TIRAMISU: LAYERING CONSENSUS PROTOCOLS FOR SCALABLE AND SECURE BLOCKCHAINS

Anurag Jain
Machine Learning Laboratory
International Institute of Information Technology
Hyderabad, TN, India 500032
anurag.jain@research.iiit.ac.in

Sanidhay Arora
University of Oregon
Eugene, US
sanidhay@uoregon.edu

Sankarshan Damle
Machine Learning Laboratory
International Institute of Information Technology
Hyderabad, TN, India 500032
sankarshan.damle@research.iiit.ac.in

Sujit Gujar
Machine Learning Laboratory
International Institute of Information Technology
Hyderabad, TN, India 500032
sujit.gujar@iiit.ac.in

March 25, 2022

ABSTRACT

Cryptocurrencies are poised to revolutionize the modern economy by democratizing commerce. These currencies operate on top of blockchain-based distributed ledgers. Existing permissionless blockchain-based protocols offer unparalleled benefits like decentralization, anonymity, and transparency. However, these protocols suffer in performance which hinders their widespread adoption. In particular, high time-to-finality and low transaction rates keep them from replacing centralized payment systems such as the Visa network. Permissioned blockchain protocols offer attractive performance guarantees, but they are not considered suitable for deploying decentralized cryptocurrencies due to their centralized nature. Researchers have developed several multi-layered blockchain protocols that combine both permissioned and permissionless blockchain protocols to achieve high performance along with decentralization. The key idea with existing layered blockchain protocols in literature is to divide blockchain operations into two layers and use different types of blockchain protocols to manage each layer. However, many such works come with the assumptions of honest majority which may not accurately reflect the real world where the participants may be self-interested or rational. These assumptions may render the protocols susceptible to security threats in the real world, as highlighted by the literature focused on exploring game-theoretic attacks on these protocols. We generalize the "layered" approach taken by existing protocols in the literature and present a framework to analyze the system in the BAR Model and provide a generalized game-theoretic analysis of such protocols. Using our analysis, we identify the critical system parameters required for a distributed ledger's secure operation in a more realistic setting.

Keywords Blockchain · Scalability · Game Theory

1 Introduction

Bitcoin Nakomoto promised to transform the financial system by proposing a decentralized peer-to-peer currency. *Miners* maintain this system by honestly following the underlying consensus protocol through appropriate incentives/rewards. In Bitcoin, and similar cryptocurrencies such as Ethereum, every miner validates transactions and tries to append a set of valid transactions to the set of validated transactions, i.e., append a block on the blockchain. These miners compete

against one another in a "mining" game. Typically, these miners have to produce "Proof of Work" (PoW) to write new transaction data to the blockchain.

As of December 2021, Bitcoin's and Ethereum's network processes an average of 4 and 15 transactions per second (TPS), respectively blo [2020], eth [2020]. In contrast, Visa's global payment system handles a reported 1,700 TPS and claims to be capable of handling more than 24,000 TPS vis [2021]. Besides, using such cryptocurrencies for real-time payments, like buying a cup of coffee is not feasible as the protocols only provide *eventual consistency*. It means that each transaction requires a certain number of block confirmations to be confirmed; Bitcoin needs at least 60 mins for confirming a transaction.[1]

Researchers have proposed several multi-layered protocols that to improve blockchain technology's practical performance for better applicability Abraham et al. [2018], Timo Hanke and Williams [2018], Decker et al. [2016], Pass and Shi [2017]. These protocols combine both permissioned and permissionless protocols to achieve scalability while maintaining decentralization. We take inspiration from these works and abstract out this layered approach and formalize these layers for different functions of the system. These essentials lead us to a general framework we term as "Tiramisu".

We highlight that these works do not provide a game-theoretic analysis to ensure incentive compatibility, i.e., all the above protocols assume miners are either honest or byzantine. However, the miners could also be *strategic players*. That is, they may deviate from the prescribed protocol to gain additional *rewards* because it may not be a strategic miner's *best response* to follow the protocol honestly. Selfish mining, petty mining, and undercutting are some of the strategies that may lead to greater rewards for the miners Carlsten et al. [2016], Sapirshtein et al. [2016]. In the game-theory literature, a system is said to be *incentive compatible* if it rewards each player greater for playing truthfully as compared to all other possible strategies. We remark that designing such incentive compatible blockchain protocols that are robust to strategic deviations is a significant challenge. It is also a new area of research with limited prior work Siddiqui et al. [2020].

Motivation. Overall, we observe that building a scalable, consistent, and fully decentralized practical blockchain protocol remains elusive. Cryptocurrencies like EOS eos [2020], DFINITY Timo Hanke and Williams [2018], Solida Abraham et al. [2018], etc. have a committee-based blockchain protocol. Among these, PeerCensus Decker et al. [2016] and Hybrid Consensus Pass and Shi [2017] inherently use a layered approach that demonstrates its potential benefits by achieving higher throughput and faster block confirmations. We highlight none of these works provide a game-theoretic analysis to ensure incentive compatibility and thus, security in the presence of *rational* miners – the ones maximing their rewards. Hence, the scalability guarantees provided by these protocols might not hold if the protocol is not incentive compatible.

Tiramisu: Overview. We believe that with a layered-approach, one can create protocols that offer the best of both worlds: throughput and decentralization. We dub this approach as "Tiramisu" after the layered dessert. At a high-level, we illustrate Tiramisu with Figure 1.

More concretely, Tiramisu consists of two layers, namely: Access Control Layer (ACL), and Consensus Later (CSL). In ACL, we run a permissionless consensus protocol to obtain authorized nodes. Then, with CSL, these nodes form a committee to run a BFT consensus protocol on the state of the system. As our framework is general, one can use any permissioned blockchain protocol in CSL, as long as it satisfies certain conditions (refer Section 5). As standard, our security analysis assumes that ACL is secure. This assumption is based on inherent security guarantees of the underlying consensus protocol that must be carefully employed in the protocol design. Hence, we do not discuss security attacks on it and analyze the security of Tiramisu through CSL (Section 5.2). We also identify three conditions in our game-theoretic analysis (Section 5.1) that ensure incentive-compatible implementation of any Tiramisu protocol.

1.1 Contributions.

In summary, our contributions are as follows:

- We abstract out the essentials of a layered approach to combine multiple protocols for improved performance of blockchain ecosystem and propose a framework, namely *Tiramisu* (Section 4).

- Along similar lines, we present general conditions on protocol parameters, rewards, and the fraction of honest nodes (Section 5.1) to ensure an incentive compatible instantiation of a Tiramisu framework. Our analysis considers three types of nodes: rational, honest, and Byzantine (Theorem 1).

- Under the reasonable assumption that the protocol designer uses a secure protocol in ACL, we provide a rigorous security analysis formalism for Tiramisu such that the overall system is secure (Section 5.2).

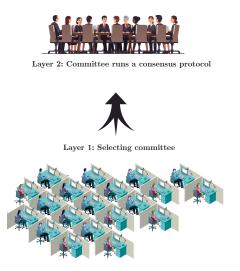[1]It is also referred to as time to finality.

