

DGZT Protocol: Consensus and Security

Oort Founding Team *

Abstract

In this paper, we build up the mathematical foundations of DGZT Protocol, which advanced the directed acyclic graph (DAG) for storing transactions. DGZT has superior performance such as high throughput and almost-zero transaction fees. The nature of its original architecture enables the flexibility and capability of chain interoperability with all other layer-1 blockchains, which is the foundation for building a Web3 data infrastructure. We provide thorough and rigorous analysis on our consensus mechanism which depends on non-anonymous reputable entities, called committees. Our scheme allows committees to be replaced to achieve higher level of decentralization. The security of DGZT-Protocol against malicious behaviors is guaranteed.

1 Introduction

The concept of blockchain as an independent technology began to surge in 2015. Prior to this, it was known as the data structure of Bitcoin. In Nakamoto's white paper [1], the two words "block" and "chain" appear together, but it only refers to "a series of blocks." With the popularity of Bitcoin, the technology and concepts in Bitcoin is often classified as Blockchain 1.0. With Ethereum [2] running as a platform for distributed applications, people began to classify Ethereum as Blockchain 2.0. Now the market is vying for the fundamental structure for a new paradigm of Internet infrastructure, interoperability and scalability, i.e., Blockchain 3.0. Many people think that directed acyclic graph (DAG) structure is one of the best candidates.

In traditional blockchain technology represented by Bitcoin and Ethereum, blocks and transactions are two separate concepts. A transaction is confirmed by the miners and packed into a block, and the throughput in terms

*Oort was formerly called "Computecoin Networks". Please visit our official website to stay up-to-date on our progress, and to view the latest version of this technical paper.

of transactions per second (TPS) is limited by the block size and the block generation speed. In addition, miners in the blockchain system have the right to decide the content of the block. The profit-seeking behavior of the miners can easily lead to excessive concentration of power or voting rights, thus losing the decentralization characteristics. DAG-based distributed ledger technology (DLT) was created to solve these problems. Compared to traditional blockchain technology, DAG-based DLT has the following advantages: 1) Strong scalability (high TPS); 2) Fast transaction speed; 3) (Almost) no transaction fee and friendly to small payments; 3) No requirement for special miners to participate.

The idea of using DAGs in the cryptocurrency space has been around for a while. DAGLabs has proposed a series of consensus protocols, such as Inclusive [3], SPECTRE [4] and PHANTOM [5]. The general idea behind them is to utilize a DAG of blocks. Also the miners in the system still compete for transaction fees, and new tokens may be created by these miners. Instead, some cryptocurrencies depend on a DAG of individual transactions other than blocks. IOTA [6] and Byteball¹ [7] are among the oldest and most representative projects. They both have the same advantages using a DAG structure, but have quite different design details in order to cater to different audiences. IOTA assigns a certain weight to each transaction, and the transaction is generated through the proof of work (PoW) mechanism. Instead of utilizing PoW, Byteball prevents junk transactions by charging a small fee, and introduces votes from committees to determine valid transactions.

Similar to IOTA and Byteball, transactions in DGZTare stored and or-ganized in a DAG structure. However, we impose some additional rules, which results in a special DAG called DGZTdirected acyclic graph (DGZT-DAG). Consensus in our DGZT-DAG is achieved through committees, which are non-anonymous reputable entities. It is a Byzantine Fault Tolerant(BFT) consensus protocol which can tolerate malicious behaviors. Since the FLP impossibility result [8] has demonstrated the impossibility of distributed consensus in an asynchronous environment, we assume one of the two forms of partial synchrony defined in [9]. That is, the upper bound on the time required for a message to be delivered is fixed but not known a priori. The main advantage of our consensus algorithm, compared with the state-of-the-art BFT protocols such as PBFT [10] and Tendermint [11], is the exclusion of additional messages for voting purpose. It significantly reduces the communication overhead, which in turn alleviates the scaling

Byteball project has been renamed as Obyte.




Figure 1: Example of consensus in DGZT-DAG structure

issues to achieve higher TPS.

The remainder of the paper is organized as follows. The DGZT-DAG structure is presented in Section 2.

The proposed consensus algorithm is described in Section 3.

Section 4 rigorously proves the correctness of our consensus protocol, including both safety and liveness properties.

2 DGZT -DAG

In DGZT, each block represents one transaction, which contains references to previous blocks (called parents) through their hashes. Blocks and their parent-child links are the vertices and edges of the DAG, respectively. As depicted in Fig. 1, our DGZT-DAG structure has two layers, namely the con-sensus layer and the transaction layer.

All blocks in the consensus layer are composed by some non-anonymous reputable people or companies, called committees, who might have a long established reputation, or great benefits in keeping the network healthy. Each block in the consensus layer can reference multiple blocks from both the consensus layer and the transaction layer. committees are expected to post transactions frequently and behave honestly. However, it is unreasonable to totally trust any single committee. Our proposed scheme allows committees to be replaced without jeopardizing the consensus and security in the network. Details on how to change committees will be elaborated

in Section 3. Transactions in the consensus layer is for the sole purpose of achieving consensus in the network, while real transactions happen in the transaction layer. In the transaction layer, each account has its own chain of blocks, which records the transaction history of this account. In addition, each block in the transaction layer is referenced by blocks in the consensus layer.

The consensus in the computecoin network is achieved via total ordering of all blocks. Each node starts by finding out the “stable” main chain within the consensus layer of its local DAG. The rigorous definition of stable main chain will be described later in Section 3.1. Each node then numbers all blocks included by blocks on the stable main chain as follows. It first defines indices for blocks that lie directly on the stable main chain. The genesis block has index 0, the next block on the stable main chain that is a child of the genesis block has index 1, and so on. By traveling forward along the stable main chain, it assigns indices to blocks that lie on the stable main chain. For any block that does not lie on the stable main chain, its index is assigned by the index of the block on the stable main chain that first references it directly or indirectly. Now each node can determine the order for any two blocks B_1 and B_2 with assigned indices using the following rule \mathcal{O} : B_1 precedes B_2 if and only if

- a) B_1 has lower index than B_2 ; or
- b) B_1 and B_2 have the same indices, but B_1 is referenced by B_2 directly or indirectly; or
- c) B_1 and B_2 have the same indices, and there is no reference relationship between B_1 and B_2 , but B_1 has lower hash than B_2 .

