**

$$\forall e_{rsp}(r), e_{rsp}(r') \in E^2, \text{ if } e_{rsp}(r) \mapsto e_{inv}(r'), \text{ then } \text{score}(e_{rsp}(r) : bc) \leq \text{score}(e_{rsp}(r') : bc').$$

- **Strong prefix:**

$$\forall e_{rsp}(r), e_{rsp}(r') \in E^2, (e_{rsp}(r') : bc' \sqsubseteq e_{rsp}(r) : bc) \vee (e_{rsp}(r) : bc \sqsubseteq e_{rsp}(r') : bc').$$

- **Ever growing tree:** $\forall e_{rsp}(r) \in E(a^*, r^*), s = \text{score}(e_{rsp}(r) : bc)$ then

$$|\{e_{inv}(r') \in E \mid e_{rsp}(r) \nearrow e_{inv}(r'), \text{score}(e_{rsp}(r') : bc') \leq s\}| < \infty.$$

Figure 2 shows a concurrent history H admissible by the BT Strong consistency criterion. In this example the score is the length l of the blockchain and the selection function f selects the longest blockchain, and in case of equality, selects the largest based on the lexicographical order. For ease of readability, we do not depict the `append()` operation. We assume the block validity property is satisfied. The Local monotonic read is easily verifiable as for each couple of read blockchains one prefixes the other. The first `read()` r operation, enclosed in a black rectangle, is taken as reference to check the consistency criterion (the criterion has to be iteratively verified for each `read()` operation). Let l be the score of the blockchain returned by r . We can identify two sets, enclosed in rectangles defined by different patterns: (i) the finite sets of `read()` operations such that the score associated to each blockchain returned is smaller than or equal to l , and (ii) the infinite set of `read()` operations such that the score is greater than l . We can iterate the same reasoning for each `read()` operation in H . Thus H satisfies the Ever growing tree property.




Figure 2: Concurrent history that satisfies the BT Strong consistency criterion. In such scenario f selects the longest blockchain and the blockchain score is length l .

BT Eventual consistency. The BT Eventual consistency criterion is the conjunction of the block validity, the Local monotonic read and the Ever growing tree of the BT Strong consistency criterion together with the Eventual prefix which states that for each blockchain returned by a `read()` operation with s as score, then eventually all the `read()` operations will return blockchains sharing the same maximum common prefix at least up to s . Say differently, let H be a history with an infinite number of `read()` operations, and let s be the score of the blockchain returned by a read r , then the set of `read()` operations r' , such that $e_{rsp}(r) \nearrow e_{inv}(r')$, that do not return blockchains sharing the same prefix at least up to s is finite.

Definition 3.3 (Eventual prefix property). Given a concurrent history $H = \langle \Sigma, E(a, r^*), \Lambda, \mapsto, \prec, \nearrow \rangle$ of the system that uses a BT-ADT, we denote by s , for any read operation $r \in \Sigma$ such that $\exists e \in E(a, r^*), \Lambda(r) = e$, the score of the returned blockchain, *i.e.*, $s = \text{score}(e_{rsp}(r) : bc)$. We denote by E_r the set of response events of read operations that occurred after r response, *i.e.*

$E_r = \{e \in E \mid \exists r' \in \Sigma, r' = \text{read}, e = e_{rsp}(r') \wedge e_{rsp}(r) \not\prec e_{rsp}(r')\}$. Then, H satisfies the Eventual prefix property if for all $\text{read}()$ operations $r \in \Sigma$ with score s ,

$$|\{(e_{rsp}(r_h), e_{rsp}(r_k)) \in E_r^2 \mid h \neq k, \text{mpcs}(e_{rsp}(r_h) : bc_h, e_{rsp}(r_k) : bc_k) < s\}| < \infty$$

The Eventual prefix properties captures the fact that two or more concurrent blockchains can co-exist in a finite interval of time, but that ly all the participants adopts a same branch for each cut of the history. This cut of the history is defined by a read that picks up a blockchain with a given score.

Based on this definition, the BT Eventual consistency criterion is defined as follows:

Definition 3.4 (BT Eventual consistency criterion *EC*). A concurrent history $H = \langle \Sigma, E, \Lambda, \mapsto, \prec, \succ \rangle$ of the system that uses a BT-ADT verifies the *BT Eventual consistency criterion* if it satisfies the Block validity, Local monotonic read, Ever growing tree, and the Eventual prefix properties.




Figure 3: Concurrent history that satisfies the Eventual BT consistency criterion. In such scenario f selects the longest blockchain and the blockchain score is the length l . In case (a) and case (b) the concurrent history is the same but different sets are outlined.

Figure 3 shows a concurrent history that satisfies the Eventual prefix property but not the Strong prefix one. Strong Prefix is not satisfied as blockchain¹ $b_0 \widehat{1}$ returned from the first $\text{read}()$ at process j is not a prefix of blockchain $b_0 \widehat{2} \widehat{4}$ returned from the first read at process i . Note that we adopt the same conventions as for the example depicted in Figure 2 regarding the score, length and $\text{append}()$ operations. We assume that the Block validity property is satisfied. The

¹For ease of readability we extend the notation $b_i \widehat{b}_j$ to represent concatenated blocks in a blockchain.

Local monotonic read property is easily verifiable. In both Figures 3a and 3b, the first $\text{read}()$ r operation at i , enclosed in a black rectangle, is taken as reference to check the consistency criterion (the criterion has to be iteratively verified for each $\text{read}()$ operation). Let l be the score of the blockchain returned by r . In Figure 3b we can identify two sets, enclosed in rectangles defined by different patterns: (i) the finite set of $\text{read}()$ operations sharing a maximum common prefix score (mcps) smaller than l (the set to check for the satisfiability of the Eventual Prefix property), and (ii) the infinite set of $\text{read}()$ operations such that for each couple of them bc, bc' , $\text{mcps}(bc, bc') \geq l$. We can iterate the same reasoning for each $\text{read}()$ operation in H . Thus H satisfies the Eventual Prefix property. Figure 4 shows a history that does not satisfy any consistency criteria defined so far.



Figure 4: Concurrent history that does not satisfy any BT consistency criteria. In such scenario f selects the longest blockchain and the blockchain score is the length l .

Relationships between EC and SC . Let us denote by \mathcal{H}_{EC} and by \mathcal{H}_{SC} the set of histories satisfying respectively the EC and the SC consistency criteria.

Theorem 3.1. Any history H satisfying SC criterion satisfies EC and $\exists H$ satisfying EC that does not satisfy SC , i.e., $\mathcal{H}_{SC} \subset \mathcal{H}_{EC}$.

Proof. $EC \leq SC$ implies that $\mathcal{H}_{SC} \subset \mathcal{H}_{EC}$, and $\mathcal{H}_{SC} \subset \mathcal{H}_{EC}$ implies that $\forall H \in \mathcal{H}_{SC} \Rightarrow H \in \mathcal{H}_{EC}$. By hypothesis, H verifies the Ever Growing Tree property, thus $\forall e_{rsp}(r) \in E(a^*, r^*)$ with $s = \text{score}(e_{rsp}(r) : bc)$ then set $\{e_{inv}(r') \in E | e_{rsp}(r) \nearrow e_{inv}(r'), \text{score}(e_{rsp}(r'') : bc) \leq s\}$ is finite, and thus, there is an infinite set $\{e_{inv}(r') \in E | e_{rsp}(r) \nearrow e_{inv}(r'), \text{score}(e_{rsp}(r'') : bc) > s\}$. The Strong prefix property guarantees that $\forall e_{rsp}(r), e_{rsp}(r') \in H, (e_{rsp}(r) : bc \sqsubseteq e_{rsp}(r') : bc) \vee (e_{rsp}(r) :$

$bc \sqsubseteq e_{rsp}(r') : bc'$), thus in this infinite set, all the `read()` operations return blockchains sharing the same maximum prefix whose score is at least $s+1$, which satisfies the Eventual prefix property. The Eventual Prefix property demands that for each $\forall e_{rsp}(r) \in E(a, r^*)$ with $s = \text{score}(e_{rsp}(r) : bc)$ there is an infinite set defined as $\{(e_{rsp}(r_h), e_{rsp}(r_k)) \in E_r^2 | h \neq k, \text{mpcs}(e_{rsp}(r_h) : bc_h, e_{rsp}(r_k) : bc_k) \geq s\}$ where E_r denotes the set of response events of read operations that occurred after r response. To conclude the proof we need to find a $H \in \mathcal{H}_{EC}$ and $H \notin \mathcal{H}_{SC}$. Any H in which at least two `read()` operations return a blockchain sharing the same prefix but diverging in their suffix violate the Strong prefix property, which concludes the proof. \square

Let us remark that the BlockTree allows at any time to create a new branch in the tree, which is called a *fork* in the blockchain literature. Moreover, an append is successful only if the input block is valid with respect to a predicate. This means that histories with no append operations are trivially admitted. In the following we will introduce a new abstract data type called Token Oracle that when combined with the BlockTree will help in (i) validating blocks and (ii) controlling forks. We will first formally introduce the Token Oracle in Section 3.2 and then we will define the properties on the BlockTree augmented with the Token Oracle in Section 3.4.

