Unveiling Decentralisation in the Algorand Blockchain: A Detailed Inquiry through Full-Node Parsing and Exogenous Artefacts

Rizzatama Nurrokhman Santosa

Master of Science
Cyber Security, Privacy and Trust
School of Informatics
University of Edinburgh
2023

Abstract

Serving as the groundwork for the Edinburgh Decentralisation Index (EDI), this dissertation delves into the nuanced realm of decentralisation in permissionless blockchain networks, spotlighting the Algorand and its Pure Proof-of-Stake (PPoS) protocol. A specialized parser, combined with off-ledger data and three key metrics—Gini Coefficient, Herfindahl-Hirschman Index, and Nakamoto Coefficient—were used to evaluate seven tangible layers: *Hardware*, *Software*, *Network*, *Consensus*, *Tokenomics*, *Governance*, and *Geography*. In spite of the protocol's soundness against forks, its resistance to potential foundational exits or network partition challenges, and remarkable community engagement, a propensity towards centralization over time within Algorand was also uncovered, reflecting comparable realities observed in other blockchains notably Bitcoin and Ethereum. Beyond the noted monetary dynamics, this concerning pattern is influenced by factors like inventive governance procedures and a muted involvement in development. The methodologies, tools, and findings presented in this dissertation are anticipated to assist relevant stakeholders in initiating due diligence and facilitate subsequent inquiries into decentralisation in blockchain.

Research Ethics Approval

This project was planned in accordance with the Informatics Research Ethics policy. It did not involve any aspects that required approval from the Informatics Research Ethics committee.

Declaration

I declare that this thesis was composed by myself, that the work contained herein is my own except where explicitly stated otherwise in the text, and that this work has not been submitted for any other degree or professional qualification except as specified.

(*Rizzatama Nurrokhman Santosa*)

Acknowledgements

The successful culmination of this dissertation owes much to the unwavering academic, emotional, and spiritual support bestowed upon me by numerous individuals. Primarily, I wish to express my heartfelt gratitude to my chief project supervisor, whose unflagging guidance has been instrumental in navigating the research methodologies, the labyrinthine corridors of my thoughts, and the technological obstacles confronted during the execution of this dissertation.

I was indeed fortunate to cross paths with Dr. Woods, who graciously inducted me into the Edinburgh Decentralisation Index Project and its underlying ethos. Additionally, Ms. Christina Ovezik deserves my sincere appreciation for her lucid and thorough elucidations concerning the technical citations. Their fervent mentorship, combined with an uncanny understanding of their students' needs, facilitated a conducive environment for my intellectual expansion through a process of guided inquiry. Moreover, their insightful feedback proved to be invaluable, ensuring that I remained steadfast and true to my individual academic journey.

In my heart, a special recognition is reserved for Mita Dwi Rahmawati and Muhammad Arsyad Arrafif, whose consistent and patient support has been unwavering throughout my studies. Their generosity and selflessness have been truly remarkable. Similarly, Wifaq Santosa and Istiari, despite their roles as parents imparting unconditional love and understanding, exceeded their responsibilities by devoting countless hours discussing my study objectives, irrespective of their expertise.

An endeavor of this dissertation is intrinsically a collaborative enterprise, with intellectual discourse forming its core. I am indebted to my colleagues, all of whom are engaged in diverse blockchain environments, for their spirited discussions, innovative insights, and support, all of which were freely extended and have assisted in the development and evolution of the ideas presented in this dissertation. In particular, my thanks are extended to Shengying Li, Zhi Jian Tow, Shuren Miao, and Yu Cen for sharing their insightful experiences.

The role of friendship in my life is immeasurable and I am deeply thankful for their continuous presence. I wish to extend my appreciation to Chendika Bayu Kurniawan, Nugrahani Sulistyowati, Muhammad Andi Miftachul Huda, and Wawa Uswatun Hasanah. Their camaraderie, shared during the trials and tribulations of our journey as international students, has been a source of joy and a platform for lifelong friendships.

The completion of this dissertation owes much to the unwavering support and precious shared experiences, and consistent encouragement from my esteemed colleagues at the Central Bank of Indonesia. It is with profound gratitude that I acknowledge the invaluable contributions of Utami Dyah, Diva Amelia, and Andi Ari. Notwithstanding our diverse backgrounds, we jointly undertook this scholarly endeavor after deliberate contemplation. Their collective insights, born out of our shared institutional legacy, have been instrumental in fortifying impetus for the fruition of this work. I am profoundly indebted to them for their enduring commitment to my personal development.

Finally, I owe an immense debt of gratitude to Dr. Agusman. As my professional career mentor in the Central Bank of Indonesia, his influence in my academic journey extended beyond any formal social or academic obligations. He was the catalyst in encouraging me to pursue further studies, even after nearly a decade of professional engagement. His compassionate heart has profoundly influenced me, for which and for much more, he merits my utmost respect and gratitude.

Table of Contents

Chapter 1

Introduction

1.1 Motivation

Centralized control often results in protracted transaction processing times and substantial fees, particularly in the context of financial settlements. In the digital era, the ability to swiftly transfer data and update ledgers between financial intermediaries is expected and almost intuitive. However, the practical realization of this potential is often hampered by existing centralized structures, thereby creating a compelling case for the adoption of decentralized systems that have been demonstrated to yield substantial influence over crucial market efficiency and outcomes [17, 22, 23, 114, 139, 153].

Decentralisation, acting as the *sine qua non* of blockchain technology, fundamentally alters the power dynamics within a system by devolving authority away from a centralized entity [26, 80, 120, 148]. While transformative and suggesting a future financial landscape that is both more democratized and resilient [80], the decentralized nature of blockchain introduces distinctive challenges pertaining to trust, security, and stability that persistently draw critical attention [77]. This transition paves the way for the potential creation of non-custodial platforms, ushering in an unprecedented level of transparency [78, 127, 147].

Discerning the extent of decentralisation intrinsic to public blockchain initiatives, together with the concomitant risks, has emerged as a concern of paramount significance. A thorough and accurate measure of decentralisation could not only aid in risk evaluation but also foster the preservation of blockchain's distinctive characteristics - transparency, resilience, and integrity - traits that are inherently tied to its decentralized principles.

1.2 Rationale and Significance

The ramifications of decentralisation are profound. However, the extant degree of its practical application and the effects of varying decentralisation levels on blockchain platforms remain nebulous and insufficiently explored [38, 155, 156]. The seminal observations suggesting a structural relationship within the blockchain domain underscore the pressing need for more exhaustive research [7, 27, 37, 38, 74, 99, 119, 157]. This view ostensibly contradicts the distributed network model that blockchain espouses, suggesting that the practical application of blockchain may deviate from its theoretical foundations.

Conversely, regardless the ostensive transparency of public blockchain, the collection and processing of a comprehensive data set present substantial challenges. Operating a full node within a blockchain network and retrieving block-related data require navigating a unique array of technical complexities to ensure completeness, precision, and reliability as it necessitates maintaining a complete replica of the ledger. Additionally, to extract valuable insights that can inform real-world governance situations, it is crucial to select a contextually appropriate and interpretable unit of decentralisation and inequality measurement for the blockchain under consideration.

1.3 Aims and Objectives

The primary objective of this dissertation is to critically appraise the decentralisation layers of the Algorand blockchain. The study will explore the complexities involved in managing a full node within the ecosystem and extracting pertinent data. Algorand was chosen as the central focus of this inquiry owing to its ground-breaking PPoS consensus protocol. The protocol is being heralded as an innovative resolution to widespread challenges prevalent in alternative blockchain networks: scalability, security, and efficiency.

The projected outcomes of this dissertation are anticipated to yield insights into effective full node management, strategies to surmount technical challenges, and considerations for parsing block data. Serving as an instrumental resource for a myriad of stakeholders, this dissertation endeavors to elucidate the intricate aspects of blockchain, thereby facilitating informed decision-making and bolstering assurance within the decentralized ecosystem.

1.4 Novelty

This dissertation was launched to promote the characterization of blockchains and lay the solid groundwork for the advancement of the Edinburgh Decentralisation Index (EDI), which marks an inaugural index dedicated to decentralisation within publicly accessible blockchain. Notwithstanding the array of past research dissecting and assessing well-known networks such as Bitcoin and Ethereum, comparable examinations of the Algorand network have not materialized up to this point.

1.5 Document Outline

The structure of the dissertation unfolds into seven distinct chapters subsequent to the Introduction, in the following manner:

Chapter 2: Related Work, presents a thorough review of related work and theoretical foundations drawn from the domain of blockchain decentralisation measurement.

Chapter 3: Methodology, lays out the approaches employed in this dissertation, embracing the structured procedure for gleaning pertinent information from the blockchain, alongside the integration of off-chain data sources to establish a robust foundation for measurement.

Chapter 4: Validation, presents the outcomes of the analysis concerning the computational requirements and efficacy for full node operation, while demonstrating the functionality of the developed tools in acquiring the necessary on-chain data.

Chapter 5: Decentralisation Measurement, confer the results of the decentralisation measurements for each examined strata, grounded in the interaction and interpretability of the collected on-chain and off-chain data.

Chapter 6: Discussion, recognizes the inherent limitations and constraints of the work, delves into the implications of the findings through the lens of analytical comparison, and draws analogies with real-world occurrences.

Chapter 7: Conclusion, delineates the contributions of this work in response to the motivations and objectives, lays a foundational basis for informed decision-making for pertinent stakeholders, and suggests potential avenues for future work informed by the central discoveries, strengths, and observed shortcomings.

Chapter 2

Related Work

The vast body of literature indicates that numerous efforts have been directed towards analyzing public blockchains. Although each methodology and instrument presents its own distinctive approach and competitive edge, they are not without limitations and potential pitfalls, primarily stemming from inherent ontological intricacies.

2.1 Initiatives Aimed at Analysing Public Blockchains

2.1.1 Observation via Direct Access to Blockchain Data

Enthusiasts and researchers have proposed and developed specific frameworks for the aggregation and administration of Bitcoin blockchain data [16, 24, 112, 129, 158]. Though these studies have harnessed the ledger to probe salient economic dynamics, the intricate structure necessitates profound technical acumen to access and analyze the available data. Predominantly used frameworks largely draw from traditional and general-purpose databases; however, they are outperformed by highly-tailored, clustered, specialized in-memory methodologies [89, 94, 126, 130].

A pioneering blockchain parsing tool was spearheaded by a pseudonymous enthusiast known as `Znort987` [158]. The introduced BlockParser provides a framework for developers to integrate analysis code while parsing Bitcoin. However, its stateful attributes restrict its effectiveness, impeding parallel processing. Drawing inspiration from [158], Spagnuolo et al. [129] unveiled BitIodine. By adapting the previous developed parser, BitIodine demonstrated the ability to sift, consolidate, and visually exhibit crucial blockchain data, while also enabling the grouping, categorisation, and labelling of addresses prior to database storage.

In a different vein, adhering to the proposed future enhancements put forth by [129], Möser and Böhme [112] leveraged the Neo4J graph database for processing Bitcoin blockchain. The integration of the Cypher query language with the low-level version of the Java API showcases empirical adeptness in evaluating the blocks. Nonetheless, beyond the architecture's inability to support parallel processing, the lack of transparency concerning the infrastructure hinders a comprehensive comprehension of the tool's performance.

Kalodner et al. [89] expanded on the previous foundational works by introducing BlockSci, a multifaceted platform capable of executing swift query and analysing a variety of graph-structured-type blockchains. BlockSci stands out due to its in-memory analytical capabilities and its versatility to adapt to blockchains beyond just Bitcoin. While it can process at speeds 15-600 times faster than [129], it is not ideally optimized to manage voluminous transactions and requires a high-specification infrastructure.

In response to the limitations imposed by specific types of blockchain networks, Bartoletti et al. [16] presented a platform grounded in Scala, which mixtures in-database blockchain data restructuring techniques. However, even though it can operate on consumer-grade computing, such a configuration requires significant time for both data importing and processing.

Simultaneously, Bragagnolo et al. [24] adopted a query-based methodology for analyzing Ethereum. The researchers employed more advanced techniques, specifically leveraging parallelization prevalent in Big Data platforms. However, the research observed that increased parallelization does not consistently yield proportional performance improvements. Coupled with the requirement for an extensive computing infrastructure, the study faced challenges related to storage space constraints stemming from the indexing process, even for a relatively limited number of blocks.

Subsequent scholarly investigations, led by Rubin [126], Kiliç et al. [94] and Su et al. [130], embarked on tackling the deficiencies previously pinpointed, by utilizing cloud-based infrastructure. BTCSpark [126], emerges as one of the first to employ a distributed analysis platform, integrating the Python to C interpreter and Apache Spark as integral components of its architecture. Kiliç et al. [94] exploited the scalable computational and storage resources provided by Amazon Cloud. Although they might exhibit high performance, both necessitate substantial infrastructure owing to the utilization of parallel architecture.

In contrast, Su et al. [130] demonstrated the feasibility of extracting comprehensive interactions from the Ethereum network by employing a fully automated tool hosted on Google BigQuery. While it showcases notable scalability and storage efficiency, the data ingestion pace aligns closely with the speed of raw data acquisition.

Cumulatively, to optimize performance and ensure accurate analysis, aforementioned studies in return calls for a sturdy and reliable infrastructure. Nevertheless, they largely overlook the expenditure evaluation essential for the deployment of the proposed solutions, with the exception of [94]. A comparative overview of the infrastructure required by each study can be found in the Appendix A.

2.1.2 Employment of Public Blockchain Data Services

Subsequent studies have emphasized the prospect of examining blockchain transactions derived from accessible public data repositories. A compilation of well-known public blockchain explorers, along with their provided classifications of information, is available in Appendix B.

Although the significance of data extraction is indisputable, the breadth of these inquiries is frequently limited. In terms of Application Programming Interfaces (APIs), a level of uniformity is generally observed. These APIs are deficient in accommodating highly tailored queries, and also constrain the quantity of data that can be accessed per request. This limitation detrimentally affects the feasibility of thorough and detailed examinations, and some of these services have been subject to criticism.

Among the pioneers to adopt the aforementioned methodology were Reid and Harrigan [121] as well as Ron and Shamir [125]. Their research objectives were achieved by combining data from crawled nodes and transactions with specific details extracted using proprietary algorithms. However, due to the absence of performance metrics and undisclosed code, these studies have been criticized for their over-reliance on blockchain explorers, which subsequently narrows the potential for comprehensive evaluations in subsequent studies [73].

2.2 Decentralisation Measurement Basis

Blockchain was envisioned as a decentralized trust system, primarily aimed at fostering wealth creation and economic balance [25, 26]. However, a thorough analysis of its foundational and functional framework uncovers its inherent complex multidimensionality [156]. This multi-faceted nature is further accentuated by the absence of a clear-cut definition for ”decentralisation” and the elements influencing it.

2.2.1 Attempts Concentrated on Economic Equilibrium

Scholarly discourse suggests that transaction decentralisation is integral to blockchain's essence and can be viewed from two vantage points. On one hand, the decentralisation of transactions embodies the principles of peer-to-peer (P2P) networks and the digital payment system landscape [108, 117, 131]. On the other hand, centralized transactions pose risks of manipulative control. With this context in mind, this research delves into three fundamental aspects: governance, consensus, and wealth distribution.

The pivotal roles of governance and consensus in safeguarding the blockchain ecosystem gain prominence when considering that even the slightest alteration in the ”code-is-law” domain can ignite a destabilization of the ecosystem [93]. Despite the involvement of a diverse range of entities within the ecosystem aspiring to bolster the principles of decentralized control and wealth distribution, the tangible outcomes may not necessarily correspond to these initial expectations. This incongruity can be especially pronounced in the spheres of voting power, transaction validation, and system upgrades [150]. The crux of this divergence can often be traced to situations wherein blockchains are predominantly overseen by an elite group of entities that orchestrate the system's proceedings [10, 32, 133, 140, 141].

In light of the above-mentioned context, Pelt et al. [141] have recontextualized the definition of Open Source Software (OSS) governance to better align with blockchain nuances.Similarly, Beck et al. [18] have derived dimensions of blockchain governance by referencing established IT governance frameworks. In contrast, studies by Allen and Berg [4], Campbell-Verduyn [28], and Hsieh et al. [84] adopt a more traditional stance on governance, segregating it into its intrinsic and extrinsic facets. Adding layers of granularity, McMullen's subsequent research [104] dissects Bitcoin's governance into two pivotal domains: technical influence, underpinned by the protocol, and social influence, steered by the community and other stakeholders.

Building upon the insights provided by [18, 141], Khan et al. [92] delved deeper into the intricacies of blockchain governance, paying particular attention to the varied actors and strategies involved in platform decision-making processes. [92] found that certain platforms, mirroring traditional governance structures, experienced a significant limitation: the inability of participants to change their voting preferences between successive elections. An in-depth exploration of delegated voting mechanisms has been undertaken by Gersbach et al. [75]. While these mechanisms offer potential solutions, their implementation should be approached with caution and discernment.

Subsequently, a systematic review of blockchain governance was initiated by Kiayias and Lazos [93]. In accordance with the proposed seven fundamental properties, the authors highlighted that a degree of trade-off between these properties is inevitable. This observation is grounded in the fact that each blockchain platform is uniquely engineered to accommodate specific objectives. Therefore, universally applicable decision-making processes are unlikely to be supported by the specific protocols.

