

Blockchain Consensus Layer as a Plug-in Tools for Real-time Analysis and Monitoring

Guilherme João Bidarra Breia Lopes

Dissertação para obtenção do grau de Mestre em

Engenharia informática
(2º ciclo de estudos)

Orientador: Prof. Doutor Simão Melo de Sousa

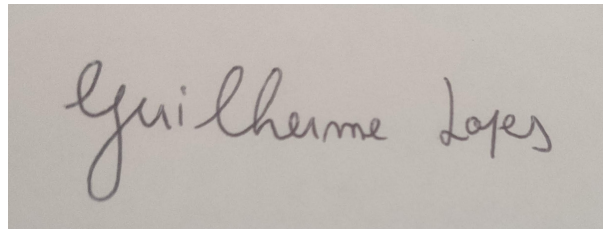
Outubro 2023

Declaração de Integridade

Eu, Guilherme João Bidarra Breia Lopes, que abaixo assino, estudante com o número de inscrição M12155 do Mestrado em Engenharia Informática da Faculdade de Engenharia, declaro ter desenvolvido o presente trabalho e elaborado o presente texto em total consonância com o Código de Integridades da Universidade da Beira Interior.

Mais concretamente afirmo não ter incorrido em qualquer das variedades de Fraude Académica, e que aqui declaro conhecer, que em particular atendi à exigida referenciação de frases, extratos, imagens e outras formas de trabalho intelectual, e assumindo assim na íntegra as responsabilidades da autoria.

Universidade da Beira Interior, Covilhã 09/10/2023

A rectangular box containing a handwritten signature in dark ink. The signature is written in a cursive style and reads "Guilherme Lopes".

(assinatura conforme Cartão de Cidadão ou preferencialmente assinatura digital no documento original se naquele mesmo formato)

Acknowledgments

First and foremost, I extend my deepest gratitude to my family, who have always been my cornerstone. To my friends and colleagues at the Release Lab, thank you for making the long hours of coding and debugging bearable, and for the endless intake of coffee, no matter its quality, from the bar downstairs.

A special thank you goes to my girlfriend, who has been my steadfast companion throughout the majority of this project and thesis journey. You were often the only audience to my thesis monologues, listening attentively even when the jargon got too dense. Your presence made the journey not just bearable, but enjoyable.

I am immensely grateful to my coordinator, Professor Simão Melo de Sousa. Your guidance has been invaluable, and your friendship has enriched this academic journey, making it more than just a series of deadlines and milestones.

I would also like to express my gratitude to Vincent Botbol. Your timely assistance helped me overcome one of the most challenging hurdles in this project. Your help was the missing semicolon in my long string of code.

Last but not least, a nod to my cats. Thank you for being there during my most challenging times, even if your idea of support was occasionally walking across my keyboard or napping on my lap.

Resumo

Esta dissertação aborda os desafios complexos e soluções associados à implementação e teste de algoritmos de consenso em redes *blockchain*. A pesquisa é baseada na plataforma *blockchain Tezos*, escolhida pela sua adaptabilidade e versatilidade. O objetivo principal é avançar no campo da tecnologia *blockchain*, desenvolvendo ferramentas e métodos que facilitem a troca fácil, teste e análise comparativa de diferentes algoritmos de consenso em um ambiente de rede ao vivo.

O documento está organizado em várias seções cruciais. Começa com uma introdução que delinea o problema e a solução, enfatizando as lacunas existentes no desenvolvimento e teste de algoritmos de consenso, particularmente a ausência de ferramentas para testes ao vivo e os desafios na troca de algoritmos. Segue-se uma revisão abrangente do estado da arte em algoritmos de consenso, oferecendo uma análise detalhada da sua eficiência, limitações e adequação para várias aplicações.

Um estudo de caso preliminar é então apresentado, focando na implementação de um algoritmo de consenso de Prova de Trabalho (PoW na plataforma Tezos. Isso serve como uma prova de conceito e fornece insights valiosos sobre o funcionamento interno dos algoritmos de consenso e sua integração com nós de *blockchain* existentes.

A contribuição central desta dissertação é o desenvolvimento de duas ferramentas: uma voltada para facilitar o desenvolvimento de algoritmos de consenso e outra projetada para seus testes ao vivo. Essas ferramentas destinam-se a servir tanto a comunidade acadêmica quanto a industrial interessadas em tecnologia *blockchain*.

Por fim, a dissertação apresenta um cenário do mundo real que emprega ambas as ferramentas. Um algoritmo de Prova de Autoridade (PoA) é implementado usando a ferramenta de desenvolvimento e depois testado fazendo uso da ferramenta de teste ao vivo. Os resultados são comparados com o algoritmo PoW previamente implementado, fornecendo evidências empíricas da eficiência e limitações de ambos os algoritmos.

Em resumo, esta dissertação oferece um *framework* robusto para o desenvolvimento e teste de algoritmos de consenso, substanciado por evidências empíricas da plataforma Tezos. Tem implicações significativas para ambos os pesquisadores acadêmicos e *stakeholders* da indústria, estabelecendo uma base sólida para trabalhos futuros neste campo em rápida evolução.

Palavras-chave

Blockchain, Algoritmos de Consenso, Plataforma Tezos, Proof of Work (PoW), Proof of Authority (PoA), Testes ao Vivo, Desenvolvimento de Protocolo

Abstract

This thesis addresses the complex challenges and solutions associated with the implementation and testing of consensus algorithms in blockchain networks. The research is anchored on the Tezos blockchain platform, selected for its adaptability and versatility. The primary aim is to advance the field of blockchain technology by developing tools and methods that facilitate the easy swapping, testing, and comparative analysis of different consensus algorithms in a live network environment.

The document is organized into several pivotal sections. It begins with an introduction that outlines the problem and solution, emphasizing the existing gaps in the development and testing of consensus algorithms, particularly the absence of tools for live testing and the challenges in swapping algorithms seamlessly. This is followed by a comprehensive review of the current state-of-the-art in consensus algorithms, offering a detailed analysis of their efficiency, limitations, and suitability for various applications.

A preliminary case study is then presented, focusing on the implementation of a Proof of Work (PoW) consensus algorithm on the Tezos platform. This serves as a proof of concept and provides valuable insights into the inner workings of consensus algorithms and their integration with existing blockchain nodes.

The core contribution of this thesis is the development of two tools: one aimed at easing the development of consensus algorithms and another designed for their live testing. These tools are intended to serve both the academic and industrial communities interested in blockchain technology.

Finally, the thesis presents a real-world scenario that employs both tools. A Proof of Authority (PoA) algorithm is implemented using the development tool and then tested using the live testing tool. The results are compared with the previously implemented PoW algorithm, providing empirical evidence of the efficiency and limitations of both algorithms.

In summary, this thesis offers a robust framework for the development and testing of consensus algorithms, substantiated by empirical evidence from the Tezos platform. It holds significant implications for both academic researchers and industry stakeholders, laying a strong foundation for future work in this rapidly evolving field

Keywords

Blockchain, Consensus Algorithms, Tezos Platform, Proof of Work (PoW), Proof of Authority (PoA), Live Testing, Protocol Development

Resumo Alargado

Esta dissertação debruça-se sobre os desafios e as soluções inovadoras associadas à implementação e teste de algoritmos de consenso em redes blockchain. O trabalho é centrado na plataforma blockchain Tezos, selecionada pela sua notável adaptabilidade e versatilidade. O objetivo primordial é contribuir para o avanço da tecnologia blockchain, criando ferramentas e métodos que permitam não só a implementação, mas também o teste rigoroso e a análise comparativa de diferentes algoritmos de consenso em ambientes de rede em tempo real.

O problema que esta dissertação procura resolver é duplo: primeiro, a falta de ferramentas que permitam a implementação ágil de algoritmos de consenso; segundo, a ausência de meios para testar esses algoritmos em condições de rede ao vivo. Estas lacunas tornam difícil para os investigadores e profissionais da indústria avaliar a eficácia e as limitações de diferentes algoritmos de consenso.

O documento está organizado em várias secções críticas. A introdução fornece uma visão geral do problema e da solução proposta, sublinhando as lacunas e desafios existentes. Uma revisão aprofundada do estado da arte em algoritmos de consenso segue-se, abordando uma variedade de algoritmos e as suas respectivas eficiências, limitações e aplicações.

Um estudo de caso preliminar é detalhado, focando na implementação de um algoritmo de Prova de Trabalho (PoW) na plataforma Tezos. Este estudo não só serve como uma prova de conceito, mas também oferece *insights* profundos sobre os mecanismos subjacentes dos algoritmos de consenso e os desafios de integrá-los em nós de blockchain existentes.

A contribuição principal desta dissertação reside no desenvolvimento de duas ferramentas específicas. A primeira, denominada 'Protocol Bootstrapper', é uma ferramenta de desenvolvimento que simplifica o processo de implementação de novos algoritmos de consenso. A segunda, conhecida como 'Live Testing Tool', é uma ferramenta de teste em tempo real que permite aos utilizadores executar testes rigorosos em algoritmos de consenso em condições de rede ao vivo.

Por último, a dissertação apresenta um cenário prático, onde ambas as ferramentas são aplicadas na implementação e teste de um algoritmo de Prova de Autoridade (PoA). Os resultados deste teste são comparados com os do algoritmo PoW, fornecendo dados empíricos que destacam as eficiências e limitações de ambos os algoritmos e também servem de demonstração da utilidade destas ferramentas.

Em resumo, esta dissertação oferece um quadro abrangente e robusto para o desenvolvimento, implementação e teste de algoritmos de consenso, apoiado por dados empíricos recolhidos na plataforma Tezos. Este trabalho tem implicações significativas para a comunidade académica e para os stakeholders da indústria, estabelecendo um novo padrão para futuras investigações e desenvolvimentos neste campo em rápido crescimento.

Contents

1	Introduction	1
1.1	Motivation	1
1.2	The Problem	2
1.3	Proposed Solution	3
1.4	Main Contributions	4
1.5	Document Organization	4
2	Core Concepts and State of the Art	7
2.1	Core Concepts	7
2.2	State of the Art	14
2.3	Conclusion	23
3	Preliminary Case Study	25
3.1	Tezos Blockchain as a tool for the Experiment	25
3.2	How Tezos is structured and How it allows self-amending	26
3.3	Implementation of a Protocol	28
3.4	Tools developed to interact with the protocol	35
3.5	Conclusion	36
4	Tool Development	37
4.1	Bootstrapping Tool	37
4.2	Live Testing Tool	39
4.3	Conclusion	44
5	Use Case: Developing and Testing a New Protocol	45
5.1	Goals and Approach	45
5.2	Bootstrapping and Implementing Proof of Authority	45
5.3	Methodology	47
5.4	Evaluation and Insights: PoA vs PoW	48

5.5 Conclusion	52
6 Conclusion	55
6.1 Future Work	55
Bibliography	57

List of Figures

2.1	Blockchain Datastructure.	8
2.2	FLP and CAP theorems visualized	12
2.3	Blockchain trilemma	14
3.1	Architecture of Tezos (Source: Tezos Documentation)	27
4.1	Live Testing Tool Architecture Visualized, with 4 Tezos Nodes Running	40

Acronyms

SMR	State Machine Replication
ASIC	Application-specific integrated circuit
BFT	Byzantine Fault Tolerance
DAG	Directed Acyclic Graph
DSL	Domain-Specific Language
FLP	MJ Fischer, NA Lynch, MS Paterson -
PoA	Proof of Authority
PoS	Proof of Stake
PoW	Proof of Work
RPC	Remote Procedure Call
TEE	Trusted Execution Environments
TTC	Time-to-Consensus

Chapter 1

Introduction

Blockchains are a revolutionary technology that have the potential to transform various industries by providing a secure and transparent platform for storing and exchanging information and assets. Decentralized in nature, blockchains use cryptography and consensus algorithms to create a shared ledger system, eliminating the need for a trusted intermediary.

In recent years, the demand for secure and transparent systems has led to the growing popularity of blockchains, especially in finance, supply chain management, digital identity, and digital assets. The decentralization of blockchains reduces the risk of single points of failure and ensures the integrity of data stored on the network.

Consensus algorithms play a critical role in the functioning of blockchains, determining the process of validating transactions, adding new blocks to the chain, and other crucial aspects like the selection of the next block producer. Different consensus algorithms offer varying levels of security, scalability, and decentralization, making the choice of the right algorithm for a specific use case crucial for the success of a blockchain network.

The main objective is to develop tools and methods to easily swap, test, develop, and compare different consensus algorithms in a live environment using actual nodes on a network. This study aims to provide insight into the field of blockchains and consensus algorithms, contributing to the advancement of more secure, scalable, and efficient blockchain systems. The Tezos [1] blockchain node will serve as a testbed for this project, and the focus will be on examining the potential of consensus algorithms to be swapped and changed on demand. Through this work, we introduce innovative tools that facilitate these processes, setting the stage for future advancements in the field.

From tools and methods to test consensus algorithms to the development of algorithms and their integration into existing blockchain nodes, this thesis will leave a lasting impact on the future of blockchains.

1.1 Motivation

While blockchain technology has garnered significant attention and adoption, the field is still in its infancy when it comes to the ease of experimenting with and implementing consensus algorithms. The introduction has laid out the landscape of blockchain technology and the pivotal role that consensus algorithms play in it. However, the real-world application of these algorithms is often hindered by the complexities involved in their development, testing, and deployment in a live network.

The current state of affairs presents a paradox: on one hand, we have an ever-growing list of consensus algorithms, each promising to solve the "blockchain trilemma" of scalability, security, and decentralization. On the other hand, we lack the tools and frameworks that allow for seamless experimentation and comparison of these algorithms in a live setting. This gap not only stunts innovation but also restricts the practical evaluation of whether an algorithm delivers on its promises.

This study is motivated by the need to bridge this gap. It aims to provide the tools that make it easier to swap, test, and compare different consensus algorithms in a live environment, using the Tezos blockchain as a testbed. By doing so, this research aspires to remove the barriers that currently limit the full exploration and understanding of consensus algorithms, thereby contributing to the advancement of blockchain technology as a whole.

The following section will delve deeper into the specific problems that this research addresses, further highlighting the necessity and timeliness of the proposed solutions.