3.2 Token oracle Θ -ADT

In this section we formalize the Token Oracle Θ to capture the creation of blocks in the BlockTree structure. The block creation process requires that the new block must be closely related to an already existing valid block in the BlockTree structure. We abstract this implementation-dependent process by assuming that a process will obtain the right to chain a new block b_ℓ to b_h if it successfully gains a token tkn_h from the token oracle Θ . Once obtained, the proposed block b_ℓ is considered as valid, and will be denoted by $b_\ell^{tkn_h}$. By construction $b_\ell^{tkn_h} \in \mathcal{B}'$. In the following, in order to be as much general as possible, we model blocks as objects. More formally, when a process wants to access a generic object obj_h , it invokes the `getToken(obj_h, obj_\ell)` operation with object obj_ℓ from set $\mathcal{O} = \{obj_1, obj_2, \dots\}$. If `getToken(obj_h, obj_\ell)` operation is successful, it returns an object $obj_\ell^{tkn_h} \in \mathcal{O}'$, where (i) tkn_h is the token required to access object obj_h and (ii) each object $obj_k \in \mathcal{O}'$ is valid with respect to predicate P , i.e. $P(obj_k) = \top$. We say that a token is *generated* each time it is provided to a process and it is *consumed* when the oracle grants the right to connect it to the previous object. Each token can be consumed at most once. To consume a token we define the token consumption `consumeToken(obj_\ell^{tkn_h})` operation, where the consumed token tkn_h is the token required for the object obj_h . A maximal number of tokens k for an object obj_h is managed by the oracle. The `consumeToken(obj_\ell^{tkn_h})` side-effect on the state is the insertion of the object $obj_\ell^{tkn_h}$ in a set K_h as long as the cardinality of such set is less than k .

In the following we specify two token oracles, which differ in the way tokens are managed. The first oracle, called *prodigal* and denoted by Θ_P , has no upper bound on the number of tokens consumed for an object, while the second oracle Θ_F , called *frugal*, and denoted by Θ_F , assures controls that no more than k token can be consumed for each object.

Θ_P when combined with the BlockTree abstract data type will only help in validating blocks, while Θ_F manages tokens in a more controlled way to guarantee that no more than k forks can occur on a given block.

3.2.1 Θ_P -ADT and Θ_F -ADT definitions

For both oracles, when $\text{getToken}(obj_k, obj_h)$ operation is invoked, the oracle provides a token with a certain probability $p_{\alpha_i} > 0$ where α_i is a “merit” parameter characterizing the invoking process i .² Note that the oracle knows α_i of the invoking process i , which might be unknown to the process itself. For each merit α_i , the state of the token oracle embeds an infinite tape where each cell of the tape contains either tkn or \perp . Since each tape is identified by a specific α_i and p_{α_i} , we assume that each tape contains a pseudorandom sequence of values in $\{tkn, \perp\}$ depending on α_i .³ When a $\text{getToken}(obj_k, obj_h)$ operation is invoked by a process with merit α_i , the oracle pops the first cell from the tape associated to α_i , and a token is provided to the process if that cell contains tkn .

Both oracles also enjoy an infinite array of sets, one for each object, which is populated each time a token is consumed for a specific object. When the set cardinality reaches k then no more tokens can be consumed for that object. For a sake of generality, Θ_P is defined as Θ_F with $k = \infty$ while for Θ_F a predetermined $k \in \mathbb{N}$ is specified.




Figure 5: The Θ_F abstract state. The infinite K array, where at the beginning each set is initialized as empty and the infinite set of infinite tapes, one for each merit α_i in \mathcal{A} .

We first introduce some definitions and notations.

- $\mathcal{O} = \{obj_1, obj_2, \dots\}$, infinite set of generic objects uniquely identified by their index i ;
- $\mathcal{O}' \subset \mathcal{O}$, the subset of objects valid with respect to predicate P , i.e. $\forall obj'_i \in \mathcal{O}', P(obj'_i) = \top$.
- $\mathfrak{T} = \{tkn_1, tkn_2, \dots\}$ infinite set of tokens;
- $\mathcal{A} = \{\alpha_1, \alpha_2, \dots\}$ an infinite set of rational values;
- \mathcal{M} is a countable not empty set of mapping functions $m(\alpha_i)$ that generate an infinite pseudo random tape $tape_{\alpha_i}$ such that the probability to have in a cell the string tkn is related to a specific α_i , $m \in \mathcal{M} : \mathcal{A} \rightarrow \{tkn, \perp\}^*$;
- $K[\]$ is a infinite array of sets (one per object) of elements in \mathcal{O}' . All the sets are initialized as empty and can be fulfilled with at most k elements, where $k \in \mathbb{N}$ is a parameter of the oracle ADT;
- $pop : \{tkn, \perp\}^* \rightarrow \{tkn, \perp\}^*$, $pop(a \cdot w) = w$;
- $head : \{tkn, \perp\}^* \rightarrow \{tkn, \perp\}^*$, $head(a \cdot w) = a$;

²The merit parameter can reflect for instance the hashing power of the invoking process.

³We assume a pseudorandom sequence mostly indistinguishable from a Bernoulli sequence consisting of a finite or infinite number of independent random variables X_1, X_2, X_3, \dots such that (i) for each k , the value of X_k is either tkn or \perp ; and (ii) $\forall X_k$ the probability that $X_k = tkn$ is p_{α_i} .




Figure 6: A possible path of the transition system defined by the Θ_F and Θ_P -ADTs. We use the following syntax on the edges: operation/output.

- $\text{add} : \{K\} \times \mathbb{N} \times \mathcal{O}' \rightarrow \{K\}$, $\text{add}(K, i, obj_\ell^{tkn_h}) = K : K[i] = K[i] \cup \{obj_\ell^{tkn_h}\}$ if $|K[i]| < k$; else $K[i] = K[i]$;
- $\text{get} : \{K\} \times \mathbb{N} \rightarrow \mathbb{N}$, $\text{get}(K, i) = K[i]$;

Definition 3.5. (Θ_F -ADT Definition). The Θ_F Abstract Data type is the 6-tuple Θ_F -ADT $= \langle A = \{\text{getToken}(obj_h, obj_\ell), \text{consumeToken}(obj_\ell^{tkn_h}) : obj_h, obj_\ell^{tkn_h} \in \mathcal{O}', obj_\ell \in \mathcal{O}, tkn_h \in \mathfrak{T}\}$, $B = \mathcal{O}' \cup \text{Boolean}$, $Z = m(\mathcal{A})^* \times \{K\} \times k \cup \{\text{pop}, \text{head}, \text{dec}, \text{get}\}$, $\xi_0, \tau, \delta \rangle$, where the transition function $\tau : Z \times A \rightarrow Z$ is defined by

- $\tau(\{\text{tape}_{\alpha_1}, \dots, \text{tape}_{\alpha_i}, \dots\}, K, k, \text{getToken}(obj_h, obj_\ell)) = (\{\text{tape}_{\alpha_1}, \dots, \text{pop}(\text{tape}_{\alpha_i}), \dots\}, K, k)$ with α_i the merit of the invoking process;
- $\tau(\{\text{tape}_{\alpha_1}, \dots, \text{tape}_{\alpha_i}, \dots\}, K, k, \text{consumeToken}(obj_\ell^{tkn_h})) = (\{\text{tape}_{\alpha_1}, \dots, \text{tape}_{\alpha_i}, \dots\}, \text{add}(K, h, obj_\ell^{tkn_h}))$, if $tkn_h \in \mathfrak{T}$; $\{\{\text{tape}_{\alpha_1}, \dots, \text{tape}_{\alpha_i}, \dots\}, K, k\}$ otherwise.

and the output function $\delta : Z \times A \rightarrow B$ is defined by

- $\delta(\{\text{tape}_{\alpha_1}, \dots, \text{tape}_{\alpha_i}, \dots\}, K, k, \text{getToken}(obj_h, obj_\ell)) = obj_\ell^{tkn_h} : obj_\ell^{tkn_h} \in \mathcal{O}', tkn_h \in \mathfrak{T}$, if $\text{head}(\text{tape}_{\alpha_i}) = tkn$ with α_i the merit of the invoking process; \perp otherwise;
- $\delta(\{\text{tape}_{\alpha_1}, \dots, \text{tape}_{\alpha_i}, \dots\}, K, k, \text{consumeToken}(obj_\ell^{tkn_h})) = \text{get}(K, h)$.

Definition 3.6. (Θ_P -ADT Definition). The Θ_P Abstract Data type is defined as the Θ_F -ADT with $k = \infty$.

Figure 6 shows a possible path of the transition system defined by the Θ_F and Θ_P -ADTs.