In the domain of consensus, substantial emphasis has been accorded to transactional indices such as mining power, and network indices such as bandwidth. Detailed analyses of these factors, particularly as they relate to Bitcoin and Ethereum, have been documented by Gencer et al. [74] and Wang et al. [145]. Notably, well-established metrics such as the Gini and Nakamoto coefficients, along with Shannon entropy, have been utilized to facilitate these analyses, as demonstrated in the further studies [11, 77, 79, 97, 99, 149, 156].

2.2.2 Expanded Facets for Measuring Decentralisation

Venturing into the ontological dimensions of blockchain, Bellaj et al. [19] present a model that classifies the foundation into four core layers: data, distributed consensus protocols, execution, and application layers. This method primarily takes a holistic view, encompassing distinct properties that might be studied individually for a more detailed comprehension. The consensus layer is illustrative of this, merging both the network and governance models, enriching the overall layered framework.

Broadening the conversation, Zhang et al. [156] present an elaborate taxonomy, incorporating additional dimensions: network and transactions. Rooted in the principle of *ceteris paribus*, their research proposes an index that introduces certain ambiguities. Particularly, there are questions regarding how this methodology evaluates transactions of disparate utilities or values across ledgers with varied throughputs.

In their vanguard study, Karakostas et al. [90] employ the Open Systems Interconnection (OSI) model's architectural principles as a point of departure. They dissect the intricate structure of blockchain into eight distinct, well-articulated layers. Every layer introduces its own set of unique challenges that need to be navigated to fully leverage the intrinsic merits of decentralisation. Central to [90]'s stratified framework is a meticulous investigation of each layer, aimed at identifying areas of decentralisation that could potentially be susceptible to the encroachment of centralizing influences, encapsulating vital facets like *security*, *liveness*, *privacy*, and *stability*.

Chapter 3

Methodology

The approach was constructed with awareness of the particular characteristics of the Algorand ecosystem and the associated development kit. Following data acquisition, a quantitative analysis employing the pertinent metrics is conducted.

3.1 Algorand Ecosystem

3.1.1 Network and Nodes Structure

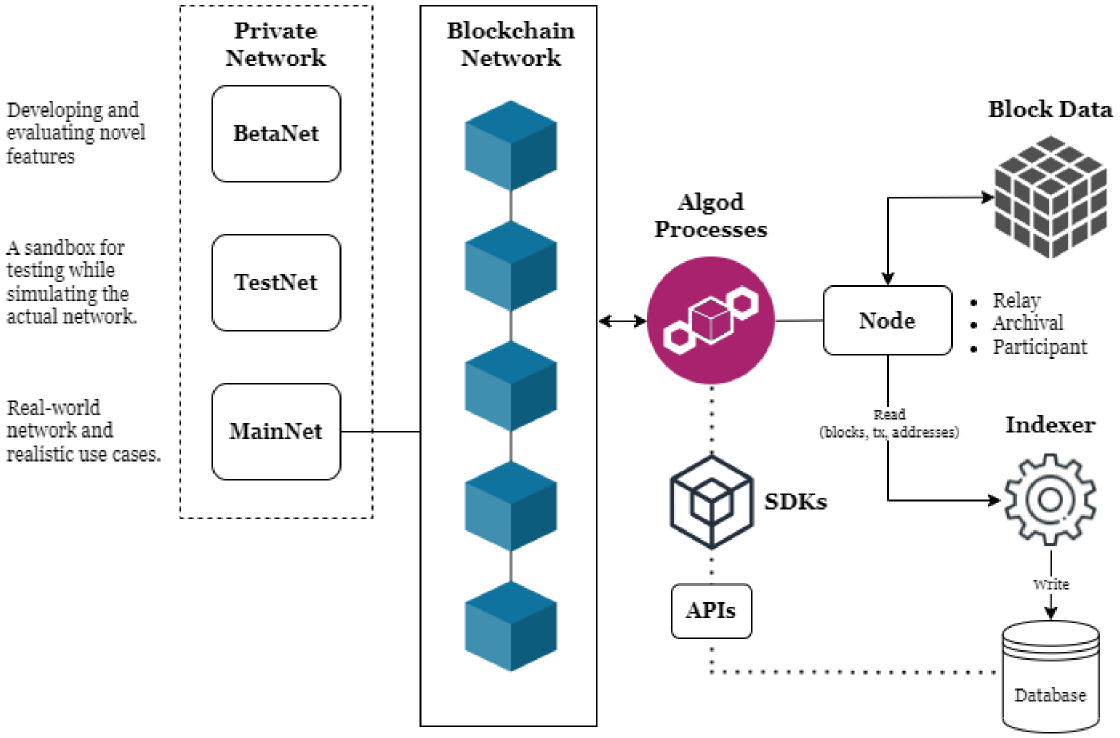

Figure 3.1: Global environment of Algorand ecosystem

The Algorand network is composed of three public networks, each serving a distinct purpose: the MainNet, TestNet, and BetaNet. The MainNet is the primary network where genuine use cases are conducted. On the other hand, the TestNet and BetaNet are used predominantly for developmental purposes.

Beyond their primary roles pertaining to the relay or non-relay of MainNet communications, Algorand offers open stance and versatile customization for nodes based on distinct ledger data preservation requirements. Nodes can be categorized into two main types: archival and non-archival. Archival nodes bear the onus of storing the comprehensive history of ledger data, dating back to the network's origin. In contrast, non-archival nodes store only the latest blocks until they synchronize with the broader network. This inclusive approach allows any category of node to actively contribute to the network without facing any prescribed constraints.

3.1.2 Exposed Development Tools

Within the technical architecture of the Algorand network, three principal technical tools are available for block data retrieval: the Algod daemon, the Algorand Indexer, and the Algorand Software Development Kits (SDKs).

The Algod daemon holds a central position in the Algorand network, overseeing vital protocol phases, enabling communication between nodes, and recording the blockchain on individual nodes. In addition, it presents a series of Representational State Transfer (REST) APIs, creating an avenue for interaction with the network and acquisition of essential data. Working alongside Algod, the Algorand Indexer facilitates adaptable search functionalities within the locally-preserved blockchain. The Algorand SDKs, at the other end of the spectrum, serve as a conduit for seamless engagement with the network.

3.2 Problem Settings

Establishing a full node introduces an added degree of flexibility in conducting an in-depth analysis in measuring blockchain decentralisation [16, 73]. Nevertheless, the aspiration often demands a formidable infrastructure complemented by considerable computational capabilities [16, 24, 89, 94, 126].

3.2.1 Network Size

The size of a blockchain ecosystem typically experiences a gradual increase. Recent observations indicate that the Algorand network is witnessing a daily growth of 2 GB.

Node	MainNet	BetaNet	TestNet
Archival	1.190	651	490

Table 3.1: Algorand network size as of 1 June 2023[3], measured in Gigabytes (GB)

3.2.2 Computational and Financial Constraint

A solution is deemed suitable if it can operate within acceptable financial bounds and exhibit reasonable performance. A number of studies present strategies to tackle a significant hurdle in processing a full-node blockchain: managing storage space. These include implementing distributed computing and parallelization [16, 24, 94, 126, 129], as well as using in-memory databases [89]. While these methods increase storage efficiency, they bring the requirement for many server instances with adequate memory. It can be postulated that the implementation of such solutions might necessitate significant financial outlays.

3.3 Proposed Architecture: Unmediated parse via Algod Daemon and PyAlgorand SDK

The expansive decentralisation evaluation scope required for this dissertation mandates the establishment of a MainNet full node on a dependable infrastructure. On a technical note, Algorand specifies minimum hardware requirements that render personal computing unsuitable for implementation.

Considerations regarding financial feasibility remain pivotal. Drawing from the demonstrated viability in [94, 126, 130], the dissertation elects to utilize cloud computing services with judicious auto-scale management to circumvent the financial strain associated with its operation. From the methodological perspective, informed by the findings in [89], the work has embraced a single-threaded strategy. This approach has been demonstrated to outperform parallel or distributed computing configurations, especially considering the innate graph-structured architecture of the blockchain.

Consistent with the goals of the dissertation, the node has been configured in an archival, non-indexer, and non-relay mode. This setup ensures the preservation of historical data without the obligation of partaking in the communication relay. The choice to forgo the use of the Indexer stems from its demand for a distinct instance, which essentially multiplies the required investment, as illustrated by the simulation presented in Appendix C.

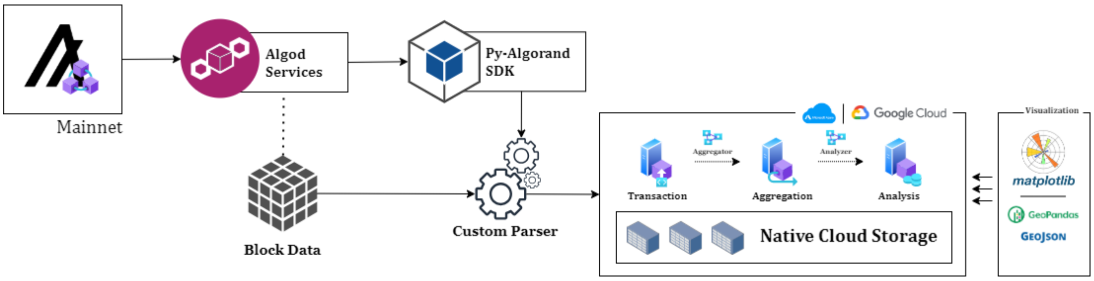

Figure 3.2: Deciphering block data facilitated by Algod and Py-Algorand SDK

The architectural framework synergizes the Algod daemon with the Algorand SDK. Investigations indicate that the Algod daemon offers a range of API calls proficient in retrieving the majority of necessary data from synchronized blocks. These calls can be accessed by the custom parser through the integration of the pertinent SDK version. Notwithstanding the comprehensive functionalities both tools offer, it has been ascertained that specific intrinsic details within the block remain inaccessible.

In light of the most recent version of Algod daemon and PyAlgorand SDK, their capacity to extract proposer attributes from the blocks is revoked. This specific data is pivotal for assessing the PPoS protocol. Upon thorough observation, it was discerned that such information is exclusively retrievable from the encoded block certifier section, which exists in tandem with the block header and block content sections. To address this limitation, the custom parser was engineered as well to directly extract and decode the necessary data from the saved blocks.

3.4 Decentralisation Measurement Basis

A plethora of methodologies have been proposed to quantify decentralisation across various taxonomies owing to the wide-ranging interpretations. To render a holistic understanding of the scrutinized blockchain, the dissertation's measurement of decentralisation is anchored to literature that offers a significant depth of granularity.

Leveraging essential features such as *safety* (which ensures system integrity), *liveness* (which guarantees consistent system responsiveness through transactions and updates), *privacy* (which shields real-world identities), and *stability* (related to operational resilience and market sustainability), the measurement will explore seven of the eight dimensions as outlined by [90], namely:

Hardware, the analysis will delve into the variety of hardware apparatus used to facilitate the network and its consensus.

Software, investigate the progression of the fundamental blockchain components, in addition to the availability, application, and governance of testing and sandboxing environments.

Network, analyze the topology and the processes integral to its bootstrapping, while also evaluating resilience to potential threats to its sustainability.

Consensus, assess the manner in which enables the integration of new nodes into the consensus and pinpoint any potential focal points that may pose safety vulnerabilities within the protocol.

Tokenomics, focus on the economic dimensions and market liquidity of the blockchain, as well as the consequential control exerted over the native tokens of the ledger by specific entities.

Governance, evaluation towards the democratic facets of the ecosystem's improvement proposals, the strategies for resolving conflicts, and the distribution of resources for research and development efforts.

Geography, aimed at evaluating jurisdictional distribution and the socio-technical-economic factors connected to the network's sustainability and resilience.

The API layer, while intimately related to the aspect of safety, lacks adequate information concerning native token wallets and the associated data necessary for a thorough evaluation. The only information that can be found pertains to the decentralized application (dApp), which likewise provides limited details [42]. In order to circumvent the potential issue of arriving at an inconclusive resolution, this dissertation opts to omit the discussion pertaining to this layer.

3.5 Used Metrics

Established quantitative indices including the Gini and Nakamoto coefficients, and the Herfindahl-Hirschman Index (HHI) are employed contingent upon the relevance of the layer under assessment, guaranteeing valid and insightful outcomes.

3.5.1 Gini Coefficient

$$G = \frac{\sum_{A_i \in A, A_j \in A} \left[NB_{A_i} - NB_{A_j} \right]}{2\,|A| \sum_{NB_{A_j} \in NB} NB_{A_j}}$$

Traditionally employed for assessing wealth distribution disparities within a populace, the Gini coefficient can also be adapted to represent the unequal distribution of power among participants (NB_{A_i}) in a blockchain ecosystem [97, 99]. A reduced coefficient indicates the resilience of a blockchain infrastructure against potential collusion threats.

3.5.2 Herfindahl-Hirschman Index (HH Index)

$$HHI = \sum_{i=1}^{N} (S_i)^2$$

The index serves as a metric that captures the extent of potential anti-competitive conduct by evaluating the relative size of each entity within the domain [96, 123, 143]. Although it overlooks the possible ramifications of semi-cooperative actors [103], an increase in the index value is typically viewed as a sign of centralization, which might undermine competitive dynamics.

3.5.3 Nakamoto Coefficient

$$N = \min \left\{ k \in [1, \cdots, K] : \sum_{i=1}^{k} \rho_i \geq 0.51 \right\}$$

The coefficient is acknowledged as a more definitive metric relevant to security [11, 99]. It quantifies the proportion of entities that, when collaborating in a collusive manner, could threaten the system's integrity. A higher coefficient implies greater vulnerability, *per se*.

3.6 Data Sets

For a precise measurement of the intrinsic degree of decentralisation, it is incumbent on the fusion of both on-chain and off-chain data sources, thereby offering a multitude of informational perspectives.

3.6.1 On-Chain Data

On-chain data refers to information directly recorded on the blockchain. Unlike prior studies which chose to omit the blockchain's initial adoption phase because of its limited popularity [99, 130], this dissertation initiates the analysis from the Algorand's genesis to provide comprehensive and long-term insights into both network and agent dynamics. Given its direct relevance to the measurement of decentralisation, the data extracted from the blockchain comprises:

Addresses, representing the daily tally of entities participating in the blockchain's operations, which sets the foundational populace for the ecosystem.

Protocol Participation, pinpointing instances wherein designated addresses pledge to contributing to the protocol's safeguard.

Block Proposal and Block Reward Distribution, unveiling the quantity of blocks produced, the degree of influence exerted by certain addresses, and the monetary advantage derived from this influence.

Balance, Payments and Asset Transfers, pertaining to monetary volume, transaction frequency, and distribution patterns to deduce the breadth of economic activity on the blockchain.

3.6.2 Off-Chain Data

On-chain data does not fully capture the real-world factors that are integral to the operation of the blockchain. Neglecting to integrate off-chain data into the evaluation of decentralisation within relevant layers could lead to fallacious conclusions. For instance, even though a system might be theoretically decentralized based on on-chain data, if one organization maintains the authority to unilaterally modify the protocol, the degree of decentralisation is called into question. Likewise, any demographic concentration of nodes could jeopardize assertions of decentralisation.

Apart from the fundamental technical specifications of Algorand, which are further enriched by its business and governance documents and will be factored into analyses across all layers, the specific off-chain data utilized for the assessment of decentralisation includes:

Strata	**Related Information**
Hardware	Profiles of identified relay nodes' hosting providers
Network	- Documented network issues across continents and regions - Global underwater cable network connectivity map
Tokenomics	- Strategy and execution of token allocation - Secondary market's capitalization and exchanges platforms
Governance	- Periodic governance projects and community involvement program - Historical data on the project's codebase on GitHub - Details on project progression and related costs
Geography	- SRV records of the `mainnet.algorand.network` domain - Positional data of the network's relay nodes - Worldwide report on stability concerning disasters - Regulatory measures on blockchain and digital currencies by country - Gross Domestic Product (GDP) by country - Index of political stability - Grade of investment and cost of starting business by countries

Table 3.2: Utilized off-chain data on the measurement

Chapter 4

Validation

Harnessing the unique characteristics and functionalities of the Algorand ecosystem, the suggested architectural framework together with the custom parser deployed around prudent cloud configurations demonstrates its aptitude in acquiring the necessary data sets of on-chain data.

4.1 Node Installation and Bootstrapping

The Algorand network offers two viable methods for node installation: one through the deployment of publicly distributed binaries and the other via an updater shell script. The first approach, though relatively user-friendly, necessitates manual interventions. Due to its adaptability to cater to a broad spectrum of purposes, the second method is perceived as the versatile option. For ease of reference and to streamline potential subsequent work, a fully automated shell script proven to set up the node and commence the bootstrapping procedure with a single input is provided in Appendix D.

Compellingly, the bootstrapping process exhibited a discernible decline in performance over time. Blocks from the genesis were processed expeditiously, with a rate of approximately 8.4M blocks in the inaugural week. However, in subsequent weeks, there was a consistent decrease in this rate, amounting to a 26% weekly reduction as delineated in Appendix E. Due to the temporal limitations tied to the dissertation, synchronization was discontinued at the outset of the sixth week, resulting in a final tally of 17.4M synchronized blocks. Given the absence of evidence suggesting over-utilization of the specified cloud computing resources, this phenomenon is hypothesized to be linked to the ledger synchronization performance.

