

# When Is Spring Coming? A Security Analysis of Avalanche Consensus

Ignacio Amores-Sesar ✉ 

University of Bern, Switzerland

Christian Cachin ✉ 

University of Bern, Switzerland

Enrico Tedeschi ✉ 

The Arctic University of Norway, Tromsø, Norway

---

## Abstract

Avalanche is a blockchain consensus protocol with exceptionally low latency and high throughput. This has swiftly established the corresponding token as a top-tier cryptocurrency. Avalanche achieves such remarkable metrics by substituting proof of work with a random sampling mechanism. The protocol also differs from Bitcoin, Ethereum, and many others by forming a directed acyclic graph (DAG) instead of a chain. It does not totally order all transactions, establishes a partial order among them, and accepts transactions in the DAG that satisfy specific properties. Such parallelism is widely regarded as a technique that increases the efficiency of consensus.

Despite its success, Avalanche consensus lacks a complete abstract specification and a matching formal analysis. To address this drawback, this work provides first a detailed formulation of Avalanche through pseudocode. This includes features that are omitted from the original whitepaper or are only vaguely explained in the documentation. Second, the paper gives an analysis of the formal properties fulfilled by Avalanche in the sense of a generic broadcast protocol that only orders related transactions. Last but not least, the analysis reveals a vulnerability that affects the liveness of the protocol. A possible solution that addresses the problem is also proposed.

**2012 ACM Subject Classification** Theory of computation → Cryptographic protocols; Software and its engineering → Distributed systems organizing principles

**Keywords and phrases** Avalanche, security analysis, generic broadcast

**Digital Object Identifier** 10.4230/LIPIcs.OPODIS.2022.10

**Related Version** *Full Version:* <https://arxiv.org/abs/2210.03423>

**Funding** This work has been funded by the Swiss National Science Foundation (SNSF) under grant agreement Nr. 200021\_188443 (Advanced Consensus Protocols).

## 1 Introduction

The Avalanche blockchain with its fast and scalable consensus protocol is one of the most prominent alternatives to first-generation networks like Bitcoin and Ethereum that consume huge amounts of energy. Its AVAX token is ranked 14th according to market capitalization in August 2022 [9]. Avalanche offers a protocol with high throughput, low latency, excellent scalability, and a lightweight client. In contrast to many well-established distributed ledgers, Avalanche is not backed by proof of work. Instead, Avalanche bases its security on a deliberately metastable mechanism that operates by repeatedly sampling the network, guiding the honest parties to a common output. This allows Avalanche to reach a peak throughput of up to 20'000 transactions per second with a latency of less than half a second [29].

This novel mechanism imposes stricter security constraints on Avalanche compared to other networks. Traditional Byzantine fault-tolerant consensus tolerates up to a third of the parties to be corrupted [24] and proof-of-work protocols make similar assumptions in



© Ignacio Amores-Sesar, Christian Cachin, and Enrico Tedeschi;  
licensed under Creative Commons License CC-BY 4.0

26th International Conference on Principles of Distributed Systems (OPODIS 2022).

Editors: Eshcar Hillel, Roberto Palmieri, and Etienne Rivière; Article No. 10; pp. 10:1–10:22

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

terms of mining power [13, 12]. Avalanche, however, can tolerate only up to  $O(\sqrt{n})$  malicious parties. Furthermore, the transactions in the “exchange chain” of Avalanche (see below) are not totally ordered, in contrast to most other cryptocurrencies, which implement a form of atomic broadcast [6]. As the protocol is structured around a directed acyclic graph (DAG) instead of a chain, it permits some parallelism. Thus, the parties may output the same transactions in a different order, unless these transactions causally depend on each other. Only the latter must be ordered in the same way.

The consensus protocol of a blockchain is of crucial importance for its security and for the stability of the corresponding digital assets. Analyzing such protocols has become an important topic in current research. Although Bitcoin appeared first without formal arguments, its security has been widely understood and analyzed meanwhile. The importance of proving the properties of blockchain protocols has been recognized for a long time [8].

However, there are still protocols released today without the backing of formal security arguments. The Avalanche whitepaper [29] introduces a family of consensus protocols and offers rigorous security proofs for some of them. Yet the Avalanche protocol itself and the related Snowman protocol, which power the platform, are not analyzed. Besides, several key features of this protocol are either omitted or described only vaguely.

In this paper, we explain the Avalanche consensus protocol in detail. We describe it abstractly through pseudocode and highlight features that may be overlooked in the whitepaper (Sections 3–4). Furthermore, we use our insights to formally establish safety properties of Avalanche. Per contra, we also identify a weakness that affects its liveness. In particular, Avalanche suffers from a vulnerability in how it accepts transactions that allows an adversary to delay targeted transactions by several orders of magnitude (Section 5), which may render the protocol useless in practice. The problem results from dependencies that exist among the votes on different transactions issued by honest parties; the whitepaper does not address them. The attack may be mounted by a single malicious party with some insight into the network topology. Finally, we suggest a modification to the Avalanche protocol that would prevent our attacks from succeeding and reconstitute liveness of the protocol (Section 6). This version, which we call *Glacier*, restricts the sampling choices in order to break the dependencies, but also eliminates the parallelism featured by Avalanche.

The vulnerability has been acknowledged by the Avalanche developers. However, the deployed version of the protocol implements another measure that prevents the problem.

## 2 Related work

Despite Avalanche’s tremendous success, there is no independent research on its security. Recall that Avalanche introduces the “snow family” of consensus protocols based on sampling [29, 4]: Slush, Snowflake, and Snowball. Detailed proofs about liveness and safety for the snow-family of algorithms are given. The Avalanche protocol for asset exchange, however, lacks such a meticulous analysis. The dissertation of Yin [32] describes Avalanche as well, but does not analyze its security in more detail either.

Recall that Nakamoto introduced Bitcoin [23] without any formal analysis. This has been corrected by a long line of research, which established the conditions under which it is secure (e.g., by Garay, Kiayias, and Leonardos [13, 14] and by Eyal and Sirer [12]).

The consensus mechanisms that stand behind the best-known cryptocurrencies are meanwhile properly understood. Some of them, like the proof-of-stake protocols of Algorand [15] and the Ouroboros family that powers the Cardano blockchain [17, 10], did apply sound design principles by first introducing and analyzing the protocols and only later implementing them.

Many others, however, have still followed the heuristic approach: they released code first and were confronted with concerns about their security later. This includes Ripple [3, 1] and NEO [31], in which several vulnerabilities have been found, or Solana, which halted multiple times in 2021–2022. Stellar comes with a formal model [21], but it has also been criticized [18].

Protocols based on DAGs have potentially higher throughput than those based on chains. Notable examples include PHANTOM and GHOSTDAG [27], the Tangle of IOTA ([www.iota.org](http://www.iota.org)), Conflux [20], and others [16]. However, they are also more complex to understand and susceptible to a wider range of attacks than those that use a chain. Relevant examples of this kind are the IOTA protocol [22], which has also failed repeatedly in practice [30] and PHANTOM [27], for which a vulnerability has been shown [19] in an early version of the protocol.

## 3 Model

### 3.1 Avalanche platform

We briefly review the architecture of the Avalanche platform [4]. It consists of three separate built-in blockchains, the *exchange* or *X-Chain*, the *platform* or *P-Chain*, and the *contract* or *C-Chain*. Additionally there are a number of subnets. In order to participate in the protocols and validate transactions, a party needs to stake at least 2'000 AVAX (about 50'000 USD in August 2022 [9]).

The *exchange chain* or *X-Chain* secures and stores *transactions* that trade digital assets, such as the native AVAX token. This chain implements a variant of the *Avalanche consensus protocol* that only partially orders the transactions and that is the focus of this work. All information given here refers to the original specification of Avalanche [29].

The *platform chain* or *P-Chain* secures *platform primitives*; it manages all other chains, allows parties to join the network, designates parties to become validators or removes them again from the validator list, and creates or deletes wallets. The P-Chain implements the *Snowman* consensus protocol: this is a special case of Avalanche consensus that always provides total order, like traditional blockchains. It is not explained in the whitepaper and we do not describe it further here.

The *C-Chain* hosts *smart contracts* and runs transactions on an Ethereum Virtual Machine (EVM). It also implements the Snowman consensus protocol of Avalanche and totally orders all transactions and blocks.

### 3.2 Communication and adversary

We now abstract the Avalanche consensus protocol and consider a static network of  $n$  parties  $\mathcal{N} = \{p, q, \dots\}$  that communicate with each other by sending messages. An adversary may *corrupt* up to  $f$  of these parties and cause them to behave *maliciously* and diverge arbitrarily from the protocol. Non-corrupted parties are known as *honest*, messages and transactions sent by them are referred to as *honest*. Analogously, corrupted parties send *malicious* transactions and messages. The parties may access a low-level functionality for sending messages over authenticated point-to-point links between each pair of parties. In the protocol, this functionality is accessed by two events *send* and *receive*. Parties may also access a second low-level functionality for broadcasting messages through the network by gossiping, accessed by the two events *gossip* and *hear* in the protocol. Both primitives are subject to network and timing assumptions. We assume the same network model as in the original

Avalanche whitepaper [29]. Messages are delivered according to an exponential distribution, that is, the amount of time between the sending and the receiving of a message follows an exponential distribution with unknown parameter to the parties. However, messages from corrupted parties are not affected by this delay and will be delivered as fast as the adversary decides. This model differs from traditional assumptions like partial synchrony [11], because the adversary does not possess the ability to delay honest messages as it pleases.

### 3.3 Abstractions

The *payload transactions* of Avalanche are submitted by users and built according to the *unspent transaction output (UTXO)* model of Bitcoin [23]. A payload transaction  $tx$  contains a set of *inputs*, a set of *outputs*, and a number of digital signatures. Every input refers to a position in the output of a transaction executed earlier; this output is thereby *spent* (or *consumed*) and distributed among the outputs of  $tx$ . The balance of a user is given by the set of unspent outputs of all transactions (UTXOs) executed by the user (i.e., assigned to public keys controlled by that user). A payload transaction is valid if it is properly authenticated and none of the inputs that it consumes has been consumed yet (according to the view of the party executing the validation).