Figure 1: A high-level view of Tiramisu

- Analysis with Tiramisu framework also helps in identifying the optimal parameters for a multi-layered blockchain protocol to ensure maximal performance within safe operation.

Organisation. The rest of the paper is structured as follows: First, we discuss and compare our work with the previous related work in Section 2. Then, in Section 3, we outline preliminary concepts for understanding our approach. Next, we propose the Tiramisu approach in Section 4 and present our analysis of Tiramisu in Section 5. Finally, in Section 7, we conclude our paper and discuss future work.

2 Related Work

We now discuss the existing literature involving committee-based blockchains, layered approach protocols, and game-theoretic analysis in blockchain protocols, and other scalability solutions.

Committee-based approach. Committee-based blockchains like EOS eos [2020] employ byzantine agreement protocols for a scalable blockchain protocol. EOS claims to scale up to 5000 TPS but still has open security concerns He et al. [2020].

PeerCensus Decker et al. [2016] proposes a protocol that employs a layered approach, which fits as an example in our framework. In PeerCensus, the nodes in the network gain authorized identities by solving a Proof-of-Work puzzle. The nodes controlling these identities then gain privileges to join a *committee* which runs PBFT to produce and validate blocks. We observe that once authorized, these nodes can join and leave this committee during the same session of the protocol. However, we observe that it is not required for a node to leave the committee at any time, and hence these identities are ever-increasing. This observation realizes the following explained drawback for the protocol. The network latency for the exchange of messages in PBFT increases with a message complexity of $O(n^2)$, where n denotes the number of nodes. Now since these nodes are ever-increasing, PeerCensus becomes infeasible to scale because of network latency and time complexity of PBFT limiting the real-world deployment of the protocol.

Garay et al. Garay et al. [2015a] showed that deferring the resolution of blockchain forks is inefficient and strong consistency provides greater applicability. One might observe that running a committee with only a fraction of the total nodes in the system significantly improves the speed of the protocol due to less number of participants needed for consensus. We use this observation to our advantage when designing our framework.

Layered approach. Existing consensus protocols most inspire our work in literature that uses a layered approach in some manner. Some of these consensus protocols are used in Algorand Chen and Micali [2019], OHIE Yu et al. [2020], DFINITY Timo Hanke and Williams [2018], Solida Abraham et al. [2018], PeerCensus Decker et al. [2016], and Hybrid Consensus Pass and Shi [2017]. Algorand uses two layers with smart contracts to provide secure and fast common-case transactions and off-chain contracts in layer two for more customization and programming. OHIE composes as many parallel instances of Bitcoin's original backbone protocol as needed to achieve excellent throughput. Dfinity divides its consensus algorithm into four layers tasked to provide identities and registries, random beacon, blockchain, and

notarization. Solida is a decentralized blockchain protocol based on a re-configurable Byzantine consensus augmented by proof-of-work. Hybrid Consensus is probably the closest to an instantiation of our approach. It uses a proof-of-stake protocol to manage a fixed-size committee that runs a consensus protocol on transactions.

We take inspiration from these works and aim to build our framework that could capture similar protocols described above. Specifically, we incorporate the notion of dividing tasks into different layers for better abstraction and simplification of both implementation and understanding. The concept of layers also provides independence between different tasks. We believe that this independence may overcome the drawbacks of multiple techniques used in a single layer.

Game-theoretic Considerations. Solidus Abraham et al. [2016] is the first to consider rational participants by proposing an incentive-compatible Byzantine-Fault-Tolerant protocol for blockchain. However, it does not provide a game-theoretic analysis. Biais et al. Biais et al. [2019] model Bitcoin as a coordination game considering rational participants. Yackolley et al. in Amoussou-Guenou et al. [2020] provide a game-theoretic analysis in consensus-based blockchains considering rational and Byzantine players and model a dynamic game. Manshaei et al. in Manshaei et al. [2018] consider a non-cooperative static game approach for an intra-committee protocol where they show that rational players can free-ride if rewards are equally shared.

Contrary to previous literature, we consider the cost of validating a block in the player's utilities. This consideration makes our work more realistic. We follow the BAR model Aiyer et al. [2005] and study rational behavior in a non-cooperative setting, as shown by Halpern et al. in Halpern and Vilaça [2016]. In the BAR model, three types of participants are considered in the system, i.e., honest, rational, and Byzantine.

Scalability Solutions. Our framework is different from off-chain platforms that use a blockchain like the Bitcoin Lightning network Poon and Dryja [2015], Ethereum's Plasma Joseph Poon [2017], and Truebit Jason Teutsch [2017]. Lightning networks and Plasma are decentralized networks using smart contract functionality in the blockchain. Plasma works by creating a tree-like structure of numerous smaller chains. The main idea is to establish a framework of secondary chains that will communicate and interact as sparingly as possible with the main chain. Truebit guarantees correctness in two layers yielding a unanimous consensus layer. In this layer, anyone can object to faulty solutions and an on-chain mechanism that incentivizes participation and ensures fair remuneration. These platforms enhance overall performance and enable instant payments across a network of participants.

These networks are relatively independent and do not always require synchrony with the blockchain, raising additional security and liveness concerns. However, we provide a general framework that splits a blockchain network's operation into its abstracted components and not a synchrony-independent network. Thus, we leverage the advantages from both layers in our framework.

3   Preliminaries

In this section, we outline few important preliminary concepts that will be useful for understanding this paper. First, we give a brief of two types of distributed ledgers (DLs). Second, we discuss the properties that distributed ledgers should follow. Next, we discuss BFT protocols. We then discuss permissioned and permissionless protocols. Further, we describe the performance metrics of a given blockchain. And finally, we explain our network model.

3.1   Permissioned and Permissionless Blockchain Protocols

Permissionless blockchains are the blockchains that onboard new participants without requiring the permission of existing participants. In this paper, we refer to these participants as *nodes*. Thus, no particular entity can be said to have complete control over the functioning of the system. They are also known as public blockchains. Popular blockchains such as Bitcoin, Ethereum, Litecoin, Dash, and Monero fall under this category. These blockchains allow anyone to transact and join as a validator. The data on these blockchains is publicly available, and complete copies of the ledgers are stored across the network.

Permissioned blockchains can be considered as an additional blockchain security system. They maintain an access control layer to allow pre-defined actions to be performed only by specifically identifiable participants. For this reason, these blockchains differ from public and private blockchains.

The main distinction between them is that a permissioned blockchain needs prior approval before use, whereas a permissionless blockchain lets anyone participate in the system. Ethereum Wood [2014], and Hyperledger Androulaki et al. [2018] are widely adopted permissionless and permissioned blockchain protocols, respectively.

3.2 Properties of Distributed Ledgers

We now formally list some important and necessary properties of distributed ledgers. Garay et al. Garay et al. [2015b] propose important properties of *Liveness* and *Safety* for blockchain protocols to meet to create a secure distributed ledger.