3.3 BT-ADT augmented with Θ Oracles

In this section we augment the BT-ADT with Θ oracles and we analyze the histories generated by their combination. Specifically, we define a refinement of the $\text{append}(b_\ell)$ operation of the BT-ADT with the oracle operations which triggers the $\text{getToken}(b_h \leftarrow \text{last_block}(f(bt)), b_\ell)$ operation as long as it returns a token on b_k , i.e., $b_\ell^{tkn_h}$ which is a valid block in \mathcal{B}' . Once obtained, the token is consumed and the append terminates, i.e. the block $b_\ell^{tkn_h}$ is appended to the block h in

the blockchain $f(bt)$ ($\{b_0\} \frown f(bt)|_h \frown \{b_\ell\}$). Notice that those two operations and the concatenation occur atomically.

We say that the *BT-ADT* augmented with Θ_F or Θ_P oracle is a *refinement* $\mathfrak{R}(BT-ADT, \Theta_F)$ or $\mathfrak{R}(BT-ADT, \Theta_P)$ respectively.

Let us define the following auxiliary function:

- *evaluate*: $\mathcal{B} \times B^\ominus \rightarrow bool$. $evaluate(b, \delta_b \circ \delta_a^*) = \text{true}$ if $(\exists h : b^{tkn_h} \in \delta_b \wedge (\exists X : b^{tkn_h} \in X \wedge X \in \delta_a^*))$; **false** otherwise.

Definition 3.7. [$\mathfrak{R}(BT-ADT, \Theta_F)$ refinement] Given the BT-ADT $= \langle A, B, Z, \xi_0, \tau, \delta \rangle$, and the Θ_F -ADT $= \langle A^\ominus, B^\ominus, Z^\ominus, \xi_0^\ominus, \tau^\ominus, \delta^\ominus \rangle$, we have $\mathfrak{R}(BT-ADT, \Theta_F) = \langle A' = A \cup A^\ominus, B' = B \cup B^\ominus, Z' = Z \cup Z^\ominus, \xi'_0 = \xi_0 \cup \xi_0^\ominus, \tau', \delta' \rangle$, where the transition function $\tau' : Z' \times A' \rightarrow Z'$ is defined by

- $\tau_a = \tau'(\{\{tape_{\alpha_1}, \dots, tape_{\alpha_i}, \dots\}, K, k, bt, f, P\}, \text{getToken}(b_k \leftarrow \text{last_block}(bt), b_\ell)) = (\{tape_{\alpha_1}, \dots, pop(tape_{\alpha_i}), \dots\}, K, k, bt, f, P)$;
- $\tau_b = \tau'(\{\{tape_{\alpha_1}, \dots, tape_{\alpha_i}, \dots\}, K, k, bt, f, P\}, \text{consumeToken}(b_\ell^{tkn_h})) = (\{tape_{\alpha_1}, \dots, tape_{\alpha_i}, \dots\}, add(K, h, b_\ell^{tkn_h}), k, \{b_0\} \frown f(bt)|_h \frown \{b_\ell\}, f, P)$ if $tkn_h \in \mathfrak{T} \wedge b_\ell^{tkn_h} \in get(K, l)$; $(\{tape_{\alpha_1}, \dots, tape_{\alpha_i}, \dots\}, K, k, bt, f, P)$ otherwise;
- $\tau'(\{\{tape_{\alpha_1}, \dots, tape_{\alpha_i}, \dots\}, K, k, bt, f, P\}, \text{append}(b)) = \tau_b \circ \tau_a^*$
where $\tau_b \circ \tau_a^*$ is the repeated application of τ_a until $\delta_a(\{\{tape_{\alpha_1}, \dots, tape_{\alpha_i}, \dots\}, K, k, bt, f, P\}, \text{getToken}(b_k \leftarrow \text{last_block}(bt), b_\ell)) = b_\ell^{tkn_h}$ concatenated with the τ_b application;
- $\tau'(\{\{tape_{\alpha_1}, \dots, tape_{\alpha_i}, \dots\}, K, k, bt, f, P\}, \text{read}()) = bt$.

and the output function $\delta' : Z' \times A' \rightarrow B'$ is defined by:

- $\delta_a = \delta'(\{\{tape_{\alpha_1}, \dots, tape_{\alpha_i}, \dots\}, K, k, bt, f, P\}, \text{getToken}(b_k \leftarrow \text{last_block}(bt), b_\ell)) = b_\ell^{tkn_h} : b_\ell^{tkn_h} \in \mathcal{B}', tkn_h \in \mathfrak{T}$, if $head(tape_{\alpha_i}) = tkn$ with α_i the merit of the invoking process; \perp otherwise;
- $\delta_b = \delta'(\{\{tape_{\alpha_1}, \dots, tape_{\alpha_i}, \dots\}, K, k, bt, f, P\}, \text{consumeToken}(obj_\ell^{tkn_h})) = get(K, h)$;
- $\delta'(\{\{tape_{\alpha_1}, \dots, tape_{\alpha_i}, \dots\}, K, k, bt, f, P\}, \text{append}(b)) = evaluate(b, \delta_b \circ \delta_a^*)$, where $\delta_b \circ \delta_a^*$ is the repeated application of δ_a until $\delta_a(\{\{tape_{\alpha_1}, \dots, tape_{\alpha_i}, \dots\}, K, k, bt, f, P\}, \text{getToken}(\text{last_block}(bt), b)) = b_\ell^{tkn_h}$ concatenated with the δ_b application;
- $\delta'(\{\{tape_{\alpha_1}, \dots, tape_{\alpha_i}, \dots\}, K, k, bt, f, P\}, \text{read}()) = \{b_0\} \frown f(bt)$;
- $\delta'(\{\{tape_{\alpha_1}, \dots, tape_{\alpha_i}, \dots\}, K, k, bt_0, f, P\}, \text{read}()) = b_0$.

Definition 3.8 ($\mathfrak{R}(BT-ADT, \Theta_P)$ refinement). Same definition as the $\mathfrak{R}(BT-ADT, \Theta_F)$ refinement.

Definition 3.9 (k-Fork Coherence). A concurrent history $H = \langle \Sigma, E, \Lambda, \mapsto, \prec, \nearrow \rangle$ of the BT-ADT composed with Θ_F -ADT satisfies the *k-Fork Coherence* if there are at most k $\text{append}()$ operations that return \top for the same token.




Figure 7: Refinement of the `append()` operation. We use the following syntax on the edges: operation/output.

Theorem 3.2 (*k-Fork Coherence*). Each concurrent history $H = \langle \Sigma, E, \Lambda, \mapsto, \prec, \nearrow \rangle$ of the BT-ADT composed with a Θ_F -ADT satisfies the *k-Fork Coherence*.

Proof. We prove the theorem by considering the defined refinement (Definition 3.7) where (i) there are a infinite number of `getToken()` invocations for object *obj* and (ii) given a valid block as input parameter, the `consumeToken()` operation successfully terminates if it has been invoked less than k times for the same token. From the properties of the pseudo random sequences of tapes, if there are an infinite number of `getToken()` invocations for object *obj* then there exists at least one response for which `getToken()` operation returns a token t , which, when passed as input of the `consumeToken()` operation it successfully terminates if at most $k - 1$ tokens t have been already consumed. \square

Let us notice, the Θ_F -ADT guarantees by construction the safety property (Theorem 3.2). Liveness properties (i.e., the Termination) for Θ_F -ADT and Θ_P -ADT depend on the communication model and failure model in which those are implemented.

3.4 Hierarchy

In this section we define a hierarchy between different BT-ADT satisfying different consistency criteria when augmented with different oracle ADT. We use the following notation: BT-ADT_{SC} and BT-ADT_{EC} to refer respectively to BT-ADT generating concurrent histories that satisfies the *SC* and the *EC* consistency criteria. When augmented with the oracles we have the following four typologies, where for the *frugal* oracle we explicit the value of k : $\mathfrak{R}(\text{BT-ADT}_{SC}, \Theta_{F,k})$, $\mathfrak{R}(\text{BT-ADT}_{SC}, \Theta_P)$, $\mathfrak{R}(\text{BT-ADT}_{EC}, \Theta_P)$, $\mathfrak{R}(\text{BT-ADT}_{EC}, \Theta_{F,k})$.




Figure 8: $\mathfrak{R}(\text{BT-ADT}, \Theta)$ Hierarchy.

In the following we want study the relationship among the different refinements. Without loss of generality, let us consider only the set of histories $\hat{\mathcal{H}}^{\mathfrak{R}(\text{BT-ADT}, \Theta)}$ such that each history $\hat{H}^{\mathfrak{R}(\text{BT-ADT}, \Theta)} \in \hat{\mathcal{H}}^{\mathfrak{R}(\text{BT-ADT}, \Theta)}$ is purged from the unsuccessful `append()` response events (i.e., such that the returned value is \perp). Let $\hat{\mathcal{H}}^{\mathfrak{R}(\text{BT-ADT}, \Theta_{F,k})}$ be the concurrent set of histories generated by a BT-ADT refined with $\Theta_{F,k}$ -ADT and let $\hat{\mathcal{H}}^{\mathfrak{R}(\text{BT-ADT}, \Theta_P)}$ be the concurrent set of histories generated by a BT-ADT refined with Θ_P -ADT.