As a concrete example shown in Fig. 1, a node is trying to decide the order of two blocks B_1 and B_2 marked in blue. The stable main chain it finds out is marked in bold arrows. And the numbers inside each block are indices assigned according to the stable main chain. Now block B_1 has index 3 and block B_2 has index 4. Therefore, the node will determine that B_1 precedes B_2 since B_1 has lower index than B_2 .

3 Consensus in DGZT

In this section, we will focus on the consensus layer of our DGZT-DAG structure, and explain in detail how a node finds out the stable main chain of its local graph. The remainder of this section is organized as follows. The

key terms which will be used intensively throughout the paper are described in Section 3.1. In Section 3.2, we list the key assumptions we rely on in order to guarantee that the computecoin network is secure. Based on the definitions and assumptions, Section 3.3 presents the consensus algorithm which is implemented in the computecoin mainnet.

3.1 Definitions

At any time, each node in the network would observe slightly different graph due to network delay. Let $G_n(t)$ denote the graph node n has observed at time t . In this section, we drop n and t and use G to represent a general DAG which satisfies that if a block B is in G , all B 's parents are also in G . In the following, we describe some key terms which will be used intensively in the subsequent sections.

- D1 Graph inclusion relation: We use $G \subseteq G^*$ to represent that G^* contains all blocks in G , and G^* satisfies the condition that if a block B is in G^* , all B 's parents are also in G^* .
- D2 Block inclusion relation: We say a block B_1 includes another block B_0 if $B_1 = B_0$ or B_1 references B_0 directly or indirectly.
- D3 Block comparison: Suppose each block in G has its epoch, level and hash, where the definitions of epoch and level will be discussed in D6 and D7, respectively. For any pair of blocks B_0 and B_1 , we call B_1 is better than B_0 if and only if B_1 has larger epoch, or larger level if B_0 and B_1 have the same epoch, or larger hash in the case that B_0 and B_1 have the same epoch and the same level. We denote this comparison rule as \mathcal{R} .
- D4 Best Parent: The best parent of a block is one of its parents, which is the best under block comparison rule \mathcal{R} . The best parent of a block B is denoted by $\text{bp}(B)$.
- D5 Block height: The height of a block B , denoted by $\text{h}(B)$, refers to the length of the path from B to the genesis block through best parent links. Note that the height of the genesis block is 0.
- D6 Epoch: The system moves through a succession of configurations called epochs. In each epoch, there is a different set of committees, denoted by \mathcal{W}_i . Let N_i denote the number of committees in \mathcal{W}_i and $K_i = \lfloor \frac{2}{3}N_i \rfloor + 1$. We represent the set of all nonnegative integers as a union

of disjoint consecutive integer sequences, i.e., $\mathbb{N} \cup \{0\} = \bigcup_{i=1}^{\infty} \mathcal{I}_i$, where \mathcal{I}_i is a consecutive integer sequence ranging from a_i to b_i . Here, all the numbers in \mathcal{I}_j is larger than those in \mathcal{I}_i for any $j > i$, i.e., $a_j > b_i$. The epoch a block B belongs to is determined by which interval the height of the last stable block (defined later in D10) of B 's best parent falls in. Specifically, if the height of the last stable block of $\text{bp}(B)$ is in \mathcal{W}_i , the epoch of block B , denoted by $\text{ep}(B)$, is i .

D7 Block level: The level of a block B , denoted by $\text{lv}(B)$, is defined as follows:

$$\text{lv}(B) = \begin{cases} 0, & \text{if } B \text{ is the genesis block,} \\ 1, & \text{if } \text{ep}(B) > \text{ep}(\text{bp}(B)), \\ \text{lv}(\text{bp}(B)) + 1, & \text{if } \text{ep}(B) = \text{ep}(\text{bp}(B)). \end{cases} \quad (1)$$

D8 Main chain: The main chain of graph G is defined as the path starting from the best tip block in G under block comparison rule \mathcal{R} to the genesis block through best parent links. Here, tip blocks refer to blocks without any child.

D9 Stable block: A block on the main chain of G is called a stable block of G if it is guaranteed to be contained in the main chain of any graph G^* that includes G , i.e., $G \subseteq G^*$.

D10 Last stable block: The last stable block of the genesis block is itself. Now for a block B_1 , given that the last stable block of its best parent is defined, the last stable block of B_1 is determined by the following procedure. For any two blocks B and B^* , we use $B^* \rightarrow B$ to denote that B^* includes B through parent links and all blocks in the path (including both B^* and B) must be in the same epoch. Similarly, we use $B^* \xrightarrow{b} B$ to denote that B^* includes B through best parent links and all blocks in the path need not be in the same epoch. The degenerated case of $B = B^*$ is regarded true, i.e., $B^* \rightarrow B$ and $B^* \xrightarrow{b} B$. For any block B_0 such that $B_1 \xrightarrow{b} B_0$, let $C(B_0, B_1)$ denote the set of blocks from B_1 to B_0 through best parent links, which includes B_1 but not B_0 . Assume $\text{ep}(B_1) = i$. Start with $B_0 = \text{lsb}(\text{bp}(B_1))$, and check whether the following condition holds

$$\text{lv}(B_1) > \max_{B \in S(B_0, B_1)} \text{lv}(B) + 2(K_i - 1), \quad (2)$$




Figure 2: One step in finding out the last stable block of B_1 . Solid and dashed lines represent parent-child links and ancestor-descendant links, respectively. Bold and regular lines represent \xrightarrow{b} and \rightarrow relations, respectively.

where $S(B_0, B_1) = \{B \mid B \xrightarrow{b} B_0, B_1 \rightarrow B, C(B_0, B) \cap C(B_0, B_1) = \emptyset\}$. If $S(B_0, B_1) = \emptyset$, the maximal value over $S(B_0, B_1)$ in (2) is set to be 0. If the condition (2) holds, update B_0 to be its child on $C(B_0, B_1)$ and go back to check the condition (2) again, and so on. We repeatedly advance B_0 till $h(B_0) \in \mathcal{I}_{i+1}$ or B_1 does not satisfy the condition (2) with respect to B_0 . The block B_0 we stop at is the last stable block of B_1 , denoted by $\text{lsb}(B_1)$. One advancement of B_0 described above is depicted as the blue arrow in Fig. 2.