Definition 1 (Liveness Garay et al. [2015b]). Liveness *with parameters u, κ states that if all parties attempt to include a valid transaction in the ledger at time t, then the transaction would be accepted by every honest party with security parameter κ at a time $t + u$ with high probability.*

Definition 2 (Safety Garay et al. [2015b]). Safety *with parameter κ states that if an honest party has accepted a transaction with a security parameter κ at time t, then at any time after $t + \delta$ the transaction must be present in the longest chain available with every honest party.*

In contrast, consensus requires the following properties:

- *Agreement* - All non-faulty nodes decide on the same value.
- *Termination* - All non-faulty nodes terminate in finite time.
- *Validity* - If all non-faulty nodes propose the same value, then no other value can be decided.

We now describe consensus protocols that also incorporate byzantine fault tolerance.

3.3 Byzantine Agreement Protocols

A Byzantine fault is a condition of a distributed system, where components (or nodes) may fail, and there is imperfect information on whether a component has failed or not. Byzantine Fault Tolerance (BFT), derived from Byzantine Generals' Problem, is the property of a distributed network to reach consensus even with some amount of byzantine fault in the system. The objective of a BFT mechanism is to safeguard against system failures. Byzantine agreement protocols or BFT protocols strive to provide the property of Byzantine Fault Tolerance.

The most widely known and used BFT agreement protocol is the Practical Byzantine Fault Tolerance Castro and Liskov [1999], commonly known as PBFT. One might observe that these protocols can reach consensus with absolute finality fairly quickly if there are fewer participants. However, their speed reduces drastically with every increasing node in the system.

3.4 Network Model

Here, we formally describe our network model. We consider a partially synchronous peer-to-peer network consisting of participants who control identities in the network. These identities are denoted by their public-private (pk, sk) key pair and hence, are only pseudonyms since they do not leak any real-world information about the participant[2]. Participants connect by a broadcast network over which they can send messages to everyone.

3.4.1 Agent Model

We follow the BAR Model introduced by Aiyer et al. [2005] that consists of Byzantine, Altruistic (or Honest) and Rational Parties.

Definition 3 ($\mathcal{H}$ Honest Party). *A party is said to be* honest *if and only if it strictly follows the protocol.*

Definition 4 ($\mathcal{R}$ Rational Party). *A party is said to be* rational *or* self-interested *if it may strategically deviate from the protocol if the deviation is expected to yield a higher reward.*

Definition 5 ($\mathcal{B}$ Byzantine Party). *A party is said to be an* adversarial *or* byzantine *if it may or may not follow the protocol. The adversary's goal is to disrupt the operation of the protocol, and it does **not** try to maximize its reward.*

With this background, we now present Tiramisu in the next section.

4 Tiramisu: A Layered Approach

In Tiramisu, we split the operation of the blockchain into two independent layers. The first and the second layers are called Access Control Layer (ACL) and Consensus Layer (CSL) respectively. ACL manages nodes in the network via

---

[2]For brevity, we use identities in place of pseudonymous identities.

Notation	Meaning
$\alpha_{participant}$	Fraction of network resources owned by a participant
α_A	Fraction of network resources owned by adversary A
Π_{ACL}	Underlying Access Control Layer Protocol
Π_{CSL}	Underlying Consensus Layer Protocol
$\Pi_{\mathcal{TM}}$	Tiramisu Protocol: $\{\Pi_{ACL}, \Pi_{CSL}\}$
n_{CSL}	Total number of committee nodes in CSL.
n_f	Minimum nodes needed to compromise Π_{CSL}.

Table 1: Protocol specific notations

Term	Definition
$\mathcal{N}$	Set of nodes, $\{1, 2, \ldots, n_{CSL}\}$
τ_i	Node i's type, i.e.,$\tau_i \in \{\mathcal{H}, \mathcal{R}, \mathcal{B}\}$
s_i^1	Node i signs the block without validating
s_i^2	Node i signs the block only if its valid
s_i^3	Node i signs and proposes only invalid blocks
S_i	Set of pure strategies of nodes i, $\{s_i^1, s_i^2, s_i^3\}$
TR_i	Total rewards of node i for a transaction block.
u_i	Utility of node i; $u_i : \tau_i \times TR_i \times S \to \mathbb{R}$

Table 2: Game constituents

a permissionless consensus protocol. Whereas, CSL employs a permissioned consensus protocol among the authorized nodes from the first layer. These nodes will run a Byzantine agreement protocol to verify transactions and reach a consensus on the state of the system.

Let $\Pi_{\mathcal{TM}} = \{\Pi_{ACL}, \Pi_{CSL}\}$ be a Tiramisu protocol where Π_{ACL} and Π_{CSL} refer to the underlying consensus protocols running in the Access Control Layer and the Consensus Layer respectively.

4.1   Access Control Layer (ACL)

ACL is responsible for providing sybil-resistant node identities to the Consensus Layer. The identity of these *nodes* are simply a derivative of the public key of a public-secret-key pair $\langle pk, sk \rangle$, owned by a participant. Each node is identified by its public address.

In this layer, participants maintain a separate blockchain just to obtain sybil-resistant identities. Participants in this layer run any permissionless blockchain consensus protocol, by choice of protocol design, denoted by Π_{ACL}. Note that Π_{ACL} must satisfy some pre-defined conditions. Towards stating these pre-defined conditions, we first describe a *democratized resource*. A democratized resource is essential to run the protocol Π_{ACL}. It can be captured as a real-world resource and must be directly responsible for: first, the basis of sybil-resistance in Π_{ACL}; and second, for obtaining control over specific network resources. These network resources must be directly proportional to the fraction of control that they gain over the protocol Π_{ACL}. Typically Π_{ACL} can be a proof-of-X base protocol like Proof-of-Work. We assume that ACL is secure based on the inherent security of Π_{ACL}.

Formally, Π_{ACL} must satisfy the following properties:

- There must exist a protocol parameter that is directly able to express the democratized resource present in the network.
- The probability of a node adding a new block to its blockchain must be linearly dependent on the democratized resource.

For instance, Π_{ACL} could be a Proof-of-Work protocol where the democratized resource is *computational power* and the protocol parameter could be the *hash-rate* of the network. Another Π_{ACL} could be a Proof-of-Stake protocol where both the resource and this parameter are the stake invested in the system itself.

Operation. Π_{ACL} is run to determine the identities of the nodes that will be participating in the Consensus Layer. Each block in the blockchain of this layer must contain a single public address representing the identity of the node that wishes to join the CSL. Once a block reaches finality (Section 3), ACL uses a pre-defined interface to interact with CSL and propose the identity present in this block to join the committee in CSL. Observe that any participant

Term	Definition
c_{mine}	The cost of mining a block in ACL.
c_{val}	The cost to validate one byte of data.
n_{TX}	Number of transaction blocks for a given committee session.
κ_i	Cost incurred by a rational node i if an invalid transaction block is accepted.
$\kappa_{\mathcal{R}}$	κ_i for the rational node with minimum stake invested.
ϕ	Number of bytes of shared state to be validated for a single block.
TR	Minimum TR_i for the transaction block.
$P_{invalid}$	Belief probability of invalid block being accepted, $P_{invalid} : \mathcal{N}^{n_{CSL}} \times S_i \to [0,1]$.