Blockchain protocols are generally formalized as *atomic broadcast*, since every party running the protocol outputs the same ordered list of transactions. However, the transaction sequences output by two different parties running Avalanche may not be exactly the same because Avalanche allows more flexibility and does not require a total order. Avalanche only orders transactions that causally depend on each other. Thus, we abstract Avalanche as a *generic broadcast* according to Pedone and Schiper [25], in which the total-order property holds only for *related* transactions as follows.

► **Definition 3.1.** *Two payloads  $tx$  and  $tx'$  are said to be related, denoted by  $tx \sim tx'$ , if  $tx$  consumes an output of  $tx'$  or vice versa.*

Our generic broadcast primitive is accessed through the two events  $broadcast(tx)$  and  $deliver(tx)$ . Similar to other blockchain consensus protocols, it defines an “external” validity property and introduces a predicate  $V$  that determines whether a transaction is valid [7].

► **Definition 3.2.** *A payload  $tx$  satisfies the validity predicate of Avalanche if all the cryptographic requirements are fulfilled and there is no other delivered payload with any input in common with  $tx$ .*

For the remainder of this work, we fix the external validation predicate  $V$  to check the validity of payloads according to the logic of UTXO mentioned before.

Since Avalanche is a randomized protocol, the properties of our broadcast abstraction need to be fulfilled only with all but negligible probability.

► **Definition 3.3.** *A protocol solves validated generic broadcast with validity predicate  $V$  and relation  $\sim$  if it satisfies the following conditions, except with negligible probability:*

**Validity.** *If a honest party broadcasts a payload transaction  $tx$ , then it eventually delivers  $tx$ .*

**Agreement.** *If a honest party delivers a payload transaction  $tx$ , then all honest parties eventually deliver  $tx$ .*

**Integrity.** *For any payload transaction  $tx$ , every honest party delivers  $tx$  at most once, and only if  $tx$  was previously broadcast by some party.*

**Partial order.** *If honest parties  $p$  and  $q$  both deliver payload transactions  $tx$  and  $tx'$  such that  $tx \sim tx'$ , then  $p$  delivers  $tx$  before  $tx'$  if and only if  $q$  delivers  $tx$  before  $tx'$ .*

**External validity.** *If a honest party delivers a payload transaction  $tx$ , then  $V(tx) = \text{TRUE}$ .*

Note that different instantiations of the relation  $\sim$  transform the generic broadcast primitive into well-known primitives. For instance, when no pair of transactions are related, generic broadcast degenerates to reliable broadcast. Whereas when every two transactions are related, generic broadcast transforms into atomic broadcast. In our context, broadcasting corresponds to submitting a payload transaction to the network, whereas delivering corresponds to accepting a payload and appending it to the ledger.

The Avalanche protocol augments payload transactions to *protocol transactions*. A protocol transaction additionally contains a set of *references* to previously executed protocol transactions, together with further attributes regarding the execution. A protocol transaction in the implementation contains a batch of payload transactions, but this feature of Avalanche is ignored here, since it affects only efficiency. Throughout this paper, *transaction* refers to a protocol transaction, unless the opposite is indicated, and *payload* means simply a payload transaction.

A transaction references one or multiple previous transactions, unlike longest-chain protocols, in which each transaction has a unique parent [23]. An execution of the Avalanche protocol will therefore create a directed acyclic graph (DAG) that forms its ledger data structure.

Given a protocol transaction  $T$ , all transactions that it references are called the *parents* of  $T$  and denoted by  $parents(T)$ . The parents of  $T$  together with the parents of those, recursively, are called the *ancestors* of  $T$ , denoted by  $ancestors(T)$ . Analogously, the transactions that have  $T$  as parent are called the *children* of  $T$  and are denoted by  $children(T)$ . Finally, the children of  $T$  together with their recursive set of children are called the *descendants* of  $T$ , denoted by  $descendants(T)$ .

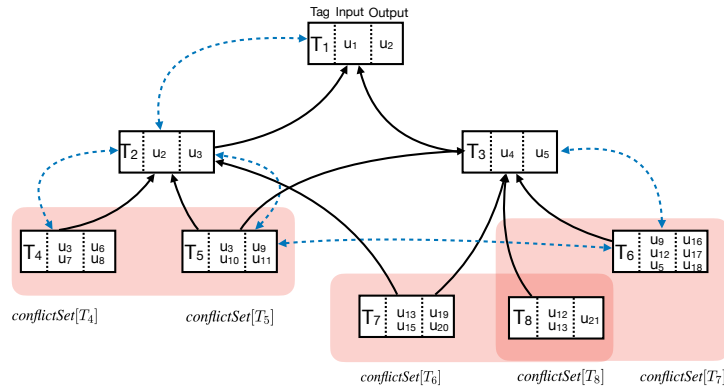
Note that two payload transactions  $tx_1$  and  $tx_2$  in Avalanche that consume the same input are not related, unless the condition of Definition 3.1 is fulfilled. However, two Avalanche payloads consuming the same output *conflict*. For each transaction  $T$ , Avalanche maintains a set  $conflictSet[T]$  of transactions that conflict with  $T$ .

## 4 A description of the Avalanche protocol

Avalanche’s best-known quality is its efficiency. Permissionless consensus protocols, such as those of Bitcoin and Ethereum, are traditionally slow, suffer from low throughput and high latency, and consume large amounts of energy, due to their use of proof-of-work (PoW). Avalanche substitutes PoW with a random sampling mechanism that runs at network speed and that has every party adjust its preference to that of a (perceived) majority in the system. Avalanche also differs from more traditional blockchains by forming a DAG of transactions instead of a chain.

### 4.1 Overview

Avalanche is structured around its *polling* mechanism. In a nutshell, party  $u$  repeatedly selects a transaction  $T$  and sends a *query* about it to  $k$  randomly selected parties in the network. If a majority of them send a positive reply, the query is successful and the transaction contributes to the security of other transactions. Otherwise, the transaction is still processed but does not contribute to the security of any other transactions. Then the party selects a new transaction and repeats the procedure. A bounded number of such polls may execute concurrently. Throughout this work the terms “poll” and “query” are interchangeable.



■ **Figure 1** The UTXO model, conflicting transactions, and related transactions in Avalanche. The eight transactions are labeled  $T_1, \dots, T_8$ . Each transaction is divided into three parts: the left part is a tag  $T_i$  to identify the transaction, the middle part is its set of inputs, and the right part is its set of outputs. The solid arrows indicate the references added by the protocol, showing the parents of each transaction. For instance,  $T_5$  references  $T_2$  and  $T_3$  and has them as parents. The dashed double-arrows indicate related transactions. For example,  $T_5$  and  $T_2$  are related because  $u_3$  is created by  $T_2$  and consumed by  $T_5$ . The conflict sets are denoted by the shaded (red) rectangles. As illustrated, conflict sets can be symmetric, as for  $T_4$  and  $T_5$ , where the conflict sets are identical ( $\text{conflictSet}[T_4] = \text{conflictSet}[T_5]$ ) or asymmetric, as for  $T_6$ ,  $T_7$ , and  $T_8$  where  $\text{conflictSet}[T_6] \cup \text{conflictSet}[T_7] = \text{conflictSet}[T_8]$ .

In more detail, the protocol operates like this. Through the *gossip* functionality, every party is aware of the network membership  $\mathcal{N}$ . A party locally stores all those transactions processed by the network that it knows. The transactions form a DAG through their references as described in the previous section.

Whenever a user submits a payload transaction  $tx$  to the network, the user actually submits it through a party  $u$ . Then,  $u$  randomly selects a number of leaf nodes from a part of the DAG known as the *virtuous frontier*; these are the leaf nodes that are not part of any conflicting set. Party  $u$  then extends  $tx$  with references to the selected nodes and thereby creates a transaction  $T$  from the payload transaction  $tx$ . Next,  $u$  sends a QUERY message with  $T$  to  $k$  randomly, according to stake, chosen parties in the network and waits for their replies in the form of VOTE messages. When a party receives a query for  $T$  and if  $T$  and its ancestors are *preferred*, then the party replies with a positive vote. The answer to this query depends exclusively on the status of  $T$  and its *ancestors* according to the local view of the party that replies. Moreover, the definition of *preferred* is non-trivial and will be explained further below. If the polling party receives more than  $\alpha > \frac{k}{2}$  positive votes, the poll is defined to be successful.

Every party  $u$  running the Avalanche protocol sorts transactions of its DAG into conflict sets.

► **Definition 4.1.** *The conflict set  $\text{conflictSet}[T]$  of a given transaction  $T$  is the set of transactions that have an input in common with  $T$  (including  $T$  itself).*

Note that even if two transaction  $T$  and  $T'$  consume one common transaction output and thus conflict, their conflict sets  $\text{conflictSet}[T]$  and  $\text{conflictSet}[T']$  can differ, since  $T$  may consume outputs of further transactions. (In Figure 1, for example,  $T_8$  conflicts with  $T_6$  and  $T_7$ , although  $T_7$  conflicts with  $T_8$  but not with  $T_6$ .)



Decisions on accepting transactions are made as follows. For each of its conflict sets, a party selects one transaction and designates it as *preferred*. This designation is parametrized by a *confidence value*  $d[T]$  of  $T$ , which is updated after each transaction query. If the confidence value of some conflicting transaction  $T^*$  surpasses  $d[T]$ , then  $T^*$  becomes the preferred transaction in the conflict set.

It has been shown [28, 29] that regardless of the initial distribution of such confidence values and preferences of transactions, this mechanism converges. For the transactions of one conflict set considered in isolation, this implies that all honest parties eventually prefer the same transaction from their local conflict sets. (The actual protocol has to respect also dependencies among the transactions; we return to this later.)