Theorem 3.3. $\hat{\mathcal{H}}^{\mathfrak{R}(\text{BT-ADT}, \Theta_F)} \subseteq \hat{\mathcal{H}}^{\mathfrak{R}(\text{BT-ADT}, \Theta_P)}$.

Proof. The proof follows from Theorem 3.2 considering that $\mathfrak{R}(\text{BT}, \Theta_P)$ can generate histories with an infinite number of `append()` operations that successfully terminate while $\mathfrak{R}(\text{BT}, \Theta_F)$ can generate history with at most k `append()` operations that successfully terminate. \square

Theorem 3.4. If $k_1 \leq k_2$ then $\hat{\mathcal{H}}^{\mathfrak{R}(\text{BT-ADT}, \Theta_{F,k_1})} \subseteq \hat{\mathcal{H}}^{\mathfrak{R}(\text{BT-ADT}, \Theta_{F,k_2})}$.

Proof. The proof follows from Theorem 3.2 applying the same reasoning as for the proof of Theorem 3.3 with $k_1 \leq k_2$. \square

Finally, from Theorem 3.1 the next corollary follows.

Corollary 3.4.1. $\hat{\mathcal{H}}^{\mathfrak{R}(\text{BT-ADT}_{SC}, \Theta)} \subseteq \hat{\mathcal{H}}^{\mathfrak{R}(\text{BT-ADT}_{EC}, \Theta)}$.

Combining Theorem 3.1 and Theorem 3.3 we obtain the hierarchy depicted in Figure 8.

4 Implementing BT-ADTs

4.1 Implementability in a concurrent model

In this Section we show that $\Theta_{F,k=1}$ has consensus number ∞ and that Θ_P has consensus number 1.

We consider a concurrent system composed by n processes such that up to f processes are faulty (stop prematurely by crashing), $f < n$. Moreover, processes can communicate through atomic registers.

<pre> (1) consumeToken($b_\ell^{tkn_h}$) : (2) $previous_value \leftarrow K[h]$; (3) if ($previous_value == \{\} \wedge tkn_h \in \mathfrak{T}$)then; (4) $K[h] \leftarrow K[h] \cup \{b_\ell^{tkn_h}\}$; (5) endIf (6) return $K[h]$ </pre>	<pre> (1) compare&swap($register, old_value, new_value$) : (2) $previous_value \leftarrow register$; (3) if ($previous_value == old_value$)then; (4) $register \leftarrow new_value$; (5) endIf (6) return $previous_value$ </pre>
---	--

Figure 9: Compare&Swap() and consumeToken() in the case of $\Theta_{F,k=1}$.

4.1.1 Frugal with $k = 1$ at least as strong as Consensus

In the following we prove that there exists a wait-free implementation of the Consensus [25] by the $\Theta_{F,k=1}$ Oracle object. In particular, in this case $\Theta_{F,k=1} = \langle A = \{\text{getToken}(b_h, b_\ell), \text{consumeToken}(b_\ell^{tkn_h})\} : b_h, b_\ell^{tkn_h} \in \mathcal{B}', b_\ell \in \mathcal{B}, tkn_h \in \mathfrak{T}\rangle, B = \mathcal{B}' \cup \text{Boolean}, Z = m(\mathcal{A})^* \times \{K\} \times k \cup \{\text{pop}, \text{head}, \text{dec}, \text{get}\}, \xi_0, \tau, \delta\rangle$. We explicitly consider blocks and valid blocks (\mathcal{B} and \mathcal{B}') rather than objects and valid objects (\mathcal{O} and \mathcal{O}'). Moreover, we consider a version of the Consensus problem for the blockchain. Thus, we consider the Validity property as in [11] such that the decided block b satisfies the predicate P .

Definition 4.1. Consensus \mathcal{C} :

- **Termination.** Every correct process eventually decides some value.
- **Integrity.** No correct process decides twice.
- **Agreement.** If there is a correct process that decides a value b , then eventually all the correct processes decide b .
- **Validity[11].** A decided value is valid, it satisfies the predefined predicate denoted P .

To this aim, we first prove that there exists a wait-free implementation of Compare&Swap() object by consumeToken() object in the case of $\Theta_{F,k=1}$, implying that consumeToken() has the same Consensus number as Compare&Swap() which is ∞ (see [21]). Finally we compose the consumeToken() with the getToken() object proving that there exist a wait-free implementation of \mathcal{C} by $\Theta_{F,k=1}$.

Figure 9 describes consumeToken() (CT), as specified by the Θ -ADT, along with the Compare&Swap() (CAS). Compare&Swap() takes three parameters as input, the *register*, the *old_value* and the *new_value*. If the value in *register* is the same as *old_value* then the *new_value* is stored in *register* and in any case the operation returns the value that was in *register* at the beginning of the operation. In comparison with consumeToken($b_\ell^{tkn_h}$) we have that $b_\ell^{tkn_h}$ is the *new_value*, *register* is $K[h]$ and the implicit *old_value* is $\{\}$. That is, $\text{add}(K, h, b)$ stores b in $K[h]$ if $|K[h]| < k = 1$, then if $K[h] = \{\}$. In any case the operation returns the content of $K[h]$ at the end of the operation itself. Figure 10 describes an algorithm that reduces CAS to consumeToken().

Theorem 4.1. If input values are in \mathcal{B}' then there exists an implementation of CAS by CT in the case of $\Theta_{F,k=1}$.

Proof. The proof simply follows by construction. Let us consider the algorithm in Figure 10. When the Compare&Swap() operation is invoked, if $K[h]$ is empty, then when consumeToken() is invoked

```

(1) compare&swap( $K[h], \{\}, b_\ell^{tkn_h}$ ) :
(2)  $returned\_value \leftarrow consumeToken(b_\ell^{tkn_h});$ 
(3) if ( $returned\_value == b_\ell^{tkn_h}$ )then;
(4)   return  $\{\}$ ;
(5)   else return  $returned\_value$ ;
(6)   endIf

```

Figure 10: An implementation of CAS by CT in the case of $\Theta_{F,k=1}$.

```

upon event propose( $b$ ):
(1)  $validBlock \leftarrow \perp$ ;
(2)  $validBlockSet \leftarrow \emptyset$ ; % since  $k = 1$  then it contains only one element.
(3) while ( $validBlock = \perp$ ):
(4)    $validBlock \leftarrow getToken(b_0, b)$ ;
(5)  $validBlockSet \leftarrow consumeToken(validBlock)$ ; % it can be different from validBlock
(6) trigger decide( $validBlockSet$ );

```

Figure 11: The Protocol \mathcal{A} that reduces the Consensus problem to the Frugal Oracle with $k = 1$.

with $b_\ell^{tkn_h}$ (valid by hypothesis) $K[h]$ is populated with $b_\ell^{tkn_h}$. Such value is later returned by the $consumeToken()$ operation in $returned_value$. Since it is the same value as $b_\ell^{tkn_h}$ (line 3) then the **Compare&Swap** returns the value of $K[h]$ at the beginning of the operation, $\{\}$. If the condition at line (line 3) does not hold, then this means that $K[h]$ did not change during the operation and its value, in $returned_value$ is returned. \square

Figure 11 describes a simple implementation of Consensus by $\Theta_{F,k=1}$. When a correct process p_i invokes the **propose**(b) operation it loops invoking the $getToken(b_0, b)$ operation as long as a valid block is returned (lines 3-4). In this case the $getToken()$ operation takes as input some block b_0 and the proposed block b . Afterwards, when the valid block has been obtained p_i invokes the $consumeToken(validBlock)$ operation whose result is stored in the $tokenSet$ variable (line 5). Notice, the first process that invokes such operation is able to successfully consume the token, i.e., the valid block is in the Oracle set corresponding to b_0 , which cardinality is $k = 1$, and such set is returned each time the $consumeToken()$ operation is invoked for a block related to b_0 . Finally, (line 6) the decision is triggered on such set (with contains one element).

Theorem 4.2. $\Theta_{F,k=1}$ Oracle has Consensus number ∞ .

Proof. The proof proceeds by construction, let us consider the implementation in Figure 11. All correct processes performing the Consensus are looping on the $getToken(b_0, b)$ operation. From the properties of the pseudo random sequences of tapes, if there are an infinite number of $getToken()$ invocations for an block b_0 then there exists at least one response for which $getToken()$ operation returns a valid block b^{tkn_0} . Thus, all correct process i can invoke the $consumeToken(b^{tkn_0})$ operation with valid values. Since all the processes invoke such operation with valid values with can apply Theorem 4.1 which concludes the proof considering that CAS has Consensus number ∞ ([21]). \square

```

(1) consumeTokenk(tkn) :
(2)   Rh,m ← update(Rh,m, tknm)
(3)   returned_value ← scan(Rh,1, Rh,2, ..., Rh,m, ..., Rh,n)
(4)   return returned_value;
```

Figure 12: An implementation of CT by Atomic Snapshot in the case of Θ_P .