Table 3: Relevant Game Notations

Name	State
POWCHAIN	Proof of Work Blockchain
T	List of transaction validators
COMCHAIN	Committee Blockchain
TXCHAIN	Transaction Blockchain
O	Operation Log
B	Account Balances

Table 4: Shared state constituents

Name	Operation
powBlock(b)	Adds PoW block b to POWCHAIN
txBlock(t)	Adds transaction block t to T
comBlock(c)	Adds committee block c to COMCHAIN

Table 5: Shared state operations

will only invest their network resources for the identities of the nodes that they wish to get promoted in the Consensus Layer. The key insight behind the sybil-resistant nodes is that the democratized resources are hard to obtain and may not be scaled at will. Observe that the probability of a node joining the CSL is directly proportional to the amount of democratized resources owned by the participant, denoted by $\alpha_{participant}$. This is a key observation that will be used in our security analysis (Section 5).

In Section 5, we show that for the system to function correctly, the fraction of resources controlled by the adversary must be less than $\frac{n_f}{n_{CSL}}$. Here, n_{CSL} and n_f are the total number of nodes in the Consensus Layer and the minimum number of nodes needed to compromise Π_{CSL}, respectively. Then we show that α_A is a reasonable bound for practical systems. We now move on to the second layer called the Consensus Layer (CSL) in which the committee of nodes from this layer handles the transactions.

4.2 Consensus Layer (CSL)

This layer is responsible for handling transactions and reaching a consensus on the state of the system. The nodes in this layer are controlled by participants. This layer maintains a shared state of the system by running any Byzantine agreement protocol or BFT protocol, denoted by Π_{CSL}. One can use any BFT protocol suited for the permissioned blockchain setting as long as it satisfies the network model. The shared-state can only be modified by pre-determined operations. This shared-state can consist of anything related to the purpose of the protocol like running a cryptocurrency, notarization platform, smart-contract platform, etc. For a cryptocurrency, this state may contain account balances, a transaction blockchain, etc. Note that however, for the sake of simplicity in this paper, we use the term *transaction blockchain* to denote the shared state which represents the purpose of the Tiramisu protocol, $\Pi_{\mathcal{TM}}$. Similarly, we use the term *transaction* throughout this paper to denote a state change in the transaction blockchain.

We now claim that the Tiramisu protocol satisfies the properties of distributed ledgers described in Section 3.2. BFT protocols structure their execution into a sequence of views, each with a designated leader process. These protocols guarantee Safety and liveness by ensuring that all correct nodes eventually overlap in a single view, with the right leader, for enough time to reach consensus.

Claim 1. *A Tiramisu protocol session achieves Safety and liveliness if* Π_{CSL} *does so.*

At any instance in time, this layer is operated by n_{CSL} nodes, collectively called the *committee* of this layer. The committee size can be either fixed or variable depending on network parameters. Fixed-size can be achieved through various techniques. For example, by allowing only the latest n_{CSL} nodes from the ACL blockchain to join the system. Note that, the shared-state must store a copy of the ACL blockchain, or an equivalent, to ensure verification of the nodes. The nodes in the committee are decided only after a BFT agreement on the identity of these nodes. This agreement ensures that any behavior of the ACL blockchain, like a blockchain fork, is not relevant in this layer. Since we use Byzantine agreement protocols, CSL inherently satisfies the properties of Safety and liveness. This justifies Claim 1.

Now we discuss our analysis of Tiramisu, in which we also identify three conditions on protocol parameters that ensure incentive compatible implementation of a Tiramisu protocol.

5   Tiramisu: Analysis

It is important to note that each node is independently selected to be on the committee. The probability of this selection is proportional to the number of democratic resources, as required for participating in ACL which is owned by the node in the network.

The BAR model (described in Section 3.4.1) allows us to analyze the game-theoretic and security aspects of Tiramisu as presented next.

5.1   Game-theoretic Analysis

As stated in our player model, the reward structure of any blockchain consensus protocol induces a game among the participants. We highlight that we account for the computational costs of validating the state of the system. This consideration forms the basis of our game model, which also makes it more suited for realistic setting. To analyze this induced game, we define the following notations: Consider a set of nodes $\mathcal{N} = \{1, \ldots, n_{CSL}\}$. This set includes all the information about a node, including its strategy and type. Let each participant i's strategy be s_i and its type be τ_i, where $\tau_i \in \{\mathcal{H}, \mathcal{R}, \mathcal{B}\}$. Note that, the strategy will depend on the participant's type. We have $\vec{s} = (s_1, \ldots, s_m)$ as the strategy vector and $\vec{s}_{-i}$ as the vector without node i. Let TR_i denote participant i's reward for one transaction block. Note that this reward is left for design of the protocol. With this, $u_i(\tau_i, TR_i, \vec{s})$ represents a participant i's utility from its participation. We assume that rational participants do not collude with each other based on the premise that the distributed environment will make it resistant to collusion.

In Tiramisu, we game-theoretically analyze the node's behavior at equilibrium using the following notion.

Definition 6 (Pure Strategy Nash Equilibrium (PSNE)). *A strategy vector $\vec{s^*} = \{s_1^*, \ldots, s_m^*\}$ is said to be a Pure Strategy Nash Equilibrium (PSNE) if for every node i, it maximizes its utility $u_i(\vec{s^*}, \tau_i, r_i)$ i.e., $\forall i \in \mathcal{M}$, $\forall j$,*

$$u_i(\vec{s^*}, \tau_i, r_i) \geq u_i(s_i, \mathcal{S}_{-i}^*, \tau_i, r_i); \ \forall s_i. \tag{1}$$

Intuitively, PSNE states that it is the best response for a node to follow $\vec{s^*}$, given that every other node is following it.

Equilibrium. Towards this, we assume that a rational player $\mathcal{R}$ incurs a cost $\kappa_{\mathcal{R}}$ if any malicious transaction block is committed in CSL, similar to Amoussou-Guenou et al. [2020]. This assumption is based on the premise that the entire ecosystem of the currency inflicts harm when an invalid block is accepted. This cost is directly proportional to the amount of stake that the rational node has invested in the system. Note that this stake can be in the form of electricity for mining in a proof-of-work blockchain, amount locked in the blockchain for participation in a proof-of-stake blockchain, rewards earned by the miner, and likewise. For simplicity, we take this proportionality constant to be 1, making the cost equal to the stake invested. We denote this cost of a rational node i by κ_i. Let $\kappa_{\mathcal{R}}$ be κ_i of the rational player with minimum invested stake. Observe that the protocol designer can set a lower-bound on $\kappa_{\mathcal{R}}$ using a variety of well-known techniques.

Although Byzantine nodes behave arbitrarily by definition, we consider a specific behavior of Byzantine nodes. Specifically, we assume that the objective of Byzantine nodes is to minimize the utility of the rational nodes and prevent the protocol from achieving its goal, regardless of the cost they incur. Note that this assumption ensures the security of

a Tiramisu protocol in the worst possible case, and hence it ensures security for all cases, i.e. regardless of the behavior of Byzantine nodes. With these assumptions, we model a game between honest, rational and Byzantine nodes' nodes in the committee, defined as: $\Gamma = \langle N, (\tau_i), (S_i), (u_i) \rangle$.