- D11 **Stable main chain:** From the last stable blocks of all blocks in G , we pick the one with the largest height, denoted by $\text{SB}(G)$. The stable main chain of G , denoted by $\text{SC}(G)$, is then defined as the chain of blocks starting from $\text{SB}(G)$ to the genesis block through best parent links. Note that the stable main chain of G is part of the main chain that will not change as G expands.
- D12 **Main chain index (MCI):** The MCI for any block that lies directly on the stable main chain is equal to its height. For any block that does not lie on the main chain, its MCI is assigned by the MCI of the block on the stable main chain that first includes it. The MCI of a block B is denoted by $\text{mci}(B)$.

Many definitions above depend on each other. However, they can be incrementally built up as the DAG grows. To start with, the genesis block belongs to epoch 0, has level 0 and its last stable block is itself. For a new

block B added to the graph, assume that all terms for its parents are already well defined. We first find out its best parent $\text{bp}(B)$ via block comparison rule \mathcal{R} . Next, we find out its epoch $\text{ep}(B)$ by checking the height of the last stable block of $\text{bp}(B)$. B 's level $\text{lv}(B)$ can then be determined by (1). And the last step is to find out the last stable block of B , i.e., $\text{lsb}(B)$ by the procedure described in D10. After that, we will know whether the stable main chain of the graph has been extended or not.

3.2 Assumptions

The key assumptions used in DGZT consensus protocol and subsequent technical discussions are as follows:

- A1 Honest committees should generate blocks serially. In other words, each honest committee should reference (directly or indirectly) all its previous blocks in every subsequent block.
- A2 When an honest committee composes a block, he always chooses the best tip block of its local graph under block comparison rule \mathcal{R} as the best parent of this new block.
- A3 If a block is in epoch i , the issuer of this block must be in the committee set \mathcal{W}_i .
- A4 Start from any block in epoch i and traverse through best parent links, we stop as soon as we encounter K_i blocks or a block of level 1, whichever comes first. Each block we encountered (including the one we stop at) must be issued by a different committee from the committee set \mathcal{W}_i .
- A5 In each epoch i , more than $2/3$ of the committees in \mathcal{W}_i are honest. In other words, at least K_i committees are honest, where $K_i = \lfloor \frac{2}{3}N_i \rfloor + 1$ is defined in D6.
- A6 Any block will be delivered to all honest committees within some fixed but unknown amount of time. It implies that for honest committees, the graphs they eventually observe would be consistent with each other. That is to say, for any pair of honest committees i and j , the graph $\mathbf{G}_i(t_i)$ node i observed at time t_i will also be observed by node j at some time t_j , i.e., $\mathbf{G}_i(t_i) \subseteq \mathbf{G}_j(t_j)$.

The assumptions from A1 to A4 are also constraints that need to be satisfied when a committee issues a block. Among those, however, only A3

and A4 are binding. That is to say, other committees can perform certain sanity check on A3 and A4, and reject the block if either of these two conditions is not met. Note that assumption A6 is a form of partial asynchrony [9], which is a middle ground between synchrony and asynchrony.

3.3 Consensus Algorithm

Based on the definitions and assumptions above, the consensus algorithm implemented in DGZT is summarized in Algorithm 1. The key idea is on how to consistently expand the local graph when receiving a block. For consensus purpose, we only need to deal with blocks issued by committees and update the stable main chain accordingly, since only those blocks can contribute to the consensus of the system.

4 Correctness

This section provides the technical proofs to show that the consensus algorithm described in Algorithm 1 is correct. Section 4.1 provides some useful propositions that will be used in the subsequent sections. In Section 4.2, we show that the advance of last stable block defined in D10 guarantees that the last stable block is indeed stable. Section 4.3 and Section ?? are dedicated to prove that our consensus algorithm satisfies safety and liveness properties, respectively. Note that in this section, we still focus on the consensus layer of our DGZT-DAG structure.

4.1 Propositions

Recall that for any two blocks B and B^* , $B^* \rightarrow B$ denotes that B^* includes B through parent links and all blocks in the path (including both B^* and B) are in the same epoch. Similarly, $B^* \xrightarrow{b} B$ denotes that B^* includes B through best parent links and all blocks in the path are not necessarily in the same epoch. In the following, we prove some useful results which will be used in later analysis.

Proposition 1. *For any two blocks B_0 and B_1 , if $B_0 = \text{bp}(B_1)$, we have $\text{lsb}(B_1) \xrightarrow{b} \text{lsb}(B_0)$, and $\text{ep}(B_1) = \text{ep}(B_0)$ or $\text{ep}(B_1) = \text{ep}(B_0) + 1$.*

Proof. It can be directly inferred from how the last stable block is determined as described in D10. To find the last stable block of B_1 , we start with $B^* = \text{lsb}(B_0)$, and update B^* to be its child in $\mathcal{C}(B^*, B_1)$ in each step

Algorithm 1 DGZTConsensus Algorithm

1: *Input:* Local graph $G = \{G\}$ for some node, where G is the genesis block
2: *Initialization:* Set $\text{ep}(G) = 0, \text{lv}(G) = 0, \text{lsb}(G) = G$.
3: *Main iterations:*
4: **for all** received block B_1 **do**
5: **if** B_1 does not pass the sanity checks **then**
6: Reject block B_1 .
7: Continue
8: **end if**
9: **if** At least one of B_1 's parent is not in G **then**
10: Add block B_1 into a buffer for future consideration.
11: Continue
12: **end if**
13: **if** B_1 is not issued by a committee **then**
14: Continue
15: **end if**
16: Determine B_1 's best parent $\text{bp}(B_1)$ by block comparison rule \mathcal{R} .
17: Determine B_1 's epoch $\text{ep}(B_1)$ by checking which interval the height of $\text{lsb}(\text{bp}(B_1))$ falls in. Assume the interval is \mathcal{I}_i , i.e., $\text{ep}(B_1) = i$.
18: **if** Assumptions A3 or A4 is not satisfied **then**
19: Reject block B_1 .
20: Continue
21: **end if**
22: Add B_1 to G , and determine B_1 's level $\text{lv}(B_1)$ according to (1).
23: Set $B_0 = \text{lsb}(\text{bp}(B_1))$.
24: **while** The condition (2) holds **do**
25: Update B_0 to be its child in $C(B_0, B_1)$.
26: **end while**
27: Set $\text{lsb}(B_1) = B_0$.
28: **if** $\text{lsb}(B_1)$ has larger height than the tip block of the existing stable main chain **then**
29: Update the stable main chain $\text{SC}(G)$ to end with $\text{SB}(G) = \text{lsb}(B_1)$.
30: **end if**
31: Find out MCIs of all blocks that are included by any block on $\text{SC}(G)$.
32: **end for**
33: *Output:* Linear ordering of all blocks that are included by any block on $\text{SC}(G)$ using rule \mathcal{O} .