4.1.2 Prodigious not stronger than an Atomic Register

In order to show that the Prodigious oracle Θ_P has consensus number 1, it suffices to find a wait-free implementation of the oracle by an object with consensus number 1. To this end we present a straightforward implementation of the Prodigious oracle by Atomic Snapshot [7].

Let us firstly simplify the notation of the consume token operation. Let us consider a consume token invoked for a given block b_h , denoted as $\text{consumeToken}_h(\text{tkn}_m)$, which simply writes a token from the set $\mathfrak{T} = \{\text{tkn}_1, \text{tkn}_2, \dots, \text{tkn}_m, \dots\}$ in the set $K[h]$. Without loss of generality let us assume that: (i) tokens are uniquely identified, (ii) cardinality of \mathfrak{T} is n finite but not known and (iii) the set $K[h]$ is represented by a collection of n atomic registers $\mathfrak{R}[h] = \{R_{h,1}, R_{h,2}, \dots, R_{h,m}, \dots, R_{h,n}\}$, where $R_{h,m}$ is assigned to the tkn_m token, i.e. $R_{h,m}$ can contain either \perp or tkn_m .

It can be observed that the $\text{consumeToken}_h(\text{tkn}_m)$ in the case of k infinite, always allows to write the token tkn_m in $R_{h,m}$, i.e. there always exists a register $R_{h,m}$ for the proposed token tkn_m . By the oracle definition, moreover, the $\text{consumeToken}_h(\text{tkn}_m)$ returns a read of the n registers that includes the last written token. Figure 12 shows a trivial implementation of commit token CT using Atomic Snapshot that offers $\text{update}(R_i, \text{value})$, $\text{scan}(R_1, R_2, \dots, R_n)$ operation to update a particular register and perform an atomic read of input registers, respectively.

Theorem 4.3. Θ_P Oracle has Consensus number 1.

Proof. The proof trivially follows from implementation in Figure 12 of the consensus token operation of the Prodigious oracle by the Atomic Snapshot object and from [7]. \square

4.2 Implementability in a message-passing system model

We consider a message-passing system composed of an arbitrary large but finite set of n processes, $\Pi = \{p_1, \dots, p_n\}$. The passage of time is measured by a fictional global clock (*e.g.*, that spans the set of natural integers). Processes in the system do not have access to the fictional global time. Each process of the distributed system executes a single instance of a distributed protocol \mathcal{P} composed of a set of algorithms, i.e., each process is running an algorithm. Processes can exhibit a Byzantine behavior (i.e., they can arbitrarily deviate from the protocol \mathcal{P} they are supposed to run). A process affected by a Byzantine behavior is said to be faulty, otherwise we refer to such process as non-faulty or correct. We make no assumption on the number of failures that can occur during the system execution. Processes communicate by exchanging messages via communication channels. We say that a communication channels are asynchronous if there is no upper bound on the message delivery delay. Contrarily, communication channels are synchronous if messages sent by correct processes at time t are delivered by correct processes by time $t + \delta$. Finally, communication channels are weakly synchronous if there exist an unknown a priori time τ after which the communication channels behave as synchronous. We specify time to time the channels synchrony assumption considered, when left untold we consider asynchronous channels.

The BlockTree being now a shared object replicated at each process, we note by bt_i the local copy of the BlockTree maintained at process i . To maintain the replicated object we consider histories made of events related to the **read** and **append** operations on the shared object, i.e. the **send** and **receive** operations for process communications and the **update** operation for BlockTree updates. We also use subscript i to indicate that the operation occurred at process i : $\text{update}_i(b_g, b_i)$ indicates that i inserts its locally generated valid block b_i in bt_i with b_g as a predecessor. Updates are communicated through **send** and **receive** operations. An update related to a block b_i generated on a process p_i , sent through $\text{send}_i(b_g, b_i)$, and received through a $\text{receive}_j(b_g, b_i)$, takes effect on the local replica bt_j of p_j with the operation $\text{update}_j(b_g, b_i)$.

We assume a generic implementation of the **update** operation: when process i locally updates its BlockTree bt_i with the valid block b_i (returned from the `consumeToken()` operation), we write $\text{update}_i(b, b_i)$. When a process j execute the $\text{receive}_j(b, b_i)$ operation, it locally updates its BlockTree bt_j by invoking the $\text{update}_j(b, b_i)$ operation.

In the remaining part of the work we consider implementations of BT-ADT in a Byzantine failure model where the set of events is restricted as follows.

Definition 4.2. The execution of the system that uses the BT-ADT $= (A, B, Z, \xi_0, \tau, \delta)$ in a Byzantine failure model defines the concurrent history $H = \langle \Sigma, E, \Lambda, \mapsto, \prec, \nearrow \rangle$ (see Definition 2.4) where we restrict E to a countable set of events that contains (i) all the BT-ADT `read()` operations invocation events by the *correct* processes, (ii) all BT-ADT `read()` operations response events at the *correct* processes, (iii) all **append**(b) operations invocation events such that b satisfies the predicate P and, (iv) *send*, *receive* and *update* events generated at correct processes.

In this Section we consider a message passing system model and we show the

(i) impossibility to achieve Strong Prefix without Consensus and impossibility to achieve Eventual Prefix if at least one message sent by a correct process is lost.

TBC: (ii) Eventual Prefix is impossible in an asynchronous system (iii) Eventual Prefix is impossible if the interval between the generation of two successive blocks is less than the upper bound on the message delay. (iv) Impossible to solve Strong Prefix without the Frugal oracle with $k = 1$.

4.3 Communication Abstractions

We now define the properties that each history H generated by a BT-ADT satisfying the Eventual Prefix Property has to satisfy and then we prove their necessity.

Definition 4.3 (Update Agreement). A concurrent history $H = \langle \Sigma, E, \Lambda, \mapsto, \prec, \nearrow \rangle$ of the system that uses a BT-ADT satisfies the Update Agreement if satisfies the following properties:

- R1. $\forall \text{update}_i(b_g, b_i) \in H, \exists \text{send}_i(b_g, b_i) \in H$;
- R2. $\forall \text{update}_i(b_g, b_j) \in H, \exists \text{receive}_i(b_g, b_j) \in H$ such that $\text{receive}_i(b_g, b_j) \mapsto \text{update}_i(b_g, b_j)$;
- R3. $\forall \text{update}_i(b_g, b_j) \in H, \exists \text{receive}_k(b_g, b_j) \in H, \forall k$.

Figure 13 depicts a concurrent history that satisfies the Update Agreement properties.

In the following, for ease of notation we consider that the selection function $f \in \mathcal{F}$ returns directly also the genesis block.

Lemma 4.4. Property R1 or Property R2 are necessary conditions for any protocol \mathcal{P} to implement a BT-ADT generating histories H satisfying the Eventual Prefix property.




Figure 13: Example of concurrent history that satisfies R1,R2 and R3, the Update Agreement properties.

Proof. Let us assume that there exists a protocol \mathcal{P} implementing a BT-ADT that generates histories H satisfying Eventual Prefix property but not Property R1 or Property R2. Thus, in H there is some `update` u that is not sent to the other processes (R1) or once received, u is not locally applied (R2). Let us consider the following history where R1 is not verified and process i issues the first update event in H .

Let us construct the following execution history H . i issues the `update $_i(b_0, b'_i)$` (thus $bt_i = b_0 \widehat{b}'_i$) but not the `send $_i(b_0, b'_i)$` event. It follows that if there is no `send $_i(b_0, b'_i)$` event in H then in H are no present any `receive $_j(b_0, b)$` events, $j \neq i$ and thus not process $j \neq i$ can issue `update $_j(b_0, b'_j)$` (on the other side, if R2 is not satisfied, even if the `receive $_j(b_0, b)$` event occur then `update $_j(b_0, b'_j)$` may not occur), thus $\forall j \neq i, bt_j = b_0$. Let us assume that i performs a `read()` operation, the selection function $f \in \mathcal{F}$ is applied on $bt_i = b_0 \widehat{b}'_i$. By the score function definition it follows that $\text{score}(b_0 \widehat{b}'_i) > \text{score}(b_0)$. Thus if i issues a `read()` operation after `update $_i(b_0, b'_i)$` it returns a blockchain such that $\text{score}(b_0 \widehat{b})$ and the possible infinite `read()` operations issued by other processes always return blockchain such that $\text{score}(b_0)$, violating the Eventual Prefix property. The construction of H can be completed iterating the same reasoning for an infinite number of `append()` operation issued by i , thus H violates the Eventual Prefix Property leading to a contradiction. \square

Lemma 4.5. Property R3 is a necessary condition for any protocol \mathcal{P} to implement a BT-ADT generating histories H satisfying the Eventual Prefix property.

Proof. Let us assume that there exists a protocol \mathcal{P} implementing a BT-ADT that generates histories H satisfying Eventual Prefix property but not Property R3. Thus, in H there is some `update $_i(b, b'_i)$` u at some process i such that the `receive $_j(b, b'_j)$` events do not occur at all processes $j \neq i$.