4.2 Architecture Evaluation

The efficacy of the secondary approach utilized by the custom parser, characterized by its direct decoding of the locally replicated ledgers, has been demonstrated to exceed the swiftness and efficiency of the API calls. Such superiority is ascribed to the avoidance of possible congestion and the overheads related to data complexity.

4.2.1 API Calls Performance Through PyAlgorand SDK

The PyAlgorand SDK provides the custom parser a capability to engage with the ledger using the Algod daemon's REST API calls via the local HTTP protocol. This results in a significant utilization of computational resources over a prolonged duration.

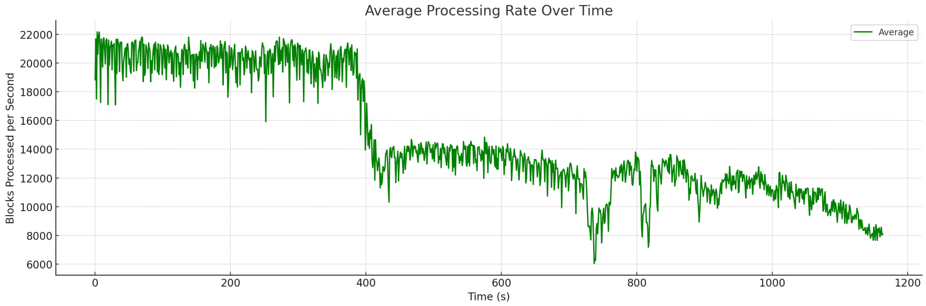

Figure 4.1: Block parser performance using Algod daemon and PyAlgorand SDK

4.2.2 Direct Decoding of the Block Data

Compared to the previous method, the following diagram highlights the marked advantage of this approach, exhibiting a processing rate that is roughly five times higher.

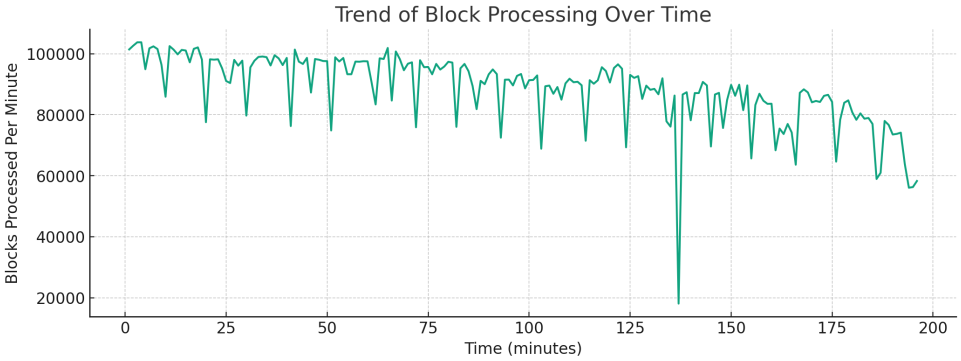

Figure 4.2: Block parser performance through direct access

A comparative analysis between data extraction through API calls and direct access to the encoded blockchain data in the Archival Node is made available in Appendix F. This significant difference in pace might potentially be traced back to the inherent operational characteristics. Direct decoding of the synchronized ledger enables local interaction, circumventing procedures such as serialization and deserialization of data thus promoting accelerated acquisition of required on-chain data.

However, it is essential to underscore that the current decoder integrated into the custom parser is operational only for the block header and certifier sections. This limitation arises from the specialized encoding protocol employed in the block content section, leaving the block data decoder unearthed to date.

4.3 Parsed On-Chain Data

The developed custom parser adeptly parsed a vast array of on-chain information as the intended datasets. Given the time constraints previously outlined, data sets were collected across two distinct frames to provide overarching insights into Algorand. The first stems from the network's inception phase, beginning on June 12, 2019, and continuing until November 17, 2021. The subsequent period focuses on the network's current status as of July 31, 2023. These combined datasets necessitated a storage capacity of 94GB, originating from 1.03TB of synchronized data blocks. In conjunction with the computational functionalities, the monthly operational expenditures were recorded at GBP 240, consistent with the model detailed in Appendix C.

Data	**First Time Frame**			**Second Time Frame**		
	Block round		**Records**	**Block round**		**Records**
	Start	**End**		**Start**	**End**	
Addresses & Balance	0	17.471.695	15,197,214	30,930,248	30,931,249	15,198,564
Payment			48,860,963			3,346
Participation			855.712			0
Asset Transfer			364.450.896			7.359
Block Proposal			17.471.695			1.001
Block Reward			12.905.515			0

Table 4.1: Collected on-chain data sets using the developed custom parser

Chapter 5

Decentralisation Measurement

The assessment of Algorand blockchain's decentralisation primarily entails scrutinizing the proposed layers [90] *vis-à-vis* empirical data harvested from the blockchain. In exceptional instances, the amalgamation of off-chain data with is utilized to yield a finely tuned measurement basis. In conjunction with the pertinent metrics, salient insights concerning Algorand's decentralisation have been brought to light.

5.1 Hardware

In a homogeneous hardware landscape, regardless of its technicality or administrative orientation, there exists the potential to compromise the integrity of the network. On the other hand, a variegated hardware ecosystem underpins a more robust network by ensuring there's no single point of failure within the system as a whole.

In the Algorand network, the detectability of participating nodes is considerably masked. Contrary to other widely adopted PoS blockchains, there exists no service nor public repository providing intricate details about the nodes in question. As a result, the evaluation of this layer is principally concentrated on the relay nodes.

Relay nodes in the Algorand network operate as the primary infrastructure, demanding a high degree of connectivity and efficient propagation of messages across the nodes. Stimulated by this requisite, Algorand strategically positions its relay nodes amidst a diverse range of services, each of which are contracted under commercial terms. The Algorand Foundation, aware of its pronounced impact on operational sustainability [68], regularly revisits these agreements.

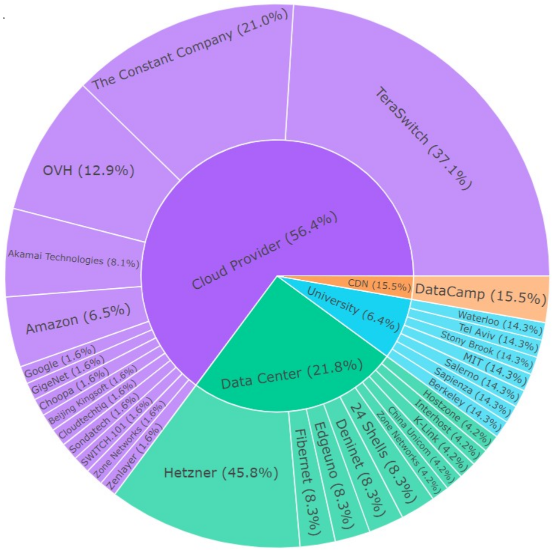

Figure 5.1: Hosted relay nodes

The distribution and management strategies of these nodes reveal a marked degree of centralization and permissioned features. Initial analysis revealed that a scant five companies controlled 60% of the 110 nodes. Moreover, although these nodes do not partake in the consensus or functioning as staking pools, they function as the hubs that weave the consensus. Should the majority of the nodes under the command of a company act antagonistically, transactions from participating nodes could conceivably censored [20, 95].

On the contrary, while the protocol theoretically allows the community to establish relay nodes, this option is often neglected due to the steep computational demands coupled with absence of incentivization. Furthermore, since relay nodes serve as the custodians of information on the blockchain and fall under the control of the foundation, they become particularly susceptible to compliance pressures from regulatory mandates or censorship [135, 144]. This includes potential interjections from bodies such as the United States Securities and Exchange Commission (SEC) or Office of Foreign Assets Control (OFAC) [20] as well as European Securities and Market Authority (ESMA) [9, 128]. Under these conditions, *qua ratione*, the foundational infrastructure tends gravitates towards centralization, which poses risks to its security and liveliness.

5.2 Software

Algorand protocol is fundamentally anchored in a software layer that supports the democratization of control. The architecture conceived serves to simultaneously allow iterative growth of the MainNet. Nevertheless, the implicit barriers to participating in the woven fabric of the network's communication, as well as the homogeneity in software choices, might inadvertently hinder the progression towards decentralisation.

5.2.1 Participation in Upholding the Integrity of the Protocol

In a departure from other PoS protocols, Algorand provides a unique approach toward decentralisation by not mandating the operation of full nodes as validators. The `fast-catchup` mechanisms not only drastically curbing the substantial computational demands characteristic of full nodes but also significantly shortens the time required for a node to start participating in the network.

This reduction in entry barriers, in terms of capital and hardware requirements, essentially encourages a wider pool of participants to help secure the protocol. As of July 31, 2023, the network boasted 1,216 participating nodes [69]. However, crucial details required for an in-depth grasp of the background technicalities involved in its establishment are agnostic and remain enigmatic, with no discernible information from off-chain sources that is contrary to Ethereum [50]. This opacity limits the visibility into the actual distribution of control within the network.

Despite the low entry barriers, it is of significance to highlight that unlike other networks [2], Algorand currently does not feature an active incentive program for its participation nodes. As of May 2022, the program has been completely supplanted by the Governance reward [66]. Although this could be a deliberate strategy to prevent potential disparities, it inadvertently exposes the network to the 'lazy validator' problem [91], as the participants may opt to leave the network due to diminishing interest—a rational behaviour of *'homo economicus'*.

Furthermore, Algorand mandates Algod daemon as a uniformity regarding the software needed for network interaction. The reliance on a single software introduces to complete centralization of this context [39, 146]. Paired with the lack of incentives that could result in a decline in the blockchain's health, this monopolistic scenario of software distribution raises concerns about the system's overall safety and liveliness.

5.2.2 Facilities for Isolated Experimentation

To foster inclusivity within its ecosystem, Algorand has established TestNet and BetaNet to cater to varied testing actors and backgrounds. In opposition to a plethora of prevalent blockchain systems [136], the testing architecture hinges on the experiment and final rehearsal stage for enhancements destined for the MainNet. This methodology closely mirrors practices found in entities such as Cardano [70] and Filecoin [51]. As unambiguously suggested by their designations, the tokens transacted within these environments do not carry any monetary value.

Both the BetaNet and TestNet environments continue to be operational to date, with ongoing maintenance undertaken by the Algorand developer. Dynamic telemetry information procured from these testing landscapes are examined at a grand scale, irrespective of the tester's underlying motives, whether adversarial or well-intentioned. The insights gleaned from this data then feed back into the development trajectory, aiding in the fine-tuning and augmentation of the primary network. These frameworks fundamentally provides an auxiliary layer of provision, assuring the safety, liveness, and stability of the blockchain under a variety of real-world circumstances.

5.3 Network

With a goal of facilitating global adoption, the Algorand architecture institutes participation nodes and relay nodes to guarantee swift information propagation across the network and timely finality. Despite the distinct roles, the heterogeneity of the nodes and their intrinsic technical details are essential in ensuring safety and liveness.

5.3.1 Participation Nodes

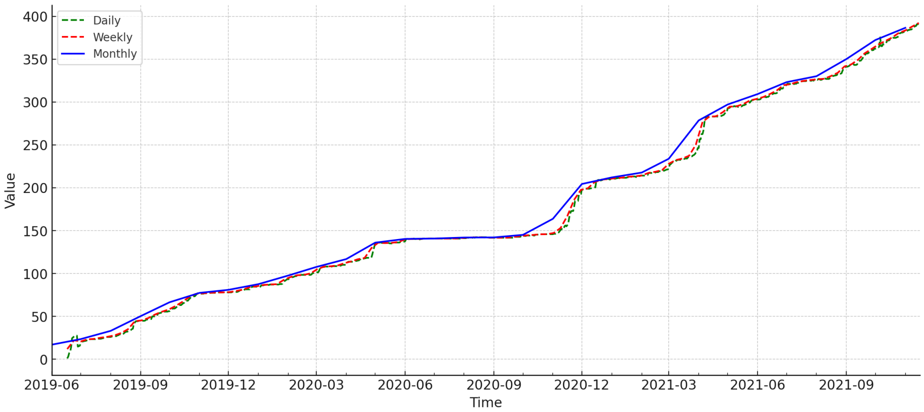

Figure 5.2: Nakamoto coefficient of participation nodes over time

The sparser the number of participation nodes in operation, the more susceptible the Algorand network becomes to an adversary overwhelming its security protocols. Upon quantitative assessment, a pronounced upward trend emerges in the count of participation nodes with an appreciable growth rate of 38.45%. This signifies an advancing decentralisation while enhancing safety, as underscored by the maximum Nakamoto coefficient of 393.32 which could be attributed to the reasonable computational demands required to operate.

5.3.2 Relay Nodes

Algorand has devised a holistic strategy to enhance the robustness against network partitioning stemming from natural disturbances or deliberate adversaries. On the physical plane, to reinforce the continuity and steady functionality of the system, relay nodes are strategically placed across diverse locations as detailed in Geography section.

At the foundational level, three pivotal technical mechanisms have been established. First, the consensus protocol implements the Partition Recovery Mode (PRM), which is triggered when there is a disruption in network progression that extends beyond a specified timeframe [31, 107, 142]. Second, the agreement protocol unveils a singular non-forking ledger [6, 31, 56, 85, 115]. Lastly, the architectural design ensures that the community has the capacity to sustain network operations [59].

Hypothetically, within the PRM state, the relay nodes are designed to identify alternative routes to bolster network resilience. When confronted with disturbances, such as those related to regional network incidents [36, 52, 86, 113, 132] and network partitioning attacks [8, 81, 102], affected nodes pursuit to redirect towards the closest functional relay nodes. Simultaneously, these nodes will continuously send out recovery signals and maintain a state of preparedness to swiftly re-establish connections once they undergo protocol validation. While this approach invariably results in an increase in transaction latency, it is vital in averting complete unavailability [29].

In scenarios of intensified gravity, the network is engineered to prevent the emergence of a fork, irrespective of consensus disputes or adversarial interventions. Formal verification substantiates that the Algorand network's likelihood of encountering a fork is practically non-existent (10^{-18}) [6, 105, 115]. When faced with analogous conditions, Bitcoin or Ethereum are more susceptible to experience a network fork [8, 81, 102]. The appearance of a divergent ledger jeopardizes safety and stability.

Altera parte, should the foundation decide to discontinue the project, rendering the relay nodes non-functional [68], Algorand's architectural design ensures its persistence irrespective of the foundation's continuity. This endurance can be achieved either by establishing private relay nodes or by instituting a communal SRV record registered within the Domain Name System (DNS) for the MainNet.

Within the rationale previously outlined, the CAP theorem continues to be upheld [137], and Algorand conforms to this principle. While the blockchain network does provide a level of liveliness, its fundamental emphasis is on ensuring safety to maintain the consistency of the blocks.

5.4 Consensus

Algorand's genesis was driven by a dedication to heightened security. It notably challenges the conventional miner-user dichotomy [56, 59]. Concurrently, the protocol empowers nodes with the discretion to choose their peers, unaffected by their individual stakes, while preserving the capacity for every node to access messages [31, 76]. These factors collectively rendering the notion that every node is crucial to the safety.

5.4.1 Level of Unpredictability and Disparity in the Block Proposal

A high degree of block formation randomness enhances the blockchain resistance to manipulative or malicious intents. HH index is utilized as a quantitative measure to assess the degree of uncertainty associated with the election of a node to serve as the block proposer within the consensus.

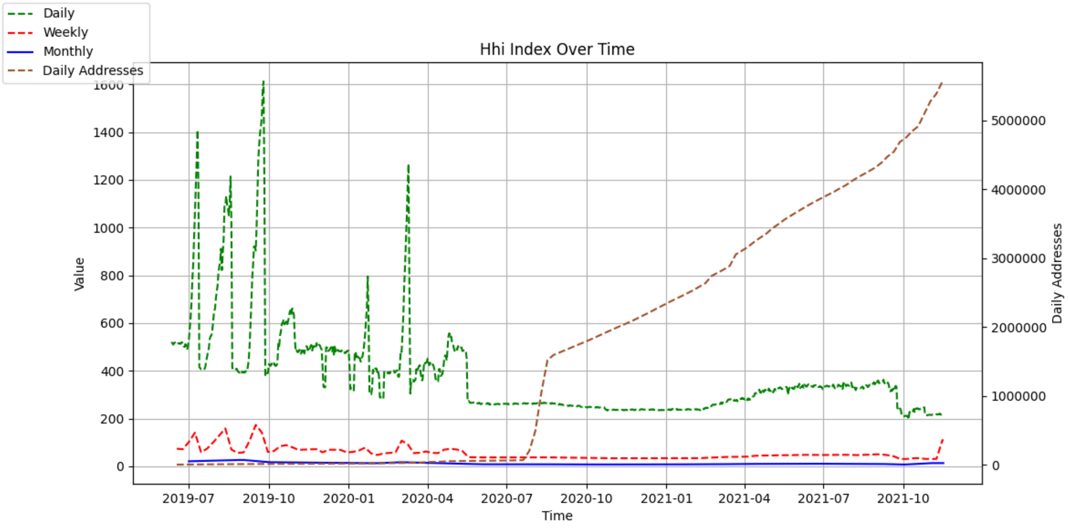

Figure 5.3: Block proposal unpredictability rate following to HH Index

The temporal dynamism inherent in the resulting index is intriguing. It was ascertained that the nascent stages of Algorand's deployment were marked by pronounced concentration fluctuations. Such variations strikingly parallel the evolutionary trajectory observed in Bitcoin [99]. Throughout the inaugural six months, there were marked daily oscillations, with the apex reaching an index of 1,612, a figure that denotes a moderate level of centralization [30, 45, 118, 134]. By the onset of its sophomore year, the volatility began to taper, transitioning the blockchain into a state typified by elevated randomness.