To illustrate this phenomenon, assume that there exist only two transactions  $T$  and  $T'$  and that half of the parties prefer  $T$ , whereas the other half prefers  $T'$ . This is the worst-case scenario. Randomness in sampling breaks the tie. Without loss of generality, assume that parties with preferred transaction  $T$  are queried more often. Hence, more parties consider  $T$  as preferred as a consequence. Furthermore, the next time when a party samples again, the probability of hitting a party that prefers  $T$  is higher than hitting one that prefers  $T'$ . This is the “snowball” effect that leads to ever more parties preferring  $T$  until every party prefers  $T$ .

This preferred transaction is the candidate for acceptance and incorporation into the ledger. The procedure is parametrized by a confidence counter for each conflict set, which reflects the probability that  $T$  is the preferred transaction in the local view of the party. The party increments the confidence counter whenever it receives a positive vote to a query on a descendant of  $T$ ; the counter is reset to zero whenever such a query obtains a negative vote. When this counter overcomes a given threshold,  $T$  is accepted and its payload is added to the ledger. We now present a detailed description of the protocol and refer to the pseudocode in Algorithm 1–4.

## 4.2 Data structures

The information presented here has been taken from the whitepaper [29], the source code [5], or the official documentation [4].

**Notation.** We introduce the notation used in the remaining sections including the pseudocode. For a variable  $a$  and a set  $\mathcal{S}$ , the notation  $a \stackrel{R}{\leftarrow} \mathcal{S}$  denotes sampling  $a$  uniformly at random from  $\mathcal{S}$ . We frequently use hashmap data structures: A hashmap associates keys in a set  $\mathcal{K}$  with values in  $\mathcal{V}$  and is denoted by  $HashMap[\mathcal{K} \rightarrow \mathcal{V}]$ . For a hashmap  $\mathcal{F}$ , the notation  $\mathcal{F}[K]$  returns the entry stored under key  $K \in \mathcal{K}$ ; referencing an unassigned key gives a special value  $\perp$ .

We make use of timers throughout the protocol description. Timers are created in a stopped state. When a timer has been started, it produces a *timeout* event once after a given duration has expired and then stops. A timer can be (re)started arbitrarily many times. Stopping a timer is idempotent.

**Global parameters.** We recall that we model Avalanche as run by an immutable set of parties  $\mathcal{N}$  of size  $n$ . There are more three global parameters: the number  $k$  of parties queried in every poll, the majority threshold  $\alpha > \frac{k}{2}$  for each poll, the acceptance parameters  $\beta_1$  and  $\beta_2$ , and the maximum number *maxPoll* of concurrent polls.

**Local variables.** Queried transactions are stored in a set  $\mathcal{Q}$ , the subset  $\mathcal{R} \subset \mathcal{Q}$  is defined to be the set of *repollable* transactions, a feature that is not explained in the original paper [29]. The number of active polls is tracked in a variable *conPoll*. The parents of a transaction are

selected from the *virtuous frontier*,  $\mathcal{VF}$ , defined as the set of all *non-conflicting* transactions that have no known descendant and whose ancestors are preferred in their respective conflict sets. A transaction is non-conflicting if there is no transaction in the local DAG spending any of its inputs. For completeness, we recall that conflicting transactions are sorted in  $\text{conflictSet}[T]$  formed by transactions that conflict with  $T$ , i.e., transactions which have some input in common with  $T$ .

Transactions bear several attributes related to queries and transaction preference. A *confidence value*  $d[T]$  is defined to be the number of positive queries of  $T$  and its descendants. Given a conflict set  $\text{conflictSet}[T]$ , the variable  $\text{pref}[\text{conflictSet}[T]]$ , called *preferred transaction*, stores the transaction with the highest confidence value in  $\text{conflictSet}[T]$ . The variable  $\text{last}[\text{conflictSet}[T]]$  denotes which transaction was the preferred one in  $\text{conflictSet}[T]$  after the most recent update of the preferences. The preferred transaction is the candidate for acceptance in each conflict set, the acceptance is modeled by a counter  $\text{cnt}[\text{conflictSet}[T]]$ . Once accepted, a transaction remains the preferred one in its conflict set forever.

### 4.3 Detailed description

Each transaction goes through three phases during the consensus protocol: query of transactions, reply to queries, and update of preferences. All of the previous phases call the same set of functions.

**Functions.** The function  $\text{updateDAG}(T)$  sorts the transactions in the corresponding conflict sets. The function  $\text{preferred}(T)$  (L 98) outputs TRUE if  $T$  is the preferred transaction in its conflict set and FALSE otherwise. The function  $\text{stronglyPreferred}(T)$  (L 100) outputs TRUE if and only if  $T$ , and everyone of its ancestors is the preferred transaction in its respective conflict set.

The function  $\text{acceptable}(T)$  (L 102) determines whether  $T$  can be accepted and its payload added to the ledger or not. Transaction  $T$  is considered accepted when one of the two following conditions is fulfilled:

- $T$  is the unique transaction in its conflict set, all the transactions referenced by  $\mathcal{T}$  are considered accepted, and  $\text{cnt}[\text{conflictSet}[T]]$  is greater or equal than  $\beta_1$ .
- $\text{cnt}[\text{conflictSet}[T]]$  is greater or equal than  $\beta_2$ .

Finally, the function  $\text{updateRepollable}()$  (L 106) updates the set of repollable transactions. A transaction  $T$  is repollable if  $T$  has already been accepted; or all its ancestors are preferred, a transaction in its conflict set has not already been accepted, and no parent has been rejected

**Transaction query.** A party in Avalanche progresses only by querying transactions. In each of these queries, party  $u$  selects a random transaction  $T$  (L 38), from the set of transactions that  $u$  has not previously queried by  $u$ . Then, it samples a random subset  $\mathcal{S}[T] \subset \mathcal{N}$  of  $k$  parties from the set of parties running the Avalanche protocol and sends each a [QUERY,  $T$ ] message. In the implementation of the protocol, party  $u$  performs up to  $\text{maxPoll}$  simultaneous queries. The repoll functionality (L 33–48) consists of performing several simultaneous transactions. When  $u$  does not know of any transaction that has not been queried,  $u$  queries a transaction that has not been accepted yet. The main idea behind this functionality is to utilize the network when this is not saturated. The repoll functionality (L 33–48) constitutes one of the most notable changes from Avalanche’s whitepaper [29].



**Query reply.** Whenever  $u$  receives a query message with transaction  $T$ , it replies with a message  $[\text{VOTE}, u, T, \text{stronglyPreferred}(T)]$  containing the output of the binary function  $\text{stronglyPreferred}(T)$  according to its local view (L 100).

**Update of preferences.** Party  $u$  collects the replies  $[\text{VOTE}, v, T, \text{stronglyPreferred}(T)]$ , and counts the number of positive votes. On the one hand, if the number of positive votes overcomes the threshold  $\alpha$  (L 53), the query is considered successful. In this case party  $u$  loops over  $T$  and all its ancestors  $T'$ , increasing the confidence level  $d[T']$  by one. If  $T'$  is the preferred transaction in its conflict set, then party  $u$  increases the counter for transaction  $\text{cnt}[\text{conflictSet}[T']]$  by one. Subsequently,  $u$  checks whether  $T'$  has also previously been the preferred transaction in its conflict set. And when  $T'$  is not the preferred transaction according to the most recent query, party  $u$  will set the counter to one (L 53–67), in order to ensure that  $\text{cnt}[\text{conflictSet}[T']]$  correctly reflects the number of consecutive successful queries of descendants of  $T'$ .

On the other hand, if  $u$  receives more than  $k - \alpha$  negative votes, party  $u$  loops also over  $T$  and its ancestors, and sets their counters  $\text{cnt}[\text{conflictSet}[T']]$  to zero as if to indicate that  $T'$  and the other transactions should not be accepted yet. (L 68–73). Party  $u$  only waits until  $\alpha$  positive votes or  $k - \alpha$  votes in total are received, since  $u$  can then determine the outcome of the query.

**Acceptance of transactions.** Party  $u$  accepts  $T$  when its counter  $\text{cnt}[\text{conflictSet}[T]]$  reaches a certain threshold  $\beta_1$  or  $\beta_2$ . If  $T$  is the only transaction in its conflicting set and all its parents have already been accepted, then  $u$  accepts  $T$  if  $\text{cnt}[\text{conflictSet}[T]] \geq \beta_1$ , otherwise  $u$  waits until the counter overcomes a higher value  $\beta_2$ .

**No-op transactions.** The local DAG is modified whenever a poll is finalized. In particular, only the queried transaction and its ancestors are modified. Avalanche makes use of *no-op transactions* to modify all the transactions in the DAG. After finalizing a poll, party  $u$  queries the network with all the transactions in the virtuous frontier whose state has not been modified, in a sequential manner.

#### 4.4 Life of a transaction

We follow an honest transaction  $T$  through the protocol. The user submits the payload transaction  $tx$  to some party  $u$ , then  $u$  adds references  $refs$  to the payload transaction, creating a transaction  $T = (tx, refs)$ . These references point to transactions in the virtuous frontier  $\mathcal{VF}$ . Transaction  $T$  is then *gossiped* through the network and added to the set of known transactions  $\mathcal{T}$  (L 22–28). Party  $u$  may also *hear* about new transactions through this gossip functionality. Whenever this is the case,  $u$  add the transaction to its set of known transactions  $\mathcal{T}$  (L 29–32).

Party  $u$  eventually selects  $T$  to be processed. When this happens,  $u$  *samples*  $k$  random parties from the network and stores them in  $\mathcal{S}[T]$ . Party  $u$  *queries* parties in  $\mathcal{S}[T]$  with  $T$  and starts a timer  $\text{timeout}[T]$ .  $T$  is added to  $\mathcal{Q}$  (L 33–48).

Parties queried with  $T$  reply with the value of the function  $\text{stronglyPreferred}(T)$  (L 100). This function answers positively (TRUE) if  $T$  is *strongly preferred*, i.e., if  $T$  and all of its ancestors are the preferred transaction inside each respective conflict set. A negative answer (FALSE) is returned if either  $T$  or any of its ancestors fail to satisfy these conditions.