Definition 7 (Nash Incentive Compatible). *We say that a Tiramisu protocol $\Pi_{\mathcal{TM}}$ is Nash Incentive Compatible (NIC) if $\vec{s^*} = (s_1^2, \cdots, s_{n_\mathcal{R}}^2)$ is a PSNE for all rational nodes and at this PSNE all rational nodes obtain non-negative utility.*

For the analysis, let $P_{invalid}(\mathcal{N}, s_i^*)$ be node i's belief that an invalid block is accepted after it follows s_i^*. Honest nodes will always follow the protocol, i.e. s_i^2 strategy. Byzantine nodes will always follow s_i^3 because this is the strategy that best aligns with our assumption on the objective of Byzantine nodes. Hence we only consider rational nodes' behavior and therefore their equilibrium strategies.

The expected utility of a rational node, when following s_i^1 strategy for a single transaction block in CSL, is total rewards obtained for that block, minus the average cost of mining a block in ACL for that committee session, and minus the expected κ_i. Similarly, to calculate the utility of a rational node for strategy s_i^2, we need to subtract the cost of validating a block as well. The expected utilities of a rational node i for one round are as follows:

$$u_i(\cdot, \cdot, s_i^1) = TR_i - \frac{c_{mine}}{n_{TX}} - P_{invalid}(\mathcal{N}, s_i^1) \cdot \kappa_i \tag{2}$$

$$u_i(\cdot, \cdot, s_i^2) = TR_i - \phi \cdot c_{val} - \frac{c_{mine}}{n_{TX}} - P_{invalid}(\mathcal{N}, s_i^2) \cdot \kappa_i \tag{3}$$

For n_{CSL} nodes, let n_H, n_R, and n_B be number of honest, rational and Byzantine nodes respectively. Trivially, $n_{CSL} = n_H + n_R + n_B$. We use PSNE for analysis, as we have a static game with asymmetric information.

Claim 2. *For every rational node i, the probability of an invalid block being accepted will be more if it signs a block without validating, as compared to when it validates and then signs, i.e.,*

$$\delta = P_{invalid}(\mathcal{N}, s_i^1) - P_{invalid}(\mathcal{N}, s_i^2) > 0.$$

Intuitively, Claim 2 follows from the fact that the chance of an invalid block being accepted is more if a rational node decides to sign the block as valid without verifying it. Now we present the formal proof for the claim.

Proof. Without loss of generality, for the proof we consider a rational player i. Let

$$P_{invalid}(\mathcal{N}, s_i^1) = k_1$$

and

$$P_{invalid}(\mathcal{N}, s_i^2) = k_2$$

s.t.

$$0 < k_1, k_2 < 1.$$

In the event when node i plays the strategy s_i^1 i.e. signs the block as valid without validating, its belief regarding an invalid block being accepted can only increase. This follows by observing that the size of the committee, i.e., n_{CSL} is finite. Thus, node i not validating a block will directly imply that the chance of an invalid block being accepted will be more, i.e., $k_1 > k_2$. Now,

$$P_{invalid}(\mathcal{N}, s_i^1) - P_{invalid}(\mathcal{N}, s_i^2) = k_1 - k_2 = \delta > 0$$

$\square$

We denote the minimum possible value of δ for a protocol session running with specific security guarantees as δ_{min}. Along with the node's belief and some protocol parameters, δ_{min} is dependent on the CDF F which is appropriately defined in security analysis under proof of Proposition 1.

NIC Conditions. For all rational nodes, there exists a PSNE under certain conditions. We identify three such conditions to ensure that a Tiramisu protocol session $\Pi_{\mathcal{TM}}$ is NIC (Definition 7). We call these set of conditions as *NIC-Conditions* which are based on the following requirements. First, to ensure that Π_{CSL} is secure. Second, to ensure that rational nodes attain non-negative utility when they deviate from following the honest strategy. And third, to ensure that each rational node must obtain positive utility. The respective NIC-Conditions are as follows:

1. **Faithful Fault Tolerance Condition** $n_B < n_f$
2. **Maximum Payload Condition** $\phi \leq \frac{\kappa_\mathcal{R} \cdot \delta_{min}}{c_{val}}$

3. **Minimum Reward Condition** $TR \geq \phi \cdot c_{val} + \frac{c_{mine}}{n_{TX}}$

Note that NIC-Conditions provide constrains on protocol parameters which are dependent on protocol design, implementation, and live-network data[3]. Hence, satisfying these conditions rely on the careful implementation by protocol designer.

Theorem 1. *In Tiramisu, Faithful Fault Tolerance Condition, Maximum Payload Condition and Minimum Reward Condition are sufficient to ensure that the protocol is Nash Incentive Compatible (NIC). Formally, if Faithful Fault Tolerance Condition, Maximum Payload Condition and Minimum Reward Condition are true, then $\Pi_{\mathcal{TM}}$ is NIC according to Definition 7.*

Proof. To ensure that $\Pi_{\mathcal{TM}}$ is NIC, first requirement is that Π_{CSL} must work correctly. Hence the condition $n_B < n_f$ should be true, justifying the requirement of Faithful Fault Tolerance Condition.

Second, $s_i^* = s_i^2$ must be a PSNE for every rational node i. To ensure this, $u_i(R, s_i^1) \leq u_i(R, s_i^2)$ should satisfy according to Equation 1. On solving this condition using equations 2 and 3 we get,

$$\kappa_i \cdot (P_{invalid}(\mathcal{N}, s_i^1) - P_{invalid}(\mathcal{N}, s_i^2)) \geq \phi \cdot c_{val}.$$

Substituting $(P_{invalid}(\mathcal{N}, s_i^1) - P_{invalid}(\mathcal{N}, s_i^2))$ as δ, where $\delta \in (0, 1]$ following from Claim 2, we get our result for the value of ϕ as:

$$\phi \cdot c_{val} \leq \kappa_i \cdot \delta.$$

Now, since $\kappa_{\mathcal{R}} \leq \kappa_i$ and $\delta_{min} \leq \delta$, we get the following condition as:

$$\phi \cdot c_{val} \leq \kappa_{\mathcal{R}} \cdot \delta_{min},$$

justifying the requirement of Maximum Payload Condition.

Finally in third, for the above mentioned PSNE, utilities gained by all rational nodes must be non-negative, i.e.
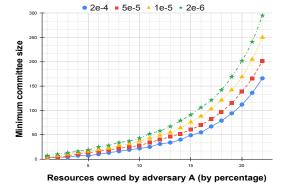
$$TR_i - \phi \cdot c_{val} - \frac{c_{mine}}{n_{TX}} - P_{invalid}(\mathcal{N}, s_i^2) \cdot \kappa_i \geq 0.$$

Notice that in this PSNE, since all rational nodes will validate the blocks, $P_{invalid}(\mathcal{N}, s_i^2) = 0$. Also, $TR_i \leq TR$. From this we get the condition as:

$$TR \geq \phi \cdot c_{val} + \frac{c_{mine}}{n_{TX}},$$

justifying the Minimum Reward Condition. □

However, in a special case where $n_H > n_{CSL} - n_f$, $S^* = (s_i^1)$ is the PSNE for rational nodes where only valid blocks are accepted and is as special case of a NIC $\Pi_{\mathcal{TM}}$.
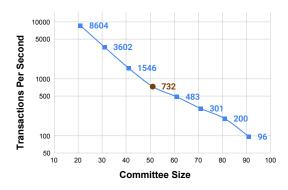

Figure 2: For a given α_A, n_{CSL}^{min} for different values of ϵ.