Let us consider a system composed by three processes, i, j and k . The system execution generates the following history H where R3 is not verified. In particular, in H are present the `update $_i(b_0, b'_i)$` , `receive $_j(b_0, b'_j)$` events but there is no any `receive $_k(b_0, b'_k)$` event. It follows that $bt_i = bt_j = b_0 \widehat{b}'_i$ and $bt_k = b_0$. We apply the same argument as for Lemma 4.4. Let us assume that j and k perform `read()` operations. Such operation returns the result of $f(bt_j)$ and $f(bt_k)$ respectively. By the score function definition it follows that $\text{score}(b_0 \widehat{b}'_i) > \text{score}(b_0)$. If j issues a `read()` operation after `update $_j(b_0, b'_j)$` it returns a blockchain with $\text{score}(b_0 \widehat{b})$ and the other `read()` operations issued by k

will always return blockchain with $\text{score}(b_0)$. The construction of H can be completed iterating the same reasoning for an infinite number of $\text{append}()$ operation issued by i , thus H violates the Eventual Prefix Property leading to a contradiction. \square

Theorem 4.6. The update agreement property is necessary to construct concurrent histories $H = \langle \Sigma, E, \Lambda, \mapsto, \prec, \nearrow \rangle$ generated by a BT-ADT that satisfy the BT Eventual Consistency criterion.

Proof. The proof follows directly from Lemma 4.4, Lemma 4.5 and the definition of Eventual BT consistency criterion. \square

Considering Theorem 4.6 and Theorem 3.1 the next Corollary follows.

Corollary 4.6.1. There not exists a concurrent history $H = \langle \Sigma, E, \Lambda, \mapsto, \prec, \nearrow \rangle$ of the system that uses a BT-ADT that satisfies the Strong BT consistency criterion but not the Update Agreement.

In the following we consider a communication primitive that is inspired by the Liveness properties of the reliable broadcast [9]. We will prove that this abstraction is necessary to implement Eventual BT Consistency.

Definition 4.4 (Light Reliable Communication (LRC)). A concurrent history H satisfies the properties of the LRC abstraction if and only if:

- (Validity): $\forall \text{send}_i(b, b_i) \in H, \exists \text{receive}_i(b, b_i) \in H$;
- (Agreement): $\forall \text{receive}_i(b, b_j) \in H, \forall k \exists \text{receive}_k(b, b_i) \in H$

In other words, if a correct process i sends a message m then i eventually receives m and if a message m is received by some correct process (e.g., i itself), then m is eventually received by every correct process.

Theorem 4.7. The LRC abstraction is necessary to for any BT-ADT implementation that generates concurrent histories that satisfies the BT Eventual Consistency criterion.

Proof. The proof done by generating a concurrent history H that violates the LRC properties and showing that H also violate the Update Agreement properties. For Theorem 4.6 the Update Agreement properties are necessary condition to implement BT-ADT that generates concurrent histories that satisfies the BT Eventual Consistency criterion.

Let us consider H where at process n occurs the event $\text{update}_n(b, b_n)$ and $\text{send}_n(b, b_n)$ and where the LRC2 property is not satisfied. If LRC2 is violated then in H we can have that there exist some process i at which occurs the $\text{receive}_i(b, b_n)$ event and some process j at which never occurs the $\text{receive}_j(b, b_n)$ event. Since at process n occurred the event $\text{update}_n(b, b_n)$, then, for the R3 property, for each process k $\text{update}_n(b, b_n)$ has to occur. For R2 the $\text{update}_m(b, b_n)$ event at some process m has to be preceded by a $\text{receive}_m(b, b_n)$ event at the same process m . Since by hypothesis not at all processes m the $\text{receive}_m(b, b_n)$ occurs then the property is violated, violating the Update Agreement properties, which are necessary conditions to implement BT-ADT that generates concurrent histories that satisfies the BT Eventual Consistency criterion, which concludes the proof. \square

Finally, from Theorem 3.1 and Theorem 4.7 the next Corollary follows.

Corollary 4.7.1. The LRC abstraction is necessary to for any BT-ADT implementation that generates concurrent histories that satisfies the BT Strong Consistency criterion.

4.4 System model and hierarchy

Observation. Following our Oracle based abstraction (Section 3.4) we assume by definition that the synchronization on the block to append is oracle side and takes place during the append operation. It follows that when a process takes the token to append a block it can only use the LRC communication abstraction.

Theorem 4.8. There does not exist an implementation of $\mathfrak{R}(\text{BT-ADT}_{SC}, \Theta)$ with $\Theta \neq \Theta_{F,k=1}$ that uses a LRC primitive and generates histories satisfying the BT Strong consistency.

Proof. Let us assume that there exist a BT-ADT implementation that satisfies the BT Strong consistency criterion refined with a Θ -ADT different from $\Theta_{F,k=1}$, which implies that forks in the bt can occur. Let us now construct the following history H generated by the system execution at two correct processes i and j . At the beginning $bt_i = bt_j = b_0$. At the same time instant t_0 both processes invoke $\text{append}(b_1)$ and $\text{append}(b_2)$ operations respectively and $b_1, b_2 \in \mathcal{B}'$. By definition, the $\text{append}()$ operation applies a selection function $f \in \mathcal{F}$ to select the block from the BlockTree to which the new block has to be appended, in this case such block is $f(bt_i) = f(bt_j) = f(b_0) = b_0$. By construction, $b_i, b_j \in \mathcal{B}'$, let us assume that a fork occurs and both $\text{append}()$ operations take place and update events are triggered. Since an LRC primitive is used, each update is sent to the other processes. Since synchronous channels are employed, then by time $t_0 + \delta$ the update events are delivered by i and j . Let us consider that H contains the following ordered events: $\text{update}_i(b_0, b_j) \mapsto \text{update}_i(b_0, b_i)$ and $\text{update}_j(b_0, b_i) \mapsto \text{update}_j(b_0, b_j)$. It follows that at a time instant $t < t_0 + \delta$ it can occur that $bt_i = b_0 \widehat{b}_j$ and $bt_j = b_0 \widehat{b}_i$. Let us finally assume that at time t both i and j issue a $\text{read}()$ operation. By definition it returns the result of the selection function f to the BlockTree. For both processes the BlockTree is a blockchain, thus the $\text{read}()$ operations returns $b_0 \widehat{b}_j$ at i and $b_0 \widehat{b}_i$ at j violating the Strong Prefix property leading to a contradiction. Thus, there no exists an implementation of a BT-ADT refined with a Θ -ADT different from $\Theta_{F,k=1}$ that generates histories satisfying the BT Strong consistency even in a fault-free environment. \square

From Theorem 4.8 the next Corollary follows.

Corollary 4.8.1. $\Theta_{F,k=1}$ is necessary for any implementation of any $\mathfrak{R}(\text{BT-ADT}_{SC}, \Theta)$ that generates histories satisfying the BT Strong consistency.

Thanks to Theorem 4.2 the next Corollary also follows.

Corollary 4.8.2. Consensus is necessary for any implementation of a BT-ADT that generates histories satisfying the BT Strong consistency.

As direct implication of the Theorem 4.8 we can eliminate from the hierarchy in Figure 8 both $\mathfrak{R}(\text{BT-ADT}_{SC}, \Theta_P)$ and $\mathfrak{R}(\text{BT-ADT}_{SC}, \Theta_{F,k>1})$, since in both cases the Θ -ADT employed allows forks, thus such enriched ADTs can not generate histories that satisfies the BT Strong consistency criterion. The resulting hierarchy is depicted in Figure 14.

5 Mapping with existing Blockchain-like systems

This section completes this work by illustrating the mapping between different existing systems and the specifications and abstractions presented in this paper. The following table summarizes

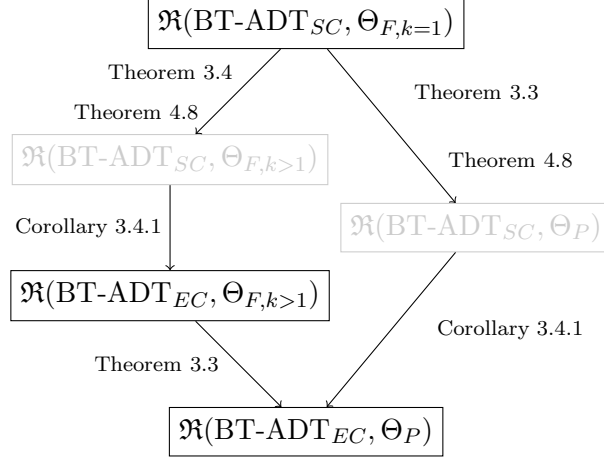


Figure 14: $\mathfrak{R}(\text{BT-ADT}, \Theta)$ Hierarchy. In gray the combinations impossible in a message-passing system

the mapping between different existing systems and these abstractions. More details are given in the following sections. In those sections we refer to a permissionless system as a system where the cardinality of the process set is not a-priori known and each process can read and append into the blockchain. When we do not consider permissionless systems we explicitly state the differences.