as long as B_1 satisfies the condition (2) with respect to B^* . It guarantees that in every step, the new B^* references the old one through the best parent link. Therefore, we have $\text{lsb}(B_1) \xrightarrow{b} \text{lsb}(B_0)$. Assume $\text{ep}(B_0) = i$, i.e., $\text{h}(\text{lsb}(\text{bp}(B_0))) \in \mathcal{I}_i$. To find the last stable block of B_0 , the block we stop at, i.e., $\text{lsb}(B_0)$ must satisfy that $\text{h}(\text{lsb}(B_0))$ is still in \mathcal{I}_i or in \mathcal{I}_{i+1} . It follows that $\text{ep}(B_1) = i$ or $i + 1$, which leads to $\text{ep}(B_1) = \text{ep}(B_0)$ or $\text{ep}(B_1) = \text{ep}(B_0) + 1$. \square

Proposition 2. *For any two blocks B_0 and B_1 , if B_1 includes B_0 , we have $\text{ep}(B_1) \geq \text{ep}(B_0)$.*

Proof. The statement is true for the trivial case $B_0 = B_1$. Now we assume that $B_0 \neq B_1$. First, we show that if B_0 is a parent of B_1 , $\text{ep}(B_1) \geq \text{ep}(B_0)$ holds. Consider the following two cases.

- 1) B_0 is the best parent of B_1 : We have $\text{lsb}(B_0) \xrightarrow{b} \text{lsb}(\text{bp}(B_0))$ by Proposition 1. It follows that $\text{h}(\text{lsb}(B_0)) \geq \text{h}(\text{lsb}(\text{bp}(B_0)))$. Thus, there exists $i \geq j$ such that $\text{h}(\text{lsb}(B_0)) \in \mathcal{I}_i$ and $\text{h}(\text{lsb}(\text{bp}(B_0))) \in \mathcal{I}_j$. Therefore, $\text{ep}(B_1) = i \geq j = \text{ep}(B_0)$.
- 2) $B_2 \neq B_0$ is the best parent of B_1 : Similarly as in the previous case, we have $\text{ep}(B_1) \geq \text{ep}(B_2)$. According to the definition of best parent, B_2 is better than B_0 under block comparison rule \mathcal{R} . It implies that $\text{ep}(B_2) \geq \text{ep}(B_0)$. Therefore, we have $\text{ep}(B_1) \geq \text{ep}(B_2) \geq \text{ep}(B_0)$.

For the general case that B_1 does not directly reference B_0 , we can apply the chain rule to show that $\text{ep}(B_1) \geq \text{ep}(B_0)$. \square

Proposition 3. *For any two blocks B_0 and B_1 , if $B_1 \rightarrow B_0$, we have $\text{lv}(B_1) \geq \text{lv}(B_0)$.*

Proof. The statement is true for the trivial case $B_0 = B_1$. Now we assume that $B_0 \neq B_1$. First, we show that if B_0 is a parent of B_1 , $\text{lv}(B_1) \geq \text{lv}(B_0)$ holds. Consider the following two cases.

- 1) B_0 is the best parent of B_1 : Since B_0 and B_1 are in the same epoch by the definition of $B_1 \rightarrow B_0$, we have $\text{lv}(B_1) = \text{lv}(B_0) + 1 > \text{lv}(B_0)$ by (1).
- 2) $B_2 \neq B_0$ is the best parent of B_1 : According to the definition of best parent, B_2 is better than B_0 under block comparison rule \mathcal{R} . It implies that $\text{ep}(B_2) \geq \text{ep}(B_0)$. It follows that

$$\text{ep}(B_2) \geq \text{ep}(B_0) \stackrel{(a)}{=} \text{ep}(B_1) \stackrel{(b)}{\geq} \text{ep}(B_2), \quad (3)$$

where (a) is by the definition of $B_1 \rightarrow B_0$ and (b) is by Proposition 2. Thus, the following condition holds: $\text{ep}(B_0) = \text{ep}(B_1) = \text{ep}(B_2)$. Therefore, we have

$$\text{lv}(B_1) \stackrel{(a)}{=} \text{lv}(B_2) + 1 \stackrel{(b)}{\geq} \text{lv}(B_0) + 1 > \text{lv}(B_0), \quad (4)$$

where (a) is by (1) and (b) is due to the fact that $\text{lv}(B_2) \geq \text{lv}(B_0)$ since B_2 is better than B_0 under \mathcal{R} but $\text{ep}(B_0) = \text{ep}(B_2)$.

For the general case that B_1 does not directly reference B_0 , we can apply the chain rule to show that $\text{lv}(B_1) \geq \text{lv}(B_0)$. \square

The following is a direct corollary of Proposition 2 and Proposition 3.

Corollary 1. *For any two blocks B_0 and B_1 , if B_1 includes B_0 and $\text{ep}(B_1) = \text{ep}(B_0)$, we have $B_1 \rightarrow B_0$ and $\text{lv}(B_1) \geq \text{lv}(B_0)$.*

4.2 Advance of Last Stable Block

Let \mathbb{G}^B denote the induced graph from a block B in \mathbb{G} which consists of all blocks that B includes. In this section, we will analyze the procedure to determine the last stable block of B , i.e., $\text{lsb}(B)$. Our main goal is to show that $\text{lsb}(B)$ is a stable block of graph \mathbb{G}^B . Recall that from Assumption A4, if we start from block B in epoch i , traverse through best parents links, and stop as soon as K_i blocks or a block of level 1 has been visited, all blocks encountered must be issued by different committees from the committee set \mathcal{W}_i . Let $\mathbb{T}(B)$ and $\mathbb{W}(B)$ denote the set of blocks encountered and the set of committees who issue these blocks, respectively. Note that all blocks in set $\mathbb{T}(B)$ are in the same epoch as B . In the following, we first prove three lemmas which are crucial for the proof of our claim.