1.2 The Problem

The development of blockchain technology has led to the creation of a large number of consensus protocols, each with its unique characteristics, from the selection of the leader, to the type of network, data structure, block structure, time between each block, and multiple ways to have a chain, to name just a few. Newer and innovative protocols are designed "everyday" with totally different characteristics compared to previous protocols. There's a lot of "fine tuning" to be done, even when comparing protocols of the same type. These small changes are enough to start a totally different network. For example, just increasing the original block size of Bitcoin was enough to start a new network, Bitcoin Cash [2].

This abundance of choices has made it difficult for individuals or organizations looking to start their own blockchain to determine which protocol is the best fit for their needs, as there's a lot to account for.

Furthermore, there is no easy way to test and compare these consensus protocols in a live, non-simulated environment. There is no blockchain node software that allow for a seamless swap of the consensus algorithm with testing and benchmarking, making it challenging to properly evaluate the performance of each protocol. This presents a major obstacle for researchers and developers looking to experiment with and improve upon existing consensus protocols.

Additionally, the process of designing and coding a new consensus algorithm is also challenging and requires a significant amount of expertise and resources. There is currently no platform that allows for the easy design, swap, and testing of new consensus algorithms.

All these problems contribute to the lack of standardization in the field of consensus protocols, a lot of attributes and characteristic to take in, making it difficult for individuals and organizations to adopt blockchain technology effectively. It is crucial to find a solution that

allows for the easy and efficient comparison and testing of consensus protocols, as well as the design and implementation of new protocols. This will pave the way for the further development and widespread adoption of blockchain technology.

1.3 Proposed Solution

The proposed solution aims to tackle the challenges mentioned in the previous section by providing a method of easily swapping consensus algorithms, so they can then be tested. To achieve this, the solution leverages the advantages of a pre-existing blockchain that is suitable for both swapping and testing.

By using an already built blockchain, the focus of this solution is solely on the consensus protocol/layer, the crucial aspect of any blockchain. The need for implementing other parts of the blockchain node such as the Peer-to-Peer layer, client, validation layer, storage of the blockchain state, etc., is eliminated. This benefits the solution by making use of already tested and industrial-grade blockchain software, freeing up time and resources to concentrate solely on the consensus.

One of the ideas for implementing this solution is to use the Tezos blockchain. The reason for using Tezos is its inherent adaptability, which allows for the removal of its current consensus algorithm and the implementation of a new one. This approach provides a layer of adaptability to the Tezos blockchain, making it suitable for testing and swapping various consensus algorithms with ease.

As of the writing of this document, there's no knowledge of projects or documents that describe a similar solution to this one proposed. There exists the concept of consensus testing when talking about a blockchain node implementation in specific, that is, for example, tools to test the Ethereum's consensus protocol, or the Tezos node, but there are no tools to test specifically the consensus part of these projects and compare them, or even, a tool to compare and test consensus protocols in general.

In summary, this thesis presents a comprehensive approach to address the challenges outlined in the problem statement. We have developed tools that facilitate the easy swapping, testing, and comparison of consensus algorithms in a live environment, using the Tezos blockchain as a base. This work serves as a pioneering step in standardizing the evaluation and development of consensus protocols, thereby making blockchain technology more accessible and customizable. The following section will elaborate on the main contributions of this research, detailing the tools and protocols developed to realize this solution.

1.4 Main Contributions

The work presented in this thesis addresses several challenges in the development and testing of blockchain protocols, leading to the following main contributions, that can be found here [here](#):

- **Development of Bootstrapper:** A tool that simplifies the process of creating new blockchain protocols by providing a template-based approach, reducing the amount of code that needs to be written.
- **Creation of Live Testing Tool:** A tool designed for real-time testing and debugging of blockchain protocols, offering various metrics for performance evaluation.
- **Implementation of a Proof of Work Protocol:** A fully functional PoW protocol developed from zero, serving as a practical example of the tool's capabilities.
- **Implementation of a Proof of Authority Protocol:** A PoA protocol developed to showcase the versatility of the Bootstrapper and Live Testing Tool, featuring a round-robin mechanism for validator selection.
- **Open Source Availability:** All tools and protocols developed are open-source and publicly available, encouraging further research and development in the field.

1.5 Document Organization

The following list describes the structure of this document:

1. **Introduction** - Introduces the research problem, objectives, and scope. Highlights the focus on blockchain consensus algorithms.
2. **Core Concepts and State of Art** - Offers a comprehensive review of the existing literature and technologies related to blockchain and consensus algorithms.
3. **Preliminary Case Study** - Details an experiment conducted on the Tezos blockchain to implement and test a Proof of Work consensus protocol. This chapter serves as a precursor to the tool development phase.
4. **Tool Development** - Describes the development of tools designed to facilitate the swapping and testing of consensus algorithms on the Tezos blockchain.
5. **Use Case: Developing and Testing a new Protocol** - Presents the use of the tools, results of the experiments using the tools, including metrics such as Time-to-Consensus, and discusses the implications of these findings.
6. **Conclusion** - Summarizes the research findings, discusses their significance, and suggests avenues for future research.

This is the organization of the document, where firstly we describe how this type of system works, how they can differ and then we elaborate on the problem and solution that will be worked on in this project.

Chapter 2

Core Concepts and State of the Art

The consensus mechanism is a critical component of any blockchain network. It determines the rules of the system, like validating transactions and adding them to the ledger. In order to understand the current development and innovation in this area, it is important to have a solid understanding of the core concepts of blockchains as well as the current state of the art in consensus algorithms. This chapter aims to provide an overview of both of these key topics.

2.1 Core Concepts

This section covers State Machine Replication (SMR), Blockchain Data Structure and Networks, Network Models, and the Concept of Consensus. Understanding these concepts is crucial for understanding how blockchains networks function in their core, the different trade-offs and design considerations involved in creating and operating a blockchain network. It starts with an overview of SMR and how it provides consistency and availability in decentralized systems.

Then, the concept of blockchain data structure and its application in decentralized networks is explored. The different Network Models in distributed systems, such as synchronous, asynchronous, and partially-synchronous are also discussed. Finally, the section concludes with a exploration of the topic of consensus and it's importance in blockchain networks.

State Machine Replication

State Machine Replication is the concept of replicating data in a distributed system, where nodes in a system maintain the same state of a machine. That is, where multiple entities maintain the same information.

In this approach, a node broadcasts an update to its state using a total order broadcast, and all other nodes receive the same updates in the same order and apply it to their own state.

Total order broadcast refers to notion of broadcasting messages in a specific order, such that all nodes in the network receive the messages in the same order, even if the messages are sent from different nodes at different times.

The main advantages of using SMR, like mentioned before, are that it can provide both **consistency** and **availability** (in some cases it only provides availability), making it a popular

choice for use in blockchain technology, databases, file systems, and other decentralized systems.

Consistency, because it guarantees that every node contains the same state of the machine (not necessarily at the same time), and availability is also ensured because, in case of a node failure, there are other nodes in the system that have the information.

Overall, the concept of State Machine Replication is a fundamental component of decentralized systems that plays a crucial role in ensuring the consistency and availability of shared data. In case of blockchains, to ensure that every node has the same chain and that there's no single point of failure of the whole system. This makes these networks reliable, since in case multiple nodes fail in the network, it can still operate, and that no single node has the power over the whole information, making the network decentralized.

Blockchain Data Structure and Networks

Blockchain is a data structure made of blocks that is often compared to a linked list (to which blocks are the nodes of the list), that are connected to a previous block in the chain by a reference to it, where usually is the hash of the block it is pointing to.

That is, every block, except for the first block (also called "Genesis" block), contains the hash of its previous (or parent) block.

Once a block is added to the blockchain, the block it's pointing to cannot be altered or deleted, making the ledger tamper-resistant and immutable and ensures data integrity, since, in order to change a block's information, one would need to change every child's block hash.

It's also possible to say that, the older the block, the more "immutable" it is, or the bigger the number of child blocks a block has, the harder is the tampering of such block. That is because the hash of a block depends on the hash of the previous block, and so on.

Figure 2.1 illustrates the blockchain datastructure, where it's possible to observe that the last block (Block N) points to it's parent, and so on. This data structure is often separated by two parts, the header of the block, that contains it's metadata, such as the previous block hash, and the data part, where, in most blockchain networks, this data is where the transactions are stored and visible to everyone that has access to it.

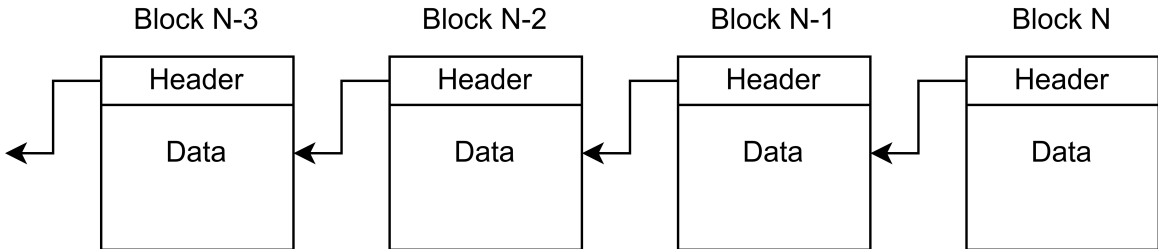


Figure 2.1: Blockchain Datastructure.

Since it is used to record transactions, a blockchain can be considered a “ledger”, and it’s transparent, as all transactions are recorded.

There are networks, like Bitcoin [3] and Ethereum [4] that store and replicate the blockchain (where the chain is the State Machine), with no central authority or entity controlling the network, like mentioned in the previous section. This makes the state of the blockchain highly available, since there are multiple nodes that contain its information.

The network reaches “common consent/decision” on the contents of a new block through the use of consensus mechanisms, such as Proof of Work or Proof of Stake, which will be described in the next section (2.2). The blockchain is then a distributed ledger, as all nodes in the network have a copy of the ledger, ensuring availability of the state.

The blockchain can be considered the state of the machine, where multiple copies of the state of this machine are stored in multiple nodes, and newer states, that is, the process of appending a new block to the chain, are replicated to every machine in the network.

Beyond cryptocurrency, blockchains have a wide range of potential applications in various industries, like supply chain management, voting systems, and identity verification. The transparent nature of blockchains ensures that all transactions are publicly accessible and verifiable.

In conclusion, blockchain is a data structure that leverages the concepts of state machine replication to provide availability, leverages the concepts of cryptographic concepts to make it secure, immutable and tamper-resistant ledger where all transactions are recorded, making it transparent. The decentralized nature of blockchain networks and the use of consensus mechanisms ensure that all nodes in the network have the same state of the chain, providing the availability to access of this information.

Network Models

Network models are fundamental to the design and implementation of distributed systems. The three main types- synchronous, asynchronous, and partially-synchronous — each have unique properties that influence the behavior of algorithms and protocols. Understanding these differences is essential for creating effective and efficient systems.

Synchronous networks are characterized by the fact that the entities in it have a common notion of time and upper bounds on message delay. That is, in the Synchronous model, a finite time limit Δ is established and known. The adversary can only delay the delivery of a message sent by at most Δ time. This means that the network can guarantee that messages will be delivered within a certain amount of time and that all nodes in the network have a consistent view of the current time. This type of network is relevant when strict timing constraints are necessary, such as in real-time systems.

Asynchronous networks [5], on the other hand, neither have a common notion of time

nor guarantee upper bounds on message delay. In this type of network, messages may be delayed indefinitely, and nodes in the network may have a different view of the current time. Asynchronous networks are relevant when timing constraints are not strict. The unpredictability of message delays makes asynchronous networks more challenging to design and implement, but they are also more robust and can handle failures and network partitions more effectively.

Partially-synchronous networks [6] are a hybrid of synchronous and asynchronous models. They have a common notion of time but only provide partial guarantees on message delay. In this type of network, some messages may be delivered within a certain amount of time, while others may be delayed indefinitely. Partially-synchronous networks are relevant in scenarios where some timing constraints are necessary, but not all and is often used in blockchain systems, such as Tendermint [7](used in Cosmos Network[8] and others), and Tezos' blockchain take on this protocol, Tenderbake [9], which balance the need for speed and reliability with the need for resilience and robustness.

In conclusion, the concept of network types is critical in the design and implementation of distributed systems. Understanding the differences between the three networks can help protocol designers make informed decisions about the type of network that best suits their needs and can ensure the success of their systems.

Consensus

Consensus in Distributed Systems refers to the process of achieving agreement/common consent among all participants in a network on the state of a shared information.

While State Machine Replication is the notion of broadcasting the update of a state and multiple entities maintaining a copy of it, consensus is the concept of how the nodes replicate and decide on a replicated value.

One way to implement total order broadcast, is by sending messages via a designated leader node, but if that leader becomes unavailable, the approach fails.

So, nodes have to reach an agreement together, in a decentralized manner, on what is the next state of the replicated machine (or who's going decide that next state), and should not be decided by a single node.

Consensus protocols are **critical** to the blockchain context because they are the mechanisms that allow the network participants to agree on the state of the distributed ledger. In this type of network, transactions are submitted and processed by nodes in a decentralized manner, and the consensus protocol ensures that all nodes have the same view of the ledger and agree on which transactions are valid and should be included in the next state. The construction of a block is used to make the system more efficient, by "packaging" multiple transactions together, but it's possible to reduce this idea to a single transaction being the update to the state. Also, the algorithms enable the selection of the next leader, that is, the node that takes the transactions, builds the block and broadcasts it to the networks.

This is important for maintaining the integrity and reliability of the blockchain, as well as for enabling trust in the network among participants who may not necessarily trust each other. A well-designed consensus protocol is essential for ensuring that the blockchain is able to process transactions efficiently and securely, even in the face of network failures, attacks, or other challenges.

Theorems

There are several important theorems related to consensus algorithms that are relevant to the context of Blockchain Networks, as these are used to attribute characteristics and properties to this type of networks (not only), which include the **FLP** Theorem [10] and the **CAP** Theorem [11].

The **FLP** Theorem states that, in an asynchronous network where messages may be delayed but not lost, it is impossible to achieve both agreement and termination in the presence of process failures. In other words, it's only possible to guarantee two of the following properties at the same time:

- **Finality or Agreement**, that is, if (functioning) the participants must agree on a single specific value.
- **Fault Tolerance or Integrity**, the system still functions in case of node failures.
- **Termination or Liveness**, where all the functioning nodes decide on value.