Figure 3: Transactions Per Seconds vs. Committee Size

---

[3]Here, live-network data at least includes the stored data of a protocol session along with its consolidated statistics.

5.2   Security Analysis

Setting. Our security analysis assumes that protocol Π_{ACL} is secure, i.e. ACL runs securely based on inherent security of Π_{ACL} and intelligent protocol design. We thus analyse the security of a $\Pi_{\mathcal{TM}}$ session by just ensuring the security of CSL. Notice that Tiramisu's security relies on securing Π_{CSL} and the inherent security guarantees of Π_{ACL}. We achieve this security just by quantifying the committee size in CSL based on an appropriately defined function in this analysis. Next, we consider this security analysis for a NIC Tiramisu protocol session. This is followed from Theorem 1 in our game-theoretic analysis. Notice that in this setting, since all rational players will follow the honest strategy, it is justified to consider all rational nodes as honest. Hence, we consider only honest and Byzantine nodes in our security analysis of a NIC Tiramisu protocol. Towards this, we first define our adversary model.

Adversary Model. We consider an adversary A controlling α_A fraction of the democratized resources in the network. The remaining resources are controlled by a meta entity $\mathcal{H}$, denoting all other participants in the network. The goal of the adversary is to control more than n_f nodes, i.e., the minimum number of nodes needed to compromise the protocol.

Note that the probability of a node joining the CSL is directly proportional to the amount of democratized resources owned by the respective participant in ACL. We denote the total participants in CSL to be n_{CSL} at any time. Similar to Decker et al. [2016], we analyze Tiramisu in its *steady state*, i.e., the number of validators and computational resources are governed by their respective expected value. To capture this formally, we define the following *secure state*.

Definition 8 (Secure State). *A $\Pi_{\mathcal{TM}}$ session is said to be in a secure state ($\mathcal{SS}$) if the number of nodes in CSL controlled by an adversary is strictly less than minimum number of nodes needed to compromise Π_{CSL}. Formally, $n_A < n_f$.*

We highlight that Claim 1 relies on security and correct execution of $\Pi_{\mathcal{TM}}$. Since it relies on security guarantees of underlying protocols Π_{ACL} and Π_{CSL}, we say that Claim 1 is true if $\Pi_{\mathcal{TM}}$ session is in a secure state. Hence, we state that a $\Pi_{\mathcal{TM}}$ session achieves Safety and liveness only in a secure state given Π_{CSL} does so. Towards this, we introduce the secure state function F_{SS} which provides n_{CSL}^{min}, the minimum number of nodes required in CSL for a $\mathcal{SS}$ with $1 - \epsilon$ probability.

$$F_{SS} : R \times R \to Z^+,$$

and

$$n_{CSL}^{min} = F_{SS}(1 - \epsilon, \alpha_A).$$

Here, $1 - \epsilon$ denotes the required minimum probability for a $\mathcal{SS}$, and α_A denotes the fraction of network resources controlled by the adversary. Note that F_{SS} is a useful function to ensure Faithful Fault Tolerance Condition with desired probability guarantee.

Proposition 1. *A Tiramisu protocol is in a secure state for a given α_A with probability of at-least $1 - \epsilon$ if the minimum number of nodes in CSL is*

$$n_{CSL}^{min} = F_{SS}(1 - \epsilon, \alpha_A),$$

where

$$F_{SS} = argmin_{n_{CSL}}\{(1 - F(n_f; n_{CSL}, \alpha_A)) \leq \epsilon\}.$$

Here $F(n_f; n_{CSL}, \alpha_A)$ is an appropriately defined function in the proof and gives the probability of $n_A < n_f$.

Proof. The probability that a participant mines a block in ACL is directly proportional to its fraction of computational power in the system, i.e.,

$$\Pr(\text{Mine a block in ACL}) \propto \alpha_{participant}.$$

This follows by observing that the size of the committee, i.e., n_{CSL} is finite. Let X be the discrete random variable denoting the number of validators controlled by the adversary. Consider the following standard binomial distribution for deriving the probability of the number of validators controlled by the adversary A as,

$$F(k; n_{CSL}, \alpha_A) = \Pr(k; n_{CSL}, \alpha_A) = \Pr(X = k)$$

$$\text{s.t. } \Pr(X = k) = \binom{n_{CSL}}{k}(\alpha_A)^k(1 - \alpha_A)^{(n_{CSL}-k)}$$

Now, let F be the corresponding CDF defined as,

$$F(x; n_{CSL}, \alpha_A) = \Pr(X \leq x).$$

$F(n_f; n_{CSL}, \alpha_A)$ gives the probability that the system is in a secure state, i.e. $n_A < n_f$, which depends on the value of n_{CSL} and α_A from the binary distribution.

With this we set

$$F_{SS}(1 - \epsilon, \alpha_A) = argmin_{n_{CSL}}\{(1 - F(n_f; n_{CSL}, \alpha_A)) \leq \epsilon\}.$$

$\square$

Selecting Safe Committee Size. ϵ denotes the probability of a system not in a secure state, i.e. $\epsilon = \Pr(\Pi_{\mathcal{TM}}$ not in $\mathcal{SS})$. Observe that the probability of a node being added to CSL is directly proportional to $\alpha_{participant}$, i.e., the resources it holds in the system. Since, each node in the CSL can either belong to the rational participants with probability $1 - \alpha_A$ or to a byzantine adversary with probability α_A, we can model this distribution as a binomial distribution. We know that, from Proposition 1, the protocol will not be in a secure state if the number of nodes controlled by the adversary in the committee are equal to or more than n_f, i.e. $\alpha_A \geq n_f$. Thus, ϵ becomes the binomial distribution of n_A successes in n_{CSL} independent experiments, with $\frac{\alpha_A}{n_{CSL}}$ as the probability of a each success. Based on different values of α_A and probability guarantees of secure state, we obtain the values of n_{CSL}^{min} [4]. We consider this analysis for probability guarantees, for system being in a secure state, ranging from 2×10^{-4} to 2×10^{-6}. Note that in Bitcoin, for 10% computational power with adversary, the system is said to be secure with probability of 2×10^{-4} i.e., 1 in 5000. We consider the values of α_A, the resources controlled by the adversary, in the range 1% to 18%. This range follows from observation that the most computational resources owned by a single mining pool in Bitcoin is 17%.