As a supplementary to evaluating randomness, the Gini coefficient is invoked as a statistical metric to assess inequality. Consistent with the aforementioned observations, an examination of the inaugural year reveals marked variations in inequality.

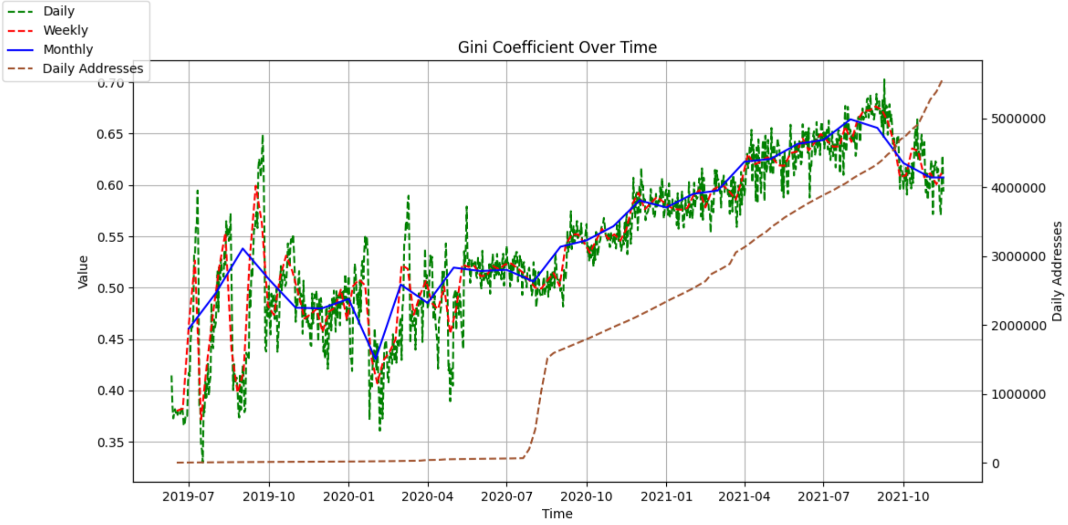

Figure 5.4: Gini variability in the first year of MainNet deployment

The figure alludes to a middling degree of inequality. The range, spanning from a minimum coefficient of 0.33 to a maximum of approximately 0.70, displays a diverse array of wealth distribution patterns. However, an unmistakable upward trajectory is observable, intensifying the level of inequality as the platform evolved. These findings illustrate the system's elevated vulnerability to risks concerning safety and stability.

The HH Index and the Gini Coefficient manifest a discernible linkage. A diminution in randomness typically equates to a heightened level of inequality. This correlation is apparent in the last 1,000 blocks as of July 31, 2023, which displayed an HH Index of 280.62 and a Gini coefficient of 0.41. The opposing evidence observed over the extended time frame, in tandem with the upward trajectory in address registrations, might hint at the emergence of dormant or ephemeral addresses.

5.4.2 Degree of Resiliency

Blockchain users necessitate assurance that any valid message they initiate will be integrated into a block and subsequently verified through consensus. In scenarios where a set of consensus nodes succumbs to compromise or partakes in orchestrated malicious actions, it could strive to hinder the network's consensus achievement on new blocks.

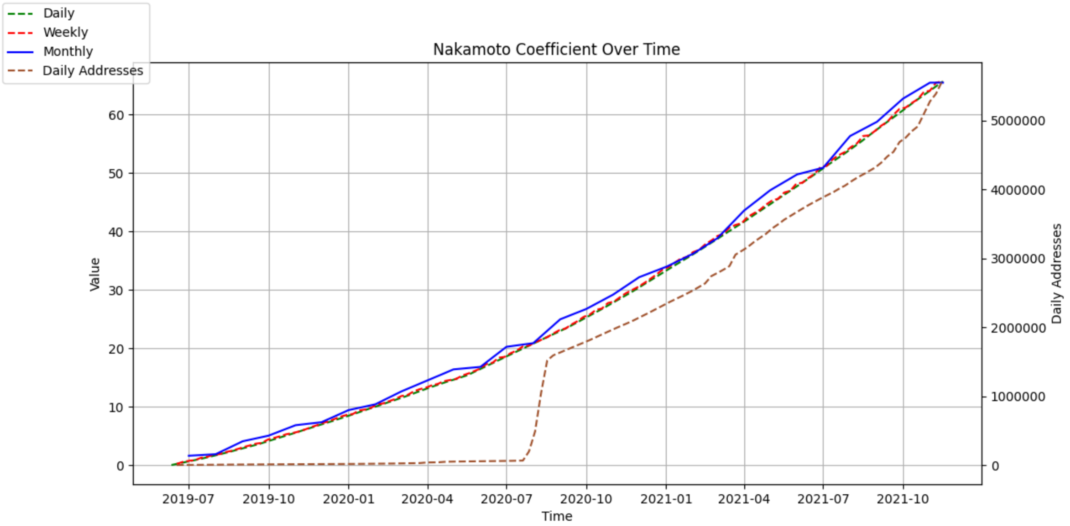

Figure 5.5: Resiliency rate of block proposal stage depicted by Nakamoto coefficient

Upon measuring the block proposer information with its respective proportion of total proposed blocks on each of the operation day, it has been unearthed that the Nakamoto coefficient is incremental days by days. The mean value for the coefficient was found to be approximately 26.4. At one extremity the blockchain attained a pinnacle of decentralisation, corroborated by coefficient of 65.8. This peak performance surpasses that of other PoS blockchains [71, 151]. These include Solana (31), Avalanche (27), Thorchain (26), and Terra (15). The most recent Algorand blocks as of July 31, 2023 also indicate relatively high coefficients at 16.

The temporal line plot of the Nakamoto coefficient clarifies the network's propensity towards amplified decentralisation. This overarching trend, presuming no single participant holds multiple addresses within the network, signals a continual shift in network governance from a confined cohort of participants towards a more distributed arrangement, thereby augmenting its safety and stability.

5.5 Tokenomics

At its core, Algorand economic structure designed to have a finite supply and it can limit the issuance of tokens through mechanisms that promote scarcity. The founding entities also device schemes that confer specific services, voting entitlements, or governance privileges to the holder, fostering incentives for retaining rather than trading these tokens on exchanges. This paradigm might inadvertently obstruct the pathway to decentralisation.

5.5.1 Token Distribution through Primary Market

Evaluating decentralisation in the context of the initial token distribution model requires an appreciation of the inherent diversity among investors. Obtaining this crucial information is often thwarted by privacy concerns, as detailed records of investor identity often remain undisclosed [83]. This opacity holds true for Algorand as well.

Another pressing concern pertains to centralization tendencies stemming from the confined scope of the primary market. The market functions within an indistinct regulatory framework [82, 152], thereby leading to a minting process susceptible to potential abusive action [154]. Manifestations of this concern include structured sales, inducements offered to early network contributors, and compensations reserved for the founding team.

Period	**Injected**	**Private Transaction / Minting**					
		ACT	**SS**	**NRR**	**CONT**	**COMP**	**PER**
Jun 19 - Nov 19	440M	25M	-	330M	-	75K	75%
Nov 19 - Feb 20	126M	-	17M	7.4M	-	125K	19%
Feb 20 - Sep 20	549M	-	169M	191M	-	200K	66%
Oct 20 - Mar 21	1.34B	-	45.3M	1.06B	50M	395K	87%
Apr 21 - Sep 21	1.37B	-	47M	1.07B	55.3M	960K	85%
Oct 21 - Mar 22	582M	-	42.2M	328M	-	1.3M	64%
Apr 22 - Sep 23	311M	-	16.5M	14M	-	1.1M	10%
Oct 22 - Mar 23	338M	-	101M	-	-	2.3M	31%

Table 5.1: Distribution of token injection through Auction (ACT), Structured Selling (SS), Node Runner Rewards (NRR), Contingency Reward (CONT), Compensation (COMP)

The table offers empirical evidence highlighting the persistent emergence of private bulk token issuance in the primary market, with manifold implications concerning the potential for centralization. The structured sales were designed to preclude front-running manoeuvres that might amplify market volatility. However, even though both structured sales and auctions are overseen by an intermediary, only auctions are explicitly delineated as public sales [60–68, 122]. This opaque nature might impede more extensive participation and, in a severe context, could inadvertently allowing certain individuals to leverage transactions underpinned by privileged information.

Furthermore, in the preliminary stages of token distribution, there appears to be a pronounced bias towards the foundational entities and early proponents. Designed as a strategy to sustain network functionality, coupled with the conditional rewards aimed at mitigating trends or demands that could hinder the ecosystem's growth, these incentives could conceivably be extended to investors as a gesture of appreciation for their backing. Concurrently, the upward trajectory observed in token remuneration granted to the members of the foundation's board and advisors with the justification that it is linked to their heightened level of involvement in the foundation invites further scrutiny.

This scenario intensifies the propensity towards centralization, granting certain investors the capacity to exert significant influence, or even potentially initiate conflicts of interest within the foundational governance. Such influence could permeate areas of consensus, developmental trajectory, and valuation of the blockchain, particularly if these investors elect to retain their tokens for participatory consensus rather than liquidate them in the secondary market. Inevitably, this scenario heightens potential vulnerabilities, risking the blockchain's safety and stability.

5.5.2 Token Ownership

Detractors contend that the PoS system intrinsically results in the concentration of wealth and control [116], as rewards are directly proportional to one's wealth, and staking rewards exhibit exponential growth in correlation with an expanding network. A broad and equitable token distribution is imperative to cultivate an ecosystem that embodies not only robustness, efficiency, and trustworthiness, but also upholds democratic values.

Position	Addresses	Top-100	Gini (10K)	Gini Overall
9 August 2022	15,197,214	73.80%	0.955	0.997
31 July 2023	15,198,564	81.49%	0.945	0.999

Table 5.2: Algorand's wealth distribution within two periods

An analysis of wealth distribution within Algorand, based on the concluding balances across two time points, has provided significant revelations. These findings corroborate initial apprehensions arising from primary market token allocation. It exhibits a consistent trajectory towards centralization, in which the 100 wealthiest accounts exhibit a marked degree of token centralization within the network. These extremity align coherently with the Gini coefficient, registering 0.955 for the top 10,000 affluent addresses and nearing perfect inequality with a coefficient of 0.997 across the ledger. Although the equality among the former group seemed to improve by 31 July 2023, the overall coefficient, on the other hand, was deteriorating. For comparative purposes, the top three blockchains with the highest wealth inequality were Ethereum Classic (0.988), Dogecoin (0.986), and a tie between Ethereum and Litecoin (0.978) [90].

To offer a more nuanced understanding of the disparity, the Gini coefficient was also utilized from the standpoint of reward distribution. This methodology has been adopted in light of prior economic studies which ascertain the significant applicability of the coefficient in assessing income distribution [21, 33, 138], and the reality that any address holding a balance is eligible to receive staking reward. In the present context, the 'income' corresponds to the reward from the network rather than the concluding address balance.

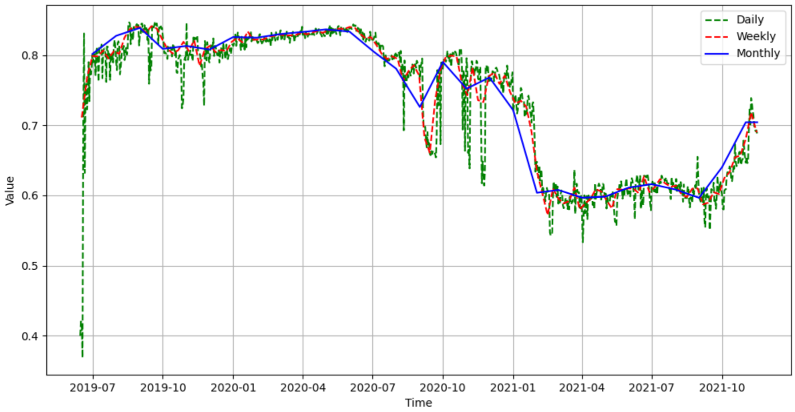

Figure 5.6: Reward distribution inequality over time

A chronological evaluation of the indicated coefficient reveals substantial homogeneity in block reward distribution. This is manifested by an average coefficient value of 0.73, accompanied by a moderate variability, as indicated by a standard deviation of 0.1. Such findings suggest that a limited number of addresses are accumulating a progressively larger share of the block rewards. Consistent with prior assessments, this phenomenon could amplify potential threats to the system's safety and stability.

The validation of these findings for July 31, 2023 is constrained. No reward distribution was observed on this date, consistent with the ledger's policy that schedules allocation phases at intervals of every 500,000 blocks. [53].

5.5.3 Secondary Markets

The secondary market, also known as exchanges for token trading, is crucial for determining decentralisation in the publicly accessible domain. Subsequent to this, a diagram offers a synopsis of this market for 60 PoS blockchains, representing data as recent as July 27, 2023 [35]. To facilitate a balanced comparison and maintain legibility, the data were normalized using logarithmic scale. Detailed values corresponding to the metrics for each blockchain are provided in the Appendix G.

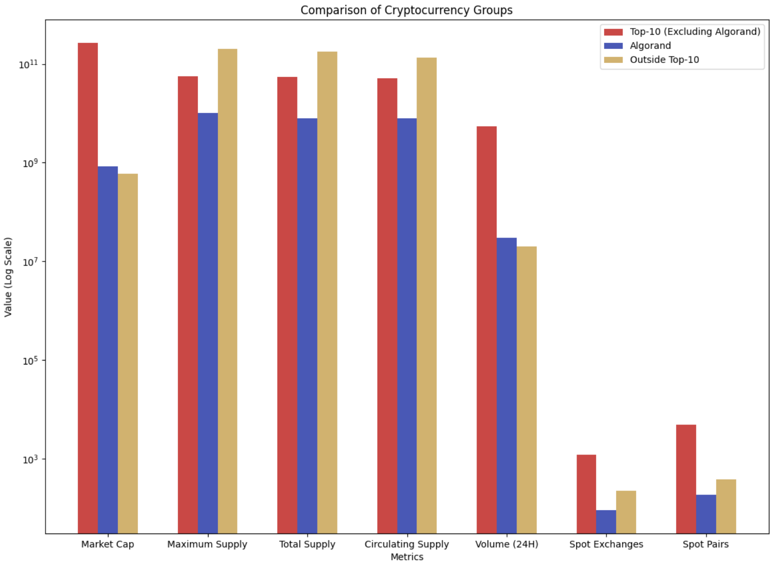

Figure 5.7: Comparison of Algorand to Top 10 PoS Blockchain and the rests

Algorand emerges with a market capitalization that exceeds platforms outside the top echelons. This prominence is accentuated by its evident growth potential. The increase in its total and circulating supply not only suggests a widespread dispersion throughout the network but also hints at its expanding influence. On the other hand, the platform faces challenges. Its 24-hour trading volume, indicative of market activity and liquidity, lags notably. This may signal limited market inclusivity and decreased trader interest. Furthermore, Algorand's lean figures of maximum supply may indicate both a potential network constraint or intrinsic scarcity value.

Regarding market accessibility, Algorand's presence is notably restricted to a limited array of exchanges. This limitation is compounded by the deficiency in the diversity of its trading pairs in comparison to its peers. Such a scenario, while underscoring a short level of accessibility and variety in its token transactions, simultaneously hints at a missed opportunity. A broader distribution could potentially augment stability as well as advancing both privacy and liveness.

5.6 Governance

Algorand employs a dual-governance system, promoting active community participation in directing both the platform's trajectory and tangible strategic ventures. Financially, the foundation dedicates a budget and designates an account on the network for the sustenance of improvement projects.

5.6.1 Ecosystem Initiatives and Resolution of Disputes

The Algorand Community Governance (ACG) [106] follows a democratic approach in resolving disputes and envisioning potential improvements. This program welcomes all eager participants and unfolds within a forum-style context. Contrasting with another prevalent blockchain and in an effort to ensure continuous growth, the program runs quarterly, provides governance rewards, and culminates with a transparency report detailing activities, reward allotments, and improvement decisions.

Governance Period	Issue(s)	Governors	Commitment	Rewards
Sep 21 – Dec 21	1	51.7K	1.76B	60M
Dec 21 – Mar 22	1	32.8K	2.81B	70.5M
Mar 22 – Jun 22	2	35.6K	3.54B	70.5M
Jun 22 – Sep 22	2	27.5K	3.65B	70.5M
Sep 22 – Dec 22	5	28.8K	3.76B	70.5M
Dec 22 – Mar 23	2	22.5K	3.82B	68.2M
Mar 23 – Jun 23	6	28.3K	2.55B	56.2M

Table 5.3: Algorand community's engagement with the ACG

The extent of community engagement that intensified by the incentive program as illustrated in the table, stands unmatched among other Layer-1 blockchains [88]. Conversely, insights from the other end of the commitment spectrum is enlightening. Declines in the number of governors often result from eligible governors' inability to sustain their balance above the requisite commitment threshold throughout the cycle [47]. This can be due to market upheavals or other factors, including missing the voting deadline, which concludes at 6 AM [109]. Although the displayed commitment indirectly strengthens the fundamental decentralisation characteristics, such as safety and stability, ample opportunities remain to further amplify the engagement.