On the other hand, the **CAP** Theorem states that, in an asynchronous network where messages may be lost, it is impossible for a distributed system to simultaneously provide the following:

- **Consistency** - all nodes see the same data at the same time
- **Availability** - every request receives a response, without guarantee that it contains the most recent version of the data
- **Partition Tolerance** - the system continues to operate despite arbitrary partitioning due to network failures.

Figure 4.1 illustrates these theorems/properties, where it is possible to observe that the area where the 3 parts meet it is impossible to achieve.

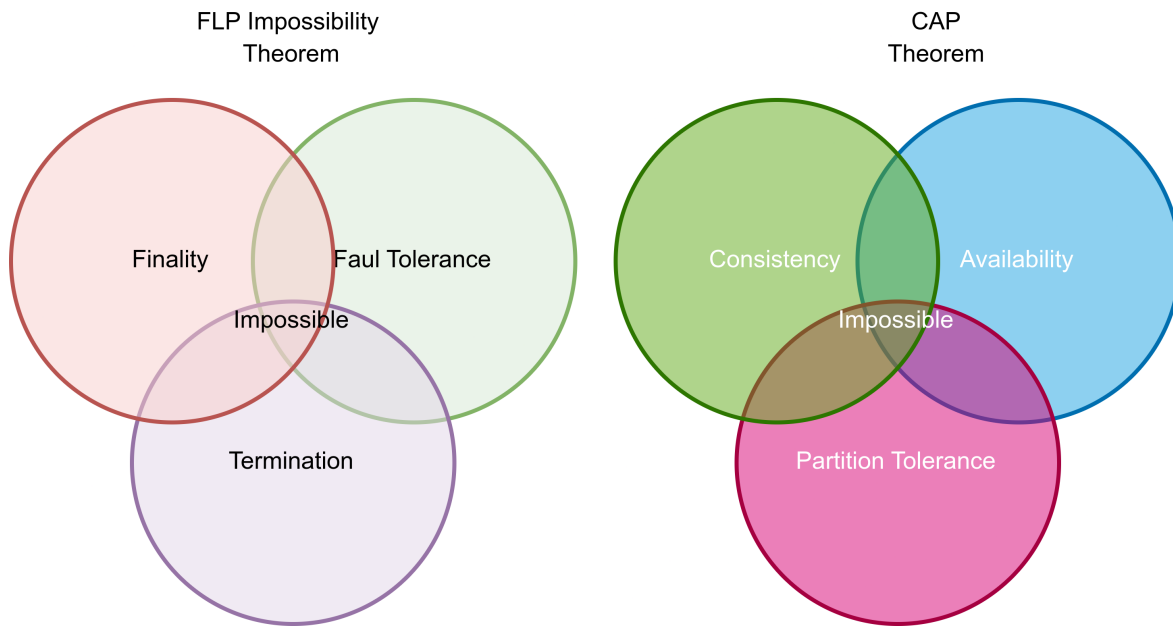


Figure 2.2: FLP and CAP theorems visualized

These (and other) theorems provide limits and trade-offs in decentralized systems, and serve as a useful reference for designers and researchers in the field. In the following sections, more insight about these properties will be provided and why they are relevant to the topic of blockchain networks, such as the fact that networks like Bitcoin chooses Availability over Consistency (while being Partition Tolerant) and chooses Termination over Finality (while being fault-tolerant and having probabilistic finality).

Permissioned and Permissionless Blockchains

Blockchains are a type of decentralized system that uses consensus algorithms to maintain the integrity of its data. In the context of blockchains, the terms “permissionless” and “permissioned” refer to the way in which participants in the network are allowed to participate in the validation of transactions and the creation of new blocks. These terms are critical in understanding the trade-offs between security, scalability, and decentralization in blockchains. In this section, we will delve deeper into what “permissionless” and “permissioned” mean and the advantages and disadvantages of each approach.

Permissionless blockchains, also known as public blockchains, are open to anyone and anyone can participate as a node. No central authority or entity is in control, making it a completely decentralized system. Permissionless blockchains are the most popular type of network and are most often used for cryptocurrencies such as Bitcoin, Tezos and others, where anyone can participate in the network and validate transactions, contributing to the security and reliability of the network.

Permissioned blockchains, are restricted to a set of trusted participants. Only approved entities are allowed to participate as nodes and validate transactions, making it a partially

decentralized system. This type of blockchain is often used in business environments where the participants are known and trusted. Usually, these networks are faster and more secure (yet lack decentralization), since the number of participating nodes is way smaller and the fact that nodes are “handpicked”, they are already trusted.

Known examples are Hyperledger-Fabric [12], an enterprise-grade network and Binance Smart Chain [13], a cryptocurrency network that’s a fork of Ethereum.

Processes of Block creation

In the context of blockchain networks, there’s a need for nodes selecting the next update to the state, or the next block to add to the blockchain structure. At each round/heartbeat of the network, a proposer decides on the structure of the block (in a cryptocurrency blockchain, the proposer selects the transactions to include in the block).

In this subsection, we will explore the two main ways of reaching consensus/selecting a leader in blockchain networks: proof-based consensus and committee-based consensus.

Proof-based consensus protocols rely on the concept of proof of leadership. Nodes are selected to generate new blocks through a cryptographic random algorithm. The selection of the new leader/proposer is similar to the idea of a lottery. The node that wins such lottery has the right to propose a new block. Nodes validate transactions, generate a Merkle Tree Root of these, and package them into a new block. The leader node broadcasts the new block and its proof of leadership to the network, that validates the new block, and upon validation, it appends it to the blockchain.

In **committee-based consensus** protocols, nodes vote to decide the next block to be appended to the blockchain. The proposer multicasts a preparation block request to other participants, who reply with their status. If the proposer receives a sufficient number of ready messages, it enters the pre-commit phase. Participants then broadcast their votes to commit the proposed block. If the number of commit responses agreeing to the new block exceeds a threshold, the block is appended to the blockchain.

Blockchain “trilemma”

Another subject that is usually used to compare different blockchain networks, with much discussion, is the topic of the “Blockchain trilemma” [14] that refers to the trade-off between scalability, security, and decentralization. This topic was first introduced by Ethereum’s founder *Vitalik Buterin*.

In other words, it is hard, if not impossible to achieve all these three characteristics at the “same time” in a public/permissionless blockchain network:

- **Decentralization:** refers to the distribution of power and decision-making authority across all participants in the network. In blockchain, this means that there is no central

authority or single point of control, allowing for a more equal and democratic system.

- **Security:** refers to the robustness and reliability of the network, ensuring that data and transactions are protected from tampering, hacking, or other forms of malicious activity.
- **Scalability:** refers to the ability of the network to handle a growing number of transactions and users, without sacrificing performance or speed.

Even though that in the context of most distributed systems the topic of decentralization isn't relevant, when talking about blockchain networks it is important, as it's one of the main reasons why people use these systems in the first place.

For example, Bitcoin is a highly decentralized and secure blockchain, as there are many nodes (more than 40 thousand [15], as of the writing of this document), yet lacks scalability, since the creation of a block takes 10min and a transaction takes about 1 to 1.5 hours to complete/be confirmed [16].

Therefore, when designing a blockchain network, it is important to understand the trade-offs between scalability, security, and decentralization and to make trade-offs based on the specific goals and requirements of the network.

Figure 2.3 illustrates this trilemma. It's common to assign blockchain network as a point in the area of this triangle. For example, since Bitcoin is secure and decentralized, yet lacks scalability, one could say that it stands near the line between "Secure" and "Decentralized".

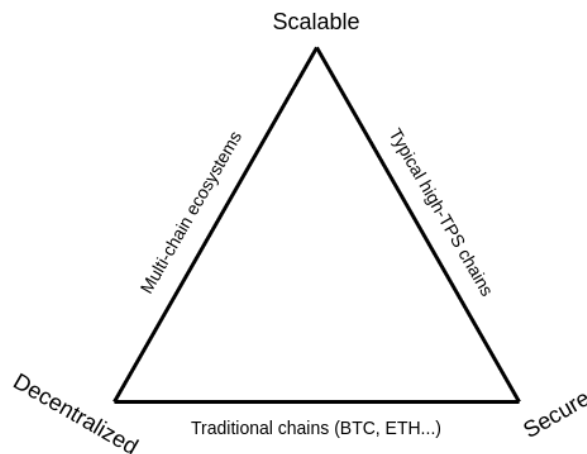


Figure 2.3: Blockchain trilemma

2.2 State of the Art

In this section, we focus on key consensus algorithms that underpin blockchain networks. These algorithms are essential for network integrity and have evolved to address the "Blockchain trilemma" and unique network requirements. We'll explore their strengths, weaknesses, and notable implementations.

This will further expose that there's way too many blockchain consensus protocols to choose from when creating a new network. The reader can obtain additional information about the topics covered in this section at the reference [17] and [18].

Proof of Work

Proof of work (PoW) [19] was originally introduced as a solution to the problem of spamming and abuse on the Internet. In a proof of work system, a computational problem is designed such that it requires a significant amount of computational effort to solve, but its solution can be easily verified. The idea is that, by requiring a party to perform a certain amount of computational work in order to perform a certain action, such as sending an email or posting a message, the system can ensure that the action is performed in a responsible and controlled manner.

The first use of proof of work was in the context of a system for fighting email spam, designed by computer scientist Adam Back in 1997. In this system, an email sender was required to perform a certain amount of computational work, in the form of solving a cryptographic puzzle, to send an email. The idea was that the cost of performing the work would make it unfeasible for spammers to send large quantities of unsolicited email, while still allowing legitimate users to send email without significant difficulty.

Proof of Work was later adapted as a consensus algorithm used by many blockchain networks, Bitcoin [3] being the most popular network to do so. Bitcoin or “Nakamoto Consensus” was also the first instance of the Proof of Work Consensus and of a blockchain network.

We will explain in detail this protocol since it was the first one to be used in a blockchain structure, because it is still one of the most relevant protocols used and we implemented it during the course of this project.

The idea behind PoW (applied to blockchain networks) is that, in order for a block (containing a list of transactions) to be added to the blockchain (and replicated), at least one node in the network has to package together a list of transactions into a block and solve a cryptographic hard problem in order for it to be valid. That is, the block has to contain a stamp or a proof that actual computational work was put in to create a block. That work/computational power is provided by nodes called “miners”, and these miners compete to be the first to solve the cryptographic puzzle, where the first to solve it is selected as the leader and has the ability to add the their created block to the blockchain.

In a permissionless environment, it's complex to make a democratic system where each entity could cast a vote, since a malicious entity could create many fake identities or nodes to manipulate the network, known as a “sybil attack”. In blockchain networks, entities make use of Public-key cryptography to be identified, so a malicious entity could create as many pairs of Public-Private keys as it wished to cast as many votes as it wanted.

The “Nakamoto Consensus” (where PoW is used) was the first to be tolerant to sybil attacks in a permissionless setting. It prevents these attacks by making it computationally expensive for

an attacker to try to disrupt the network, since, instead of voting power being concentrated on the number of votes for an election, it's replaced by the idea of computational power, where an attacker would have to do as much work as the rest of the network (assuming that it's not malicious) in order to gain a majority, and that could imply problems to the malicious entity, like large electricity costs in case of mining. By requiring a significant amount of computational power to mine blocks, PoW ensures that only legitimate nodes with real computational resources will be able to participate in the network. Mining is what is called coequally to the process of using computational power to solve cryptographic puzzles and adding blocks to the blockchain.

In other words, instead of entities voting for a block to be added to the blockchain, the "vote" is done by solving puzzles, where the first to solve such puzzle is the one deciding the block to be added. Another analogy can be that the winner of the "lottery" gets the prize of deciding the next block (solving the puzzle is actually a random chance process, where increasing the mining power also increases the chance of winning).

Mining has several purposes in PoW, such as leader selection and preventing sybil attacks, like mentioned before, but also enforcing block timing with difficulty adjustment and, in case of cryptocurrency blockchains, to add new value/mint new currency into the system and rewarding those who do the computational work.

The difficulty adjustment ensures that the time between blocks is consistent (in Bitcoin, that is 10min), and the computational difficulty to mine new blocks adjusts accordingly. When in a epoch (in Bitcoin, that is 2016 blocks), the median time taken to mine a block was lower than it should've been, the difficulty is increased proportionally to make it harder to mine, so it actually takes the pretended time, and vice versa.

The mining also introduces/mints new currency into the system, since itself doesn't have any value in it initially (it's a closed economic system, no more currency can be put into it), and that currency is awarded to the participants of the mining process.

Properties of Proof of Work Consensus

The properties of PoW consensus can be analyzed through the FLP theorem and the CAP theorem, which were discussed in the previous section.

Regarding the FLP theorem, PoW exhibits Probabilistic Finality [20]. This means that all functioning nodes eventually agree on a single block, but the process is not deterministic and involves a certain degree of probability. The waiting for confirmations adds an additional layer of security to the consensus process. Like mentioned before in this example (2.1) in Bitcoin this takes about 6 confirmations, or about 1 hour, since a block usually takes 10 minutes to be created. This happens because mining is a process independent from the state of the network and independent from other nodes, and two or more nodes can solve the cryptographic puzzle/build a block "at the same time" (until the whole network agrees on the same block). When this happens, a node can receive more than one valid candidate ex-

tending the same chain, and each protocol handles this in a specific way. This is what is called a “Fork”. Once again, taking Bitcoin as an example, a node handles this case by using the Longest Chain Rule, where, when a node has two competing blocks, where each makes a branch, it waits until one of the branch is bigger than the other.

This is done because, the branch with the most computational power has a greater chance of being extended.

PoW also demonstrates Fault Tolerance, making it a form of Byzantine Fault Tolerance (BFT). This is achieved through the complex and expensive mining process, which makes it economically infeasible for a malicious node to alter the data in a block and catch up with the rest of the network. The incentive to follow the longest blockchain also adds to the fault tolerance of the consensus mechanism.

Finally, PoW also exhibits Termination, meaning that every functioning node reaches a decision, even if it’s not the final one (because of the Probabilistic Finality).

Regarding the CAP theorem, PoW may not provide Consistent results, as a node might not have the latest block when requested, and like mentioned before, forks can happen, and so different nodes can contain different chains (for a period). However, it does provide Availability, as miners are continuously trying to mine new blocks, the difficulty level of mining is adjusted to ensure the steady addition of blocks to the blockchain and there are multiple nodes in the network. PoW is also Partition Tolerant, meaning that even if a portion of the nodes stop functioning or the network splits, the blockchain can still function and recover.