Lemma 1. *If $B_1 \xrightarrow{b} B_0$, all blocks in $\mathbb{C}(B_0, B_1)$ are in epoch i and none of them is issued by an honest committee from a set $\mathcal{W} \subseteq \mathcal{W}_i$ which consists of K_i committees, then $\mathbb{C}(B_0, B_1)$ contains at most $K_i - 1$ blocks, i.e., $|\mathbb{C}(B_0, B_1)| \leq K_i - 1$.*

Proof. Since all blocks in $\mathbb{C}(B_0, B_1)$ are issued by committees from set \mathcal{W}_i and none of them is issued by an honest committee from \mathcal{W} , they can only be issued by $N_i - K_i$ committees outside \mathcal{W} and malicious committees inside \mathcal{W} , which is at most $N_i - K_i$ by Assumption A5. Thus, due to $K_i > \frac{2}{3}N_i$




Figure 3: The case $\text{ep}(B_0) = i$. Solid and dashed lines represent parent-child links and ancestor-descendant links, respectively. Bold and regular lines represent \xrightarrow{b} and \rightarrow relations, respectively.

in assumption A5, the number of distinct committees which have issued at least one block in $\mathcal{C}(B_0, B_1)$ is at most

$$2(N_i - K_i) < \frac{2}{3}N_i < K_i. \quad (5)$$

It then follows from Assumption A4 that $|\mathcal{C}(B_0, B_1)| < K_i$, which is equivalent to $|\mathcal{C}(B_0, B_1)| \leq K_i - 1$. It completes the proof of Lemma 1. \square

Lemma 2. *If $B_1 \xrightarrow{b} B_0$, $\text{ep}(B_1) = i$ and B_1 satisfies the condition (2) with respect to B_0 , for any block B_2 such that $\text{ep}(B_2) = i$, $B_2 \xrightarrow{b} B_0$ and $\mathcal{C}(B_0, B_2) \cap \mathcal{C}(B_0, B_1) = \emptyset$, we have $\text{lv}(B_2) < \text{lv}(B_1)$.*

Proof. Since $\text{ep}(B_0) \leq \text{eq}(B_1) = i$ by Proposition 2, in the following we consider two cases, namely $\text{ep}(B_0) = i$ or $\text{ep}(B_0) < i$.

First, consider the case $\text{ep}(B_0) = i$. It means that $\mathcal{S}(B_0, B_1) \neq \emptyset$ since $B_0 \in \mathcal{S}(B_0, B_1)$. We start from B_2 , traverse through best parent links till B_0 , and stop as soon as a block in $\mathcal{S}(B_0, B_1)$ is encountered. Let B_3 denote the block we stop at, i.e.,

$$B_3 = \arg \max_{B \in (\mathcal{C}(B_0, B_2) \cup \{B_0\}) \cap \mathcal{S}(B_0, B_1)} \text{lv}(B). \quad (6)$$

We show that no block in $\mathcal{C}(B_3, B_2)$ is issued by any honest committee from set $\mathcal{W}(B_1)$. It is proved by contradiction. Assume there are blocks in $\mathcal{C}(B_3, B_2)$ issued by honest committees from $\mathcal{W}(B_1)$. Among those, let B_4 denote the one with the smallest height. As shown in Fig. 3, let B_5 denote the block in set $\mathcal{T}(B_1)$ which comes from the same committee as B_4 . Since B_4 and B_5 come from the same honest committee, by Assumption A1, either B_4 includes B_5 or B_5 includes B_4 . Since B_2 includes B_3 and $\text{ep}(B_2) = \text{ep}(B_3) = i$, we have $\text{ep}(B_4) = i$ by Corollary 1. Similarly, we have $\text{ep}(B_5) = \text{ep}(B_1) = i$. Therefore, by Corollary 1, either $B_4 \rightarrow B_5$ or $B_5 \rightarrow B_4$ holds. However, by the definition of B_3 in (6), which is the first block included by B_1 when traversing from B_2 through best parent links, it is impossible that $B_5 \rightarrow B_4$. Thus, we have $B_4 \rightarrow B_5$. Let B_6 and B_7 be parents of B_4 such that $B_4 \xrightarrow{b} B_6$ and $B_7 \rightarrow B_5$, respectively. Since $\text{ep}(B_2) = \text{ep}(B_3) = i$, all blocks in $\mathcal{C}(B_3, B_6)$ are in epoch i by Corollary 1. By the definition of B_4 , no block in $\mathcal{C}(B_3, B_6)$ is issued by any honest committee from $\mathcal{W}(B_1)$. In addition, the cardinality of $\mathcal{W}(B_1)$ is K_i since B_1 satisfies the condition (2), which implies that $\text{lv}(B_1) > K_i$. Therefore, by Lemma 1, we have $|\mathcal{C}(B_3, B_6)| \leq K_i - 1$, which leads to

$$\text{lv}(B_6) \leq \text{lv}(B_3) + (K_i - 1). \quad (7)$$

Now the following chain of inequalities hold

$$\text{lv}(B_7) \stackrel{(a)}{\geq} \text{lv}(B_5) \stackrel{(b)}{\geq} \text{lv}(B_1) - (K_i - 1) \stackrel{(c)}{>} \text{lv}(B_3) + (K_i - 1) \stackrel{(d)}{\geq} \text{lv}(B_6), \quad (8)$$

where (a) is by Proposition 3, (b) is due to $B_5 \in \mathcal{T}(B_1)$, (c) is by the fact that $B_3 \in \mathcal{S}(B_0, B_1)$ and B_1 satisfies the condition (2) with respect to B_0 , and (d) is by (7). It contradicts with the fact that $\text{lv}(B_6) \geq \text{lv}(B_7)$ since B_6 is the best parent of B_4 and $\text{ep}(B_6) = \text{ep}(B_7) = i$. It completes the proof that no block in $\mathcal{C}(B_3, B_2)$ is issued by any honest committee from $\mathcal{W}(B_1)$. In addition, $B_2 \xrightarrow{b} B_3$ and all blocks in $\mathcal{C}(B_3, B_2)$ are in epoch i , by Lemma 1 we have $|\mathcal{C}(B_3, B_2)| \leq K_i - 1$, which leads to

$$\text{lv}(B_2) \leq \text{lv}(B_3) + (K_i - 1). \quad (9)$$

It follows that

$$\text{lv}(B_1) \stackrel{(a)}{>} \text{lv}(B_3) + 2(K_i - 1) \stackrel{(b)}{\geq} \text{lv}(B_2) + (K_i - 1) \geq \text{lv}(B_2), \quad (10)$$

where (a) is by the fact that $B_3 \in \mathcal{S}(B_0, B_1)$ and B_1 satisfies the condition (2) with respect to B_0 , and (b) is by (9). It completes the proof that $\text{lv}(B_2) < \text{lv}(B_1)$ if $\text{ep}(B_0) = i$.