5.6.2 Development Activities

The analysis of Algorand's codebase activities before its inception until 29 June 2023 unveils discernible patterns. During this timeframe, significant fluctuations were observed both in the number of commits and in the affected lines of code. At the outset, the coefficient reached its conceivable peak value, indicating a marked disparity in contributions.

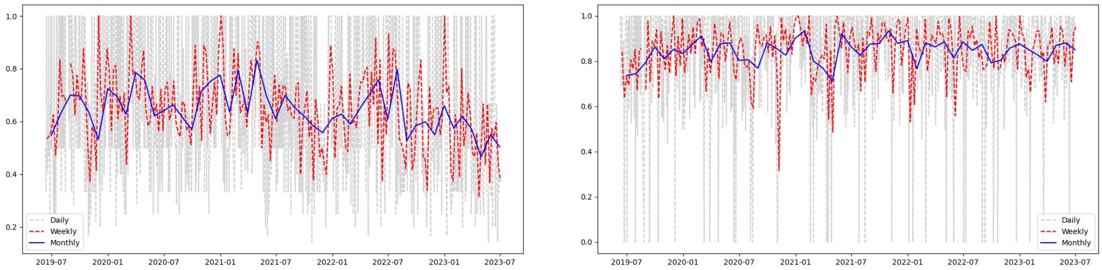

(a) Commits made

(b) Lines of code

Figure 5.8: Gini coefficient measurement on development activities

A deeper statistical analysis revealed that the average Gini coefficient for commits stands at approximately 0.65, while it rises significantly to 0.85 for lines of code. Upon closer inspection, maximum inequality was evident in the development contributions, underscoring a stark disparity among contributors to Algorand's development. Given the prevalence of modular or distributed coding techniques in modern project development, a disparity in 'commits' does not unequivocally imply a similar disparity in 'lines of code'. Nonetheless, in more evident cases, these coefficients correlate with a subdued contribution from a broader group or an undue dependence on, or a predominant role assigned to, a specific entity as identified in Appendix H.

5.6.3 Financing

Sustainable blockchain systems require a stable source of funding. Given the intrinsic goal of decentralisation that is pivotal to blockchain, the financing mechanisms for system preservation and growth should be devoid of risks stemming from centralization nor reliant on a potent minority.

The financial support of blockchain inception are generally facilitated by founding entities. Algorand conducted equity raises totalling USD 66M in 2018 to finance the development of the protocol [41, 46, 122]. Nonetheless, there exists a noticeable dearth of information pertaining the spending and the investor's ownership share transparency.

Contrariwise, Algorand earmarks a significant volume of coins as a core aspect of its preliminary coin distribution plan to emphasize and ensure long-term sustainability. Nearly half of the 10B maximum supply has been assigned to bolster continuous operations, development and stimulate inclusivity [54]. Additionally, Algorand utilizes a specific address to gather all transaction fees generated across the network. As of 29 July 2023, the balance of this address was at 1.41M Algos, which could be deployed towards ecosystem development and support [55].

Sources	**R&D**	**Ecosystem Development**	**Community Engagement**
Fiat/Stablecoins (USD)	48.2M	2.5M	37.6M
Algos	270.2M	79.2M	9.2M

Table 5.4: Investment allocation in the period subsequent to the MainNet's initiation in June 2019, extending up to March 2023[66, 68]

In line with the established ACG program, propositions regarding treasury expenditures are incorporated into the voting mechanism during recurrent governance cycles. This includes, but is not restricted to, protocol enhancements, research initiatives, funds for grants and other community projects [47, 57, 58, 106, 109–111]. The secure means of financing, when coupled with decision-making steered by the community, positively reduces risks associated with potential centralization that could unintentionally stagnate innovation or favour specific interests within the platform [5, 154], thus fostering an environment of liveliness and stability.

5.7 Geography

Geographical dispersion plays a pivotal role in sustaining the security, resilience, and inclusivity of a blockchain operation. Relay nodes, acting as the primary infrastructure of the network, are susceptible to multitude of risks based on their physical locations. These hazards include physical disruptions, natural disasters, political instability, and unsupportive regulatory environments.

The following global density map offers an in-depth visualization of the global distribution of relay nodes. Their geographical positions have been ascertained by iterating through the MainNet SRV record and employment of public IP GeoLocation API provider [87].

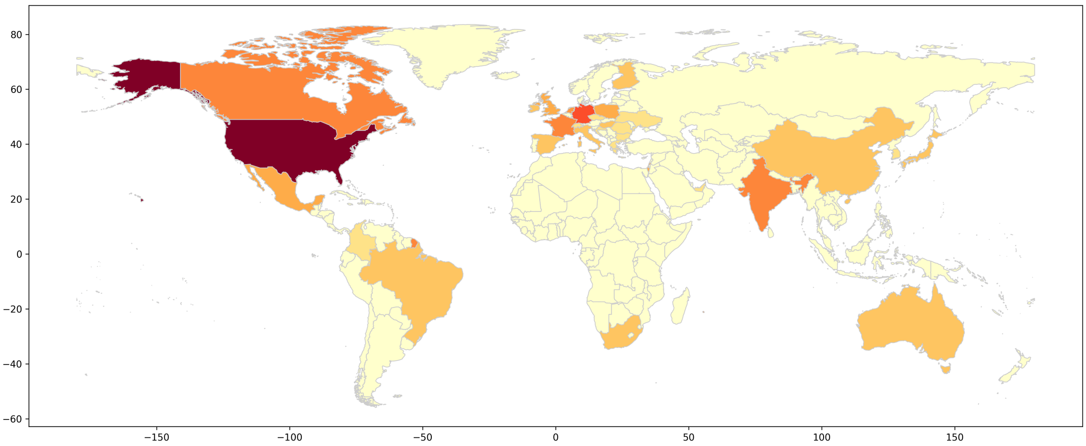

Figure 5.9: Relay nodes distribution density

The strategic positioning can be seen as a strength of the network, enhancing its liveness to potential localised issues. In the event of a disruption in one locale, the overall liveness of the network is likely to remain unaffected due to the presence of operational relay nodes in other locations as manifested in Appendix I.

The propagation of Algorand relay nodes across the United States, Germany, Canada, India, and France suggests a deliberate focus on nations that exhibit notable economic strength and varied population densities [13, 14, 43, 72]. The emphasis on nations like the U.S. and Germany, both with high GDPs, showcases the relationship between economic power and technology. Countries such as Canada demonstrate the network's adaptability to low-population, high-economic areas, while India underscores the dynamics of dense regions with emerging economic prospects.

Physical security analyses of the recent relay nodes indicate a preference for geographically stable regions [124], with Europe housing 36% of the nodes, underpinned by a strong intercontinental connectivity as elaborated in Appendix J. However, some nodes are located in disaster-prone areas, such as Los Angeles and parts of Asia. Legally, a considerable number of nodes reside in countries with blockchain-friendly regulations [34, 98], such as France and Singapore. Still, nodes in areas like China face potential regulatory challenges due to their stricter cryptocurrency stance. Regulatory details pertinent to countries hosting relay nodes are provided in Appendix K.

Socio-politically, the majority of nodes reside in stable nations [15, 49], including North America and Europe. Economically, despite higher operational costs for nodes in locations like New York or Singapore [14, 48], the network's diversified presence, including in cost-effective regions like India, helps balance these expenses.

Chapter 6

Discussion

The subsequent analysis of quantitatively measurable layers in Algorand, coupled with the discoveries presented in previous chapter, posits that the blockchain system demonstrates characteristics of centralization.

6.1 Limitation

Drawing concrete outcome about blockchain decentralisation presents challenges due to the intricate nature of address properties and their vulnerability to small value inclusions. Despite these complexities, the discussion assumes each address as a distinct holder with their respective balance.

6.2 Comparative Analysis Across Stratified Layers

Within the layered structure of a blockchain architecture, each constituent layer can be conceptualized as subsystems, providing a framework for characterizing and quantifying the system's decentralisation. Under the assumption that all these considered subsystems have equivalent significance, the ripple effect persists; a centralized subsystem has the potential to drive the overarching system towards centralization [11].

The Lorenz curve stands as a significant instrument for evaluating distribution magnitude. In conjunction with the Gini coefficient, which indicates decentralisation, and the HH Index, used for assessing concentration and competitiveness, it offers profound insights into disparities within specific subsystems. Considering the leftward skewness of the monetary dataset, the mean value is utilized to address lower outliers [44, 101]. This results in a benchmark value of USD 3M for tokens considered in the evaluation.

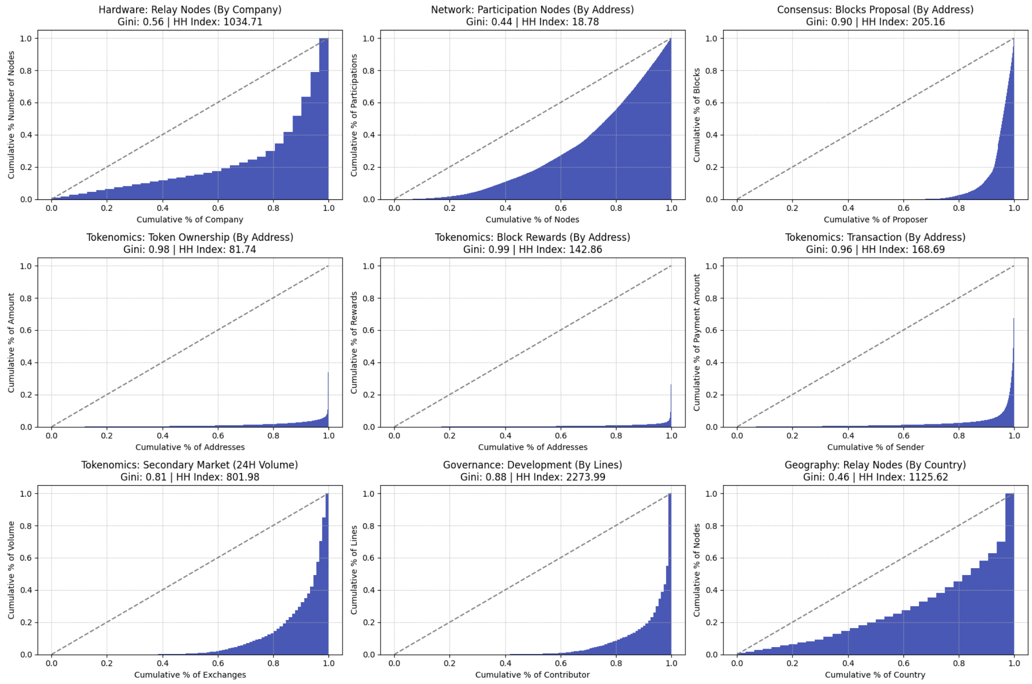

Figure 6.1: Quantitative assessment of decentralisation and inequality across layers

The panel corroborates the pronounced centralization of Algorand, echoing the empirical evidence presented in Chapter 5. Metrics show that subsystems related to consensus and tokenomics — pertaining to monetary affairs — have coefficients surpassing 0.9. Within the ecosystem, it is evident that a select cohort of entities predominantly controls the protocol's operation, manages financial transactions, and reaps most of the benefits. When contextualized in real-world terms, this value exceeds even nations with stark income inequality, such as South Africa and China [1].

This centralization pattern is reciprocated in the Governance layer, highlighted by developmental influence. In light of contemporary development techniques, the amount of affected lines of code served as the primary metric to provide a nuanced view of alterations. The documented 0.88 coefficient suggests a pronounced reliance on certain contributors. On the contrary, the Hardware, Network, and Geography layers, essential for appraising relay nodes, are perceived as fairly resilient. Their metrics indicate a balanced distribution in infrastructure setup and location.

Regarding competitiveness, despite the high Gini coefficient, all subsystems, barring Governance, display low concentration. Although the metrics are formulated for different contexts, markets with higher concentration typically exhibit greater inequality. The presence of dust balances might account for the observed divergence, necessitating further empirical studies for a comprehensive understanding.

Chapter 7

Conclusions and Future Work

Upon a thorough analysis, Algorand, despite its foundational objectives, manifests considerable centralization tendencies. Such centralization presents pressing issues for stakeholders and opens avenues for potential regulatory scrutiny, thereby offering intriguing areas for future exploration.

7.1 Conclusions

The primary aim of this dissertation was to develop a robust custom parser and to evaluate the extent of decentralisation in the Algorand blockchain. To the best of recent understanding, this dissertation represents the inaugural study that seeks to assess the decentralisation of the Algorand blockchain in a holistic manner.

On the whole, the bespoke parser devised for this dissertation has demonstrated both precision and efficiency, adeptly extracting blockchain data while optimizing resource use. Post its deployment, the harvested data has been instrumental in facilitating the examination of relevant layers, notably the Network, Consensus, and Tokenomics. By merging this with pertinent real-world information, the dataset was well-suited for evaluating seven of the eight proposed layers meant to ascertain the platform's level of decentralisation. From a quantitative standpoint, the implementation of three distinct metrics, namely, the Gini Coefficient, HH Index, and Nakamoto Coefficient, has been able to furnish vibrant perspectives of the layers under scrutiny. While each metric is designed for specific analyses, their combined use can offer a deeper insight into the complex interplay of inequality and uncertainty within decentralisation.

Over the course of its development, Algorand manifests a propensity towards centralization, tracing a path reminiscent of trends seen in other notable blockchains like Bitcoin and Ethereum. The dissertation reveals Algorand's design emphasizes resilience against network forks and ensures consistent operation, even in the absence of foundational entities or during network partitions. Besides, the global deployment of the Algorand network varies: while prevalent in developed nations with stable socio-political contexts, it adapts in developing countries to leverage economic prospects. However, the proclivity towards centralization arises are driven by inherent factors such as monetary dynamics, market incentives, and computational aspects. While this dissertation probes many layers, some parameters fostering centralization remain and require further stakeholder scrutiny.

7.2 Intuition for Stakeholders

Manifested in a real-world context, Algorand's centralization markedly surpasses those of nations with the most pronounced inequality, a phenomenon that is similarly apparent in other leading blockchain systems. Alternatively, when examined through the lens of safety, liveness, and stability, escalating values of the inequality are indicative of heightened susceptibility to adversarial intrusions. This dynamic nature demands that stakeholders rigorously evaluate it before aligning their interests.

Potential investors and regulators may have concerns regarding the distribution of holdings in Algorand. While concentrated holdings can lead to significant price fluctuations due to the trading actions of primary stakeholders, it is pertinent to acknowledge that many blockchains, including Bitcoin and Ethereum, experienced similar concerns in their early days. Nonetheless, the observed growth in extensive participant networks within Algorand may bolster its long-term credibility and trustworthiness.

In light of concerns regarding the potential implications of a centralized ledgers on decision-making, it is imperative to recognize blockchains as dynamic and evolving systems. As displayed in Algorand, governance structures are subject to evolution, and the collective community possesses the capacity to propel a more decentralized decision-making with time. From a security standpoint, while a centralized network might present vulnerabilities, the blockchain space continually grapples with and develops solutions for potential threats. It is indispensable to appraise Algorand's prospective growth in tandem with the degree of community participation and the persistent advancements in its technological and governance frameworks.

Owing to the unique complexities, prospective investors and regulators are advised to employ an enhanced level of due diligence. Key strategies for investors should encompass diversifying holdings, staying abreast of market trends, and critically assessing the governance model as well as the trajectory towards long-term viability. Regulators, serving as the custodians of the financial infrastructure, are poised with a unique opportunity to influence the emerging landscape of the market. Given the wealth concentration on a privileged group of entities and the susceptibility to market manipulation inherent in a centralized system, regulators have the potential to craft frameworks promoting prudent supervision, safeguarding investors, and upholding market integrity. As blockchain advances, striking a balance between fostering innovation and guaranteeing safety, liveness, privacy, and stability will be paramount.

7.3 Future Work

Leveraging the insights and the developed parser from this dissertation, subsequent studies might yield an expansive analysis of blockchains akin to Algorand. To further augment the present work, the ensuing suggestions are offered:

First, the establishment of a fully-synchronized node is imperative as the foundational step towards generating a holistic decentralisation measurement. In addition to cloud services' auto-scaling feature, adopting stream-based block processing can optimize computational efficiency in line with workload requirements.

Second, the observed variations in inequality and other decentralisation indicators emphasizes the evolving nature of the ecosystem. Subsequent research should strive to investigate the real-world catalysts of these shifts, along with their implications for the temporal functionality and stability of the blockchain.

Third, political philosophy suggests a reluctance to yield power, as seen in blockchain governance where foundational bodies keep voting rights despite community consensus focus. Further investigation is crucial to determine the model's resilience against potential conflicts of interest that might undermine stability.