Nakamoto Consensus

In the previous subsection, we discussed the consensus mechanism used in blockchain technology where participants in the network compete to solve complex mathematical puzzles in order to validate transactions and generate new blocks. We introduced the concept of miners, who play a crucial role in this process, and how they are incentivized through rewards for their efforts.

Building on this foundation, we now delve into the specifics of the Nakamoto Consensus, named after the pseudonym used by the unknown creator(s) of Bitcoin (Satoshi Nakamoto). The Nakamoto Consensus is a milestone in the history of blockchain technology and a major innovation in the field of consensus algorithms.

The Bitcoin network is a Peer-to-Peer network, where blocks and transactions are gossiped by the nodes and where they are also responsible for generating blocks, just like mentioned previously.

This acts as a barrier to entry for malicious actors, as it becomes extremely difficult for them to launch Sybil attacks by generating many fake nodes. The puzzle in Bitcoin involves computing a nonce (basically, a number) that meets certain conditions, as outlined by the equation $H(h_{i-1}, nonce, tx, par) < Target$, where “ h_{i-1} ” is the hash of the previous block, “ tx ”

represents the set of validated transactions, and “par” represents other parameters such as blockchain version and cryptographic parameters, and the “Target” is a value decided by the network (that encoded the current difficulty).

The usual workflow of a node is the following:

- Node gossip transactions in the network (They are temporally stored in a data structure that is called “Mempool”).
- When a node receives a block from the network, it appends it to its own blockchain stored.
- To start the mining process, a node selects multiple transactions for the transaction part of the block. There’s a size limit in bytes for this part, and the selection criteria is irrelevant, they only have to be valid.
- It creates a new block header with information, like, the hash of the previous block, a nonce (a integer number) and a Merkle Tree Root Hash of the transactions and other information (that is irrelevant to this).
- The node tries out every possible value for the nonce until the resulting hash (SHA256 hash function) of the header has a value below the current target.
- If a node receives a new block from the network before it finds itself the nonce, it stops the mining and appends such block. If it finds the nonce, it shares the block with the network and it’s rewarded.

When a entity creates a transactions, it can also offer a fee. This is used to prevent transactions spamming and also acts as a “place in the block” auction. Besides block creation reward, the miner also receives these transactions fees, so they will select transactions based on that fee, and other attributes of the transactions.

The difficulty/target of the puzzle is adjusted periodically (every 2016 blocks like mentioned before), depending on the actual block generation intervals and the expected block generation intervals (of around ten minutes in Bitcoin). This allows Bitcoin to maintain a consistent block generation time of around ten minutes, regardless of the computational power in the network.

Improving Proof of Work

The Nakamoto Consensus protocol, despite being the first consensus mechanism for blockchains and widely adopted, has limitations when it comes to scalability, security, and decentralization.

There are multiple ways of improving the Nakamoto Consensus protocol, and doing so has resulted in the creation of new blockchain networks. These changes move the consensus pro-

protocol within the triangle formed by the “trilemma”, but they also come with new challenges compared to the original consensus for blockchains.

This topic is relevant to this thesis since it’s these characteristics that differentiate consensus even further.

Improving Scalability

Scalability remains a challenge in blockchain networks due to the need for security and decentralization. Various solutions aim to enhance transaction speed and network scalability: Decoupling of Blockchain Functions, Parallel Chains, and DAG-based Protocols.

Decoupling of Blockchain Functions: This approach separates block functions into key blocks for leader election and microblocks for transactions. It boosts network throughput but has latency issues and risks like leader compromise.

Parallel Chains: Miners extend multiple chains concurrently to increase throughput. However, this method is susceptible to targeted attacks on individual chains.

DAG-based Protocols: These protocols use a Directed Acyclic Graph (DAG) instead of a linear blockchain, allowing for concurrent block generation and a unique approach to scalability.

Improving Security

Blockchain networks are vulnerable to various security attacks, such as selfish mining attacks [21] and double-spending attacks [22]. To improve the security of blockchain networks, various solutions have been proposed, including changing the incentive mechanism in the chain and how it’s shared.

The FruitChain [23] protocol is a blockchain protocol that improves on this while still providing the same consistency and liveness properties as Nakamoto consensus. With FruitChain, any honest set of nodes controlling a certain fraction of computational power is guaranteed to get at least that fraction of the block generation and rewards in any length segment of the chain, making it harder for coalitions between miners to happen.

Improving Decentralization

To improve decentralization, several solutions have been proposed, including de-incentivizing centralized activities like pool mining and incentivizing decentralized/solo mining using methods like eradicating ASIC (Application-Specific Integrated Circuits) mining.

Pool mining is a point of centralization in Proof of Work networks. The idea of pool mining is that there’s a single node (connected in the network) that communicates with multiple miners, that is, entities which their only purpose is to find a nonce. These miners join to

mine a single block as if they were only a single node, so they have more computational power together, and inherently having a bigger chance of finding a block. When a block is found by the pool (a node, as viewed from the network), the reward is shared with all the miners, depending on how much effort they put into the mining. The problem comes with the fact that most mining nowadays, like in Bitcoin, is done by groups like this, and so, the owners of these pool nodes have a lot of power over the network, as they can make decisions on the block creation, like black-listing transactions from someone, or selecting transactions based on unfair criteria.

Also, with the increase of hashing power, that is, the computational power of blockchains, miners have improved the mining process by developing highly specialized hardware, called ASICs.

By doing so, it's infeasible for anyone (without this type of hardware) to join the mining process, and centralizing the power on the owners of ASICs and their manufacturers. So there are multiple efforts to eradicate this type of mining like by using memory-hard puzzles, that are harder for ASIC hardware than for Personal Computer Hardware (like GPUs or CPUs), or using hash functions that are hard to implement as ASICs.

This is the case of Monero [24], that besides providing total privacy (the ledger isn't transparent like Bitcoin), is also designed to promote equality between miners by using an egalitarian approach based on the use of a different hash function (compared to Bitcoin, that uses SHA256) called "CryptoNight", which is said to be memory-hard, making it hard for the development of ASIC hardware and also hard for mining using GPUs.

Proof of Stake and Its Variants

Proof of Stake (PoS) has emerged as a popular alternative to Proof of Work (PoW). In PoS, nodes are selected to validate transactions and create new blocks based on their stake in the network. The core idea is similar to PoW but replaces "computational power" with "stake," making it more energy-efficient [?].

Chain-based Proof of Stake (ChainPoS): Focuses on incentivizing stakeholders to act honestly by staking assets. It offers two main approaches for block creator selection: randomized stakeholder selection and a hybrid between PoW-like selection and randomized stakeholder selection. Peercoin [25] and Casper FFG [26] are notable examples.

BFT Proof of Stake (BFT PoS): Developed to tackle Byzantine faults in decentralized systems. It operates in a multi-round consensus process involving validator selection and the use of quorums of super-majority. Tendermint and Tenderbake in the Tezos blockchain are examples [27].

Delegated Proof of Stake (DPoS): Combines elements of PoS and democratic governance. It aims to balance decentralization with efficiency by having a smaller number of delegates. Validators are selected based on reputation scores or other metrics, and stakeholders can vote for delegates.

In conclusion, PoS and its variants offer a range of options that balance security, scalability, and efficiency, each with its own set of trade-offs and benefits.

Proof of Authority and Its Variants

Proof of Authority (PoA) is a consensus algorithm commonly used in permissioned blockchain networks, where decentralization is not the primary concern. In PoA, a limited set of pre-approved validators are responsible for block creation and transaction validation. This approach offers a high degree of efficiency and scalability while eliminating the need for resource-intensive calculations. Interestingly, we also implemented a PoA, that has certain similarities with PoW, a topic that will be explored further in this document.

Clique PoA: This variant aims to improve the decentralization aspect of PoA by allowing multiple authorities to propose blocks in a round-robin manner. This reduces the risk of a single point of failure and enhances network security [28].

Aura PoA: Known for its flexibility, it allows validators to take turns in a defined order to create blocks. It also includes a mechanism for validators to communicate and come to a consensus if a block is valid or not [29].

IBFT PoA (Istanbul BFT): A Byzantine-fault-tolerant PoA algorithm that can handle malicious nodes and network splits. It uses a voting mechanism among validators to reach consensus and is known for its robustness [30].

In conclusion, PoA serves as an effective choice for networks where rapid transaction validation is crucial, and where a permissioned structure is acceptable. Its variants offer different ways to balance efficiency with security, making PoA adaptable to various specific needs.

Other Consensus Protocols

Apart from protocols mentioned before, there are other consensus algorithms that are used in the blockchain industry. However, they are not as popular as PoW and PoS and are mainly used in niche applications or not used in the popular cryptocurrency blockchains. In this section, we will explore and explain some of these alternative consensus algorithms.

Proof of Burn [31] shares some characteristics with Chain-based Proof of Stake. In this algorithm, instead of staking, the node sends some tokens to an address that is difficult to access, essentially "burning" them. The node can then mine or validate blocks proportionate to the amount of tokens that were burned. The idea behind Proof of Burn is that a node is willing to sacrifice tokens, making it less likely that they would act maliciously as they would lose their investment.

Proof of Research [32] algorithm is used in some cryptocurrencies that aim to provide an environmentally friendly alternative to Proof of Work. In this algorithm, nodes perform research tasks that can range from solving mathematical problems to annotating data. The

first node to complete the task is allowed to mine the block, providing a solution to the computational waste that is inherent in PoW.

Proof of Capacity/Space/Storage [33] shares some characteristics with Proof of Work. In this algorithm, nodes are required to dedicate a certain amount of disk space for storage. The more disk space a node dedicates, the more chances it has of being selected to mine the next block. The algorithm is energy-efficient compared to PoW, but it still suffers from centralization problems as nodes with more disk space are more likely to be selected to mine blocks. These kinds of protocols are often used in conjunction with other blockchain networks (Proof of Work or Proof of Stake) to store large amounts of information that wouldn't be feasible to store directly in those blockchains.

Trusted Computing [34] algorithms use the secure hardware provided by Trusted Execution Environments (TEEs) to secure the consensus mechanism. One example of such an algorithm is POET, used in the Hyperledger blockchain platform. In POET, a trusted party generates a wait time for each node and the first node to complete the wait time is allowed to mine the next block. The use of TEEs ensures that the wait time cannot be tampered with, providing a secure mechanism for consensus.

There are many alternative consensus algorithms that are used in the blockchain industry, each with its strengths and weaknesses. While some, like Proof of Burn and Proof of Research, offer environmentally friendly alternatives, others, like Proof of Capacity/Space/Storage, offer a totally different concept of proof. Trusted Computing Algorithms, like POET, provide secure consensus mechanisms, but rely on trusted hardware that may not be readily available and making it more centralized, as it relies on specific companies to build this kind of hardware. Ultimately, the choice of consensus algorithm will depend on the specific use case and requirements of the blockchain platform.

Lupin Language

The Lupin Language is a Domain-Specific-Language (DSL) that is currently in its early stages of development. The process of compilation of this language results in OCaml language code. It simplifies the complex task of implementing consensus algorithms, saving expertise and resources.

Not only that, but it can also be used to perform rigorous formal analysis and verification on these algorithms, bringing a whole new level of capability to analyzing consensus algorithms and will be an important tool in the future.

This project will provide feedback for the development of Lupin and may also serve as a layer of execution for the results of its compilation, and is in collaboration with the Lupin Language project and its results will be further exposed in a later section.

2.3 Conclusion

This chapter provides a comprehensive overview of the core concepts of blockchains and the state of the art in consensus algorithms. With so many different consensus protocols available, it is essential to have a clear and concise way to compare and evaluate them. As the field of blockchain continues to evolve, the development of more effective consensus algorithms will be a critical factor in determining the success of blockchain networks. By gaining a deeper understanding of the current state of the art, researchers and developers can work towards developing new consensus algorithms that are more secure, efficient, and scalable than ever before.

Chapter 3

Preliminary Case Study

In this chapter, we focus on a hands-on experiment conducted on the Tezos blockchain. We chose Tezos for its inherent flexibility and adaptability, qualities that make it an ideal testbed for exploring the implementation of various consensus protocols. The primary aim of this experiment is twofold: to gain a deeper understanding of how consensus protocols can be generalized and tested, and to delve into the specifics of Tezos' architecture and how it accommodates different protocols.

Our experiment focuses on implementing Bitcoin-like Proof of Work in Tezos. This hands-on approach helps us explore the intricacies of adding new consensus rules to an established blockchain. We'll outline the requirements, integration steps, and tools involved.

Key components like the client and baker are also discussed. The client acts as a network information hub, while the baker handles block creation and mining. Special attention is given to the mining mechanism, specifically nonce discovery to meet hash conditions.

In short, this chapter showcases Tezos as a flexible platform for testing various consensus models. Our Proof of Work implementation serves both as a proof of concept and an indicator of its adaptability to diverse consensus methods.

3.1 Tezos Blockchain as a tool for the Experiment

The choice of Tezos as the blockchain network for this experiment is a deliberate one, as it aligns well with the goals of the solution presented in the previous section.

To understand the choice of this blockchain network, it's crucial to delve into its unique features.

Launched in 2018, Tezos was designed as a self-amending crypto-ledger, capable of evolving without disrupting its community. This sets it apart from other networks that face challenges when upgrading their technology.

The platform is not only secure but also flexible and upgradable, designed to allow protocol changes without requiring a hard fork. While it originally operates on a Delegated Proof of Stake consensus, this is irrelevant for our experiment as we're replacing it with different consensus rules.

A hard fork, for context, is a permanent split in a cryptocurrency's blockchain, often due to changes in network rules or disagreements within the community. For example, a dispute over Bitcoin's block size led to the creation of Bitcoin Cash 1.2.

Compared to other options, this network is an excellent fit for our experiment. It's a well-tested, industry-grade platform with an upgradable architecture. The protocol layer is stateless and independent of other node components like the Peer-to-Peer layer and client, allowing us to focus solely on the consensus rules.

In summary, the choice of this platform was strategic, offering the features needed for a successful experiment. Its upgradable nature and stateless protocol layer allow us to concentrate on consensus rules. The experiment's results will shed light on the platform's adaptability and potential for supporting various consensus models.

3.2 How Tezos is structured and How it allows self-amending

Tezos has a unique structure, which separates the protocol, or how the project calls it, "Economic Protocol" from the rest of the node, known as the shell. The diagram 3.1 illustrates the node architecture, and it's explained in greater detail here [35].