Party  $u$  then stores the answer from party  $v$  to the query in the variable  $\text{votes}[T][v]$  and proceeds according to them.

■ **Algorithm 1** Avalanche (party  $u$ ), state.

---

<b>Global parameters and state</b>	
1: $\mathcal{N}$	// set of parties
2: $maxPoll \in \mathbb{N}$	// maximum number of concurrent polls, default value 4
3: $k \in \mathbb{N}$	// number of parties queried in each poll, default value 20
4: $\alpha \in \{\lceil \frac{k+1}{2} \rceil, \dots, k\}$	// majority threshold for queries, default value 15
5: $\beta_1 \in \mathbb{N}$	// threshold for early acceptance, default value 15
6: $\beta_2 \in \mathbb{N}$	// threshold for acceptance, default value 150
7: $\mathcal{T} \leftarrow \emptyset$	// set of known transactions
8: $\mathcal{Q} \subset \mathcal{T} \leftarrow \emptyset$	// set of queried transactions
9: $\mathcal{R} \subset \mathcal{Q} \leftarrow \emptyset$	// set of repollable transactions
10: $\mathcal{D} \subset \mathcal{T} \leftarrow \emptyset$	// set of no-op transactions to be queried
11: $\mathcal{VF} \subset \mathcal{Q} \leftarrow \emptyset$	// set of transactions in the virtuous frontier
12: $conPoll \in \mathbb{N} \leftarrow 0$	// number of concurrent polls performed
13: $conflictSet : \text{HashMap}[\mathcal{T} \rightarrow 2^{\mathcal{T}}]$	// conflict set
14: $\mathcal{S} : \text{HashMap}[\mathcal{T} \rightarrow \mathcal{N}]$	// set of sampled parties to be queried with a transaction
15: $votes : \text{HashMap}[\mathcal{T} \times \mathcal{N} \rightarrow \{\text{FALSE}, \text{TRUE}\}]$	// variable to store the replies of queries
16: $d : \text{HashMap}[\mathcal{T} \rightarrow \mathbb{N}]$	// confidence value of a transaction
17: $pref : \text{HashMap}[2^{\mathcal{T}} \rightarrow \mathcal{T}]$	// preferred transaction in the conflict set
18: $last : \text{HashMap}[2^{\mathcal{T}} \rightarrow \mathcal{T}]$	// preferred transaction in the last query
19: $cnt : \text{HashMap}[2^{\mathcal{T}} \rightarrow \mathbb{N}]$	// counter for acceptance of the conflict set
20: $accepted : \text{HashMap}[\mathcal{T} \rightarrow \{\text{FALSE}, \text{TRUE}\}]$	// indicator that a transaction is accepted
21: $timer : \text{HashMap}[\mathcal{T} \rightarrow \{\text{timers}\}]$	// timer for the query of transactions

---

- If  $u$  receives more than  $\alpha$  positive votes,  $u$  runs over all the ancestors of  $T$ . If the ancestor  $T'$  was the most recent (or “last”) preferred transaction in its conflict set, its counter is increased by one. Otherwise,  $T'$  becomes the most recent preferred transaction and its counter is reset to one (L 53–67).
- If  $u$  receives at least  $k - \alpha$  FALSE votes,  $u$  resets the counter for acceptance of all its ancestors  $cnt[T'] \leftarrow 0$  (L 68–73).
- If timer  $timeout[T]$  is triggered before the query is completed, the query is aborted instead. The votes are reset and every party is removed from the set  $\mathcal{S}[T]$ , so no later reply can be considered (L 80–83).

In parallel to the previous procedure, party  $u$  may perform up to  $conPoll$  concurrent queries of different transactions.

Once  $T$  has been queried, it awaits in the local view of party  $u$  to be accepted. Since by assumption  $T$  is honest,  $conflictSet[T] = \{T\}$ . Hence  $T$  is accepted when  $cnt[conflictSet[T]]$  reaches  $\beta_1$ , if its ancestors are already accepted, or  $\beta_2$  otherwise (L 102–104). We recall that  $cnt[conflictSet[T]]$  is incremented whenever a query involving a descendant of  $T$  is successful. However, when a non-descendant of  $T$  is queried, it may trigger a no-op transaction (L 35) that is a descendant of  $T$ .

If there is no new transaction waiting to be queried, i.e.,  $\mathcal{T} \setminus \mathcal{Q}$  is empty, the party proceeds with a *repollable* transaction (L 40–42). A *repollable* transaction is one that has not been previously accepted but it is a candidate to be accepted (L 106–110).

---

**Algorithm 2** Avalanche (party  $u$ ), part 1.

---

```

22: upon broadcast( $tx$ ) do
23:   if  $V(tx)$  then
24:      $T \leftarrow (tx, \mathcal{VF})$  // up to a maximum number of parents
25:      $\mathcal{T} \leftarrow \mathcal{T} \cup \{T\}$ 
26:      $accepted[T] \leftarrow \text{FALSE}$ 
27:      $updateDAG(T)$ 
28:     gossip message [BROADCAST,  $T$ ]

29: upon hearing message [BROADCAST,  $T$ ] do
30:   if  $T \notin \mathcal{T}$  do
31:      $\mathcal{T} \leftarrow \mathcal{T} \cup \{T\}$ 
32:      $accepted[T] \leftarrow \text{FALSE}$ 

33: upon  $conPoll < maxPoll$  do
34:    $conPoll \leftarrow conPoll + 1$ 
35:   if  $\mathcal{D} \neq \emptyset$  then // prefer no-op transactions
36:      $T \leftarrow$  least recent transaction in  $\mathcal{D}$ 
37:   else if  $\mathcal{T} \setminus \mathcal{Q} \neq \emptyset$  then // take any not yet queried transaction
38:      $T \stackrel{R}{\leftarrow} \mathcal{T} \setminus \mathcal{Q}$ 
39:      $d[T] \leftarrow 0$ 
40:   else // all transaction queried already, take one of them
41:      $updateRepollable()$ 
42:      $T \stackrel{R}{\leftarrow} \mathcal{R}$ 
43:      $\mathcal{S}[T] \leftarrow sample(\mathcal{N} \setminus \{u\}, k)$  // sample  $k$  parties randomly according to stake
44:     send message [QUERY,  $T$ ] to all parties  $v \in \mathcal{S}[T]$ 
45:      $\mathcal{D} \leftarrow \mathcal{D} \cup \{(\perp, \mathcal{VF} \setminus \{T\})\}$  // create a no-op transaction
46:     start timer[ $T$ ] // duration  $\Delta_{query}$ 
47:      $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{T\}$ 
48:      $updateDAG(T)$ 

49: upon receiving message [QUERY,  $T$ ] from party  $v$  do
50:   send message [VOTE,  $u, T, stronglyPreferred(T)$ ] to party  $v$ 

51: upon receiving message [VOTE,  $v, T, w$ ] such that  $v \in \mathcal{S}[T]$  do //  $w$  is the vote
52:    $votes[T, v] \leftarrow w$  //  $w \in \{\text{FALSE}, \text{TRUE}\}$ 

```

---

## 5 Security analysis

Avalanche deviates from the established PoW protocols and uses a different structure. Its security guarantees must be assessed differently. The bedrock of security for Avalanche is random sampling.

### 5.1 From Snowball to Avalanche

The Avalanche protocol family includes Slush, Snowflake, and Snowball [29] that implement single-decision Byzantine consensus. Every party *proposes* a value and every party must eventually *decide* the same value for an instance. The Avalanche protocol itself provides a “payment system” [29, Sec. V]; we model it here as generic broadcast.

## 10:12 When Is Spring Coming? A Security Analysis of Avalanche Consensus

■ **Algorithm 3** Avalanche (party  $u$ ), part 2.

---

```

53: upon  $\exists T \in \mathcal{T}$  such that  $|\{v \in \mathcal{S}[T] \mid \text{votes}[T, v] = \text{TRUE}\}| \geq \alpha$  do // query successful
54:   stop  $\text{timer}[T]$ 
55:    $\text{votes}[T, *] \leftarrow \perp$  // remove all entries in  $\text{votes}$  for  $T$ 
56:    $\mathcal{S}[T] \leftarrow []$  // reset  $\mathcal{S}$  for  $T$ 
57:    $d[T] \leftarrow d[T] + 1$ 
58:   for  $T' \in \text{ancestors}(T)$  do // all ancestors of  $T$ 
59:      $d[T'] \leftarrow d[T'] + 1$ 
60:     if  $d[T'] > d[\text{pref}[\text{conflictSet}[T']]]$  then
61:        $\text{pref}[\text{conflictSet}[T']] \leftarrow T'$ 
62:     if  $T' \neq \text{last}[\text{conflictSet}[T']]$  then
63:        $\text{last}[\text{conflictSet}[T']] \leftarrow T'$ 
64:        $\text{cnt}[\text{conflictSet}[T']] \leftarrow 1$ 
65:     else
66:        $\text{cnt}[\text{conflictSet}[T']] \leftarrow \text{cnt}[\text{conflictSet}[T']] + 1$ 
67:    $\text{conPoll} \leftarrow \text{conPoll} - 1$ 

68: upon  $\exists T \in \mathcal{T}$  such that  $|\{v \in \mathcal{S}[T] \mid \text{votes}[T, v] = \text{FALSE}\}| > k - \alpha$  do // query failed
69:   stop  $\text{timer}[T]$ 
70:    $\text{votes}[T, *] \leftarrow \perp$  // remove all entries in  $\text{votes}$  for  $T$ 
71:    $\mathcal{S}[T] \leftarrow []$  // reset  $\mathcal{S}$  for  $T$ 
72:   for  $T' \in \text{ancestors}(T)$  do // all ancestors of  $T$ 
73:      $\text{cnt}[\text{conflictSet}[T']] \leftarrow 0$ 

74: upon  $\exists T \in \mathcal{T}$  such that  $\text{acceptable}(T) \wedge \neg \text{accepted}[T]$  do //  $T$  can be accepted
75:    $(tx, \text{parents}) \leftarrow T$ 
76:   if  $V(tx)$  then
77:      $\text{accepted}[T] \leftarrow \text{TRUE}$ 
78:     deliver tx

80: upon timeout from  $\text{timer}[T]$  do // not enough votes on  $T$  received
81:    $\mathcal{Q} \leftarrow \mathcal{Q} \setminus \{T\}$ 
82:    $\text{votes}[T, *] \leftarrow \perp$  // remove all entries in  $\text{votes}$  for  $T$ 
83:    $\mathcal{S}[T] \leftarrow []$  // do not consider more votes from this query

```

---

The whitepaper [29] meticulously analyzes the three consensus protocols. It shows that as long as  $f = O(\sqrt{n})$ , the consensus protocols are live and safe [29] based on the analysis of random sampling [26]. On the other hand, an adversary controlling more than  $\Theta(\sqrt{n})$  parties may have the ability to keep the network in a bivalent state. For the remainder of this section we assume  $f = O(\sqrt{n})$ .