Fourth, token allocation often appears skewed in favor of early investors and the foundation. These tokens and the corresponding nodes might be traceable by adversaries due to certain communication noise levels. Given concerns about wealth concentration, future work might explore this aspect to network partition attacks.

Bibliography

[1] OECD (2023). Income inequality (indicator). *doi: 10.1787/459aa7f1-en*, 2023.

[2] Joseph Abadi and Markus K. Brunnermeier. Blockchain economics. *PSN: Technology (Topic)*, 2018.

[3] Algorand. Algorand analytics. `https://developer.algoscan.app`, 2023. Accessed: 2023-06-01.

[4] Darcy W. E. Allen and Chris Berg. Blockchain governance: What we can learn from the economics of corporate governance. *IRPN: Innovation Policy Studies (Topic)*, 2020.

[5] Darcy W. E. Allen, Chris Berg, and Aaron M. Lane. Trust and governance in collective blockchain treasuries. *Cryptocurrencies eJournal*, 2021.

[6] Musab A. Alturki, Jing Chen, Victor Luchangco, Brandon Moore, Karl Palmskog, Lucas Peña, and Grigore Roşu. Towards a verified model of the algorand consensus protocol in coq. In *Lecture Notes in Computer Science*, pages 362–367. Springer International Publishing, 2020.

[7] Ziqiao Ao, Gergely Horváth, and Luyao Zhang. Are decentralized finance really decentralized? a social network analysis of the aave protocol on the ethereum blockchain. *ArXiv*, abs/2206.08401, 2022.

[8] Maria Apostolaki, Aviv Zohar, and Laurent Vanbever. Hijacking bitcoin: Routing attacks on cryptocurrencies. *2017 IEEE Symposium on Security and Privacy (SP)*, pages 375–392, 2016.

[9] Emilios Avgouleas and Heikki Marjosola. Digital finance in europe: Law, regulation, and governance. *Digital Finance in Europe: Law, Regulation, and Governance*, 2021.

[10] Yannis Bakos, Hanna Halaburda, and Christoph Müller-Bloch. When permissioned blockchains deliver more decentralization than permissionless. *Communications of the ACM*, 64:20 – 22, 2019.

[11] Leland Lee Balaji S. Srinivasan. Quantifying decentralization. `https://news.earn.com/quantifying-decentralization-e39db233c28e`, 2017. Accessed: 2023-05-21.

[12] Anastasios Balaskas and Virginia N. L. Franqueira. Analytical tools for blockchain: Review, taxonomy and open challenges. *2018 International Conference on Cyber Security and Protection of Digital Services (Cyber Security)*, pages 1–8, 2018.

[13] World Bank. Population and gdp by country. `https://databank.worldbank.org/embed/Population-and-GDP-by-Country/id/29c4df41/`, 2021. Accessed: 2023-08-07.

[14] World Bank. The world bank indicators. `https://data.worldbank.org/indicator`, 2023. Accessed: 2023-07-16.

[15] World Bank. Worldwide governance indicators. `https://info.worldbank.org/governance/wgi/Home/Reports/`, 2023. Accessed: 2023-07-16.

[16] Massimo Bartoletti, Stefano Lande, Livio Pompianu, and Andrea Bracciali. A general framework for blockchain analytics. In *Proceedings of the 1st Workshop on Scalable and Resilient Infrastructures for Distributed Ledgers*, SERIAL '17, New York, NY, USA, 2017. Association for Computing Machinery.

[17] Stefano Battiston, Michelangelo Puliga, Rahul Kaushik, Paolo Tasca, and Guido Caldarelli. Debtrank: Too central to fail? financial networks, the fed and systemic risk. *Scientific Reports*, 2, 2012.

[18] Roman Beck, Christoph Müller-Bloch, and John Leslie King. Governance in the blockchain economy: A framework and research agenda. *J. Assoc. Inf. Syst.*, 19:1, 2018.

[19] Badr Bellaj, Aafaf Ouaddah, Emmanuel Bertin, Noël Crespi, and Abdellatif Mezrioui. Sok: A comprehensive survey on distributed ledger technologies. *2022 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, pages 1–16, 2022.

[20] BitMEX. Ofac sanctions ethereum pos – some technical nuances. `https://blog.bitmex.com/ofac-sanctions-ethereum-pos-some-technical-nuances/`, 2022. Accessed: 2023-07-12.

[21] Robert Blotevogel, Eslem Imamoglu, Kenji Moriyama, and Babacar Sarr. Measuring income inequality and implications for economic transmission channels. *IMF Working Papers*, 2020.

[22] Jakob J. Bosma, M. Koetter, and Michael Wedow. Too connected to fail? inferring network ties from price co-movements. *Journal of Business & Economic Statistics*, 37:67 – 80, 2019.

[23] Alexandre Bovet, Carlo Campajola, Francesco Mottes, Valerio Restocchi, Nicoló Vallarano, Tiziano Squartini, and Claudio J. Tessone. The evolving liaisons between the transaction networks of bitcoin and its price dynamics. *arXiv: General Finance*, 2019.

[24] Santiago Bragagnolo, Matteo Marra, Guillermo Polito, and Elisa Gonzalez Boix. Towards scalable blockchain analysis. *2019 IEEE/ACM 2nd International Workshop on Emerging Trends in Software Engineering for Blockchain (WETSEB)*, pages 1–7, 2019.

[25] Eric Budish. The economic limits of bitcoin and the blockchain. *Ewing Marion Kauffman Foundation Research Paper Series*, 2018.

[26] Eric Budish. The economic limits of bitcoin and anonymous, decentralized trust on the blockchain new form of decentralized trust suffers from a pick-your-poison conundrum. 2022.

[27] Carlo Campajola, Raffaele Cristodaro, Francesco Maria De Collibus, Tao Yan, Nicoló Vallarano, and Claudio J. Tessone. The evolution of centralisation on cryptocurrency platforms. *ArXiv*, abs/2206.05081, 2022.

[28] Malcolm Campbell-Verduyn. Bitcoin and beyond : Cryptocurrencies, blockchains, and global governance. 2017.

[29] Gabriel Reis Carrara, Leonardo M. Burle, Dianne S. V. Medeiros, Célio Vinicius N. Albuquerque, and Diogo Menezes Ferrazani Mattos. Consistency, avail-

ability, and partition tolerance in blockchain: a survey on the consensus mechanism over peer-to-peer networking. *Annals of Telecommunications*, 75:163–174, 2020.

[30] Maria Chiara Cavalleri, Alice Eliet, Peter McAdam, Filippos Petroulakis, Ana Soares, and Isabel Vansteenkiste. Concentration, market power and dynamism in the euro area. 2019.

[31] Jing Chen, Sergey Gorbunov, Silvio Micali, and Georgios Vlachos. Algorand agreement: Super fast and partition resilient byzantine agreement. *IACR Cryptol. ePrint Arch.*, 2018:377, 2018.

[32] Yan Chen, Igor Pereira, and Pankaj C. Patel. Decentralized governance of digital platforms. *Journal of Management*, 47:1305 – 1337, 2020.

[33] Federico Cingano. Trends in income inequality and its impact on economic growth. *International Organisations Research Journal*, 10:97–133, 2014.

[34] CMS. Cms expert guide to crypto regulation. `https://cms.law/en/int/expert-guides/cms-expert-guide-to-crypto-regulation`, 2023. Accessed: 2023-07-12.

[35] CoinMarketCap. Crypto coins by market capitalization. `https://coinmarketcap.com/coins/`, 2023. Accessed: 2023-07-14.

[36] James Coker. Submarine cables at growing risk of cyber-attacks. `https://www.infosecurity-magazine.com/news/submarine-cables-risk-cyber-attacks/`, 2023. Accessed: 2023-08-07.

[37] Francesco Maria De Collibus, Alberto Partida, Matija Pikorec, and Claudio J. Tessone. Heterogeneous preferential attachment in key ethereum-based cryptoassets. In *Frontiers of Physics*, 2021.

[38] Lin Cong, Ke Tang, Yanxin Wang, and Xi Zhao. Inclusion and democratization through web3 and defi? initial evidence from the ethereum ecosystem. *SSRN Electronic Journal*, 2023.

[39] Mauro Conti, Ankit Gangwal, and Michele Todero. Blockchain trilemma solver algorand has dilemma over undecidable messages. *Proceedings of the 14th International Conference on Availability, Reliability and Security*, 2019.

[40] Ben R. Craig and Joseph Kachovec. Bitcoin's decentralized decision structure. *Economic Commentary (Federal Reserve Bank of Cleveland)*, 2019.

[41] crunchbase. Algorand investment firm. `https://www.crunchbase.com/organization/algorand`, 2023. Accessed: 2023-07-20.

[42] dAppRadar. Top algorand dapps. `https://dappradar.com/rankings/protocol/algorand`, 2023. Accessed: 2023-07-04.

[43] Our World In Data. Population density vs. gdp per capita, 2021. `https://ourworldindata.org/grapher/population-density-vs-prosperity/`, 2021. Accessed: 2023-08-07.

[44] Elcio Cruz de Oliveira, Alcir de Faro Orlando, Anderson L. S. Ferreira, and Carlos Eduardo de Oliveira Chaves. Comparison of different approaches for detection and treatment of outliers in meter proving factors determination. *Flow Measurement and Instrumentation*, 48:29–35, 2016.

[45] Sebastian Jose de Ramon and Michael Straughan. Competition indicators for the uk deposit-taking sector. *Statistical implications of the new financial landscape*, 43, 2017.

[46] DealroomCo. Dealroom intelligence - algorand. `https://app.dealroom.co/companies/algorand`, 2023. Accessed: 2023-07-15.

[47] Stephen Duignan. Algorand community governance period 1 review. `https://medium.com/algorand-foundation/algorand-community-governance-period-1-review-6a9b7e35b361`, 2021. Accessed: 2023-07-20.

[48] The Global Economy. Cost of starting a business - country rankings. `https://www.theglobaleconomy.com/rankings/cost_of_starting_business/`, 2023. Accessed: 2023-07-12.

[49] The Global Economy. Political stability - country rankings. `https://www.theglobaleconomy.com/rankings/wb_political_stability/`, 2023. Accessed: 2023-07-10.

[50] ethernodes.org. Ethereum mainnet statistics. `https://www.ethernodes.org/`, 2023. Accessed: 2023-08-12.

[51] Filecoin. Filecoin docs. https://docs.filecoin.io/networks/calibration/details/, 2023. Accessed: 2023-07-30.

[52] SubMarine Telecoms Forum. Cable faults maintenance archives. https://subtelforum.com/category/cable-faults-maintenance/, 2023. Accessed: 2023-08-08.

[53] Algorand Foundation. Algorand rewards - a technical overview. https://algorand.com/resources/blog/rewards-technical-overview, 2019. Accessed: 2023-07-18.

[54] Algorand Foundation. December 2020 algo dynamics. https://assets-global.website-files.com/62835f42aef969049eba0806/62cc37d101a3d91afc30efca_December%202020%20Algo%20Dynamics.pdf, 2020. Accessed: 2023-07-19.

[55] Algorand Foundation. General faq. https://www.algorand.foundation/general-faq, 2021. Accessed: 2023-07-24.

[56] Algorand Foundation. Algorand blockchain features specification version 1.0. https://raw.githubusercontent.com/algorandfoundation/specs/master/overview/Algorand_v1_spec-2.pdf, 2022. Accessed: 2023-07-09.

[57] Algorand Foundation. Algorand community governance: Period 5 review. https://www.algorand.foundation/community-governance-period-5-review, 2022. Accessed: 2023-07-20.

[58] Algorand Foundation. Algorand community governance: Period 6 review. https://www.algorand.foundation/algorand-community-governance-period-6-review, 2022. Accessed: 2023-07-20.

[59] Algorand Foundation. Algorand node artifacts. https://developer.algorand.org/docs/run-a-node/reference/artifacts/#phonebookjson, 2022. Accessed: 2023-07-19.

[60] Algorand Foundation. 2019 transparency report (period covering june 19th - november 1st, 2019). https://algorand.foundation/news/

`algorand-foundation-transparency-report-june-19th`, 2023. Accessed: 2023-07-11.

[61] Algorand Foundation. 2020 transparency report 1 (period covering november 2, 2019 to february 10, 2020). `https://algorand.foundation/news/feb-new-transparency`, 2023. Accessed: 2023-07-11.

[62] Algorand Foundation. 2020 transparency report 2 (period covering february 11, 2020, to september 30, 2020). `https://www.algorand.foundation/news/october2020-transparency`, 2023. Accessed: 2023-07-11.

[63] Algorand Foundation. 2021 transparency report 1 (period covering october 1, 2020 to march 31, 2021). `https://www.algorand.foundation/transparency-report-march-2021`, 2023. Accessed: 2023-07-11.

[64] Algorand Foundation. 2021 transparency report 2 (period covering april 2021 - september 2021). `https://www.algorand.foundation/transparency-report-september-2021`, 2023. Accessed: 2023-07-11.

[65] Algorand Foundation. 2022 transparency report 1 (period covering october 2021 - march 2022). `https://www.algorand.foundation/transparency-report-march-2022`, 2023. Accessed: 2023-07-11.

[66] Algorand Foundation. 2022 transparency report 2 (period covering april 2022 - september 2022). `https://www.algorand.foundation/transparency-report-oct-2022`, 2023. Accessed: 2023-07-11.

[67] Algorand Foundation. Algo auctions. `https://www.algorand.foundation/algo-auction-overview`, 2023. Accessed: 2023-07-11.

[68] Algorand Foundation. Algorand foundation transparency report october 22 - march 23. `https://assets-global.website-files.com/62d96b0e9ea60fd1c96a1b50/646636ad4581f31667e57db6_Algorand%20Foundation%20Transparency%20Report%20Q4%2022%20-%20Q1%2023_.pdf`, 2023. Accessed: 2023-07-11.

[69] Algorand Foundation. Algorand metrics - developer portal. `https://metrics.algorand.org/#/decentralization/`, 2023. Accessed: 2023-07-20.

[70] Cardano Foundation. Cardano documentation. `https://docs.cardano.org/cardano-testnet/getting-started/`, 2023. Accessed: 2023-07-30.

[71] Solana Foundation. Solana validator health report: August 2022. `https://solana.com/news/validator-health-report-august-2022`, 2023. Accessed: 2023-06-20.

[72] International Monetary Fund. World economic outlook (april 2023). `https://www.imf.org/external/datamapper/datasets/WEO/`, 2023. Accessed: 2023-08-07.

[73] Jeff Garzik. Peer review of 'quantitative analysis of the full bitcoin transaction graph'. `https://gist.github.com/jgarzik/3901921`, 2012. Accessed: 2023-07-01.

[74] Adem Efe Gencer, Soumya Sankar Basu, Ittay Eyal, Robbert van Renesse, and Emin Gün Sirer. Decentralization in bitcoin and ethereum networks. In *Financial Cryptography*, 2018.

[75] Hans Gersbach, Akaki Mamageishvili, and Manvir Schneider. Vote delegation and misbehavior. In *Algorithmic Game Theory*, 2021.

[76] Yossi Gilad, Rotem Hemo, Silvio Micali, Georgios Vlachos, and Nickolai Zeldovich. Algorand: Scaling byzantine agreements for cryptocurrencies. *Proceedings of the 26th Symposium on Operating Systems Principles*, 2017.

[77] Sarada Prasad Gochhayat, Sachin S. Shetty, Ravi Mukkamala, Peter B. Foytik, Georges A. Kamhoua, and Laurent L. Njilla. Measuring decentrality in blockchain based systems. *IEEE Access*, 8:178372–178390, 2020.

[78] Peter Gomber, Jascha-Alexander Koch, and Michael Siering. Digital finance and fintech: current research and future research directions. *Journal of Business Economics*, 87:537–580, 2017.

[79] Manas Das Gupta and Parth Gupta. Gini coefficient based wealth distribution in the bitcoin network: A case study. 2017.

[80] Campbell R. Harvey, Ashwin Ramachandran, and Joseph R. Santoro. Defi and the future of finance. *LSN: Law & Finance: Theoretical (Topic)*, 2021.

[81] Ethan Heilman, Alison Kendler, Aviv Zohar, and Sharon Goldberg. Eclipse attacks on bitcoin's peer-to-peer network. In *USENIX Security Symposium*, 2015.

[82] Robert Herian. Regulating blockchain. 2018.

[83] Mohammad Belayet Hossain. Acquiring an awareness of the latest regulatory developments concerning digital assets and anti-money laundering. *Journal of Money Laundering Control*, 2023.

[84] Ying-Ying Hsieh, Jean-Philippe Vergne, and Sha Wang. The internal and external governance of blockchain-based organizations: Evidence from cryptocurrencies. *ERN: Governance & Ownership (Topic)*, 2017.

[85] Runtime Verification Inc. Modeling and verification of the algorand consensus protocol. `https://github.com/runtimeverification/algorand-verification/blob/master/report/report.pdf`, 2019. Accessed: 2023-07-20.

[86] Insikt. The escalating global risk environment for submarine cables. `https://www.recordedfuture.com/escalating-global-risk-environment-submarine-cables`, 2023. Accessed: 2023-08-07.

[87] ip api. Ip geolocation api. `https://ip-api.com/`, 2023. Accessed: 2023-07-21.