Table 1: Mapping of existing systems. Each of these systems assumes at least a light reliable communication.

References	Refinement
Bitcoin [26]	$\mathfrak{R}(BT-ADT_{EC}, \Theta_P)$
Ethereum [31]	$\mathfrak{R}(BT-ADT_{EC}, \Theta_P)$
Algorand [19]	$\mathfrak{R}(BT-ADT_{SC}, \Theta_{F,k=1})$ <i>SC</i> w.h.p
ByzCoin [24]	$\mathfrak{R}(BT-ADT_{SC}, \Theta_{F,k=1})$
PeerCensus [12]	$\mathfrak{R}(BT-ADT_{SC}, \Theta_{F,k=1})$
Redbelly [11]	$\mathfrak{R}(BT-ADT_{SC}, \Theta_{F,k=1})$
Hyperledger [5]	$\mathfrak{R}(BT-ADT_{SC}, \Theta_{F,k=1})$

5.1 Bitcoin

Bitcoin [26] is the pioneer of blockchain systems. Any process $p \in V$ is allowed to read the BlockTree and append blocks to the BlockTree. Processes are characterized by their computational power represented by α_p , normalized as $\sum_{p \in V} \alpha_p = 1$. Processes communicate through reliable FIFO authenticated channels (implemented with TCP), which models a partially synchronous setting [14]. Valid blocks are flooded in the system. The `getToken` operation is implemented by a proof-of-work mechanism. The `consumeToken` operation returns true for all valid blocks, thus there is no bounds on the number of consumed tokens. Thus Bitcoin implements a Prodigal Oracle. The `fselect` returns the blockchain which has required the most computational work, guaranteeing that concurrent blocks can only refer to the most recently appended blocks of the blockchain returned by a `read()` operation. Garay and al [17] have shown, under a synchronous environment assumption,

that Bitcoin ensures Eventual consistency criteria. The same conclusion applies as well for the FruitChain protocol [27], which proposes a protocol similar to BitCoin except for the rewarding mechanism.

5.2 Ethereum

Ethereum [31] is a permissionless blockchain. Processes are characterized by their merit parameter represented by α_p (once normalized as $\sum_{p \in V} \alpha_p = 1$). Contrarily to Bitcoin, where this merit parameter is representative of a computational power, that is this ability to quickly compute hash functions, in Ethereum this merit is bounded by the ability to move data in memory. This proof-of-work mechanism is especially designed for commodity hardware. Any process $p \in V$ is allowed to **read** the BlockTree and **append** blocks to the BlockTree. Processes communicate through reliable FIFO authenticated channels (implemented with TCP), which models a partially synchronous setting [14]. Valid blocks are flooded in the system. The **getToken** operation is implemented by a proof-of-work mechanism. The **consumeToken** operation returns true for all valid blocks, thus there is no bounds on the number of consumed tokens. Thus Ethereum implements a Prodigal Oracle. The **f selects** returns the blockchain which has required the most work (see Section 10 of [31]), guaranteeing that concurrent blocks can only refer to the most recently appended blocks of the blockchain returned by a **read()** operation. This function is implemented through GHOST algorithm [30]. Kiayias has shown [23], under a synchronous environment assumption, that GHOST protocol enjoys both common prefix and chain growth properties. Ethereum thus ensures the Eventual consistency criteria.

5.3 ByzCoin

ByzCoin [24] is a permissionless blockchain. Processes are characterized by their computational power represented by α_p (once normalized as $\sum_{p \in V} \alpha_p = 1$). Byzcoin assumes a semi synchronous environment, that is, in every period of length b there must be a strongly synchronous period of length $s < b$. The block creation process is separated from the transaction validation one. The former one is realized by a proof-of-work mechanism (similar to the Bitcoin's one), and the latter one is achieved by a Byzantine tolerant algorithm (i.e., a variant of PBFT [10]) which creates micro blocks made of transactions.

The **getToken** operation is implemented by a proof-of-work mechanism. Due to the PoW mechanism, several key blocks can be concurrently created. The **consumeToken** operation guarantees that during the synchronous periods of the semi-synchronous setting (those synchronous periods ensure that everyone receives all the concurrent key blocks in a short period of time), a single key block will be appended to the BlockTree by relying on a deterministic function f which **selects** the key block whose digest (fingerprint) has the smallest least significant bits among the concurrent key blocks. Under those assumptions, Byzcoin is an implementation of a strongly consistent BlockTree composed with a Frugal Oracle, with $k = 1$.

Note that transactions do not belong to key blocks but to microblocks which are created by a variant of PBFT where (i) the committee members are the miners of the last w appended key blocks in the BlockTree as returned by a **read()** operation; (ii) each committee member receives a voting share for each block it has created blocks among these w ones, and (iii) committee members are organized on a tree rooted at the leader, and (iv) this leader is the process that invoked the last successful **consumeToken** operation.

5.4 Algorand

Algorand [19] is an algorithm dedicated to permissionless blockchains. Users are characterized by the quantity of coins (stake) they own, represented by α_p once normalized as $\sum_{p \in V} \alpha_p = 1$. Algorand assumes a synchronous setting (rounds) in order to ensure that (i) with overwhelming probability all users agree on the same transactions (safety property) and (ii) new transactions are added to the blockchain (liveness property). Note that safety holds even in a semi synchronous environment. Users communicate among themselves through reliable communication channels (implemented via TCP). Algorand algorithm relies on two main ingredients: a cryptographic sortition and a variant of a Byzantine agreement algorithm. The cryptographic sortition implements the `getToken` operation by selecting the block proposer. This is achieved by selecting at random a committee (that is a small fraction of users weighed by their currency balance α_p , which boils down to a proof-of-stake mechanism) and providing them a random priority, so that with high probability, the highest priority committee member will be in charge of proposing the new block for the current round. The variant of Byzantine agreement algorithm `BA*` implements the `consumeToken` operation, that is the commitment to append this new valid block in the blockchain. `BA*` guarantees that in a favorable environment (strongly synchronous environment augmented with synchronized clocks), if all honest participants have received the same valid block, then this block will be appended to the blockchain (see Lemma 2 [18]). On the other hand, if there is no agreement on that block (because the highest priority committee member is malicious or the network is not strongly synchronous), then `BA*` may create forks with probability less than 10^{-7} (Theorem 2 [18]). This makes Algorand a probabilistic implementation of a strongly consistent `BlockTree` composed with a `Frugal Oracle`, with $k = 1$.

5.5 PeerCensus

PeerCensus [12] is a permissionless blockchain. Processes are characterized by their computational power represented by α_p (once normalized as $\sum_{p \in V} \alpha_p = 1$). PeerCensus assumes a semi synchronous environment, that is, in every period of length b there must be a strongly synchronous period of length $s < b$. PeerCensus is not strictly speaking a blockchain-based algorithm (as Bitcoin or Byzcoin), in the sense that it does not store a sequence of application transactions, but provides a secure and fully distributed timestamping service. This service is implemented by a dynamic Byzantine tolerant consensus algorithm which tracks the committee members of the consensus algorithm through the creation of chained key blocks. The `getToken` operation is implemented by a proof-of-work mechanism, and the `consumeToken` operation, implemented by the Byzantine consensus, commits a single key block among the concurrent ones, that is returns true for a single token, as long as no more than a $1/3$ of the committees members are Byzantine (secure state). Theorem 1 [12] states that the secure state is reachable with high probability if the computational power owned by the adversary, α_A , is less than $1/3$. Thus under these assumptions PeerCensus implements a strongly consistent `BlockTree` composed with a `Frugal Oracle`, with $k = 1$. Note however that in [2] the authors have analyzed the probability that PeerCensus reaches a secure state by examining the composition of successive quorums, and have shown that this probability is decreasing as a function of α_A . For instance, if $\alpha_A = 1/4$, then the probability that PeerCensus reaches a secure state is only equal to $1/3$.

5.6 Red Belly

Red Belly [11] is a consortium blockchain, meaning that any process $p \in V$ is allowed to read the BlockTree but a predefined subset $M \subseteq V$ of processes are allowed to **append** blocks. Each process $p \in M$ as a merit parameter set to $\alpha_p = 1/|M|$ while each process $p \in V \setminus M$ has a merit parameter $\alpha_p = 0$. Processes are asynchronous (i.e., there is no assumption on their respective computational speed) and are connected with partially synchronous [14] (i.e., messages are delivered in unknown but finite time), reliable and authenticated communication channels. Each process $p \in M$ can invoke the `getToken` operation with their new block and will receive a token. The `consumeToken` operation, implemented by a Byzantine consensus algorithm run by all the processes in V , returns true for the uniquely decided block. Thus Red Belly BlockTree contains a unique blockchain, meaning that the selection function f is the trivial projection function from $\mathcal{BT} \mapsto \mathcal{BC}$ which associates to the BT-ADT its unique existing chain of the BlockTree. As a consequence Red Belly relies on a Frugal Oracle with $k = 1$, and by the properties of Byzantine agreement implements a strongly consistent BlockTree (see Theorem 3 [11]).