However, the Avalanche protocol itself is introduced without a rigorous analysis. The most precise statement about its is that “*it is easy to see that, at worst, Avalanche will degenerate into separate instances of Snowball, and thus provide the same liveness guarantee for virtuous transactions*” [29, p. 9]. In fact, it is easy to see that this is *wrong* because every vote on a transaction in Avalanche is linked to the vote on its ancestors. The vote on a descendant  $T'$  of  $T$  depends on the state of  $T$ .

---

**Algorithm 4** Avalanche, auxiliary functions.

---

```

84: function updateDAG( $T$ )
85:    $\mathcal{VF} \leftarrow$  set of non-conflicting leaves in the DAG
86:    $\text{conflictSet}[T] \leftarrow \emptyset$ 
87:   for  $T' \in \mathcal{T}$  such that  $T' \neq T$  and  $T'$  has a common input with  $T$  do
88:      $\text{conflictSet}[T] \leftarrow \text{conflictSet}[T] \cup \{T'\}$ 
89:      $\text{conflictSet}[T'] \leftarrow \text{conflictSet}[T'] \cup \{T\}$ 
90:   if  $\text{conflictSet}[T] = \emptyset$  then //  $T$  is non-conflicting
91:      $\text{pref}[\text{conflictSet}[T]] \leftarrow T$ 
92:      $\text{last}[\text{conflictSet}[T]] \leftarrow T$ 
93:      $\text{cnt}[\text{conflictSet}[T]] \leftarrow 0$ 
94:      $\text{conflictSet}[T] \leftarrow \text{conflictSet}[T] \cup \{T\}$ 

95: function getParents( $T$ )
96:    $(tx, \text{parents}) \leftarrow T$ 
97:   return  $\text{parents}$  // set of parents stored in  $T$ 

98: function preferred( $T$ )
99:   return  $T \stackrel{?}{=} \text{pref}[\text{conflictSet}[T]]$ 

100: function stronglyPreferred( $T$ )
101:   return  $\bigwedge_{T' \in \text{ancestors}(T)} \text{preferred}(T')$ 

102: function acceptable( $T$ )
103:   return  $(|\text{conflictSet}[T]| = 1 \wedge \text{cnt}[\text{conflictSet}[T]] \geq \beta_1) \wedge \bigwedge_{T' \in \text{parents}(T)} \text{acceptable}(T')$ 
104:    $\vee \text{cnt}[\text{conflictSet}[T]] \geq \beta_2$ 

105: function isRejected( $T$ )
106:   return  $\exists T' \in \mathcal{T}$  such that  $\forall T' \in \text{conflictSet}[T] \setminus \{T\} : \text{acceptable}(T')$ 

106: function updateRepollable()
107:    $\mathcal{R} \leftarrow \emptyset$ 
108:   for  $T \in \mathcal{T}$  do
109:     if  $\text{acceptable}(T) \vee \bigwedge_{T' \in \text{parents}(T)} \text{stronglyPreferred}(T') \wedge \neg \text{isRejected}(T')$  then
110:        $\mathcal{R} \leftarrow \mathcal{R} \cup \{T\}$ 

```

---

However, we can isolate single executions of Snowball that occur inside Avalanche. For an execution of Avalanche and a transaction  $T$ , we define an *equivalent* execution of Snowball consensus as the execution in which a party  $u$  proposes 1 if it locally prefers  $T$  in the Avalanche execution, proposes 0 if  $u$  prefers some other transaction, and does not propose otherwise. Every party also selects the same parties in each round of snowball and for a query with  $T$ , for a query with a transaction that conflicts with  $T$ , or for any query with a descendant of these two. A formal description of Snowball is provided in the full version [2].

► **Lemma 5.1.** *If party  $u$  delivers an honest transaction in Avalanche, then  $u$  decided 1 in the equivalent execution of Snowball with threshold  $\beta_1$ . Furthermore,  $u$  delivers a conflicting transaction in Avalanche, then  $u$  decides 1 in Snowball with threshold  $\beta_2$ .*

**Proof.** By construction of the Avalanche and Snowball protocols [29], the counter for acceptance of value 1 in Snowball is always greater or equal than the counter for acceptance in Avalanche. Since a successful query in Avalanche implies a successful query in Snowball,

## 10:14 When Is Spring Coming? A Security Analysis of Avalanche Consensus

if an honest transaction in Avalanche is delivered, the counter in the equivalent Snowball instance is at least  $\beta_1$ . Analogously, if a conflicting transaction in Avalanche is delivered, then the counter in Snowball is at least  $\beta_2$ . Hence, a party in Snowball would decide 1 with the respective thresholds. ◀

Looking ahead, we will introduce a modification of Avalanche that ensures the complete equivalence between Snowball and Avalanche. We first assert some safety properties of the Avalanche protocol.

► **Theorem 5.2.** *Avalanche satisfies integrity, partial order, and external validity of a generic broadcast for payload transactions under relation  $\sim$  and UTXO-validity.*

**Proof.** The proof is structured by property:

- **Integrity.** We show that every payload is delivered at most once. A payload  $tx$  may potentially be delivered multiple times in two ways: different protocol transactions that both carry  $tx$  may be accepted or  $tx$  is delivered multiple times as payload of the same protocol transaction.

First, we consider the possibility of accepting two different transactions  $T_1$  and  $T_2$  carrying  $tx$ . Assume that party  $u$  accepts transaction  $T_1$  and party  $v$  accepts transaction  $T_2$ . By definition,  $T_1$  and  $T_2$  are *conflicting* because they spend the same inputs. Using Lemma 5.1, party  $u$  and  $v$  decide differently in the equivalent execution in Snowball, which contradicts agreement property of the Snowball consensus [29].

The second option is that one protocol transaction  $T$  that contains  $tx$  is accepted multiple times. However, this is not possible either because  $tx$  is delivered only if  $accepted[T] = \text{FALSE}$ ; variable  $accepted[T]$  is set to **TRUE** when transaction  $T$  is accepted (L 74–78).

- **Partial order.** Avalanche satisfies partial order because no payload is valid unless all payloads creating its inputs have been delivered (L 74–78). Transactions  $T$  and  $T'$  are related according to Definition 3.1 if and only if  $T$  has as input (i.e., spends) at least one output of  $T'$ , or vice versa. This implies that related transactions are delivered in the same order for any party.
- **External validity.** The external validity property follows from L 74, as a payload transaction can only be delivered if it is valid, i.e., its inputs have not been previously spent and the cryptographic requirements are satisfied. ◀

Theorem 5.2 shows that Avalanche satisfies the safety properties of a generic broadcast in the presence of an adversary controlling  $O(\sqrt{n})$  parties. A hypothetical adversary controlling substantially more parties could violate safety. It is not completely obvious how an adversary could achieve that. Such an adversary would broadcast two conflicting transactions  $T_1$  and  $T_2$ . As we already discussed, and also explained in the whitepaper of Avalanche [29], such an adversary can keep the network in a bivalent state, so the adversary keeps the network divided into two parts: parties in part  $\mathcal{P}_1$  consider  $T_1$  preferred, and parties in part  $\mathcal{P}_2$  prefer  $T_2$ . The adversary behaves as preferring  $T_1$  when communicating with parties in  $\mathcal{P}_1$  and as preferring  $T_2$  when communicating with parties in  $\mathcal{P}_2$ . Eventually, a party  $u \in \mathcal{P}_1$  will query only parties in  $\mathcal{P}_1$  or queries the adversary  $\beta_2$  times in a row. Thus,  $u$  will accept transaction  $T_1$ . Similarly, a party  $v \in \mathcal{P}_2$  will eventually accept transaction  $T_2$ . Party  $u$  will *deliver* the payload contained in  $T_1$  and  $v$  the payload contained in  $T_2$ , hence violating agreement. An adversary controlling at most  $O(\sqrt{n})$  can also violate agreement, but the required behavior is more sophisticated, as we explain next.



## 5.2 Delaying transaction acceptance

An adversary aims to prevent that a party  $u$  accepts an honest transaction  $T$ . A necessary precondition for this is  $\text{cnt}[\text{conflictSet}[T]] \geq \beta_1$ . Note that whenever a descendant of  $T$  is queried,  $\text{cnt}[\text{conflictSet}[T]]$  is modified. If the query is successful (L 53), then  $\text{cnt}[\text{conflictSet}[T]]$  is incremented by one. If the query is unsuccessful,  $\text{cnt}[\text{conflictSet}[T]]$  is reset to zero. Remark, however,  $\text{cnt}[\text{conflictSet}[T]]$  cannot be reset to one as a result of another transaction becoming the preferred in  $\text{conflictSet}[T]$  (L 62) because  $T$  is honest, as there exist no transaction conflicting with  $T$ .

Furthermore, a naive adversary that aims to delay transactions by not answering the query of a transaction  $T'$  would not succeed because the timers  $\text{timeout}[T']$  would be triggered and the query would be aborted. Thus, the honest party would select new  $k$  parties to query and proceed in the protocol.