The protocol is responsible for interpreting transactions and administrative operations, as well as detecting erroneous blocks, everything else is handled by the shell.

The shell includes the part of the validation code, the peer-to-peer layer, the disk storage of blocks, and the versioned state of the ledger. The distributed database abstracts the fetching and replication of new chain data to the validator.

The economic protocol on the chain is subject to an amendment procedure, which allows for the on-chain operations to switch from one protocol to another.

Finally, the RPC (Remote Procedure Call) layer is an important part of the node, allowing clients, other nodes, third-party applications, etc. to interact with the node and introspect its state using the JSON format and HTTP protocol.

In other words, the outside world of the node communicates with the shell, via RPC layer calls, and the shell communicates with the protocol. As previously mentioned, the protocol is stateless, and is only used to preform the logic part of the consensus, that is, to verify transactions, blocks and to describe what the shell should do when receiving newer information.

The communication between the shell and the protocol is based on an interface called "Updater.PROTOCOL". The protocol component is restricted to a specific environment that restricts the access to a defined set of OCaml modules. This is to improve the security of the protocol. The protocol must implement the "Updater.PROTOCOL" interface (defined in the environment, and later explained in more detail) in order to interact with the shell. The interface requires the protocol to define protocol-specific types for operations/transactions and block header, along with encoders/decoders for the types.

The protocol must also provide functions for processing blocks and updating the context, which represents the protocol state. The context is stored as a disk-based immutable key-value store. The block header and operations also have a shell header (protocol-independent)

and a protocol-specific header. For example, a Block can have two levels, one corresponding to the overall/shell level, and other is the level of the block since the current protocol is started.

Blocks in Tezos have a generic part, understood by any protocol, and a protocol-specific part.

A node can contain multiple economic protocols, but only one of them is activated at any given time, and it always starts with the “genesis” protocol.

Some protocols are linked to the node at the time of compilation, that is, the node already has information about it when starting, and some protocols can also be registered dynamically at runtime via an RPC, allowing for protocol updates without the need of manual intervention to do so.

Protocol activation in Tezos is a two-step process.

1. A command injects an “activation block” to the blockchain. This block contains the hash of the protocol to be used, and only one operation to activate it. The activation block is the only block using the genesis protocol, as this protocol doesn’t contain any other functionality besides the activation of a different protocol.
2. The next block in the blockchain will be the first block using the activated protocol.

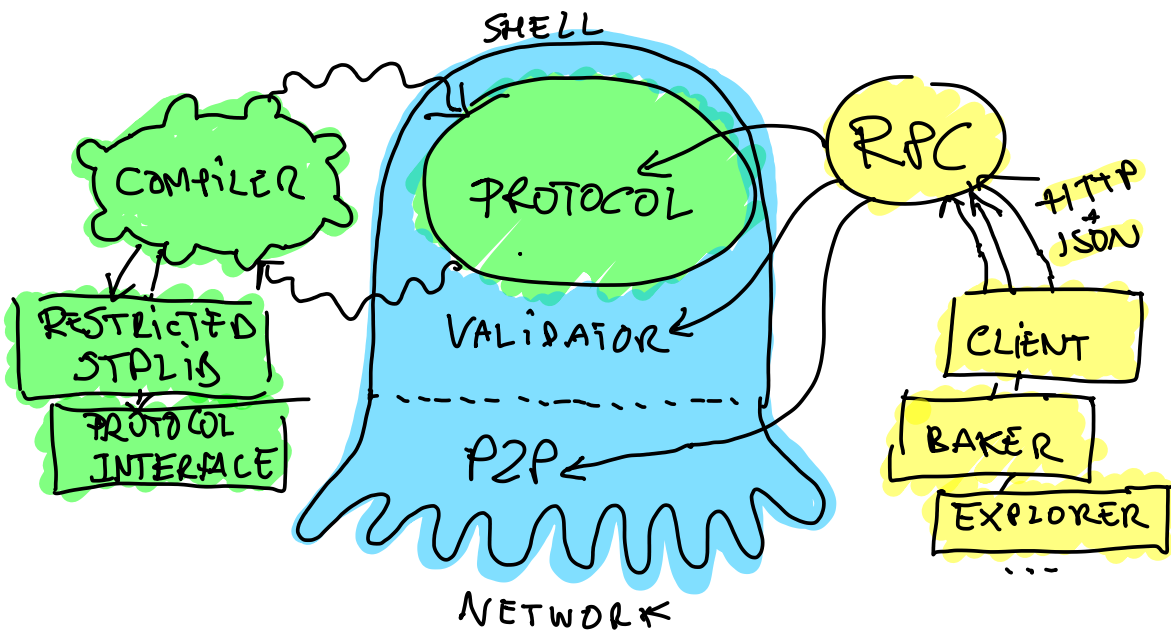


Figure 3.1: Architecture of Tezos (Source: Tezos Documentation)

In conclusion, the structure of Tezos is unique in that it separates the protocol, also known as the “Economic Protocol”, from the rest of the node, called the shell. The protocol is responsible for interpreting transactions and administrative operations, while the shell includes the validator, peer-to-peer layer, disk storage of blocks, and versioned state of the ledger. The economic protocol on the chain is subject to amendment procedures, allowing for on-chain operations to switch from one protocol to another. The RPC layer allows clients, third-party

applications, and daemons to interact with the node and its state. The communication between the shell and the protocol is based on the `Updater.PROTOCOL` interface, and the protocol is restricted to a specific environment to improve security. A Tezos node can contain multiple economic protocols, but only one of them is activated at any given time through a two-step activation process.

3.3 Implementation of a Protocol

Implementing a Proof of Work consensus algorithm as a protocol for Tezos was motivated by several reasons. Firstly, it was an opportunity for hands-on learning about how protocols are implemented and structured. This understanding is crucial for later work, such as testing and adding an adaptive layer to the DSL mentioned in previous sections. It was necessary to understand how these protocols work before attempting to create tools related to them. It also allowed for a deeper understanding of how Tezos works as a whole.

The implementation of Proof of Work was also significant because it provides a contrast to Tezos' original Proof of Stake protocol. This demonstrated that Tezos can be used to implement other types of consensus algorithms, including those that are different from the original. The implementation of Proof of Work also has the concept of mining, which is not present in the Proof of Stake protocol.

Additionally, implementing a protocol in Tezos is a big accomplishment. The fact that only a few people do so makes this achievement even more significant, especially considering that it's a different type of consensus. This demonstrates a deep understanding of consensus algorithms and the ability to put that knowledge into practice.

In conclusion, the implementation of Proof of Work as a protocol for Tezos was a valuable learning experience that allowed for a deeper understanding of how protocols are structured and executed in Tezos. It also reinforced the concept of consensus algorithms and demonstrated the ability to put that knowledge into practice by implementing a unique protocol in Tezos.

Requirements of a Protocol In Tezos

In order to implement the experiment of a Proof of Work consensus algorithm for Tezos, a `lib_protocol` module in the protocol folder had to be created, where this is the main entry to the protocol and is executed by the node as the new amendment/consensus protocol. This is the only part that has to be implemented in OCaml and has to be compiled to work with the node.

To do so, a `TEZOS_PROTOCOL` file must be included in this module, that is used to specify the version of the environment that the protocol is to be compiled against, which in this case the version 6 was the one used, the hash of the protocol folder and the set of modules

implemented in the folder.

There's the possibility to implement a newer environment for the protocol to accommodate the protocol, yet, for this experiment, this wasn't necessary.

Like said previously, the environment is an interface that provides a set of OCaml modules that the protocol can use, and the interface used to interact with the Tezos Shell.

The most important functions and types in that had to be implemented for environment "v6" were:

- *block_header_data*: This is the protocol-specific part of the header.
- *block_header*: This is the combination of the protocol header and the shell header.
- *operation_data*: This is the protocol-specific part of an operation/transaction.
- *operation*: This is the combination of the protocol operation part and the shell operation part.
- *validation_state*: This is the state of the validation of a block or operation, passed between functions.
- *init*: This function is executed to prepare the chain to start executing the new protocol.
- *begin_application*: This function is used when validating a block received from the network.
- *begin_partial_application*: This function is used when the shell receives a block more than one level ahead of the current head.
- *begin_construction*: This function is used by the shell when instructed to build a block and for validating operations as they are gossiped on the network.
- *apply_operation*: This function is called after *begin_application* or *begin_construction* and before *finalize_block* for each operation in the block or in the mempool, respectively. It validates the operation and updates the intermediary state accordingly.
- *finalize_block*: This function represents the last step in a block validation sequence.

In summary, the "TEZOS_PROTOCOL" file is used to specify the protocol environment to be used by the Tezos node, and the environment provides a set of functions and types that the protocol must implement to be able to operate within the Tezos network.

Implementation of the module "lib_protocol"

The implementation of a new protocol in the Tezos codebase is a significant challenge. In this particular case, the goal was not only to implement the new protocol, but to also learn how

existing protocols are structured and implemented within the Tezos codebase. The approach was to study and follow closely the implementation of existing protocols, rather than taking an easier approach that may also have resulted in a working protocol but would not have offered the same level of learning.

The protocol that was implemented includes a few functionalities, such as transactions between two entities, where “Tez”, the currency used in Tezos, is transferred from one account to another.

Another aspect of the protocol is the concept of mining and verifying the Proof of Work, like mentioned in a previous section about it. This is done by hashing the header whole header, both the shell part and the protocol part, this last one containing the nonce, and verifying that the hash value is lower than the target where this is one of the main mechanism behind proof of work.

The environment that the protocol is compiled against requires a definition of the protocol header, which contains three elements: “target”, “nonce”, and “miner”.

The **target** is used for comparison with the target value that the node calculated it should be (the specifics of how this is done is explained in a later point). The **nonce** is used to achieve a header hash with a value lower than the target, which is a key part of the Proof of Work mechanism. Finally, the **miner** field stores the address of the miner who found the nonce/hash, which is used to reward the miner. These elements combined make up the protocol-specific part of the header.

Implementation of Information Representation and Storage Logic

The implementation of the protocol in Tezos involves defining and representing the various types of information that are necessary for the protocol to function. This includes things like accounts, constants, headers, time, target, currency (Tez), and operations.

The information stored in the protocol is abstracted, meaning that the protocol itself is stateless. The context in which the protocol is applied maps to the actual storage of the blockchain, but this is hidden from the protocol. The protocol only sees the storage as a generic key-value store and the functions that access this are defined in the “raw_context.ml” file (in reality, they are accessed as a Tree, not just as a generic key-value). These functions allow the protocol to read, write, and update the storage, but the actual details of how this is done are not relevant to this document.

Also, as mentioned previously, the shell part of Tezos doesn’t necessarily “understand” the protocol, as the protocol is somewhat independent from it. So the way that Tezos handles this fact is by representing the protocol parts of data (either from Blocks, Storage or Operations) as bytes, then these bytes are decoded by the protocol.

The encoding and decoding process of this information is done using the “data-encoding” library in Tezos. This library is also used to encode and decode OCaml types to and/from

bytes or JSON, so it can also be sent to other nodes in the network.

The Tezos Protocol is implemented through a series of files that represent various components and functionalities. In this section, we will take a closer look at the information stored and the logic behind each component in the protocol.

The protocol contains representations/definitions and operations of the following types/information, which are stored in files that end with “repr.ml”. The Tezos documentation calls this the “representation layer” of the protocol. Unless the representations are only used within the protocol code, they also implement the encoding and decoding logic so this information can be sent/received, either from/to the network or the shell part of the node.

- **Accounts/Manager:** This represents the concept of an account in the protocol. An account is just a Public Key Hash, which is used to fetch information from the storage.
- **Constants/Parameters:** This file contains information that is constant throughout the execution of the protocol. Some constants can be defined when activating the protocol. The constants include:
 - *block_time*: Defines the time between blocks.
 - *initial_target*: Defines the first target value.
 - *difficulty_adjust_epoch_size*: Defines the number of blocks to wait to then readjust the target.
 - *halving_epoch_size*: Defines the number of blocks to wait to then halve the mining reward.
 - *reward_multiplier*: The initial reward, which then gets halved.
 - *Header*: Defines the Protocol header, as mentioned before.
 - *block_time*: Defines the time between blocks.

These parameters are crucial for achieving the project’s objective, allowing for more precise tuning of the protocol.

- **Time:** Type definition of “time” in the protocol.
- **Target:** Defines how the target, used in Proof of Work is represented, in this case, it’s a 256-bit number.
- **Tez:** Defines the type and operations for the currency.
- **Operations:** Representations and functionalities of the available operations. These include Transactions (between two entities) and Reveal (maps a Public Key to a Public Key Hash, like the ones used in the Accounts).

Some of these representations were designed by following the existing protocols implemented in the Tezos codebase. These representations ensure that the creation, conversion, encoding and decoding and other functionalities are done correctly by the protocol.

The logic behind storing information in the Tezos Protocol is an integral part of its functionality. The protocol uses an abstraction over the blockchain as its storage mechanism.

This is really important to the implementation of the protocol. While implementing a protocol, the developer doesn't have to understand how the blockchain will be stored, what optimizations for inserting and retrieving information are done, etc.

The files that are responsible for storing information have names ending with “-storage.ml”, and they store the following key pieces of information:

- **Accounts:** It stores information about each account in the protocol. It maps the account's key to the account's balance and other important information and how defines how this information should be stored, retrieved and other checks.
- **Target:** It stores the current target, which is used for comparing the value of the header hash in the proof of work process.
- **Epoch Time:** It stores the timestamp of the latest difficulty adjust epoch, which is used to determine when it's time to adjust the target value. This is an important part of maintaining the security and integrity of the network, as it allows the network to dynamically adjust the difficulty of mining to ensure a stable rate of block production.

By storing these types of information in the blockchain, the Tezos Protocol provides a secure and transparent method of tracking the state of the network, which is crucial for the proper functioning of its operations. The protocol also provides a well-defined structure for the information it stores, which ensures consistency and maintainability of the code.

From the Protocol side of view, it sees the Blockchain as a Map/Tree, where it can fetch and append information to it using the typical concept of a Map (in programming).

Implementation of the main entry points to the protocol

The implementation of the main functions of the protocol plays a crucial role in ensuring the correct execution of the protocol.

These functions are designed to take the necessary steps to apply the information and verify the validity of the information being applied. The functions presented in the previous section, such as *begin_application*, *begin_partial_application*, etc. are all a part of the main functions of the protocol.