[88] Yongpu Jia, Changqiao Xu, Zhonghui Wu, Zichen Feng, Yaxin Chen, and Shujie Yang. Measuring decentralization in emerging public blockchains. *2022 International Wireless Communications and Mobile Computing (IWCMC)*, pages 137–141, 2022.

[89] Harry A. Kalodner, Steven Goldfeder, Alishah Chator, Malte Möser, and Arvind Narayanan. Blocksci: Design and applications of a blockchain analysis platform. *ArXiv*, abs/1709.02489, 2017.

[90] Dimitris Karakostas, Aggelos Kiayias, and Christina Ovezik. Sok: A stratified approach to blockchain decentralization. *ArXiv*, abs/2211.01291, 2022.

[91] Dr. Kashif Mehboob Khan, Muhammad Abdullah Hayat, and Rana Muhammad Ibrahim. Investigating the impact of consensus algorithm on scalability in

blockchain systems. *Sir Syed University Research Journal of Engineering & Technology*, 2022.

[92] Nida Khan, Tabrez Ahmad, Anass Patel, and Radu State. Blockchain governance: An overview and prediction of optimal strategies using nash equilibrium. *ArXiv*, abs/2003.09241, 2020.

[93] Aggelos Kiayias and Philip Lazos. Sok: Blockchain governance. *Proceedings of the 4th ACM Conference on Advances in Financial Technologies*, 2022.

[94] Baran Kiliç, Can C. Özturan, and Alper Sen. Parallel analysis of ethereum blockchain transaction data using cluster computing. *Cluster Computing*, 25:1885–1898, 2022.

[95] Jae-Seok Kim, Jinmyeong Shin, Seok-Hwan Choi, and Yoon-Ho Choi. A study on prevention and automatic recovery of blockchain networks against persistent censorship attacks. *IEEE Access*, 10:110770–110784, 2022.

[96] Tarald O. Kvålseth. Cautionary note about the herfindahl-hirschman index of market (industry) concentration. *Contemporary Economics*, 2021.

[97] Yujin Kwon, Jian Liu, Minjeong Kim, Dawn Xiaodong Song, and Yongdae Kim. Impossibility of full decentralization in permissionless blockchains. *Proceedings of the 1st ACM Conference on Advances in Financial Technologies*, 2019.

[98] Freeman Law. Global cryptocurrency regulations. `https://freemanlaw.com/cryptocurrency/`, 2023. Accessed: 2023-07-12.

[99] Qinwei Lin, Chao Li, Xifeng Zhao, and Xianhai Chen. Measuring decentralization in bitcoin and ethereum using multiple metrics and granularities. *2021 IEEE 37th International Conference on Data Engineering Workshops (ICDEW)*, pages 80–87, 2021.

[100] William J. Luther and Sean Stein Smith. Is bitcoin a decentralized payment mechanism? *Journal of Institutional Economics*, 16:433 – 444, 2020.

[101] Stian Lydersen. Mean and standard deviation or median and quartiles? *Tidsskrift for den Norske laegeforening : tidsskrift for praktisk medicin, ny raekke*, 2020.

[102] Yuval Marcus, Ethan Heilman, and Sharon Goldberg. Low-resource eclipse attacks on ethereum's peer-to-peer network. *IACR Cryptol. ePrint Arch.*, page 236, 2018.

[103] Akio Matsumoto, Ugo Merlone, and Ferenc Szidarovszky. Some notes on applying the herfindahl–hirschman index. *Applied Economics Letters*, 19:181 – 184, 2012.

[104] Gr McMullen. Governance of blockchain systems: Governance of and by distributed infrastructure. 2018.

[105] Silvio Micali. Algorand's core technology (in a nutshell). `https://algorand.com/resources/blog/algorands-core-technology-in-a-nutshell`, 2019. Accessed: 2023-07-02.

[106] Silvio Micali. A proposal for decentralizing algorand governance. `https://www.algorand.com/DecentralizingAlgorandGovernance_Nov2020.pdf`, 2020. Accessed: 2023-06-30.

[107] MintingM. Algorand report – 28th january 2021. `https://mintingm.com/research/algorand-research-report/`, 2021. Accessed: 2023-07-12.

[108] Bhabendu Kumar Mohanta, Soumyashree S. Panda, and Debasish Jena. An overview of smart contract and use cases in blockchain technology. *2018 9th International Conference on Computing, Communication and Networking Technologies (ICCCNT)*, pages 1–4, 2018.

[109] Massimo Morini. Algorand community governance: Period 2 review. `https://medium.com/algorand-foundation/algorand-community-governance-victory-of-the-xgov-df7b04d547b2`, 2021. Accessed: 2023-07-20.

[110] Massimo Morini. Algorand community governance: Period 3 review. `https://medium.com/algorand-foundation/algorand-community-governance-period-3-review-7b62114b9756`, 2021. Accessed: 2023-07-20.

[111] Massimo Morini. Algorand community governance: Period 4 review. `https://medium.com/algorand-foundation/`

`algorand-community-governance-period-4-review-d11ab369a9ef`, 2021. Accessed: 2023-07-20.

[112] Malte Möser and Rainer Böhme. Trends, tips, tolls: A longitudinal study of bitcoin transaction fees. *IRPN: Innovation & Cyberlaw & Policy (Topic)*, 2014.

[113] Sebastian Moss. Saboteurs cut fiber cables in france, in second incident this year. `https://www.datacenterdynamics.com/en/news/saboteurs-cut-fiber-cables-in-france-in-second-incident-this-year/`, 2022. Accessed: 2023-08-08.

[114] Amir Pasha Motamed and Behnam Bahrak. Quantitative analysis of cryptocurrencies transaction graph. *Applied Network Science*, 4:1–21, 2019.

[115] Karl Palmskog Lucas Pena Musab Alturki, Brandon Moore. Formally verifying algorand: Reinforcing a chain of steel (modeling and safety). `https://runtimeverification.com/blog/formally-verifying-algorand-reinforcing-a-chain-of-steel-modeling-and-saf` 2023. Accessed: 2023-07-20.

[116] Cong T. Nguyen, Dinh Thai Hoang, Diep N. Nguyen, Dusit Niyato, Huynh Tuong Nguyen, and Eryk Dutkiewicz. Proof-of-stake consensus mechanisms for future blockchain networks: Fundamentals, applications and opportunities. *IEEE Access*, 7:85727–85745, 2019.

[117] Quoc Khanh Nguyen. Blockchain - a financial technology for future sustainable development. *2016 3rd International Conference on Green Technology and Sustainable Development (GTSD)*, pages 51–54, 2016.

[118] Antitrust Division U.S. Department of Justice. Herfindahl-hirschman index. `https://www.justice.gov/atr/herfindahl-hirschman-index/`, 2018. Accessed: 2023-08-07.

[119] Gustavo Ansaldi Oliva, A. Hassan, and Zhen Ming Jiang. An exploratory study of smart contracts in the ethereum blockchain platform. *Empirical Software Engineering*, 25:1864–1904, 2020.

[120] Deepak Puthal, Nisha Malik, Saraju P. Mohanty, Elias Kougianos, and Chia-Teng Yang. The blockchain as a decentralized security framework. 2018.

[121] Fergal Reid and Martin Harrigan. An analysis of anonymity in the bitcoin system. *2011 IEEE Third Int'l Conference on Privacy, Security, Risk and Trust and 2011 IEEE Third Int'l Conference on Social Computing*, pages 1318–1326, 2011.

[122] Binance Research. Algorand (algo) a permissionless pure proof-of-stake (pos) blockchain. `https://research.binance.com/en/projects/algorand`, 2023. Accessed: 2023-07-12.

[123] Stephen A. Rhoades. The herfindahl-hirschman index. *Federal Reserve Bulletin*, pages 188–189, 1993.

[124] Hannah Ritchie, Pablo Rosado, and Max Roser. Natural disasters. *Our World in Data*, 2022. https://ourworldindata.org/natural-disasters.

[125] Dorit Ron and Adi Shamir. Quantitative analysis of the full bitcoin transaction graph. In *Financial Cryptography*, 2013.

[126] Jeremy Rubin. Btcspark : Scalable analysis of the bitcoin blockchain using spark. 2015.

[127] Fabian Schär. Decentralized finance: On blockchain- and smart contract-based financial markets. *Review*, 2020.

[128] Houman B. Shadab. Regulation of Blockchain Token Sales in the United States. In *Regulating Blockchain: Techno-Social and Legal Challenges*. Oxford University Press, 06 2019.

[129] Michele Spagnuolo, Federico Maggi, and Stefano Zanero. Bitiodine: Extracting intelligence from the bitcoin network. In *Financial Cryptography*, 2014.

[130] Voon Hou Su, Sourav Sengupta, and Arijit Khan. Automating etl and mining of ethereum blockchain network. *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining*, 2022.

[131] Pinyaphat Tasatanattakool and Chian Techapanupreeda. Blockchain: Challenges and applications. *2018 International Conference on Information Networking (ICOIN)*, pages 473–475, 2018.

[132] Dina Temple-Raston. Who tried to hack hawaii's undersea cable? `https://therecord.media/who-tried-to-hack-hawaiis-undersea-cable`, 2022. Accessed: 2023-08-07.

[133] Michael Watchulonis Torsten Hoffmann. Cryptopia: Bitcoin, blockchains and the future of the internet. `https://www.imdb.com/title/tt9203586/`, 2020. Accessed: 2023-07-04.

[134] József Tóth. Bounds of herfindahl-hirschman index of banks in the european union. 2016.

[135] U.S. DEPARTMENT OF THE TREASURY. U.s. treasury sanctions notorious virtual currency mixer tornado cash. `https://home.treasury.gov/news/press-releases/jy0916`, 2022. Accessed: 2023-07-21.

[136] Triangle. Faucet - the multi-chain faucet for everyone. `https://faucet.triangleplatform.com/`, 2023. Accessed: 2023-07-31.

[137] Lewis Tseng and Moayad Aloqaily. Cryptocurrency meets cap theorem. *2023 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, pages 1–2, 2023.

[138] Margaret Uwayo, Mya Hernandez, and Denise E Ross. Income inequality. *A Scientific Framework for Compassion and Social Justice*, 2021.

[139] Nicoló Vallarano, Claudio J. Tessone, and Tiziano Squartini. Bitcoin transaction networks: An overview of recent results. In *Frontiers of Physics*, 2020.

[140] Rowan van Pelt. Blockchain governance: A framework for analysis and comparison. 2019.

[141] Rowan van Pelt, Slinger Jansen, Djuri Baars, and S. J. Overbeek. Defining blockchain governance: A framework for analysis and comparison. *Information Systems Management*, 38:21 – 41, 2021.

[142] Georgios Vlachos. Algorand's instant consensus protocol. `https://algorand.com/resources/algorand-announcements/algorands-instant-consensus-protocol`, 2018. Accessed: 2023-07-06.

[143] Philip D. Waggoner. The hhi package: Streamlined calculation and visualization of herfindahl-hirschman index scores. *J. Open Source Softw.*, 3:828, 2018.

[144] Anton Wahrstätter, Jens Ernstberger, Aviv Yaish, Liyi Zhou, Kaihua Qin, Taro Tsuchiya, Sebastian Steinhorst, Davor Svetinovic, Nicolas Christin, Mikolaj Barczentewicz, and Arthur Gervais. Blockchain censorship, 2023.

[145] Canhui Wang, Xiaowen Chu, and Yang Qin. Measurement and analysis of the bitcoin networks: A view from mining pools. In *2020 6th International Conference on Big Data Computing and Communications (BIGCOM)*, pages 180–188, 2020.

[146] Yongge Wang. Another look at algorand. *ArXiv*, abs/1905.04463, 2019.

[147] Sam M. Werner, Daniel Perez, Lewis Gudgeon, Ariah Klages-Mundt, Dominik Harz, and William John Knottenbelt. Sok: Decentralized finance (defi). *Proceedings of the 4th ACM Conference on Advances in Financial Technologies*, 2021.

[148] E. Glen Weyl, Puja Ohlhaver, and Vitalik Buterin. Decentralized society: Finding web3's soul. *SSRN Electronic Journal*, 2022.

[149] Keke Wu, Bo Peng, Hua Xie, and Zhen Huang. An information entropy method to quantify the degrees of decentralization for blockchain systems. *2019 IEEE 9th International Conference on Electronics Information and Emergency Communication (ICEIEC)*, pages 1–6, 2019.

[150] K. Wüst and Arthur Gervais. Do you need a blockchain? *2018 Crypto Valley Conference on Blockchain Technology (CVCBT)*, pages 45–54, 2018.

[151] @xenowits. Nakamoto coefficients - a measure of decentralization. `https://nakaflow.io/`, 2023. Accessed: 2023-07-31.

[152] Peter Yeoh. Regulatory issues in blockchain technology. *Journal of Financial Regulation and Compliance*, 25:196–208, 2017.

[153] Tae-Sub Yun, Deokjong Jeong, and Sunyoung Park. "too central to fail" systemic risk measure using pagerank algorithm. *Journal of Economic Behavior & Organization*, 2019.

[154] Bingsheng Zhang and Hamed Balogun. On the sustainability of blockchain funding. In *2018 IEEE International Conference on Data Mining Workshops (ICDMW)*, pages 89–96, 2018.

[155] Luyao Zhang. The design principle of blockchain: An initiative for the sok of soks. *ArXiv*, abs/2301.00479, 2023.

[156] Luyao Zhang, Xinshi Ma, and Yulin Liu. Sok: Blockchain decentralization. *ArXiv*, abs/2205.04256, 2022.

[157] Yufan Zhang, Zichao Chen, Yutong Sun, Yulin Liu, and Luyao Zhang. Blockchain network analysis: A comparative study of decentralized banks. *ArXiv*, abs/2212.05632, 2022.

[158] znort987. Blockparser. `https://github.com/znort987/blockparser`, 2023. Accessed: 2023-06-20.

Appendix A

Comparative Infrastructure Analysis

The table below describes the infrastructure utilised by previous studies that attempted to dissect and evaluate the public blockchain ledger.

Work	**Resources**				
	Type	**Nodes**	**Processor**	**RAM**	**SSD**
Spagnuolo et al. [129]	On Premise	1	Intel i7 CPU 2.7 GHz	16 GB	-
Kalodner et al. [89]	On Premise	1	**EC2 Class** 8vCPU 2.5 GHz Intel Xeon E5-2670v	61 GB	100 GB
Bartoletti et al. [16]	On Premise	1	Intel i5-4440 CPU	32 GB	2 TB
Bragagnolo et al. [24]	On Premise	20	10 vCPU 3.50 GHz Intel Xeon E3-1240	32 GB	200 GB
Rubin [126]	Cloud	10	**M3 Large** 6.5 ECUs - 2 vCPU Intel Xeon E5-2670v2 2.5GHz	8 GB	32 GB
Kilic et al. [94]	Cloud	16	**C5.4xlarge** 16 vCPU (8 HT Core)	32 GB	-
Su et al. [130]	HDFS	-	2 CPU 3.47 GHz Intel Xeon X5690	8 GB	100 GB

Table A.1: Infrastructure required by previous studies observed

Appendix B

Public Blockchain Explorers

The table below lists the digital interfaces available to the public and utilised by the observed study to extract blockchain data from a managed repository.

Nomenclature	Corpus of Data	Providers
Relationships	Transaction graph and address clustering	blockchain.info bitcoinchain.com
Obscurity	Address identifiers and tags	blockchain.info bitcointalk.org bitcoin-otc.com
Cyber-crime	DDoS attacks	bitcointalk.org
	Frauds	badbitcoin.org cryptohyips.com
Transaction fees	Conversion rate	coindesk.com
	Mining pools	blockchain.info
Market insights	Transaction graph	blockexplorer.com bitcoincharts.com block.io
	Trade records	bitcoincharts.com

Table B.1: Renowned blockchain analysis tools [12, 16, 126]

Appendix C

Simulation of Algorand Indexer

This architectural framework is modeled using the Algorand Indexer to enhance the capability of conducting in-depth analyses on local block replicas found in archival nodes.

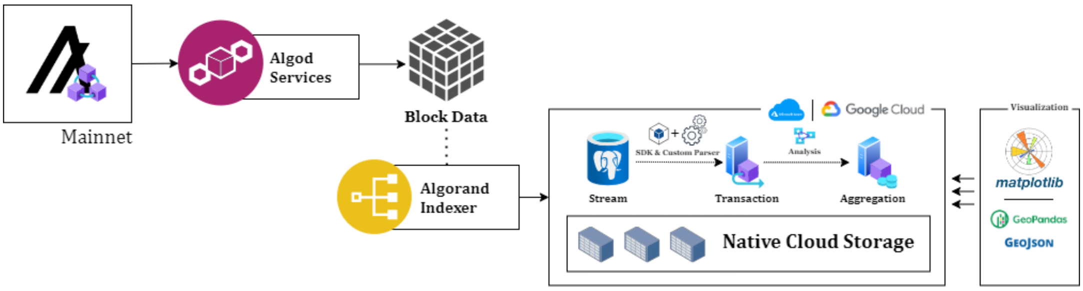

Figure C.1: of the Algorand Indexer for parsing block data

The methodological approach initiates with the collation of blocks by Algorand nodes. These are subsequently indexed by the Indexer services into a dedicated PostgreSQL database, situated on an instance provided by Google Cloud services. A specially crafted parser leverages the Indexer APIs and PyAlgorand SDKs to cull relevant details from the indexed information. This data is then channeled into two distinct tables. The Transaction table hosts the transformed data gleaned from the Indexer database, whereas the Aggregation table contains key metric values ascertained from the analytical procedure. These repositories are primed to support future analytical and visualization phases.