Our adversary proceeds by sending to  $u$  a series of cleverly generated transactions that reference  $T$ . We describe the steps that will delay the acceptance of  $T$  (see also Algorithm 5):

1. **Preparation phase.** The adversary submits conflicting transactions  $T_1$  and  $T_2$ . For simplicity, we assume that she submits first  $T_1$  and then  $T_2$ , so the preferred transaction in both conflict sets will be  $T_1$ . The adversary then waits until the target transaction  $T$  is submitted.
2. **Main phase.** The adversary repeatedly sends malicious transactions referencing the target  $T$  and  $T_2$  to  $u$ . These transactions are valid but they reference a particular set of transactions.
3. **Searching phase.** Concurrently to the main phase, the adversary looks for transactions containing the same payload as  $T$ . If some are found, she references them as well from the newly generated transactions.

■ **Algorithm 5** Liveness attack: Delaying transaction  $T$ .

---

### Initialization

```

111:   create two conflicting transactions  $T_1$  and  $T_2$ 
112:   gossip two messages [BROADCAST,  $T_1$ ] and [BROADCAST,  $T_2$ ]
113:    $\mathcal{A} \leftarrow \emptyset$ 

114: upon hearing message [BROADCAST,  $T$ ] do                                // target transaction
115:    $\mathcal{A} \leftarrow \{T\}$ 

116: upon  $\text{cnt}[\text{conflictSet}[T]] = \lfloor \frac{\beta_1}{2} \rfloor$  in the local view of  $u$  do
117:   create  $\hat{T}$  such that  $T_2 \in \text{ancestors}(\hat{T})$  and for all  $T' \in \mathcal{A}$ , also  $T' \in \text{ancestors}(\hat{T})$ 
118:   send message [BROADCAST,  $\hat{T}$ ] to party  $u$  // pretend to gossip the message

119: upon hearing message [BROADCAST,  $\tilde{T}$ ] such  $\tilde{T}$  and  $T$  contain the same payload do
120:    $\mathcal{A} \leftarrow \mathcal{A} \cup \{\tilde{T}\}$ 

```

---

For simplicity, we assume that the adversary knows the acceptance counter of  $T$  at  $u$ , so she can send a malicious transaction whenever  $T$  is close to being accepted. In practice, she can guess this only with a certain probability, which will degrade the success rate of the attack. We also assume that the query of an honest transaction is always successful, which is the worst case for the adversary.

## 10:16 When Is Spring Coming? A Security Analysis of Avalanche Consensus

After  $u$  submits  $T$ , the adversary starts the main phase of the attack. If  $u$  queries an honest transaction  $\hat{T}$ , and if  $\hat{T}$  references a descendant of  $T$ , then  $\text{cnt}[\text{conflictSet}[T]]$  increases by one. If it does not, then  $\hat{T}$  may cause  $u$  to submit a no-op transaction referencing a descendant of  $T$ . Hence, honest transactions always increase  $\text{cnt}[\text{conflictSet}[T]]$  by one, this is the worst case for an adversary aiming to delay the acceptance of  $T$ .

If  $u$  queries a malicious transaction  $\hat{T}$ , then honest parties compute  $\text{stronglyPreferred}(\hat{T})$  and reply with this value. Since  $T_2$  is an ancestor of  $\hat{T}$  and not the preferred transaction in its conflict set (as we have assumed that  $T_1$  is preferred), all queried parties return FALSE. Thus,  $u$  sets acceptance counter of every ancestor of  $\hat{T}$  to zero (L 68), in particular,  $\text{cnt}[\text{conflictSet}[T]] \leftarrow 0$ . However, since  $\hat{T}$  does not reference the virtuous frontier,  $u$  submits a no-op transaction that references a descendant of  $T$ , thus increasing  $\text{cnt}[\text{conflictSet}[T]]$  to one.

We show that when the number of transactions is low, in particular when  $|\mathcal{T} \setminus \mathcal{Q}| \leq 1$  for every party, then Avalanche may lose liveness.

► **Theorem 5.3.** *Avalanche does not satisfy validity nor agreement of generic broadcast with relation  $\sim$  with one single malicious party if  $|\mathcal{T} \setminus \mathcal{Q}| \leq 1$  for every party.*

**Proof.** We consider again the adversary described above that targets  $T$  and  $u$ .

- **Validity.** Whenever  $\text{cnt}[\text{conflictSet}[T]]$  in the local view of  $u$  reaches  $\lfloor \frac{\beta_1}{2} \rfloor$ , the adversary sends a malicious transaction to party  $u$ , who immediately queries it (since  $|\mathcal{T} \setminus \mathcal{Q}| \leq 1$ ). It follows that  $u$  sets  $\text{cnt}[\text{conflictSet}[T]]$  to zero and increases it intermediately afterwards, due to a no-op transaction. This process repeats indefinitely over time and prevents  $u$  from delivering the payload in  $T$ .
- **Agreement.** Assume that an honest party broadcasts the payload contained in  $T$ . The adversary forces a violation of agreement by finding honest parties  $u$  and  $v$  such that  $\text{cnt}[\text{conflictSet}[T]] = \beta_1 - 1$  at  $v$  and  $\text{cnt}[\text{conflictSet}[T]] < \beta_1 - 1$  at  $u$  (such parties exist because in the absence of an adversary, as  $\text{cnt}[\text{conflictSet}[T]]$  increases monotonically over time). The adversary then sends an honest transaction  $T_h$  that references  $T$  to  $v$  and a malicious transaction  $T_m$ , as described before, to  $u$ . On the one hand, party  $v$  queries  $T_h$ , increments  $\text{cnt}[\text{conflictSet}[T]]$  to  $\beta_1$ , accepts transaction  $T$ , and delivers the payload. On the other hand, party  $u$  queries  $T_m$  and sets  $\text{cnt}[\text{conflictSet}[T]]$  to one. After that, the adversary behaves as discussed before. Notice that  $v$  has delivered the payload within  $T$  but  $u$  will never do so. ◀

An adversary may thus cause Avalanche to violate validity and agreement. For this attack, however, the number of transactions in the network must be low, in particular,  $|\mathcal{T} \setminus \mathcal{Q}| \leq 1$ . In July 2022, the Avalanche network processed an average of 647238 transactions per day (<https://subnets.avax.network/stats/network>). Assuming two seconds per query, four times the value observed in our local implementation, the recommended values of 30 transactions per batch, and four concurrent polls, the condition  $|\mathcal{T} \setminus \mathcal{Q}| \leq 1$  is satisfied 88% of the time. However, the adversary still needs to know the value of the counter for acceptance of the different parties.

### 5.3 A more general attack

We may relax the assumption of knowing the acceptance counters and also send the malicious transaction to more parties through gossip. After selecting a target transaction, the adversary continuously gossips malicious transactions to the network instead of sending them only to one party as in Algorithm 5. For analyzing the performance of this attack, our figure of merit will be the number of transactions to be queried by an honest party (not counting

no-ops) for confirming the target transaction  $T$ . The larger this number becomes, the longer it will take the party until it may accept  $T$ . We assume that  $\mathcal{T} \setminus \mathcal{Q} \neq \emptyset$  and that a fraction  $\gamma$  of those transactions are malicious at any point in time<sup>1</sup>. A non-obvious implication is that the repoll function never queries the same transaction twice.

► **Lemma 5.4.** *Avalanche requires every party to query at least  $\beta_1$  transactions before accepting transaction  $T$  in the absence of an adversary.*

**Proof.** The absence of an adversary carries several simplifications. Firstly, there are no conflicting transactions, thus every transaction is the preferred one in its respecting conflict set and every query is successful. Secondly, due to the no-op transactions, the counter for acceptance of every transaction in the DAG is incremented by one after each query. Finally, a transaction  $T$  is accepted when its counter for acceptance reaches  $\beta_1$ , since the counter of the parent of any transaction reaches  $\beta_1$  strictly before  $T$  (L 102). ◀

► **Lemma 5.5.** *The average number of queried transactions before accepting transaction  $T$  in the presence of the adversary, as described in the text, is at least*

$$\beta_1 + \frac{1 + (2 + \beta_1\gamma)(1 - \gamma)^{\beta_1} - (1 - \gamma)^{2\beta_1}(1 + \beta_1\gamma)}{\gamma(1 - \gamma)^{\beta_1}(1 - (1 - \gamma)^{\beta_1})}.$$

**Proof.** We recall that in the worst-case scenario for the adversary, the query of an honest transaction increments the counter for acceptance of the target transaction  $T$  by one, while the query of a malicious transaction, effectively, resets the counter for acceptance to one, as a result of a no-op transaction.

Let a random variable  $W$  denote the number of transactions queried by  $u$  until  $T$  is accepted, and let  $X \in \{0, 1\}$  model the outcome of the following experiment. Party  $u$  samples transactions until it picks a malicious transaction or until it has sampled  $\beta_1 - 1$  honest transactions. In the first case,  $X$  takes the value zero, and otherwise,  $X$  takes the value one. By definition,  $X$  is a Bernoulli variable with parameter  $p = (1 - \gamma)^{(\beta_1 - 1)}$ . Thus, the number of attempts until  $X$  returns one is a random variable  $Y$  with geometric distribution,  $Y \sim \mathcal{G}(p)$ , with the same parameter  $p$ . We let  $W_a$  be the random variable denoting the number of queried transactions per attempt of this experiment. The expected number of failed attempts is  $E[Y] = \frac{1}{(1 - \gamma)^{\beta_1}}$ . Furthermore, the probability that an attempt fails after sampling exactly  $k$  transactions, for  $k \leq \beta_1$ , is

$$P[W_a = k | X = 0] = \frac{\gamma(1 - \gamma)^{k-1}}{1 - (1 - \gamma)^{\beta_1}}.$$

Thus, the expected number of transactions per failed attempt can be expressed as

$$E[W | X = 0] = \frac{1 - (1 - \gamma)^{\beta_1}(1 + \beta_1\gamma)}{\gamma(1 - (1 - \gamma)^{\beta_1})}. \quad (1)$$

The expected number of transaction queried during a successful attempt is at least  $\beta_1$  by Lemma 5.4. Finally, the total expected number of queried transactions can be written as the expected number of transaction per failed attempt multiplied by the expected number of failed attempts plus the expected number of transactions in the successful attempt,

$$E[W] = E[W_a | X = 0] \cdot (E[Y] - 1) + E[W_a | X = 1] \cdot 1. \quad (2)$$


<sup>1</sup> Avalanche may impose a transaction fee for processing transactions. However, since the malicious transactions cannot be delivered, this mechanism does not prevent the adversary from submitting a large number of transactions.