Following from this, observe that this upper bound on α_A, becomes closer to $\frac{n_f}{n_{CSL}}$ as we increase n_{CSL}. In general, for any Byzantine agreement protocol to function correctly it requires, at the least, that the number of participants controlled by the adversary n_A, be less than one-third of total nodes, n_{CSL}. For instance, $n_f = \frac{n_{CSL}+2}{3}$ for PBFT Castro and Liskov [1999]. Therefore, our upper bound for the fraction of network-resources controlled by an Adversary A is around $\frac{1}{3}$. In most blockchain protocols, using consensus mechanisms like Proof-of-Work, Proof-of-Stake, etc., this bound is $\frac{1}{2}$ (Bitcoin). Although, one should note that in practical scenarios, α_A does not reach over $\frac{1}{5}$ or 20%. For reference, the most computational resources owned by any miner in Bitcoin is 17%. Since our consideration is much higher than this percentage, we can say that the bounds in Tiramisu are reasonable for practical scenarios.

5.3   Discussion

We highlight that a Tiramisu protocol can be designed using any Proof-of-X consensus protocol in the first layer, inheriting its security. Similarly, any BFT protocol in the second layer can be used. In our analysis of Tiramisu, considering three types of nodes, we proved that the protocol will be Nash-Incentive-Compatible (Definition 7) for three certain conditions according to Theorem 1. We have also proved that any Tiramisu protocol is in a secure state (Definition 8) with high probability under reasonable assumptions and pre-defined conditions on protocol parameters.

One must note that PeerCensus Decker et al. [2016], Solida Abraham et al. [2018], Hybrid Consensus Pass and Shi [2017], and we believe, a few other protocols similar to Algorand Chen and Micali [2019] are a special instantiation of Tiramisu. In these special cases, we say that PeerCensus and Solida use a Proof-of-Work based protocol whereas Hybrid Consensus and Algorand use a Proof-of-Stake protocol in the ACL. However, these protocols will need formal analysis to prove so and this is left for future work. We believe our framework has further potential to be more universal and include major protocols as its instantiations.

6   ASHWAChain

ASHWAChain Arora et al. [2020] is an example of as multi-layered blockchain protocol that fits in the Tiramisu framework. ASHWAchain is a committee-based blockchain protocol in which the committee is being updated regularly. We present a formal analysis of the protocol in the full version of our paper Arora et al.. The Access Control Layer of ASHWAchain use Proof-of-Work as a method of sybil-resistance and PBFT Castro and Liskov [1999] towards achieving byzantine agreement in the Consensus Layer. As per analysis using the Tiramisu framework, we show that after selecting secure protocol parameters that ensure game-theoretic soundness the protocol can achieve a throughput of more than 700 transactions per second, more than the claim in the original paper by selecting a smaller committee size but greater Proof-of-Work difficulty. We remark that analysis with out framework is also useful in finding optimal parameters for multi-layered blockchain protocols. Figures 2 and 3 show the effect of committee size on security and performance for ASHWAChain.

7   Conclusion

In this work, we formalize a layered approach towards blockchain consensus, Tiramisu. The layered approach allows us to combine permissioned and permissionless blockchain protocols to design protocols that provide both scalability and decentralization (Section 4). We use separate layers of consensus protocols, much like a Tiramisu. Along with a security analysis, we use a game-theoretic analysis in the BAR model to show that our framework assures security in

---

[4]For additional details, we refer the reader to Proposition 1

the presence of Byzantine adversary with limited network resources while providing a better throughput than traditional permissionless blockchain protocols (Section 5). Lastly, we proved that under certain assumptions for the protocol parameters, there exists a Pure Strategy Nash Equilibrium in which the rational players validate the transactions before signing any block. We believe this abstract approach of Tiramisu is useful for blockchain practitioners to securely design layered blockchain protocols for greater performance.

References

Satoshi Nakomoto. Bitcoin: A peer-to-peer electronic cash system. `https://bitcoin.org/bitcoin.pdf`.

Blockchain Charts, 2020. `https://www.blockchain.com/charts/transactions-per-second`.

Ethereum Daily Transactions Chart | Etherscan, 2020. `https://etherscan.io/chart/tx`.

VisaNet Booklet, 2021. `https://usa.visa.com/dam/VCOM/download/corporate/media/visanet-technology/visa-net-booklet.pdf`.

Ittai Abraham, Dahlia Malkhi, Kartik Nayak, Ling Ren, and Alexander Spiegelman. Solida: A Blockchain Protocol Based on Reconfigurable Byzantine Consensus. In James Aspnes, Alysson Bessani, Pascal Felber, and João Leitão, editors, *21st International Conference on Principles of Distributed Systems (OPODIS 2017)*, volume 95 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 25:1–25:19, Dagstuhl, Germany, 2018. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. ISBN 978-3-95977-061-3. doi:10.4230/LIPIcs.OPODIS.2017.25. URL `http://drops.dagstuhl.de/opus/volltexte/2018/8640`.

Mahnush Movahedi Timo Hanke and Dominic Williams. Dfinity technology overview series consensus system. CoRR abs/1805.04548 (2018), 2018. URL `https://arxiv.org/pdf/1805.04548.pdf`.

C. Decker, J. Seidel, and R. Wattenhofer. Bitcoin meets strong consistency. In *In Proceedings of the 17th International Conference on Distributed Computing and Networking Conference (ICDCN)*. ICDCN, 2016. URL `https://tik-db.ee.ethz.ch/file/ed3e5da74fbca5584920e434d9976a12/peercensus.pdf`.

Rafael Pass and Elaine Shi. Hybrid Consensus: Efficient Consensus in the Permissionless Model. In Andréa W. Richa, editor, *31st International Symposium on Distributed Computing (DISC 2017)*, volume 91 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 39:1–39:16, Dagstuhl, Germany, 2017. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. ISBN 978-3-95977-053-8. doi:10.4230/LIPIcs.DISC.2017.39. URL `http://drops.dagstuhl.de/opus/volltexte/2017/8004`.

Miles Carlsten, Harry Kalodner, S Matthew Weinberg, and Arvind Narayanan. On the instability of bitcoin without the block reward. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 154–167, 2016.

Ayelet Sapirshtein, Yonatan Sompolinsky, and Aviv Zohar. Optimal selfish mining strategies in bitcoin. In *International Conference on Financial Cryptography and Data Security*, pages 515–532. Springer, 2016.

Shoeb Siddiqui, Ganesh Vanahalli, and Sujit Gujar. Bitcoinf: Achieving fairness for bitcoin in transaction fee only model. In Amal El Fallah Seghrouchni, Gita Sukthankar, Bo An, and Neil Yorke-Smith, editors, *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems, AAMAS '20, Auckland, New Zealand, May 9-13, 2020*, pages 2008–2010. International Foundation for Autonomous Agents and Multiagent Systems, 2020.

Documentation/TechnicalWhitepaper.md, 2020. `https://github.com/EOSIO/Documentation/blob/master/TechnicalWhitePaper.md`.

Ningyu He, Ruiyi Zhang, Lei Wu, Haoyu Wang, Xiapu Luo, Yao Guo, Ting Yu, and Xuxian Jiang. Security analysis of eosio smart contracts. 03 2020.