These functions perform various checks and updates to the context, which is an immutable representation of the state of the protocol. The context contains information such as the accounts, the target, and the epoch time, among others. The functions return some form of validation state and updated context, if applicable.

Once again, the protocol is independent from the rest, so these functions are implemented with this concept. That is, they are called by the shell part of the node, by passing information

such as the current “context” and some information about the previous block (to the one being built or validated).

- *begin_application* and *begin_partial_application*:
 - These functions prepare the current raw context, which contains information such as the context and more. The concept of raw context will be explained later.
 - Checks are performed, such as verifying if the current target is the correct one and if the header has a valid hash, that is, if the hash of the whole blockheader has a value lower than the target.
 - They update information such as the target for the next epoch, if an epoch has elapsed.
 - These functions are called when receiving a block from the network. The partial variant is used when the block received doesn’t have the successor level to the current block.
- *begin_construction*:
 - Does the same preparation to the context as *begin_application* and *begin_partial_application*.
 - Performs the same checks, but also rewards the miner if the block is valid.
 - This function is called when the node is building a node. For example, if the node is a mining node.
- *apply_operation*:
 - Verifies if the operation contains a valid signature.
 - Verifies if the operation has a positive result, for example, if the account has enough funds to transfer to another account.
 - If the operation is positive, it executes the operation on the current context.
 - This function is used after both *begin_application* and *begin_construction*.
- *finalize_block*:
 - Commits the change to the shell.
 - Returns a receipt or result log that exposes what happened with the application of either a block or operation.
- *init*:
 - Initializes everything needed to start the protocol, such as the constants and storage.
 - Also creates the first block of the protocol to be appended to the chain.

In high-level view, the way that these functions are executed is the following:

We assume that every function is independent and purely functional. Each call of any of these functions receive and return (back to the shell) an updated-copy of a value called “context”, in other words, it updates the state of the blockchain, unless something goes wrong, which then the updated context gets ignored. The shell handles the flow of execution. It starts with *begin_application*, *begin_partial_application* or *begin_construction*, depending of the conditions mentioned above. They handle the initial validation of the block header.

Then, for every operation included in the block, *apply_operation* is executed, copy-updating the context and passing it to subsequent calls of this function, until the last operation. Here, operations update the state and are validates.

After the last operation is validated, the context is finally passed to *finalize_block*, where the last updates to the context are performed and the last validations are done. After this function is called (and depending on the result of this process), the block is committed back to the shell, which is then added to the underlying blockchain.

Implementation of Alpha and Raw Contexts

The implementation of the protocol relies heavily on two core abstractions: **Alpha_context** and **Raw_context**. These two modules play crucial roles in ensuring the separation of concerns and allowing the protocol to be implemented over a generic key-value store.

Alpha_context, defined in the “*alpha_context.ml*” module, serves as the consensus view of the context. This module enforces the separation between mapping the abstract state of the ledger to the concrete structure of the key-value store, and implementing the protocol over the state. The *Alpha_context* defines a type “*t*” that represents the abstracted state of the ledger and can only be manipulated through the use of selected manipulations, which preserve the well-typed aspect and internal consistency invariants of the state.

The abstracted state is read from the disk during the validation of a block and is updated by high-level operations that preserve consistency. Finally, the low-level state is extracted to be committed to disk. This abstraction provides a well-separated structure in the code, with the code below *Alpha_context* handling the ledger’s state storage, and the code on top of it implementing the protocol algorithm using plain OCaml values.

Raw_context, defined in the “*raw_context.ml*” module, is the information view used by *Alpha_context*. It serves as the raw, storage, or non-abstract view of what is actually done in the background by the consensus/protocol. *Raw_context* provides the abstraction used by *Alpha_context* to access the raw part of the protocol, that is, the information that is actually stored and the representation layer.

In conclusion, the implementation of the protocol relies heavily on *Alpha_context* and *Raw_context* to enforce separation of concerns and to abstract away the underlying implementation details, allowing the protocol to be implemented over a generic key-value store in a

readable and well-separated manner.

Features that weren't implemented

In this feasibility study, we have focused on the implementation of consensus protocol in the Tezos node. However, there are a few features that have not been included in this study but could be considered as future work.

One of the missing features is the support for smart contracts. This was not the main focus of the experiment, but smart contracts can be added on top of the protocol to provide more functionality. Currently, the protocol only supports peer-to-peer value transactions, but with the addition of smart contracts, more complex operations can be performed.

Another missing feature is the lack of upgradability in the protocol. This was intentional as this experiment was focused on a standalone Proof of Work consensus algorithm, and there is no previous protocol nor will there ever be a next protocol that is an amendment to this one. Upgradability could be added to the protocol in the future to provide more flexible and dynamic updates, but once again, for this kind of experiment, doing so would be useless.

It should be noted that the absence of these features was not due to technical limitations, but rather a deliberate decision to focus on the implementation of a protocol in the Tezos.

3.4 Tools developed to interact with the protocol

In this study, tools related to the protocol were also implemented to enable interaction with the network. These tools could be developed independently from the Tezos project, since it's possible to communicate with the tezos node through JSON RPC. However, implementing these tools using OCaml and Tezos modules has the added benefit of being automatically integrated with the Tezos node client (and much more). Upon activation of the protocol, the client would automatically adapt to accommodate it, enabling commands that are specific to the protocol, such as "Transfer" and "Reveal", which is the case of the protocol covered here.

One of the tools implemented was the client, which can be found in the "lib_client" folder. The client is mostly used to access information on the network, such as an account's balance or the current target. It serves as a better user interface for accessing the information through JSON RPC services. The client also has the capability of injecting blocks and operations into the network, so it can perform the two operations included in the protocol.

Another tool that was implemented is the baker, which can be found in the "lib_baker" folder. The baker uses functionalities implemented in both the client module and the "Shell Services" module, this last one being already part of the Tezos codebase. The baker performs the task of baking (how block creation is called in Tezos), or mining, a block by taking multiple operations, pre-applying the block in the protocol (to check if the block header is valid), and then performing the work to find the nonce that makes the block's hash lower than the target,

similar to what a miner would do in a Proof of Work blockchain. Once the block is mined, it is pushed to the chain (and spread to the network).

3.5 Conclusion

In summary, our experiment using the Tezos blockchain has been enlightening in multiple dimensions. We chose Tezos for its inherent flexibility and adaptability, qualities that have been validated through the successful implementation of a Proof of Work consensus protocol. This experiment was not just an exercise in protocol implementation; it served as a foundational study to understand the generalizability and testability of consensus protocols on the Tezos platform.

We delved into the intricacies of protocol requirements, integration with the Tezos node, and the roles of key components like the client and baker. The insights gained from implementing and testing this protocol have been invaluable. They have not only demonstrated Tezos' capability to accommodate diverse consensus protocols but also provided us with the knowledge required to generalize and test these protocols effectively.

As we transition into the next chapter on Tool Development, the lessons learned here will be instrumental. The experiment has laid down the groundwork for developing tools that can further simplify the process of swapping, testing, and comparing consensus protocols on blockchain platforms like Tezos.

Chapter 4

Tool Development

This chapter focuses on the development of two essential tools designed to streamline the process of creating and evaluating blockchain protocols, particularly on the Tezos platform. The first tool, known as the **Bootstrapper**, automates the initial setup and coding tasks, providing a quick start for developers. The second tool, the **Live Testing Tool**, offers a real-world testing environment equipped with a variety of features for performance metrics and network control. Together, these tools aim to address the complexities and challenges inherent in blockchain protocol development, offering a more accessible and efficient pathway for both novice and experienced developers.

4.1 Bootstrapping Tool

The Bootstrapper was conceived with the aim of simplifying the development process for creating new protocols in Tezos. Our objective was to minimize the amount of code a developer needs to write to get a protocol up and running. To achieve this, we carefully selected code components that are generic enough to be applicable across multiple protocols. Additionally, we automated certain aspects of the protocol's integration with both Tezos and our Testing tool. In essence, the Bootstrapper serves as a scaffolder or starter pack for protocol development in Tezos.

Protocol Generalizations and Common Code

As previously mentioned in the document, there are specific requirements that every protocol must meet to be compatible with Tezos:

- Must reside in the 'src' folder within the Tezos codebase.
- Must contain a 'lib_protocol' folder and a 'main.ml' file.
- Requires a 'dune' file, as Tezos utilizes the Dune build system.
- Must include a 'TEZOS_PROTOCOL' JSON file with attributes:
 - 'expected_env_version': Expected environment version. Defaults to version 6 for simplicity but can be changed.
 - 'hash': Protocol's hash, serving as its identifier within the node.

- ‘modules’: Names of the OCaml Modules within the folder.

Our tool automates the creation of these required components, eliminating the need for manual setup.

Common Code Components

Here we outline the elements that are common across most protocols.

Main.ml File: Since we use environment v6, a ‘main.ml’ file is included with the essential layout. This is one of the few components that can’t be generalized, as it contains the protocol’s logic.

Account and Tez Representation: Most protocols have the concept of users/accounts and balances. Therefore, we include representations for accounts and Tez, the currency in this system.

Storage Logic, Functors, and More: We include abstractions for storage in Tezos, utilizing the concept of Maps for key-value storage, which is common to most protocols.

Block Header Representation: Every protocol needs to define how it will represent the header part of its blockchain. This is a common requirement, although the specific characteristics may vary.

Given these commonalities, our tool serves as a starter kit, providing the essential components that can be generalized across different protocols. The tool aims to fill in as much as possible, leaving only the unique logic to be implemented by the developer.

Functionalities

The Bootstrapper is a command-line program designed to streamline the initial setup of a new protocol in Tezos. It functions as an “Init Protocol” command, performing the following tasks of **folder creation**, where it automatically sets up the necessary folders, **template copy**, as it copies a predefined template (created by us) into the new folder, but are left in the folder so the developer can change it for their needs, **text replacement and integration**, as it replaces specific text in the template and ensures integration with Tezos’ Dune build system and other required code components, and finally it **includes a pre built script** to build the protocol, so there’s no need for the developer to learn how this process is done.

To further ease development and reduce boilerplate code, we enhanced a pre-existing meta-programming library, *ppx_deriving_encoding* [36].

This library allows developers to automate the definition of encoding functions, by appending a meta-programming tag [*@@deriving encoding*] to a type, the encoding function for that type is automatically generated, and this is particularly useful for type representations, as a big part of the code in protocols are these functions. We also include a block header

functor, that auto-generates the entire `block_header` file functionality by merely providing a type. This not only reduces the code volume by threefold (as experienced while improving the previously mentioned Proof of Work) but also alleviates the developer's need to delve into the intricacies of the Data-encoding library.

For the user, this is abstracted with the following command

```
./bootstrapper -name my_really_fast_protocol -env 6
```

By incorporating these functionalities, the Bootstrapper serves as a comprehensive tool for initiating and easing the development of new protocols in Tezos.

4.2 Live Testing Tool

The Live Testing Tool is central to the main objective of this thesis, which aims to facilitate the development and evaluation of blockchain protocols in a live setting. After developing a protocol in a standardized manner, particularly using Tezos as the platform, the next crucial step is to observe its real-world behavior. This tool is designed to not only run the protocol in a live environment but also to provide valuable metrics about the network's performance. Additionally, it offers the capability to control various aspects of the network, allowing for a comprehensive evaluation of the protocol under different conditions, and can also be used during the development of the protocol as part of the development workflow, as it makes it possible to write tests and get a network running in way fewer steps as with manually executing commands and operations.

Overall, the Live Testing Tool serves as an essential component for both the development and assessment phases of blockchain protocols.

Functionalities

The Live Testing Tool operates as a web server, making it language-agnostic for test development. Any programming language capable of making HTTP requests can interact with it. The server exposes a variety of endpoints that facilitate different aspects of testing, such as initiating a protocol test, controlling node operations, and gathering network metrics.

In its architecture, the server acts as a coordinator, maintaining a list of all node endpoints and ensuring seamless interaction among them. This enables the tool to have a centralized point for controlling and monitoring the network.

While executing, the tool also logs information about the tests it performed, from the protocols, to the start time of the nodes and more.

To provide real-time metrics and insights, the tool employs a monitoring mechanism for each node. Referred to as "watchers," these components continually request information from their respective nodes, focusing on blocks and transactions. This allows the tool to per-

form various calculations and assessments, offering a comprehensive view of the protocol's behavior in a live setting.

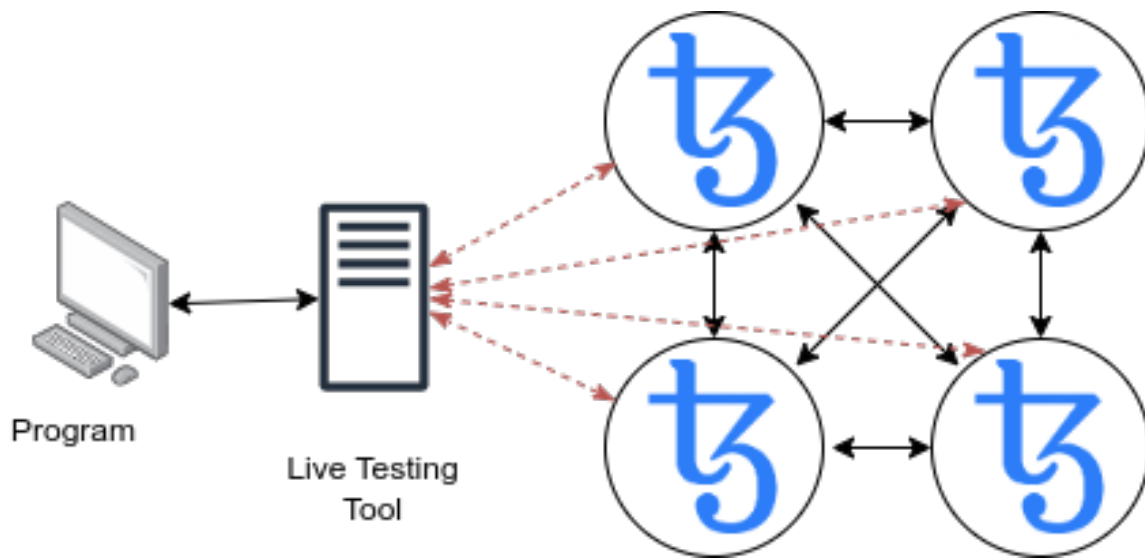


Figure 4.1: Live Testing Tool Architecture Visualized, with 4 Tezos Nodes Running

Endpoints Provided by the Tool

Generic Endpoints

- `/start-test`: Initiates a test network given the *protocol name*, *number of nodes*, and *initial parameters*.
- `/stop-test`: Halts all ongoing tests and activities.
- `/status`: Displays the current status of the test.
- `/protocols`: Lists the protocols that the tool has detected.
- `/protocol-parameters/:name`: Retrieves the parameters of a specific protocol.
- `/nodes`: Returns information about each node, including *IP*, *ports*, *process ID*, and *directory*.