Figure 4: The case $\text{ep}(B_0) < i$. Solid and dashed lines represent parent-child links and ancestor-descendant links, respectively. Bold and regular lines represent \xrightarrow{b} and \rightarrow relations, respectively.

Next, we consider the case $\text{ep}(B_0) < i$. If $\mathcal{S}(B_0, B_1) \neq \emptyset$, we can follow the same arguments as in the previous proof to show that $\text{lv}(B_2) < \text{lv}(B_1)$. Now we assume $\mathcal{S}(B_0, B_1) = \emptyset$. Since $\text{ep}(B_2) = i > \text{ep}(B_0)$, by Proposition 1, there exists a block $B_3 \in \mathcal{C}(B_0, B_2)$ such that $\text{ep}(B_3) = i$ and $\text{ep}(\text{bp}(B_3)) = i - 1$, i.e., $\text{lv}(B_3) = 1$. Similarly as in the previous case, we show that no block in $\mathcal{C}(\text{bp}(B_3), B_2)$ is issued by any honest committee from set $\mathcal{W}(B_1)$. It is also proved by contradiction. Assume there are blocks in $\mathcal{C}(\text{bp}(B_3), B_2)$ issued by honest committees from $\mathcal{W}(B_1)$. Among those, let B_4 denote the one with the smallest height. As shown in Fig. 4, let B_5 denote the block in set $\mathcal{T}(B_1)$ which comes from the same committee as B_4 . Since B_4 and B_5 come from the same honest committee, by Assumption A1, either B_4 includes B_5 or B_5 includes B_4 . Since B_2 includes B_3 and $\text{ep}(B_2) = \text{ep}(B_3)$, we have $\text{ep}(B_4) = i$ by Corollary 1. Also we have $\text{ep}(B_5) = \text{ep}(B_1) = i$. Therefore, by Corollary 1, either $B_4 \rightarrow B_5$ or $B_5 \rightarrow B_4$ holds. However, it is impossible that $B_5 \rightarrow B_4$ since it is assumed that $\mathcal{S}(B_0, B_1) = \emptyset$. Thus, we have $B_4 \rightarrow B_5$. Let B_6 and B_7 be parents of B_4 such that $B_4 \xrightarrow{b} B_6$ and $B_7 \rightarrow B_5$, respectively. If $B_4 = B_3$, we have

$$\text{ep}(B_6) = \text{ep}(\text{bp}(B_3)) = i - 1 < i = \text{ep}(B_5) \leq \text{ep}(B_7), \quad (11)$$

where the last inequality is due to Proposition 2. It contradicts with the

fact that B_6 is the best parent of B_4 . If $B_4 \neq B_3$, by the definition of B_4 , no block in $\mathcal{C}(\text{bp}(B_3), B_6)$ is issued by any honest committee from $\mathcal{W}(B_1)$. And all blocks in $\mathcal{C}(\text{bp}(B_3), B_6)$ are in epoch i . Therefore, by Lemma 1, we have $|\mathcal{C}(\text{bp}(B_3), B_6)| \leq K_i - 1$, which leads to

$$\text{lv}(B_6) \leq K_i - 1, \quad (12)$$

since $\text{lv}(B_3) = 1$. In the following, we derive a similar chain of inequalities as (8):

$$\text{lv}(B_7) \stackrel{(a)}{\geq} \text{lv}(B_5) \stackrel{(b)}{\geq} \text{lv}(B_1) - (K_i - 1) \stackrel{(c)}{>} K_i - 1 \stackrel{(d)}{\geq} \text{lv}(B_6), \quad (13)$$

where (a) is by Proposition 3, (b) is due to $B_5 \in \mathcal{T}(B_1)$, (c) is by the fact that B_1 satisfies the condition (2) which implies $\text{lv}(B_1) > 2(K_i - 1)$, and (d) is from (12). It contradicts with the fact that $\text{lv}(B_6) \geq \text{lv}(B_7)$ since B_6 is the best parent of B_4 and $\text{ep}(B_6) = \text{ep}(B_7) = i$. It completes the proof that no block in $\mathcal{C}(\text{bp}(B_3), B_2)$ is issued by any honest committee from $\mathcal{W}(B_1)$.

In addition, $B_2 \xrightarrow{b} \text{bp}(B_3)$ and all blocks in $\mathcal{C}(\text{bp}(B_3), B_2)$ are in epoch i , by Lemma 1 we have $|\mathcal{C}(\text{bp}(B_3), B_2)| \leq K_i - 1$, which leads to

$$\text{lv}(B_2) \leq K_i - 1, \quad (14)$$

since $\text{lv}(B_3) = 1$. It follows that

$$\text{lv}(B_1) \stackrel{(a)}{>} 2(K_i - 1) \stackrel{(b)}{\geq} \text{lv}(B_2) + (K_i - 1) \geq \text{lv}(B_2), \quad (15)$$

where (a) is by the fact that B_1 satisfies the condition (2) which implies $\text{lv}(B_1) > 2(K_i - 1)$, and (b) is by (14). It completes the proof that $\text{lv}(B_2) < \text{lv}(B_1)$ if $\text{ep}(B_0) < i$.