## 10:18 When Is Spring Coming? A Security Analysis of Avalanche Consensus

From equations (1) and (2) and basic algebra, we obtain

$$E[W] = \beta_1 + \frac{1 + (2 + \beta_1\gamma)(1 - \gamma)^{\beta_1} - (1 - \gamma)^{2\beta_1}(1 + \beta_1\gamma)}{\gamma(1 - \gamma)^{\beta_1}(1 - (1 - \gamma)^{\beta_1})}.$$

This expression is complex to analyze. Hence, a graphical representation of this bound is given in Figure 2. It shows the expected smallest number of transactions to be queried by an honest party (not counting no-ops) until it can confirm the target transaction  $T$ . The larger this gets, the more the protocol loses liveness. It is relevant that this bound grows proportional to  $\frac{1}{(1-\gamma)^{\beta_1}}$ , i.e., exponential in acceptance threshold  $\beta_1$  since  $(1 - \gamma) < 1$ .



■ **Figure 2** Expected delay in number of transactions needed to confirm a given transaction with acceptance threshold  $\beta_1 = 15$ , the recommended value [4], and assuming that the queries of honest transactions are successful. The (green) horizontal line shows  $\beta_1$ , the expected delay without attacker. The (blue) dotted line represents the expected confirmation delay in Avalanche depending on the fraction of malicious transactions. The (orange) squared line denotes the delay in Glacier (Section 6).

The Avalanche team has acknowledged our findings and the vulnerability. The protocol deployed in the actual network, however, differs from our formalization in a way that should prevent the problem.

### 6 Fixing liveness with Glacier

The adversary is able to delay the acceptance of an honest transaction  $T$  because  $T$  is directly influenced by the queries of its descendants. Note the issuer of  $T$  has no control over its descendants according to the protocol. A unsuccessful query of a descendant of  $T$  carries a negative consequence for the acceptance of  $T$ , regardless of the status of  $T$  inside its conflict set. This influence is the root of the problem described earlier. An immediate, but inefficient remedy might be to run one Snowball consensus instance for each transaction. However, this would greatly degrade the throughput and increase the latency of the protocol, as many more messages would be exchanged.

We propose here a modification, called *Glacier*, in which an unsuccessful query of a transaction  $T$  carries negative consequences *only* for those of its ancestors that led to negative votes and caused the query to be unsuccessful. Our protocol is shown in Algorithm 6. It

specifically modifies the voting protocol and adds to each VOTE message for  $T$  a list  $L$  with all ancestors of  $T$  that are not preferred in their respective conflict sets (L 123–127). When party  $u$  receives a negative vote like  $[\text{VOTE}, v, T, \text{FALSE}, L]$ , it performs the same actions as before. Additionally, it increments a counter for each ancestor  $T^*$  of  $T$  to denote how many parties have reported  $T^*$  as not preferred while accepting  $T$  (L 135). If  $u$  receives a positive vote, the protocol remains unchanged.

If the query is successful because  $u$  receives at least  $\alpha$  positive votes on  $T$ , then it proceeds as before (Algorithm 3, L 53). But before  $u$  declares the query to be unsuccessful, it furthermore waits until having received a vote on  $T$  from all  $k$  parties sampled in the query (L 137). When this is the case,  $u$  only resets the counter for acceptance of those ancestors  $T^*$  of  $T$  that have been reported as non-preferred by more than  $k - \alpha$  queried parties (L 141–143). If  $T^*$  is preferred by at least  $\alpha$  parties, however, then  $u$  increments its confidence level as before (L 145).

■ **Algorithm 6** Modifications to Avalanche (Algorithm 1–4) for Glacier (party  $u$ ).

---

**State**

```

121:   nonpref:  $\text{HashMap}[\mathcal{T} \times \mathcal{T} \rightarrow \mathbb{N}]$  // votes on  $T$  saying  $T'$  is not preferred

122: upon receiving message  $[\text{QUERY}, T]$  from party  $v$  do // replaces L 49
123:    $L \leftarrow []$  // contains the non-preferred ancestors of  $T$ 
124:   for  $T' \in \text{ancestors}(T)$  do
125:     if  $\neg \text{preferred}(T')$  then
126:       append  $T'$  to  $L$ 
127:   send message  $[\text{VOTE}, v, T, \text{stronglyPreferred}(T), L]$  to party  $v$ 

128: // replaces code at L 51
129: upon receiving message  $[\text{VOTE}, v, T, w, L]$  from a party  $v \in \mathcal{S}[T]$  do //  $w$  is the vote
130:    $\text{votes}[T, v] \leftarrow w$ 
131:   for  $T' \in L$  do
132:     if  $\text{nonpref}[T, T'] = \perp$  then
133:        $\text{nonpref}[T, T'] \leftarrow 1$ 
134:     else
135:        $\text{nonpref}[T, T'] \leftarrow \text{nonpref}[T, T'] + 1$ 

136: // replaces code at L 68
137: upon  $\exists T \in \mathcal{T}$  such that  $|\text{votes}[T, v]| = k \wedge \left| \{v \in \mathcal{S}[T] \mid \text{votes}[T, v] = \text{FALSE}\} \right| > k - \alpha$  do
138:   stop  $\text{timer}[T]$ 
139:    $\text{votes}[T, *] \leftarrow \perp$  // remove all entries in  $\text{votes}$  for  $T$ 
140:    $\mathcal{S}[T] \leftarrow []$  // reset the  $\text{HashMap } \mathcal{S}$ 
141:   for  $T'$  such that  $\text{nonpref}[T, T'] \neq \perp$  do // all ancestors of  $T$ 
142:     if  $\text{nonpref}[T, T'] > k - \alpha$  then
143:        $\text{cnt}[\text{conflictSet}[T']] \leftarrow 0$ 
144:     else //  $\text{nonpref}[T, T'] \leq \alpha$ 
145:        $\text{cnt}[\text{conflictSet}[T']] \leftarrow \text{cnt}[\text{conflictSet}[T']] + 1$ 
146:    $\text{nonpref}[T, *] \leftarrow \perp$ 

```

---

Considering the adversary introduced in Section 5.3, a negative reply to the query of a descendant of the target transaction  $T$  does not carry any negative consequence for the acceptance of  $T$  here. In particular, the counter for acceptance of transaction  $T$  is never reset, even when a query is unsuccessful, because  $T$  is the only transaction in its conflicting set, then always preferred. Thus, transaction  $T$  will be accepted after  $\beta_1$  successful queries, if all its parents are accepted, or  $\beta_2$  successful queries if they are not accepted. Assuming that queries of honest transactions are successful, on average  $\frac{\beta}{1-p}$  transactions are required

to accept  $T$  for  $\beta \in [\beta_1, \beta_2]$  depending on the state of the parents of  $T$ . For simplicity we assume that the parents are accepted, thus, the counter needs to achieve the value  $\beta_1$ . If this were not the case, then it is sufficient to substitute  $\beta_1$  with  $\beta$  in the upcoming expression. Avalanche requires on average  $\beta_1 + \frac{1+(2+\beta_1\gamma)(1-\gamma)^{\beta_1}-(1-\gamma)^{2\beta_1}(1+\beta_1\gamma)}{\gamma(1-\gamma)^{\beta_1}(1-(1-\gamma)^{\beta_1})}$  transactions to accept  $T$  by Lemma 5.5. The assumption that the query of honest transactions is always successful is more beneficial to Avalanche than to Glacier, since in Avalanche such a query resets the counter for acceptance of  $T$ . But in Glacier, the query simply leaves the counter as it is. The value of the acceptance threshold  $\beta_1$  is also more beneficial for Avalanche since the number of required transactions increases linearly in Glacier and exponentially in Avalanche. Figure 2 shows a comparison of both expressions.

In Glacier, the vote for a transaction is independent of the vote of its descendant and ancestors, even if a query of a transaction carries an implicit query of all its ancestors. Thus, Lemma 5.1 can be extended.

► **Lemma 6.1.** *Party  $u$  delivers a transaction  $T$  with counter for acceptance with value  $\text{cnt}[\text{conflictSet}[T]] \geq \beta_1$  in Glacier if and only if  $u$  decides 1 in the equivalent execution of Snowball with threshold  $\text{cnt}[\text{conflictSet}[T]]$ .*

**Proof.** Consider a transaction  $T$  in the equivalent execution of Snowball. The counter for acceptance of the value 1 in Snowball is always the same as the counter for acceptance of transaction  $T$  in Glacier because of the modifications introduced by Glacier. Thus, following the same argument as in Lemma 5.1, transaction  $T$  is accepted in Glacier with counter  $\text{cnt}[\text{conflictSet}[T]]$  if and only if 1 is decided with counter  $\text{cnt}[\text{conflictSet}[T]]$  in the equivalent execution of Snowball. ◀

► **Theorem 6.2.** *The Glacier algorithm satisfies the properties of generic broadcast in the presence of an adversary that controls up to  $\mathcal{O}(\sqrt{n})$  parties.*

**Proof.** Lemma 5.1 is a special case of Lemma 6.1. Theorem 5.2 shows that Lemma 5.1 and the properties of Snowball [2] guarantee that Avalanche satisfies integrity, partial order, and external validity. In the same way, Lemma 6.1 guarantees that Glacier satisfies these same properties. Thus, it is sufficient to prove that Glacier satisfies validity and agreement.