Protocol-Related Endpoints

- `/swap-protocol`: Replaces the current protocol with another.
- `/change-parameters`: Alters the existing parameters of the network's protocol.

Network Control Endpoints

- `/start-node/:id`: Starts a node with a specific *ID*, if provided.

- `/stop-node/:id`: Stops a node with a specific *ID*.

Split Brain/Fork-Related Endpoints

- `/split-network`: Splits the network into smaller networks, given a list of nodes.
- `/rejoin-network`: Rejoins the fragmented networks into a single network.

Metrics Endpoints

- `/tps`: Provides the max, min, and average *Transactions Per Second*.
- `/time-to-consensus`: Measures the time taken for the majority of the network to agree on a block.
- `/propagation-times`: Calculates the time for a block to propagate across the network.
- `/first-and-last-block-times`: Gives the time of the first and last blocks.
- `/blocks-per-second`: Offers the max, min, and average blocks per second.
- `/discarded-blocks`: Indicates the number of blocks discarded by the network.

Example Usage

In this section, we demonstrate how to use this tool with a Python script. However, the system is not language-specific, so other languages like OCaml and Go can also be used. For the sake of brevity, we won't cover the usage of every endpoint.

Listing 1 Example function that starts a test

```
def start_test(protocol_name, nodes, fitness, parameters):
    data = {
        "protocol_name": protocol_name,
        "n_nodes": int(nodes),
        "fitness": int(fitness),
        "parameters": parameters
    }
    print(data)
    response = requests.post(ADDRESS + "/start-test", json=data)
    if response.status_code == 200:
        print("Success:", response.text)
    else:
        print("Failure:", response.status_code)
        exit(1)
```

Listing 2 Multiple utility functions to fetch information about the network

```
# Function that returns the current status of the network
def get_status():
    return requests.get(ADDRESS + "/status").json()
# Function that returns information about all nodes
def get_nodes():
    return requests.get(ADDRESS + "/nodes").json()
# Function that returns the hash of the current head
def get_head_hash(node_rpc):
    return requests.get(f"http://127.0.0.1:{node_rpc}/chains/main/blocks").json()[0][0]
# Function that returns the info of a specific block, provided a hash (block_id)
def get_block_info(node_rpc, block_id):
    return requests.get(f"http://127.0.0.1:{node_rpc}/chains/main/blocks/{block_id}").json()
# Function that returns the info of the head
def get_head_info(node_rpc):
    return get_block_info(node_rpc, get_head_hash(node_rpc))
```

Listing 3 Functions that request information in a loop

```
# Loop that keeps requesting the latest info about the tps of the system
async def fetch_tps():
    while True:
        async with requests.get(ADDRESS + "/tps") as response:
            tps = await response.text()
            print(f"Current TPS: {tps}")
            await asyncio.sleep(1)
# Loop that waits until the test is running
async def wait_for_start():
    last_status = 'stopped'
    while last_status != 'running':
        status_resp = get_status()
        last_status = status_resp["status"]
        sleep(1)
    return
```

Listing 4 Main Function that tests how the TPS changes with the Block Time

```
# Main function that preforms the test
# In this test we check out what happens to TPS when we duplicate the Block Time
def test():
    blocks_to_wait = 200 # Blocks to wait until we increase the Block Time
    protocol_name = "demo" # Name of the protocol
    current_block_time = 60 # In seconds
    parameters = get_protocol_parameters(protocol_name) # Fetch parameters
    number_of_nodes = 20 # Number of nodes to start the network
    fitness = 0
    # we change the default parameters so it has the block time specified
    parameters["constants"]["block_time"] = current_block_time

    # We start the test and wait until the network is ready
    start_test(protocol_name, number_of_nodes, 0, parameters)
    wait_for_start()

    # Ommited function, that fetches information about the nodes
    nodes = get_nodes()
    first_node = nodes[0]

    # Ommited function that makes it so the first node has an baker executing
    # MINER_ADD is the address of the miner
    start_mine(first_node, MINER_ADD)

    # Ommited function that spams the network with transactions
    spam_transactions(nodes, MINER_ADD)

    # We show the current TPS in the command line
    fetch_tps()

    while True:
        last_block_info = get_head_info(first_node)
        level = last_block_info["header"]["level"]

        # In case we reached the level of blocks to wait, we increase the block_time
        if level % blocks_to_wait == 0:

            current_block_time = current_block_time * 2
            parameters["constants"]["block_time"] = current_block_time

            print("Increasing the Block Time to " + current_block_time)
            change_parameters(parameters)
            sleep(current_block_time)
```

The Live Testing Tool serves as a comprehensive platform for both deploying and evaluating blockchain protocols in a real-world setting. Its language-agnostic design and extensive set of endpoints make it a versatile tool for developers. By providing real-time metrics and control over network parameters, it offers a robust environment for rigorous protocol assessment.

4.3 Conclusion

This chapter introduced two pivotal tools aimed at simplifying the development and evaluation of blockchain protocols: the Bootstrapper and the Live Testing Tool. The Bootstrapper serves as an initial setup utility, automating much of the boilerplate code and integration tasks. On the other hand, the Live Testing Tool provides a live environment for protocol testing, complete with real-time metrics and network control. Together, these tools form a cohesive ecosystem that significantly eases the challenges of blockchain protocol development and assessment.

Chapter 5

Use Case: Developing and Testing a New Protocol

Building on the previous chapter's conclusion, the focus of this chapter is to demonstrate the practical utility of the tools we've developed—namely, the Bootstrapper and the Live Testing Tool—in a real-world development scenario. The objective is to provide a comprehensive view of how these tools can be seamlessly integrated to facilitate the development and testing phases of a blockchain protocol.

To achieve this, we chose to implement a Proof of Authority protocol. While this protocol shares some similarities with the previously implemented Proof of Work, its core consensus mechanism is fundamentally different. This serves as an excellent case study to showcase the capabilities of our tools. The Bootstrapper was used to automate the initial setup and coding tasks, while the Live Testing Tool provided a live environment for performance evaluation and fine-tuning. This chapter shows how these tools work together to make creating a protocol faster and easier.

5.1 Goals and Approach

The primary aim of this chapter is to implement a Proof of Authority protocol. We intend to start the development process using the Bootstrapper tool to set up the initial codebase. Following this, we will write tests specifically designed to evaluate the protocol's performance using our Live Testing Tool.

To meet these objectives, we have a multi-faceted approach. First, we aim to demonstrate the efficiency gains by using the Bootstrapper, highlighting the reduced number of steps needed to get the protocol up and running. Second, we plan to show how the Live Testing Tool can be an invaluable asset during the development phase, particularly for debugging and iterative testing.

Lastly, we will present example tests that can be run on the new protocol, showcasing the tool's versatility and utility in a real-world scenario.

5.2 Bootstrapping and Implementing Proof of Authority

In this chapter, we delve into the implementation of a Proof of Authority (PoA) protocol. Unlike public, permissionless protocols like Proof of Work, PoA is a permissioned protocol.

In PoA, only a select group of nodes, termed as authorities, proposers or validators, have the privilege to create new blocks. This design inherently prioritizes security and speed over decentralization. The restricted set of approved nodes enhances security by limiting block proposal to trusted entities. Additionally, the protocol operates more efficiently, minimizing forks and computational overhead.

In our discussion, the terms “validators” and “proposers” are used interchangeably, based on the assumption that participating nodes are also validating incoming blocks. However, it’s important to note that in a real-world scenario, these roles can differ. A “proposer” is a node authorized to create new blocks, while a “validator” may only have the capability to confirm the validity of these blocks. In certain rounds, a validator might only validate and not propose, depending on its designated role in that specific cycle.

Also, unlike networks using Proof of Work are considered a Synchronous network, our Proof of Authority is Partially-synchronous, as there’s an upper bound in our implementation, that is, nodes have a block time plus a tolerance to successfully create and spread a block.

Proof of Authority operates on the principle of reputation and trust. A small number of trusted nodes are given the authority to validate transactions and create new blocks. In our implementation, we replace the concept of mining with a list of these authorities. These validators are responsible for block creation and operate in a round-robin fashion.

The sequence in which these validators create blocks is determined by hashing. Specifically, for each round (or level, for simplicity), we hash the address (or Public Key hash) of each validator concatenated with the current round number. The hash function used is $H(\text{address} + \text{round})$. These hashes are then used to sort the order of the validators, to decide the order of validators for the next round.

In case a validator fails to create a block, the protocol is designed to move on to the next validator in the sorted list. This process continues until a block is successfully created. This mechanism ensures both fairness and liveness in the protocol. Every node gets an opportunity to create blocks, and the protocol continues to function even if some nodes fail.

Notice: the way we order the validators is somewhat deterministic (unless validators fail). This has real-world implications and isn’t necessarily a well designed protocol. To improve on this protocol, some sort of “randomness” should be added that is sourced from an outside entity.

For example, using previous block information to sort the validators (for example, using the hash of the operations, or timestamp) wouldn’t be a good way as the validator creating the block could have influence that information and in some way force the next block validator (to himself). This protocol was implemented just for demonstration purposes.

5.3 Methodology

To kickstart the development of our Proof of Authority protocol, we used the Bootstrapper tool with the command `./bootstrapper -name poa -env 6`. This command sets up a new protocol, ensuring it's integrated into the Tezos build system. It also creates a folder with pre-made templates for various files, saving us from starting from scratch.

Since the Bootstrapper already provides files for account logic, storage, and monetary representations, our focus was mainly on filling in the protocol-specific logic. We began by defining the block header, which in our case only includes the "validator" and "authority_list" values, along with the block's signature by the validator.

Next, we utilized the pre-provided storage file to implement the storage, logic, and representation of the validator set. This was followed by the core logic of the protocol, which is executed when a block is either received from the network or built by the node itself. This logic includes:

- Verifying if the validator address in the header is authorized.
- Confirming if the block proposer is the expected one based on our round-robin logic.
- Checking the validity of the block's signature.

When the node receives a block from the network at a level above the expected one, it only verifies if the signature is correct and if the block validator/proposer is the expected one based on the information on the header.

We also implemented the logic for round selection and created an executable for the block creator, similar to the miner executable in the previous example. That is, in order for a node to be a validator one, it must have a baker attached to it, which is just an executable that keeps track of new blocks in the network and upon new blocks it checks if the account address associate to it is of the current proposer.

By following these steps, we were able to implement the protocol with minimal effort, focusing only on the protocol-specific logic. We only had to implement the logic behind the verification of the block header as well as the round-robin mechanism.

While developing the protocol, we used the Live-testing tool with the following script, to test whether the network was executing and if the protocol works no matter the number of validators:

Listing 5 Script to live test the Proof of Authority during development

```
protocol_name = "poa"
parameters = get_protocol_parameters(protocol_name)

number_of_validators = 5 # This can be changed to account for more/less validators

validators = []
for _ in range(number_of_validators):
    account = create_account() # Returns account with address, private and public keys
    accounts.append({key: account[key] for key in ('address', 'public_key')})
return accounts

start_test_response = start_test(
    protocol_name, number_of_validators, 0, parameters)

nodes = get_nodes()

for i in range(number_of_validators):
    # Adds public and private key to node and starts baker
    add_key_and_propose(nodes[i], validators[i])
```

In summary, the Bootstrapper and Live Testing Tool significantly streamlined the development process of our Proof of Authority protocol. The Bootstrapper’s template provided a strong foundation, allowing us to focus solely on the protocol-specific logic. This was evident in the number of lines of code written; out of the 6166 lines in the complete protocol, only 1700 were written by us. That’s just 27% of the total codebase, emphasizing how much of the groundwork was already laid out.

This not only reduced the coding effort but also lowered the barrier to entry, as we didn’t have to delve deep into the intricacies of protocol development. Even though we had previous knowledge of how we could implement a protocol, we took less than one day to have a fully function PoA, where as when implementing PoW it took us more than a week. The Live Testing Tool further aided in debugging and validation, making the entire development cycle more efficient. Therefore, these tools prove to be invaluable assets for anyone looking to develop and test blockchain protocols with ease and efficiency. Before the creation of the Testing tool, either rudimentary bash scripts where used or manually starting the network was often performed, and we were able to write this script before even starting the implementation of the protocol.

5.4 Evaluation and Insights: PoA vs PoW

In this section, we aim to compare the key metrics of Proof of Authority (PoA) and Proof of Work (PoW) protocols. Specifically, we focus on the metric of Time-to-Consensus (TTC), which measures the time taken for a newly created block to be accepted by a majority of nodes or validators in the network. This comparison serves as a practical use-case for our

developed tools and provides insights into the efficiency of different consensus algorithms.

To carry out this comparison, we used a script that simulates various network conditions for both PoA and PoW. This script is designed to measure the TTC under different configurations of block time and validator set size.

The script automates Time-to-Consensus tests for different consensus protocols and network configurations. It initializes a test network, starts the baker nodes, fetches the TTC from the Live Testing Tool, and then shuts down the network. Multiple test scenarios are defined with varying numbers of nodes and validators, and the script iterates through these to gather TTC data.