By combining the two cases above, we finish the proof of Lemma 2. \square

Lemma 3. *Given $i \in \mathbb{N}$, assume $\text{lsb}(B)$ is a stable block of graph G^B for any block B with $\text{ep}(B) < i$. If $B_1 \xrightarrow{b} B_0$, $\text{ep}(B_1) = i$, $\text{h}(B_0) \in \mathcal{I}_i$ and B_1 satisfies the condition (2) with respect to B_0 , for any block B_2 such that $B_2 \xrightarrow{b} B_0$ and $\mathcal{C}(B_0, B_2) \cap \mathcal{C}(B_0, B_1) = \emptyset$, we have $\text{ep}(B_2) \leq \text{ep}(B_1)$.*

Proof. According to the procedure of determining the last stable block in D10, we have $B_2 \xrightarrow{b} \text{lsb}(B_2)$. Since $B_2 \xrightarrow{b} B_0$, either $B_0 \xrightarrow{b} \text{lsb}(B_2)$ or $\text{lsb}(B_2) \xrightarrow{b} B_0$ holds. We show that $B_0 \xrightarrow{b} \text{lsb}(B_2)$. It is proved by contradiction. Suppose $\text{lsb}(B_2) \xrightarrow{b} B_0$ and $\text{lsb}(B_2) \neq B_0$, which means that the

last stable block of B_2 has advanced past B_0 . Thus, there exists some block $B_3 \in \mathcal{C}(B_0, B_2)$ such that B_3 satisfies the condition (2) with respect to B_0 , i.e.,

$$\text{lv}(B_3) > \max_{B \in \mathcal{S}(B_0, B_3)} \text{lv}(B) + 2(K_j - 1), \quad (16)$$

where $j = \text{ep}(B_3) \leq \text{ep}(B_2) = i$ by Proposition 2. And the last stable block of B_3 has advanced past B_0 , i.e., $\text{lsb}(B_3) \in \mathcal{C}(B_0, B_2)$. Consider the following two cases.

- 1) $j < i$: Let $\mathbf{G}^* = \mathbf{G}^{B_3} \cup \mathbf{G}^{B_1}$. Since $\text{ep}(B_3) < \text{ep}(B_1)$, B_1 is the tip block of the main chain of \mathbf{G}^* . By the assumption in the statement of Lemma 3, $\text{lsb}(B_3)$ is a stable block of graph \mathbf{G}^{B_3} . Due to $\mathbf{G}^{B_3} \subseteq \mathbf{G}^*$, $\text{lsb}(B_3)$ is on the main chain of \mathbf{G}^* , i.e., $B_1 \xrightarrow{b} \text{lsb}(B_3)$. It contradicts with the fact that $\text{lsb}(B_3) \in \mathcal{C}(B_0, B_2)$ and $\mathcal{C}(B_0, B_2) \cap \mathcal{C}(B_0, B_1) = \emptyset$.
- 2) $j = i$: Since both B_1 and B_3 satisfy the condition (2) with respect to B_0 , it follows by Lemma 2 that both $\text{lv}(B_3) < \text{lv}(B_1)$ and $\text{lv}(B_1) < \text{lv}(B_3)$ hold, which is a contradiction.

Now we have shown that $B_0 \xrightarrow{b} \text{lsb}(B_2)$. In addition, we have $\text{lsb}(B_2) \xrightarrow{b} \text{lsb}(\text{bp}(B_2))$ by Proposition 1. Thus, $B_0 \xrightarrow{b} \text{lsb}(\text{bp}(B_2))$ holds. It follows that $\text{h}(\text{lsb}(\text{bp}(B_2))) \leq \text{h}(B_0)$. Since $\text{h}(B_0) \in \mathcal{I}_i$, there exists $k \leq i$ such that $\text{h}(\text{lsb}(\text{bp}(B_2))) \in \mathcal{I}_k$, which leads to $\text{ep}(B_2) = k \leq i = \text{ep}(B_1)$. It completes the proof of Lemma 3. \square

Now we can prove the following main result of this section.

Theorem 1. *For any block B_1 in graph \mathbf{G} , the last stable block of B_1 , i.e., $\text{lsb}(B_1)$ is a stable block of graph \mathbf{G}^{B_1} .*

Proof. We prove by induction. It is trivial for the case that B_1 is the genesis block. For the case $\text{ep}(B_1) = i$, we assume that for any block B such that $\text{ep}(B) < i$ or $B = \text{bp}(B_1)$, $\text{lsb}(B)$ is a stable block of \mathbf{G}^B . We will prove that $\text{lsb}(B_1)$ is a stable block of graph \mathbf{G}^{B_1} .

We first show that for any block B_0 such that B_0 is a stable block of \mathbf{G}^{B_1} , $\text{h}(B_0) \in \mathcal{I}_i$, and B_1 satisfies the condition (2) with respect to B_0 , then B_0 's child in $\mathcal{C}(B_0, B_1)$, denoted by B_0^* , is also a stable block of \mathbf{G}^{B_1} . It is equivalent to show that B_0^* is on the main chain of any graph \mathbf{G}^* such that $\mathbf{G}^{B_1} \subseteq \mathbf{G}^*$. We prove by contradiction. Assume there exists a graph \mathbf{G}^* such that $\mathbf{G}^{B_1} \subseteq \mathbf{G}^*$ and the main chain of \mathbf{G}^* does not contain B_0^* . As depicted in Fig. 5, let B_2 denote the tip block of the main chain of \mathbf{G}^* . Since B_0




Figure 5: The case where B_0^* is not a stable block of \mathbf{G}^{B_1} . Solid and dashed lines represent parent-child links and ancestor-descendant links, respectively.

is a stable block of \mathbf{G}^{B_1} and $\mathbf{G}^{B_1} \subseteq \mathbf{G}^*$, the main chain of \mathbf{G}^* must contain B_0 , i.e., $B_2 \xrightarrow{b} B_0$. Now we have $\mathcal{C}(B_0, B_2) \cap \mathcal{C}(B_0, B_1) = \emptyset$. It follows that $\text{ep}(B_2) \leq \text{ep}(B_1)$ by Lemma 3. Furthermore, if $\text{ep}(B_2) = \text{ep}(B_1) = i$, we have $\text{lv}(B_2) < \text{lv}(B_1)$ by Lemma 2. Therefore, either $\text{ep}(B_2) < \text{ep}(B_1)$ or $\text{lv}(B_2) < \text{lv}(B_1)$ when $\text{ep}(B_2) = \text{ep}(B_1)$ holds, which implies that B_1 is better than B_2 under block comparison rule \mathcal{R} . It contradicts with the fact that B_2 is the tip block of the main chain of \mathbf{G}^* which contains both B_1 and B_2 .