5.7 HyperLedger Fabric

HyperLedger Fabric [5] is a system allowing to deploy and operate permissioned blockchains. Any process $p \in V$ is allowed to read the BlockTree, however, only a subset of $M \subseteq V$ is allowed to **append** blocks to the BlockTree. Every process of M has the same merit parameter $\alpha_M = 1/|M|$ while processes of $V \setminus M$ have a null merit parameter. HyperLedger Fabric assumes eventual synchrony and reliable channels. Transactions are executed by a dedicated set of processes called endorsers. Executed transactions are then ordered through atomic broadcast primitive so as to gather them into a block. HyperLedger Fabric relies on a leader election to determine which process will generate the next block. Transactions are appended in a block until a stop condition is met. A stop condition refers either on a maximal number of transactions in a block or a maximal elapsed time since the first transaction included in the block. The block is then broadcasted and a new block is created to gather new incoming transactions. By construction, HyperLedger Fabric ensures that a unique token ($k = 1$) is consumed, thus HyperLedger Fabric implement a strongly consistent BlockTree.

6 Conclusions and Future Work

The paper presented an extended formal specification of blockchains and derived interesting conclusion on their implementability. Let us note that the presented work is intended to provide the groundwork for the construction of a sound hierarchy of blockchain abstractions and correct implementations. Future work will focus on several open issues, such as the solvability of Eventual Prefix in message-passing, the synchronization power of other oracle models, and fairness properties for oracles.

Acknowledgment

We are grateful to Mathieu Perrin and anonymous reviewers for their insightful comments on a previous version of the current paper.

References

- [1] Ittai Abraham and Dahlia Malkhi. The blockchain consensus layer and BFT. *Bulletin of the EATCS*, 123, 2017.
- [2] E. Anceaume, T. Lajoie-Mazenc, R. Ludinard, and B. Sericola. Safety Analysis of Bitcoin Improvement Proposals. In *Proceedings of the 15th IEEE International Symposium on Network Computing and Applications (NCA)*, 2016.
- [3] Emmanuelle Anceaume, Antonella Del Pozzo, Romaric Ludinard, Maria Potop-Butucaru, and Sara Tucci-Piergiovanni. Blockchain abstract data type. *arXiv preprint arXiv:1802.09877*, 2018.
- [4] Emmanuelle Anceaume, Romaric Ludinard, Maria Potop-Butucaru, and Frédéric Tronel. Bitcoin a distributed shared register. In *Stabilization, Safety, and Security of Distributed Systems - 19th International Symposium, SSS 2017, Boston, MA, USA, November 5-8, 2017, Proceedings*, pages 456–468, 2017.
- [5] E. Androulaki, A. Barger, V. Bortnikov, C. Cachin, K. Christidis, A. De Caro, D. Enyeart, C. Ferris, G. Laventman, Y. Manevich, S. Muralidharan, C. Murthy, B. Nguyen, M. Sethi, G. Singh, K. Smith, A. Sorniotti, C. Stathakopoulou, M. Vukolić, S. Weed Cocco, and J. Yellick. Hyperledger Fabric: A Distributed Operating System for Permissioned Blockchains. <https://arxiv.org/pdf/1801.10228v1.pdf>.
- [6] A. Fernández Anta, C. Georgiou, K. M. Konwar, and N. C. Nicolaou. Formalizing and implementing distributed ledger objects. *CoRR*, abs/1802.07817, 2018.
- [7] James Aspnes and Maurice Herlihy. Wait-free data structures in the asynchronous PRAM model. In *Second Annual ACM Symposium on Parallel Algorithms and Architectures*, pages 340–349, July 1990.
- [8] C. Cachin, K. Kursawe, F. Petzold, and V. Shoup. Secure and efficient asynchronous broadcast protocols. In *Advances in Cryptology - CRYPTO 2001, 21st Annual International Cryptology Conference, Santa Barbara, California, USA, August 19-23, 2001, Proceedings*, pages 524–541, 2001.
- [9] Christian Cachin, Rachid Guerraoui, and Luís E. T. Rodrigues. *Introduction to Reliable and Secure Distributed Programming (2. ed.)*. Springer, 2011.
- [10] M. Castro and B. Liskov. Practical Byzantine Fault Tolerance and Proactive Recovery. *Journal ACM Transactions on Computer Systems (TOCS)*, 2002.
- [11] T. Crain, V. Gramoli, M. Larrea, and M. Raynal. (Leader/Randomization/Signature)-free Byzantine Consensus for Consortium Blockchains. <http://csrg.redbellyblockchain.io/doc/ConsensusRedBellyBlockchain.pdf>, 2017.
- [12] C. Decker, J. Seidel, and R. Wattenhofer. Bitcoin Meets Strong Consistency. In *Proceedings of the 17th International Conference on Distributed Computing and Networking Conference (ICDCN)*, 2016.

- [13] Swan Dubois, Rachid Guerraoui, Petr Kuznetsov, Franck Petit, and Pierre Sens. The weakest failure detector for eventual consistency. In *Proceedings of the 2015 ACM Symposium on Principles of Distributed Computing*, PODC '15, pages 375–384, 2015.
- [14] C. Dwork, N. Lynch, and L. Stockmeyer. Consensus in presence of partial synchrony. *Journal of the ACM (JACM)*, 1988.
- [15] Cynthia Dwork and Moni Naor. Pricing via processing or combatting junk mail. In *Advances in Cryptology - CRYPTO '92, 12th Annual International Cryptology Conference, Santa Barbara, California, USA, August 16-20, 1992, Proceedings*, pages 139–147, 1992.
- [16] M. J. Fischer, N. A. Lynch, and M. S. Paterson. Impossibility of Distributed Consensus with One Faulty Process. *Journal of the ACM*, 1985.
- [17] J. A. Garay, A. Kiayias, and N. Leonardos. The Bitcoin Backbone Protocol: Analysis and Applications. In *Proceedings of the 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT)*, 2015.
- [18] Y. Gilad, R. Hemo, S. Micali, G. Vlachos, and N. Zeldovich. Algorand: Scaling byzantine agreements for cryptocurrencies. Technical report, MIT CSAIL, 2017. <https://people.csail.mit.edu/nickolai/papers/gilad-algorand-eprint.pdf>.
- [19] Yossi Gilad, Rotem Hemo, Silvio Micali, Georgios Vlachos, and Nickolai Zeldovich. Algorand: Scaling byzantine agreements for cryptocurrencies. In *Proceedings of the 26th Symposium on Operating Systems Principles*, pages 51–68. ACM, 2017.
- [20] A. Girault, G. Göbller, R. Guerraoui, J. Hamza, and D-A. Seredinschi. Why You Can't Beat Blockchains: Consistency and High Availability in Distributed Systems. <http://arxiv.org/abs/1710.09209>.
- [21] Maurice Herlihy. Wait-free synchronization. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 13(1):124–149, 1991.
- [22] Maurice Herlihy. Blockchains and the future of distributed computing. In *Proceedings of the ACM Symposium on Principles of Distributed Computing*, PODC '17, pages 155–155, New York, NY, USA, 2017. ACM.
- [23] A. Kiayias and G. Panagiotakos. On Trees, Chains and Fast Transactions in the Blockchain. <http://eprint.iacr.org/2016/545>, 2016.
- [24] E. Kokoris-Kogias, P. Jovanovic, N. Gailly, I. Khoffi, L. Gasser, and B. Ford. Enhancing Bitcoin Security and Performance with Strong Consistency via Collective Signing. In *Proceedings of the 25th USENIX Security Symposium*, 2016.
- [25] Leslie Lamport, Robert Shostak, and Marshall Pease. The byzantine generals problem. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 4(3):382–401, 1982.
- [26] S. Nakamoto. Bitcoin: A Peer-to-Peer Electronic Cash System. <https://bitcoin.org/bitcoin.pdf>, 2008.

- [27] Rafael Pass and Elaine Shi. Fruitchains: A fair blockchain. In *Proceedings of the ACM Symposium on Principles of Distributed Computing, PODC 2017, Washington, DC, USA, July 25-27, 2017*, pages 315–324, 2017.
- [28] M. Perrin. *Distributed Systems, Concurrency and Consistency*. ISTE Press, Elsevier, 2017.
- [29] M. Perrin, A. Mostefaoui, and C. Jard. Causal Consistency: Beyond Memory. In *21st ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, 2016.
- [30] Y. Sompolinsky and A. Zohar. Accelerating Bitcoin’s Transaction Processing. Fast Money Grows on Trees, Not Chains. <https://eprint.iacr.org/eprint-bin/getfile.pl?entry=2013/881&version=20140101:161740&file=881.pdf>, 2013.
- [31] G. Wood. Ethereum: A secure decentralised generalised transaction ledger. <http://gavwood.com/Paper.pdf>.