Listing 6 Script to execute the tests for Time-to-Consensus

```
def start_protocol_test(protocol, n_validators, n_nodes, blocks_to_wait, blocks_to_run, block_time):

    print(f"Starting {protocol} test")
    parameters = get_protocol_parameters(protocol)
    parameters["constants"]["block_time"] = str(block_time)
    baker_name = "./octez-baker-custom-"
    accounts = create_multiple_accounts(n_validators)

    if protocol == "poa":
        parameters["constants"]["initial_validator_set"] = [
            {key: account[key] for key in ('address', 'public_key')}
            for account in accounts
        ]
        baker_name = baker_name + "poa"
    else:
        baker_name = baker_name + "demo"

    print("Starting the network")
    start_test(protocol, n_nodes, 0, parameters)
    wait_until_status("running")

    nodes = get_nodes()

    print("Starting the baker")
    processes = start_bakers(protocol, baker_name, accounts, nodes)

    print("Running the test...")
    wait_for_blocks(nodes[0]['rpc'], blocks_to_run)
    ttc = fetch_time_to_consensus()
    print(f"Final TTC: {ttc}")

    print("Stopping the test...")
    stop_test()
    wait_until_status("stopped")
    stop_bakers(processes)

    print(f"{protocol} test complete.")

tests = [
    {"validators": 10, "nodes": 10, "block_time": 30},
    {"validators": 20, "nodes": 20, "block_time": 30},
    {"validators": 10, "nodes": 20, "block_time": 20},
    {"validators": 30, "nodes": 30, "block_time": 20}
]

blocks_to_wait = 5
blocks_to_run_for = 25

for config in tests:
    for protocol in ["poa", "demo"]:
        start_protocol_test(protocol, config["validators"],
                           config["nodes"], blocks_to_wait, blocks_to_run_for, config["block_time"])
```

Expectations, Results and Observations

While the primary focus of this document is not to compare protocols but rather to highlight the capabilities of the tools that enable such comparisons, it's worth noting our expectations based on the inherent characteristics of PoA and PoW. These expectations serve as compelling indicators of the tools' effectiveness and utility.

- **PoA's Consistent Time-to-Consensus:** In PoA, only a predefined set of validators are responsible for block creation. This eliminates the randomness associated with block creation in PoW, leading to a more predictable and consistent Time-to-Consensus (TTC). The TTC isn't significantly affected by block time or validator set size, as validators are known and trusted entities.
- **PoW's Increasing TTC with More Validators:** In PoW, any node can become a validator by solving a cryptographic puzzle. As the number of validators increases, the chances of multiple nodes solving the puzzle around the same time also increase. This leads to forks, requiring additional time for the network to reach consensus. Hence, TTC is expected to rise with more validators.
- **Impact of Smaller Block Times in PoW:** Smaller block times mean that validators have less time to propagate their blocks to the network. This increases the likelihood of two or more validators solving the puzzle simultaneously, leading to more frequent forks. Resolving these forks takes time, thereby increasing the TTC.
- **Need for Larger Block Times in PoW:** Shorter block times not only lead to increased TTC but also pose security risks. Frequent forks make the network more susceptible to "double-spending" attacks. Moreover, network instability can occur if blocks are propagated slower than they are created.

We present the resulting statistics in the following table:

Table 5.1: Time-to-Consensus Comparison between PoW and PoA

Parameters			Metrics (Time-to-Consensus)		
Validators	Nodes	Block Time	Max	Min	Average
Proof of Work					
10	10	30	117	10	50.959
20	20	30	143	9	70.312
10	20	20	143	10	65.229
30	30	20	191	20	72.332
Proof of Authority					
10	10	30	35	2	13.827
20	20	30	38	1	14.233
10	20	20	30	4	14.101
30	30	20	42	2	13.994

Upon a detailed examination of the results, several key insights emerge that corroborate our initial expectations:

- PoA exhibits a stable average Time-to-Consensus (approximately 14 seconds) across various configurations. This consistency is attributed to its sequential block creation mechanism, where only one block proposer is active at a given time. This eliminates the possibility of multiple forks and ensures a quicker and more reliable consensus.
- In contrast, PoW's Time-to-Consensus is highly sensitive to the number of validators and block times. This is due to its concurrent block creation approach, where multiple miners can create blocks simultaneously. This leads to the formation of multiple forks in the chain and network, necessitating a longer block time to resolve these forks and reach a stable consensus.

These findings underscore the efficiency of PoA in achieving rapid consensus, while also highlighting the limitations of PoW, particularly in scenarios with shorter block times and larger sets of validators.

It's worth noting that although not covered in this document, a specialized test suite for Proof of Work could be developed to ascertain the optimal block time that minimizes the occurrence of forks while maintaining efficient block creation.

5.5 Conclusion

In this chapter, we've demonstrated the power and utility of our developed tools—Bootstrapper and Live Testing Tool—in a real-world scenario. Through the lens of implementing and comparing two distinct consensus algorithms, Proof of Authority and Proof of Work, we've showcased how these tools can significantly streamline the protocol development and testing process.

The Bootstrapper tool allowed us to implement a fully functional PoA protocol with minimal effort, requiring us to write only 27% of the total lines of code. This not only reduces the development time but also lowers the entry barrier for newcomers in the field of blockchain protocol development.

Our Live Testing Tool provided invaluable insights into key metrics like Time-to-Consensus, enabling us to make data-driven decisions and optimizations. The tool's flexibility and adaptability were evident as we easily configured it to test different network conditions, thereby obtaining a comprehensive understanding of the protocols' performance under various scenarios.

The comparative analysis between PoA and PoW further emphasized the importance of choosing the right consensus algorithm, and how our tools can aid in making that decision based on empirical data.

In summary, our tools offer a robust, efficient, and user-friendly environment for developing and testing blockchain protocols. They not only save time and effort but also provide the

kind of deep insights that are crucial for the development of secure, efficient, and scalable blockchain networks.

Chapter 6

Conclusion

Blockchain technology has emerged as a transformative force in various sectors, from finance to healthcare and beyond. However, the development and testing of blockchain protocols remain complex and time-consuming tasks. This thesis addresses this critical issue by providing insights and tools designed to simplify and expedite these processes.

We have developed two distinct tools: the Bootstrapper and the Live Testing Tool. The Bootstrapper serves as a foundational tool for protocol development, allowing developers to focus on the unique aspects of their protocols by automating much of the boilerplate code. On the other hand, the Live Testing Tool offers a dynamic environment for testing and analyzing blockchain protocols under different network conditions. Together, these tools form a comprehensive suite that can significantly reduce the time and effort required to develop and test new blockchain protocols.

To demonstrate the utility of our tools, we implemented two consensus algorithms: Proof of Work and Proof of Authority. These implementations served as real-world test cases, allowing us to showcase the capabilities of our tools in a practical setting. The comparative analysis between these two protocols provided valuable insights into the importance of selecting the appropriate consensus algorithm for specific use-cases.

Our work contributes to the broader goal of making blockchain technology more accessible and efficient. By lowering the barriers to entry for protocol development and offering a robust testing environment, we believe our tools can play a pivotal role in the advancement of this technology.

In conclusion, this thesis presents a significant step forward in the field of blockchain protocol development and testing. The tools and insights provided herein not only facilitate the development process but also contribute to the optimization and scalability of future blockchain networks. As blockchain technology continues to evolve, the methodologies and tools developed in this thesis offer a solid foundation upon which further advancements can be built.

6.1 Future Work

While this thesis provides a comprehensive toolset for blockchain protocol development and testing, there are several avenues for future work to further enhance its utility and scope:

- **Extend to Non-Blockchain Protocols:** One of the most immediate extensions

would be to adapt the tools for non-blockchain based protocols, such as Directed Acyclic Graphs (DAGs). This would make the toolset more versatile and applicable to a broader range of decentralized technologies.

- **Additional Metrics:** The current version of the Live Testing Tool focuses on a limited set of metrics. Future work could involve the implementation of more comprehensive metrics to evaluate protocols from various angles, such as security, scalability, and fault tolerance.
- **Bootstrapper Enhancements:** The Bootstrapper tool could be evolved into a full-fledged development suite. This could include libraries, helper functions, and perhaps even a graphical user interface to assist developers in creating robust and efficient protocols more easily.
- **Integration with Lupin DSL:** There is potential for integrating the tools developed in this thesis with the Lupin Domain-Specific Language (DSL). This could streamline the development process further, allowing for more efficient code generation and testing.

By addressing these points, future work can build upon the strong foundation laid by this thesis, contributing to the development of more efficient, secure, and scalable decentralized systems.

Bibliography

- [1] L. Goodman, “Tezos—a self-amending crypto-ledger white paper,” *URL: <https://www.tezos.com/static/papers/white paper.pdf>*, vol. 4, pp. 1432–1465, 2014. 1
- [2] “Peer-to-peer electronic cash.” [Online]. Available: <https://bitcoincash.org/> 2
- [3] S. Nakamoto, “Bitcoin whitepaper,” *URL: <https://bitcoin.org/bitcoin.pdf>*-(: 17.07. 2019), 2008. 9, 15
- [4] E. Foundation, “Ethereum whitepaper.” [Online]. Available: <https://ethereum.org/en/whitepaper/> 9
- [5] G. Bracha and S. Toueg, “Asynchronous consensus and broadcast protocols,” *Journal of the ACM (JACM)*, vol. 32, no. 4, pp. 824–840, 1985. 9
- [6] C. Dwork, N. Lynch, and L. Stockmeyer, “Consensus in the presence of partial synchrony,” *Journal of the ACM (JACM)*, vol. 35, no. 2, pp. 288–323, 1988. 10
- [7] E. Buchman, “Tendermint: Byzantine fault tolerance in the age of blockchains,” Ph.D. dissertation, University of Guelph, 2016. 10
- [8] T. Inc, “Internet of blockchains.” [Online]. Available: <https://v1.cosmos.network/resources/whitepaper> 10
- [9] L. Aștefanoaei, P. Chambart, A. Del Pozzo, T. Rieutord, S. Tucci, and E. Zălinescu, “Tenderbake—a solution to dynamic repeated consensus for blockchains,” *arXiv preprint arXiv:2001.11965*, 2020. 10
- [10] M. J. Fischer, N. A. Lynch, and M. S. Paterson, “Impossibility of distributed consensus with one faulty process,” *Journal of the ACM (JACM)*, vol. 32, no. 2, pp. 374–382, 1985. 11
- [11] E. A. Brewer, “Towards robust distributed systems,” in *PODC*, vol. 7, no. 10.1145. Portland, OR, 2000, pp. 343 477–343 502. 11
- [12] E. Androulaki, A. Barger, V. Bortnikov, C. Cachin, K. Christidis, A. De Caro, D. Enyeart, C. Ferris, G. Laventman, Y. Manevich *et al.*, “Hyperledger fabric: a distributed operating system for permissioned blockchains,” in *Proceedings of the thirteenth EuroSys conference*, 2018, pp. 1–15. 13
- [13] B. Bnb-Chain, “Whitepaper/whitepaper.md at master · bnb-chain/whitepaper,” Oct 2022. [Online]. Available: <https://github.com/bnb-chain/whitepaper/blob/master/WHITEPAPER.md> 13
- [14] V. Buterin, “Why sharding is great: Demystifying the technical properties.” [Online]. Available: <https://vitalik.ca/general/2021/04/07/sharding.html> 13

- [15] Bitnodes. [Online]. Available: <https://bitnodes.io/> 14
- [16] C. Gondek, “Here is why bitcoin transactions take so long.” [Online]. Available: <https://originstamp.com/blog/here-is-why-bitcoin-transactions-take-so-long> 14
- [17] J. Xu, C. Wang, and X. Jia, “A survey of blockchain consensus protocols,” *ACM Computing Surveys*, 2023. 15
- [18] M. S. Ferdous, M. J. M. Chowdhury, M. A. Hoque, and A. Colman, “Blockchain consensus algorithms: A survey,” *arXiv preprint arXiv:2001.07091*, 2020. 15
- [19] A. Back *et al.*, “Hashcash-a denial of service counter-measure,” 2002. 15
- [20] E. Anceaume, A. Pozzo, T. Rieutord, and S. Tucci-Piergiovanni, “On finality in blockchains,” *arXiv preprint arXiv:2012.10172*, 2020. 16
- [21] C. Grunspan and R. Pérez-Marco, “On profitability of selfish mining,” *arXiv preprint arXiv:1805.08281*, 2018. 19
- [22] K. C. Chaudhary, V. Chand, and A. Fehnker, “Double-spending analysis of bitcoin,” in *Pacific Asia conference on information systems*. Association for Information Systems, 2020. 19
- [23] R. Pass and E. Shi, “Fruitchains: A fair blockchain,” in *Proceedings of the ACM symposium on principles of distributed computing*, 2017, pp. 315–324. 19
- [24] “The monero project.” [Online]. Available: <https://www.getmonero.org/> 20
- [25] P. Foundation, “The pioneer of proof-of-stake.” [Online]. Available: <https://www.peercoin.net/> 20
- [26] V. Buterin and V. Griffith, “Casper the friendly finality gadget,” *arXiv preprint arXiv:1710.09437*, 2017. 20
- [27] L. Lamport, R. Shostak, and M. Pease, “The byzantine generals problem,” in *Concurrency: the works of leslie lamport*, 2019, pp. 203–226. 20
- [28] P. Szilágyi. (2017, 3) Eip-225: Clique proof-of-authority consensus protocol. Ethereum Improvement Proposal. [Online]. Available: <https://eips.ethereum.org/EIPS/eip-225> 21
- [29] Proof of authority - poa - identity digital. Aura is one of the Blockchain consensus algorithms available in OpenEthereum (previously Parity). [Online]. Available: <https://www.poa.network> 21
- [30] H. Moniz, “The istanbul bft consensus algorithm,” 2020. 21
- [31] K. Karantias, A. Kiayias, and D. Zindros, “Proof-of-burn,” in *Financial Cryptography and Data Security: 24th International Conference, FC 2020, Kota Kinabalu, Malaysia, February 10–14, 2020 Revised Selected Papers 24*. Springer, 2020, pp. 523–540. 21

- [32] “Gridcoin wiki home.” [Online]. Available: <https://gridcoin.us/wiki/> 21
- [33] S. Dziembowski, S. Faust, V. Kolmogorov, and K. Pietrzak, “Proofs of space,” in *Advances in Cryptology—CRYPTO 2015: 35th Annual Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2015, Proceedings, Part II*. Springer, 2015, pp. 585–605. 22
- [34] M. Bowman, D. Das, A. Mandal, and H. Montgomery, “On elapsed time consensus protocols,” in *Progress in Cryptology—INDOCRYPT 2021: 22nd International Conference on Cryptology in India, Jaipur, India, December 12–15, 2021, Proceedings 22*. Springer, 2021, pp. 559–583. 22
- [35] N. Labs, “Tezos software architecture¶.” [Online]. Available: https://tezos.gitlab.io/shell/the_big_picture.html 26
- [36] Max, “ppx_deriving_encoding,” gitLab repository. [Online]. Available: https://gitlab.com/o-labs/ppx_deriving_encoding 38