We start with $B_0 = \text{lsb}(\text{bp}(B_1))$. Since $\text{ep}(B_1) = i$, we have $\text{h}(B_0) \in \mathcal{I}_i$. In addition, B_0 is a stable block of $\mathbf{G}^{\text{bp}(B_1)}$ by our assumption. And since $\mathbf{G}^{\text{bp}(B_1)} \subseteq \mathbf{G}^{B_1}$, B_0 is also a stable block of \mathbf{G}^{B_1} . Thus, by the result we have proved above, B_0 's child in $\mathcal{C}(B_0, B_1)$, denoted by B_0^* , is a stable block of \mathbf{G}^{B_1} . We set B_0 to be B_0^* , and repeat this process until $\text{h}(B_0) \notin \mathcal{I}_i$ or B_1 does not satisfy the condition (2) with respect to B_0 . The block we stop at, i.e., the last stable block of B_1 is a stable block of \mathbf{G}^{B_1} . It completes the proof of Theorem 1. \square

4.3 Safety

Recall that the local graph node i observes at time t is denoted by $\mathbf{G}_i(t)$. To determine the order of two blocks at time t , node i will first find the stable main chain of $\mathbf{G}_i(t)$, i.e., $\text{SC}(\mathbf{G}_i(t))$, and then find out the order of these two blocks by rule \mathcal{O} in Section 2 given both of them have main chain indices (defined in D12). Therefore, in order to show the safety property of our consensus algorithm, it suffices to prove that the stable main chains

different nodes observe at different time are consistent, which is stated in the following Theorem 2.

Theorem 2. *For any $i, j \in \mathbb{N}$ and $t_i, t_j \geq 0$, we have either $\text{SC}(\mathbf{G}_i(t_i)) \subseteq \text{SC}(\mathbf{G}_j(t_j))$ or $\text{SC}(\mathbf{G}_j(t_j)) \subseteq \text{SC}(\mathbf{G}_i(t_i))$.*

Proof. Recall that $\text{SB}(\mathbf{G}_i(t))$ denotes the tip block of the stable main chain node i observes at time t . We first show that $\text{SB}(\mathbf{G}_i(t))$ is a stable block of graph $\mathbf{G}_i(t)$. In fact, by the definition of stable main chain in D11, $\text{SB}(\mathbf{G}_i(t))$ can be represented as

$$\text{SB}(\mathbf{G}_i(t)) = \arg \max_{B \in \mathbf{G}_i(t)} h(\text{lsb}(B)). \quad (17)$$

For any $B \in \mathbf{G}_i(t)$, let $\mathbf{G}_i^B(t)$ denote the induced graph which consists of all blocks included by B . By Theorem 1, $\text{lsb}(B)$ is a stable block of $\mathbf{G}_i^B(t)$. For any graph \mathbf{G}^* such that $\mathbf{G}_i(t) \subseteq \mathbf{G}^*$, we have $\mathbf{G}_i^B(t) \subseteq \mathbf{G}_i(t) \subseteq \mathbf{G}^*$. It follows that $\text{lsb}(B)$ is on the main chain of \mathbf{G}^* . Thus, $\text{lsb}(B)$ is a stable block of $\mathbf{G}_i(t)$. Therefore, according to the definition in (17), $\text{SB}(\mathbf{G}_i(t))$ is a stable block of $\mathbf{G}_i(t)$.

In order to prove that either $\text{SC}(\mathbf{G}_i(t_i)) \subseteq \text{SC}(\mathbf{G}_j(t_j))$ or $\text{SC}(\mathbf{G}_j(t_j)) \subseteq \text{SC}(\mathbf{G}_i(t_i))$ holds, it is equivalent to show that $\text{SB}(\mathbf{G}_i(t_i)) \xrightarrow{b} \text{SB}(\mathbf{G}_j(t_j))$ or $\text{SB}(\mathbf{G}_j(t_j)) \xrightarrow{b} \text{SB}(\mathbf{G}_i(t_i))$. In fact, by Assumption A6, there exists some time t_j^* such that $\mathbf{G}_i(t_i) \subseteq \mathbf{G}_j(t_j^*)$. Let $T = \max\{t_j, t_j^*\}$. We have both $\mathbf{G}_i(t_i) \subseteq \mathbf{G}_j(T)$ and $\mathbf{G}_j(t_j) \subseteq \mathbf{G}_j(T)$. Since $\text{SB}(\mathbf{G}_i(t_i))$ is a stable block of $\mathbf{G}_i(t_i)$, it follows that $\text{SB}(\mathbf{G}_i(t_i))$ is on the main chain of $\mathbf{G}_j(T)$. Similarly, $\text{SB}(\mathbf{G}_j(t_j))$ is on the main chain of $\mathbf{G}_j(T)$. Therefore, due to the uniqueness of the main chain, we have either $\text{SB}(\mathbf{G}_i(t_i)) \xrightarrow{b} \text{SB}(\mathbf{G}_j(t_j))$ or $\text{SB}(\mathbf{G}_j(t_j)) \xrightarrow{b} \text{SB}(\mathbf{G}_i(t_i))$. It completes the proof of Theorem 2. \square

References

- [1] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system, 2009. <http://www.bitcoin.org/bitcoin.pdf>.
- [2] Vitalik Buterin. Ethereum whitepaper. <https://github.com/ethereum/wiki/wiki/White-Paper>.
- [3] Yoad Lewenberg, Yonatan Sompolsky, and Aviv Zohar. Inclusive block chain protocols. In *Financial Cryptography*, 2015.

- [4] Yonatan Sompolinsky, Yoad Lewenberg, and Aviv Zohar. Spectre: A fast and scalable cryptocurrency protocol, 2016. <https://eprint.iacr.org/2016/1159>.
- [5] Yonatan Sompolinsky and Aviv Zohar. Phantom: A scalable blockdag protocol, 2018. <https://eprint.iacr.org/2018/104>.
- [6] Serguei Popov. The tangle. https://iota.org/IOTA_Whitepaper.pdf.
- [7] Anton Churyumov. Byteball: A decentralized system for storage and transfer of value. <https://byteball.org/Byteball.pdf>.
- [8] Michael J. Fischer, Nancy A. Lynch, and Michael S. Paterson. Impossibility of distributed consensus with one faulty process. *J. ACM*, 32(2):374–382, April 1985.
- [9] Cynthia Dwork, Nancy Lynch, and Larry Stockmeyer. Consensus in the presence of partial synchrony. *J. ACM*, 35(2):288–323, April 1988.
- [10] Miguel Castro and Barbara Liskov. Practical byzantine fault tolerance. In *Proceedings of the Third Symposium on Operating Systems Design and Implementation*, OSDI '99, pages 173–186, Berkeley, CA, USA, 1999.
- [11] Ethan Buchman. Tendermint: Byzantine fault tolerance in the age of blockchains. Master’s thesis, The University of Guelph, Canada, 2016.