Juan Garay, Aggelos Kiayias, and Nikos Leonardos. The bitcoin backbone protocol: Analysis and applications. In Elisabeth Oswald and Marc Fischlin, editors, *Advances in Cryptology - EUROCRYPT 2015*, pages 281–310, Berlin, Heidelberg, 2015a. Springer Berlin Heidelberg. ISBN 978-3-662-46803-6.

Jing Chen and Silvio Micali. Algorand: A secure and efficient distributed ledger. *Theoretical Computer Science*, 777: 155 – 183, 2019. ISSN 0304-3975. doi:https://doi.org/10.1016/j.tcs.2019.02.001. In memory of Maurice Nivat, a founding father of Theoretical Computer Science - Part I.

H. Yu, I. Nikolić, R. Hou, and P. Saxena. Ohie: Blockchain scaling made simple. In *2020 IEEE Symposium on Security and Privacy (SP)*, pages 90–105. IEEE, 2020. doi:10.1109/SP40000.2020.00008.

Ittai Abraham, Dahlia Malkhi, Kartik Nayak, Ling Ren, and Alexander Spiegelman. Solidus: An incentive-compatible cryptocurrency based on permissionless byzantine consensus. In *CoRR abs/1612.02916v1 (2016)*, 2016.

Bruno Biais, Christophe Bisière, Matthieu Bouvard, and Catherine Casamatta. The blockchain folk theorem. In *The Review of Financial Studies (2019)*, 2019.

Yackolley Amoussou-Guenou, Bruno Biais, Maria Potop-Butucaru, and Sara Tucci-Piergiovanni. Rational vs byzantine players in consensus-based blockchains. In *In Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), Auckland, New Zealand, May 9–13, 2020, IFAAMAS, 9 pages*. IFAAMAS, 2020. URL http://ifaamas.org/Proceedings/aamas2020/pdfs/p43.pdf.

M. H. Manshaei, M. Jadliwala, A. Maiti, and M. Fooladgar. A game-theoretic analysis of shard-based permissionless blockchains. *IEEE Access*, 6:78100–78112, 2018. doi:10.1109/ACCESS.2018.2884764.

Amitanand S. Aiyer, Lorenzo Alvisi, Allen Clement, Jean Philippe Martin Michael Dahlin, and Carl Porth. Bar fault tolerance for cooperative services. In *Proceedings of the 20th ACM Symposium on Operating Systems Principles 2005, SOSP 2005, Brighton, UK, October 23-26, 2005*, page 45–58, 2005.

Joseph Y. Halpern and Xavier Vilaça. Rational consensus: Extended abstract. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing*, PODC '16, page 137–146, New York, NY, USA, 2016. Association for Computing Machinery. ISBN 9781450339643. doi:10.1145/2933057.2933088.

Joseph Poon and Tadge Dryja. Lightning network. 2015. URL https://lightning.network/lightning-network-paper.pdf.

Vitalik Buterin Joseph Poon. Plasma: Scalable autonomous smart contracts. In *White paper*, pages 1–47, 2017. URL https://www.plasma.io/plasma-deprecated.pdf.

Christian Reitwießner Jason Teutsch. A scalable verification solution for blockchains. 2017. URL https://people.cs.uchicago.edu/~teutsch/papers/truebit.pdf.

Gavin Wood. Ethereum: A secure decentralised generalised transaction ledger. *Ethereum project yellow paper*, 151:1–32, 2014.

Elli Androulaki, Artem Barger, Vita Bortnikov, Christian Cachin, Konstantinos Christidis, Angelo De Caro, David Enyeart, Christopher Ferris, Gennady Laventman, Yacov Manevich, Srinivasan Muralidharan, Chet Murthy, Binh Nguyen, Manish Sethi, Gari Singh, Keith Smith, Alessandro Sorniotti, Chrysoula Stathakopoulou, Marko Vukolić, Sharon Weed Cocco, and Jason Yellick. Hyperledger fabric: A distributed operating system for permissioned blockchains. In *Proceedings of the Thirteenth EuroSys Conference*, EuroSys '18, New York, NY, USA, 2018. Association for Computing Machinery. ISBN 9781450355841. doi:10.1145/3190508.3190538. URL https://doi.org/10.1145/3190508.3190538.

Juan A. Garay, Aggelos Kiayias, and Nikos Leonardos. The bitcoin backbone protocol: Analysis and applications. In *EUROCRYPT (2)*, pages 281–310. Springer, 2015b. doi:10.1007/978-3-662-46803-6_10. URL https://www.iacr.org/archive/eurocrypt2015/90560288/90560288.pdf.

Miguel Castro and Barbara Liskov. Practical byzantine fault tolerance. In *Third Symposium on Operating Systems Design and Implementation, New Orleans, USA, February 1999*. Laboratory for Computer Science, Massachusetts Institute of Technology, 545 Technology Square, Cambridge, MA 02139, 1999.

Sanidhay Arora, Anurag Jain, Sankarshan Damle, and Sujit Gujar. Ashwachain: A fast, scalable and strategy-proof committee-based blockchain protocol. In *Workshop on Game Theory in Blockchain at WINE 2020 (GTiB@WINE 2020)*, 2020. URL https://econcs.pku.edu.cn/wine2020/wine2020/Workshop/GTiB20_paper_9.pdf.

Sanidhay Arora, Anurag Jain, Sankarshan Damle, and Sujit Gujar. Tiramisu: Layering consensus protocols for scalable and secure blockchains. https://www.dropbox.com/s/c1am9xbvyd51ndf/Tiramisu.pdf?dl=0.

8 ASHWAChain Details

Algorithm 1: ASHWAchain's agreement process for identity p

---

begin
 if p = *Primary* **then**
  | Creates any operation and broadcasts
 else
  | Validate operation, sign and broadcast
 wait(Latency Time)
 if *signed by super-majority* **then**
  | Append *operation* to O
  | Commit *operation* according to its commit rule
 else
  | Nothing is committed, RESET

Algorithm 2: ASHWAchain: Predefined operations to modify the shared state in CSL.

Validate powBlock(b): **begin**
 $b^* \leftarrow$ latest block in $PowChain$
 if b *is child of* b^* *and* b *is valid* **then**
  | return Valid
 else
  | return Invalid

Validate comBlock(c): **begin**
 if *All* n_{CSL} *nodes in* c *match*
 with the POWCHAIN *and* c *is valid* **then**
  | return Valid
 else
  | return Invalid

Validate txBlock(t): **begin**
 if *signatures and all txs are valid* **then**
  | return Valid
 else
  | return Invalid

Commit powBlock(b): **begin**
 Append powBlock(b) in O
 Append b to POWCHAIN

Commit comBlock(c): **begin**
 Append comBlock(c) in O
 Append c to COMCHAIN

Commit txBlock(t): **begin**
 Append txBlock(t) to O
 for $tx \in t$ **do**
  | $B(from) \leftarrow B(from) - coins$
  | $B(to) \leftarrow B(to) + coins$
 Append the list of validators of t in T
 for $v \in T(t)$*;* $txFee \leftarrow$ *Total transaction fee;*
 $nSig \leftarrow$ *Total signatures* **do**
  | $B(v) \leftarrow B(v) + \frac{txFee + \text{Block Reward}}{nSig}$
 Append t to TXCHAIN