- **Validity.** Assume that an honest party broadcasts a payload  $tx$ . Because the party is honest, the transaction  $T$  containing  $tx$  is valid and non-conflicting. In the equivalent execution of Snowball, every honest party that proposes a value proposes 1. Hence, using the validity and termination properties of Snowball, every honest party eventually decides 1. Using Lemma 6.1, every honest party eventually delivers  $tx$ .
- **Agreement.** Assume that an honest party delivers a payload transaction  $tx$  contained in transaction  $T$ . Using Lemma 6.1, an honest party decides 1 in the equivalent execution of Snowball. Because of the termination and agreement properties of Snowball, every honest party decides 1. Using Lemma 6.1 again, every honest party eventually delivers payload  $tx$ .

We conclude that Glacier satisfies the properties of generic broadcast. ◀

With the modification to Glacier, Avalanche can be safely used as the basis for a payment system. Notice that the sample mechanism is not modified, thus remains the same as in the original protocol. The only possible concern with Glacier could be a decrease in performance compared to Avalanche. However, Glacier does not reduce the performance but rather improves it. Glacier only modifies the update in the local state of party  $u$  after a query has been unsuccessful. The counter of acceptance of a given transaction  $T$  in Glacier implementation is always greater or equal than its counterpart in Avalanche. This follows



because a reset of  $\text{cnt}[\text{conflictSet}[T]]$  in Glacier implies the same reset in Avalanche. Such a reset in Glacier occurs if the query of a descendant of  $T$  fails and  $T$  was reported as non-preferred by more than  $k - \alpha$  parties, whereas in Avalanche it is enough if the query of the descendant failed. In Avalanche,  $\text{cnt}[\text{conflictSet}[T]]$  is incremented if the query of a descendant of  $T$  succeeds, and the same occurs in Glacier. Thus,  $\text{cnt}[\text{conflictSet}[T]]$  in Glacier is at least as large as in Avalanche. We recall that a transaction is accepted when  $\text{cnt}[\text{conflictSet}[T]]$  reaches a threshold depending on some conditions of the local view of the DAG, but these are identical for Glacier and Avalanche. Hence, every transaction that is accepted in Avalanche is accepted in Glacier with equal or smaller latency. This implies not only that the latency of Glacier is smaller than the latency of Avalanche, but also that the throughput of Glacier is at least as good as the throughput of Avalanche.

## 7 Conclusion

Avalanche is well-known for its remarkable throughput and latency that are achieved through a metastable sampling technique. Our pseudocode captures in a compact and relatively simple manner the intricacies of the protocol. We show that Avalanche, as originally introduced, possesses a vulnerability allowing an adversary to delay transactions arbitrarily. We also address such vulnerability with a modification of the protocol, Glacier, that allows Avalanche to satisfy both safety and liveness.

The developers of Avalanche have acknowledged the vulnerability, and the actual implementation does not suffer from it due to an alternative fix. Understanding this variant of Avalanche remains open and is subject of future work.

---

## References

- 1 Ignacio Amores-Sesar, Christian Cachin, and Jovana Micic. Security analysis of ripple consensus. In *OPODIS*, volume 184 of *LIPICs*, pages 10:1–10:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
- 2 Ignacio Amores-Sesar, Christian Cachin, and Enrico Tedeschi. When is spring coming? A security analysis of avalanche consensus. *CoRR*, abs/2210.03423, 2022. [arXiv:2210.03423](https://arxiv.org/abs/2210.03423).
- 3 Frederik Armknecht, Ghassan O. Karame, Avikarsha Mandal, Franck Youssef, and Erik Zenner. Ripple: Overview and outlook. In *TRUST*, volume 9229 of *Lecture Notes in Computer Science*, pages 163–180. Springer, 2015.
- 4 Ava Labs, Inc. Avalanche documentation. <https://docs.avax.network/>.
- 5 Ava Labs, Inc. Node implementation for the Avalanche network. <https://github.com/ava-labs/avalanchego>.
- 6 Christian Cachin, Rachid Guerraoui, and Luís E. T. Rodrigues. *Introduction to Reliable and Secure Distributed Programming (2. ed.)*. Springer, 2011.
- 7 Christian Cachin, Klaus Kursawe, Frank Petzold, and Victor Shoup. Secure and efficient asynchronous broadcast protocols. In *CRYPTO*, volume 2139 of *Lecture Notes in Computer Science*, pages 524–541. Springer, 2001.
- 8 Christian Cachin and Marko Vukolic. Blockchain consensus protocols in the wild (keynote talk). In *DISC*, volume 91 of *LIPICs*, pages 1:1–1:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017.
- 9 Coinmarketcap: Today’s cryptocurrency prices by market cap. <https://coinmarketcap.com/>, 2022.
- 10 Bernardo David, Peter Gazi, Aggelos Kiayias, and Alexander Russell. Ouroboros praos: An adaptively-secure, semi-synchronous proof-of-stake blockchain. In *EUROCRYPT (2)*, volume 10821 of *Lecture Notes in Computer Science*, pages 66–98. Springer, 2018.

## 10:22 When Is Spring Coming? A Security Analysis of Avalanche Consensus

- 11 Cynthia Dwork, Nancy A. Lynch, and Larry J. Stockmeyer. Consensus in the presence of partial synchrony. *J. ACM*, 35(2):288–323, 1988.
- 12 Ittay Eyal and Emin Gün Sirer. Majority is not enough: Bitcoin mining is vulnerable. In *Financial Cryptography*, volume 8437 of *Lecture Notes in Computer Science*, pages 436–454. Springer, 2014.
- 13 Juan A. Garay, Aggelos Kiayias, and Nikos Leonardos. The bitcoin backbone protocol: Analysis and applications. In *EUROCRYPT (2)*, volume 9057 of *Lecture Notes in Computer Science*, pages 281–310. Springer, 2015.
- 14 Juan A. Garay, Aggelos Kiayias, and Nikos Leonardos. The bitcoin backbone protocol with chains of variable difficulty. In *CRYPTO (1)*, volume 10401 of *Lecture Notes in Computer Science*, pages 291–323. Springer, 2017.
- 15 Yossi Gilad, Rotem Hemo, Silvio Micali, Georgios Vlachos, and Nickolai Zeldovich. Algorand: Scaling byzantine agreements for cryptocurrencies. In *SOSP*, pages 51–68. ACM, 2017.
- 16 Idit Keidar, Eleftherios Kokoris-Kogias, Oded Naor, and Alexander Spiegelman. All you need is DAG. In *PODC*, pages 165–175. ACM, 2021.
- 17 Aggelos Kiayias, Alexander Russell, Bernardo David, and Roman Oliynykov. Ouroboros: A provably secure proof-of-stake blockchain protocol. In *CRYPTO (1)*, volume 10401 of *Lecture Notes in Computer Science*, pages 357–388. Springer, 2017.
- 18 Minjeong Kim, Yujin Kwon, and Yongdae Kim. Is stellar as secure as you think? In *EuroS&P Workshops*, pages 377–385. IEEE, 2019.
- 19 Chenxing Li, Peilun Li, Wei Xu, Fan Long, and Andrew Chi-Chih Yao. Scaling nakamoto consensus to thousands of transactions per second. *CoRR*, abs/1805.03870, 2018. [arXiv:1805.03870](https://arxiv.org/abs/1805.03870).
- 20 Chenxing Li, Peilun Li, Dong Zhou, Zhe Yang, Ming Wu, Guang Yang, Wei Xu, Fan Long, and Andrew Chi-Chih Yao. A decentralized blockchain with high throughput and fast confirmation. In *USENIX Annual Technical Conference*, pages 515–528. USENIX Association, 2020.
- 21 Marta Lohkava, Giuliano Losa, David Mazières, Graydon Hoare, Nicolas Barry, Eli Gafni, Jonathan Jove, Rafal Malinowsky, and Jed McCaleb. Fast and secure global payments with stellar. In *SOSP*, pages 80–96. ACM, 2019.
- 22 Hamed Mamache, Gabin Mazué, Osama Rashid, Gewu Bu, and Maria Potop-Butucaru. Resilience of IOTA consensus. *CoRR*, abs/2111.07805, 2021. [arXiv:2111.07805](https://arxiv.org/abs/2111.07805).
- 23 Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. Whitepaper, <https://bitcoin.org/bitcoin.pdf>, 2009.
- 24 Marshall C. Pease, Robert E. Shostak, and Leslie Lamport. Reaching agreement in the presence of faults. *J. ACM*, 27(2):228–234, 1980.
- 25 Fernando Pedone and André Schiper. Generic broadcast. In *DISC*, volume 1693 of *Lecture Notes in Computer Science*, pages 94–108. Springer, 1999.
- 26 Sheldon Ross. *Stochastic Processes*. Wiley, second edition, 1996.
- 27 Yonatan Sompolinsky, Shai Wyborski, and Aviv Zohar. PHANTOM GHOSTDAG: a scalable generalization of nakamoto consensus: September 2, 2021. In *AFT*, pages 57–70. ACM, 2021.
- 28 W. Y. Tan. On the absorption probabilities and absorption times of finite homogeneous birth-death processes. *Biometrics*, 32(4):745–752, 1976.
- 29 Team Rocket, Maofan Yin, Kevin Sekniqi, Robbert van Renesse, and Emin Gün Sirer. Scalable and probabilistic leaderless BFT consensus through metastability. e-print, [arXiv:1906.08936](https://arxiv.org/abs/1906.08936) [cs.CR], 2019.
- 30 Bozhi Wang, Qin Wang, Shiping Chen, and Yang Xiang. Security analysis on tangle-based blockchain through simulation. *CoRR*, abs/2008.04863, 2020. [arXiv:2008.04863](https://arxiv.org/abs/2008.04863).
- 31 Qin Wang, Jiangshan Yu, Zhiniang Peng, Van Cuong Bui, Shiping Chen, Yong Ding, and Yang Xiang. Security analysis on dbft protocol of NEO. In *Financial Cryptography*, volume 12059 of *Lecture Notes in Computer Science*, pages 20–31. Springer, 2020.
- 32 Maofan Yin. *Scaling the Infrastructure of Practical Blockchain Systems*. PhD thesis, Cornell University, USA, 2021.