C.1 Performance Evaluation

The architecture pivots on the utilization of Indexer, which expected to be beneficial to expedite the analysis. A distinct instance of PostgreSQL is indispensable for populating and indexing the blockchain data, enabling it to cater to an extensive variety of criteria. This procedure inherently leads to a bifurcation for the local node.

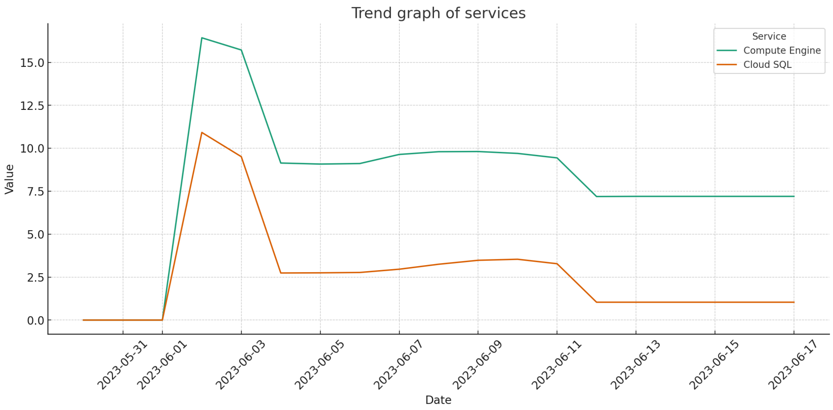

Figure C.2: GCP Computing Engine and Cloud SQL expenditure

C.2 Financial Evaluation

The infrastructure requirements is effectively doubled. Nonetheless, while conforming to the least specifications, given the variation in unit prices, the Indexer service might impose a higher cost burden. Considering the high cost necessary to deploy and operate, the proposed architecture is deemed infeasible and inefficient.

Instance	Components	Units	Monthly Unit Price (£)	Monthly Total (£)
Archival Node	vCPU	4	113.28	113.28
	Memory (GB)	16		
	Storage (GB)	1,250	0.1	125
Indexer	vCPU	4	30.14	120.56
	Memory (GB)	16	5.11	81.76
	Storage (GB)	1,700	0.17	289
				729.6

Table C.1: Simulation of Archival Node and Algorand Indexer monthly expenses

Appendix D

Node Installation Shell Script

This code facilitates the installation of a node in two distinct modes: Archival and Non-Archival `Fast-Catchup`.

An Archival Node within the Algorand network refers to a specialized type of node that is tasked with maintaining a complete record of the ledger's history, starting from its genesis block. In contrast, Algorand's `Fast-Catchup` is an optimized feature that allows a node to synchronize rapidly with the blockchain by leveraging a trusted catchpoint snapshot. Unlike an Archival Node, a node using the `Fast-Catchup` mode will only store the last 1,000 blocks of the network, marking a significant deviation in its storage requirement from the former.

```bash
#!/bin/bash

# Step 1: Opt the user to select installation mode:
   ↪ Archival or Fast-Catchup.
echo "Select installation mode: 1 for Archival, 2 for
   ↪ Fast-Catchup"
read mode

# Check if the user has entered a valid selection.
if [[ "$mode" != 1 && "$mode" != 2 ]]; then
    echo "Invalid selection. Please select 1 for Archival
        ↪ or 2 for Fast-Catchup."
    exit 1
fi
```

```
# Step 2: Establish a provisional directory to
   ↪ accommodate the installation package along with the
   ↪  associated files.
mkdir ~/node
cd ~/node

# Step 3: Procure the script responsible for updates.
wget https://raw.githubusercontent.com/algorand/go-
   ↪ algorand/rel/stable/cmd/updater/update.sh

# Step 4: Guarantee that the system recognizes the file
   ↪ as executable.
chmod 744 update.sh

# Step 5: Initiate the installer within the context of
   ↪ the node directory.
./update.sh -i -c stable -p ~/node -d ~/node/data -n

# Step 6: Adapt the node to operate in the Archival
   ↪ configuration, if selected.
if [ "$mode" -eq 1 ]; then
    echo '{
        "Archival": true
    }' > ~/node/data/config.json
fi

# Step 7: Configure the export path within the shell
   ↪ configuration files.
echo 'export ALGORAND_DATA="$HOME/node/data"' >> ~/.
   ↪ bashrc
echo 'export PATH="$HOME/node:$PATH"' >> ~/.bashrc
source ~/.bashrc

# Step 8: Start the node based on the selected mode.
if [ "$mode" -eq 1 ]; then
```

```
    goal node start
elif [ "$mode" -eq 2 ]; then
    # Fetch the latest catchpoint.
    catchpoint=$(curl -s https://algorand-catchpoints.s3.
       ↪ us-east-2.amazonaws.com/channel/mainnet/latest.
       ↪ catchpoint)
    goal node start
    goal node catchup $catchpoint
fi

# Step 9: Maintain the node progressing between two
   ↪ successive status updates.
echo "Examine the Sync Time parameter. If it exhibits a
   ↪ value of 0.0s, this signifies that the node is
   ↪ thoroughly synchronized with the network."
goal node status -w 1000
```

Appendix E

Archival and Indexer Node Performance

The following tables present the comparative computational performance of two distinct systems: the Archival Node deployed on Google Compute Engine and the Indexer Node which leverages Google Cloud SQL.

E.1 Node Synchronization

Date	Archival Node		Indexer		Check Point			
	CPU	Storage	CPU	Storage	Blocks	Storage	DB	DB Size
30-May-23	2/16	1.25TB	4/16	1.7TB	8,4M	150.1	3,391,743	3.15
1-Jun-23	4/16	1.25TB	4/16	1.7TB	8,2M	130.2	3,362,842	3.02
3-Jun-23	4/16	1.3TB	2/8	1.7TB	8,4M	150.7	3,463,881	3.18
					12,7M	299.7	Suspended	
					15,7M	429.1	Terminated	
					17,0M	580.2		
					18,1M	734.2		
					20,1M	999.4		
					22,7M	1,147		

Table E.1: Archival and Indexer synchronization performance

E.2 CPU Performance

Table E.2: Daily performance comparison between utilized cloud instances

Operation Date	Archival Node	Indexer
2023-06-02	16.43	10.92
2023-06-03	15.72	9.51
2023-06-04	9.14	2.74
2023-06-05	9.08	2.75
2023-06-06	9.11	2.77
2023-06-07	9.64	2.96
2023-06-08	9.8	3.25
2023-06-09	9.81	3.48
2023-06-10	9.7	3.54
2023-06-11	9.44	3.28
2023-06-12	7.19	1.04
2023-06-13	7.2	1.04
2023-06-14	7.2	1.04
2023-06-15	7.2	1.04
2023-06-16	7.2	1.04
2023-06-17	7.2	1.04

Appendix F

Comparative Performance of Parser Architecture

Analytical comparison between acquisition methodologies: API calls and the direct engagement with the encoded blockchain data.

F.1 Average Rates

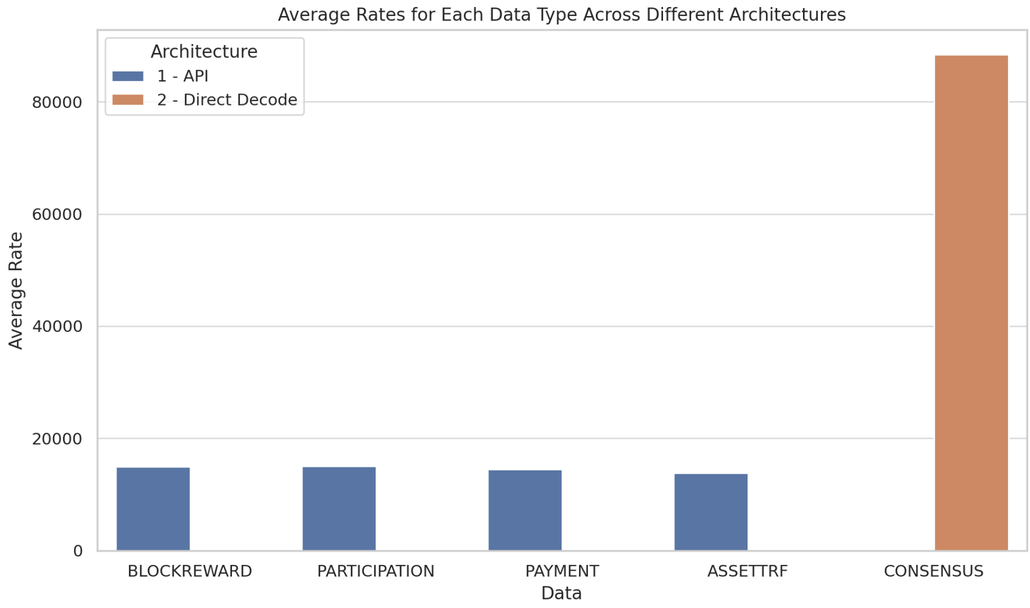

Figure F.1: API vs direct decoding average rates comparison

F.2 Distribution of Rates

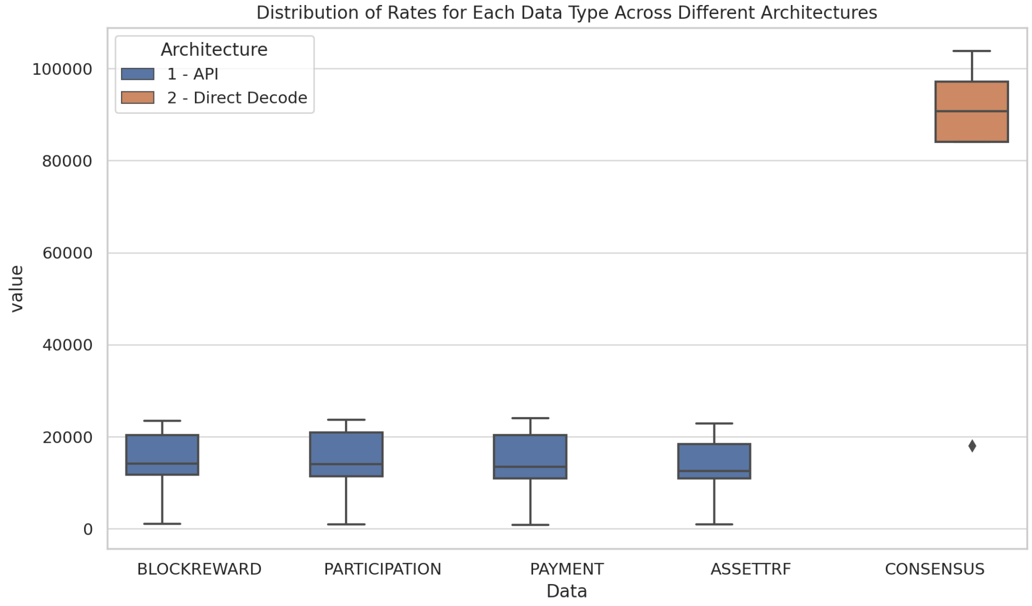

Figure F.2: API vs direct decoding distribution of rates

F.3 Completion Time

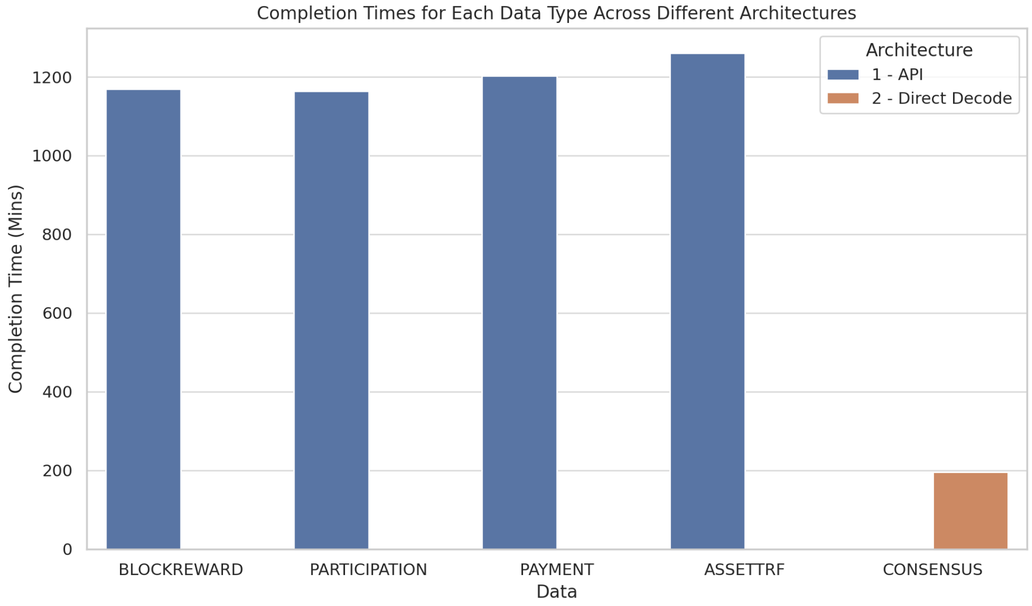

Figure F.3: API vs direct decoding time of completion comparison

Appendix G

PoS Blockchains Market Activities

Details of the market capitalization, the available supply, and the operational values that align with the measurement parameters for PoS Blockchains.

Table G.1: Market and activities values as of 27 July 2023 [35]

No	Crypto	Market	Volume (24H)	Max Supply	Total Supply	Circ. Supply	Exchs	Pairs
1	Ethereum	222.9B	4.3B	0	120.4M	120.3M	284	3038
2	Solana	13.2B	387.8M	0	553.3M	404.3M	133	253
3	Cardano	10.6B	191.6M	45.0B	36.0B	35.0B	156	480
4	Polygon	6.6B	256.1M	10.0B	10.0B	9.3B	175	336
5	Toncoin	4.8B	32.7M	0	5.0B	3.4B	40	63
6	Avalanche	4.6B	139.4M	720.0M	432.7M	345.9M	135	228
7	Cosmos	3.1B	60.0M	0	0	346.6M	123	220
8	Algorand	849.8M	30.3M	10.0B	7.8B	7.8B	91	186
9	Tezos	767.0M	12.0M	0	967.7M	946.5M	88	157
10	Mina	413.2M	5.7M	0	1.1B	940.2M	33	51
11	Thor chain	313.3M	18.8M	500.0M	486.1M	333.9M	42	87
12	Celo	243.4M	12.2M	1.0B	1.0B	505.1M	49	103
13	Osmosis	240.9M	4.1M	1.0B	587.4M	492.6M	19	62

Continued on next page

Table G.1: Market and activities values as of 27 July 2023 [35] (Continued)

No	Crypto	Market	Volume (24H)	Max Supply	Total Supply	Circ. Supply	Exchs	Pairs
14	BitShares	28.5M	575.01K	3.6B	3.0B	3.0B	19	31
15	LTO Network	26.2M	826.711K	500.0M	417.3M	417.3M	21	29
16	PIVX	16.0M	344.058K	0	75.0M	75.0M	12	15
17	Validity	9.8M	1.2M	9.0M	4.9M	4.9M	8	12
18	PRIZM	9.2M	44.869K	6.0B	3.4B	3.4B	4	4
19	Electra Protocol	4.7M	18.734K	30.0B	17.8B	17.8B	7	9
20	Oxen	4.6M	215.841K	0	64.0M	64.0M	3	7
21	Omax Coin	3.6M	39.735K	9.0B	9.0B	8.7B	4	4
22	Particl	3.3M	1.146K	0	11.8M	13.3M	4	5
23	Otocash	2.9M	0	0	38.3M	36.8M	1	1
24	GCR	2.4M	184.133K	0	107.0M	107.0M	3	5
25	Ghost	1.8M	0	0	23.1M	23.1M	0	0
26	HiCoin	1.3M	0	0	10.0B	4.4B	1	1
27	Black Coin	1.0M	0	0	62.0M	62.2M	2	2
28	PAC Protocol	989.281K	41.702K	50.0B	16.0B	17.4B	4	10
29	Metrix Coin	705.504K	0	30.0B	18.4B	18.8B	2	5
30	BlackHat	566.057K	64.57K	21.0M	10.9M	10.3M	5	9
31	Bitcoin Plus	546.769K	758	1.0M	211.804K	211.804K	5	10
32	DAPS Coin	510.049K	5	70.0B	62.7B	58.0B	1	1

Continued on next page

Table G.1: Market and activities values as of 27 July 2023 [35] (Continued)

No	**Crypto**	**Market**	**Volume (24H)**	**Max Supply**	**Total Supply**	**Circ. Supply**	**Exchs**	**Pairs**
33	Bismuth	499.651K	287	0	30.7M	29.1M	3	5
34	FYDcoin	436.506K	70.437K	650.0M	640.9M	635.0M	3	3
35	Alias	294.639K	52	0	27.2M	27.2M	4	4
36	Zennies	293.196K	0	0	1.0B	1.0B	1	1
37	NoLimit	243.938K	408	0	1.1B	490.2M	2	2
38	Freedom	129.922K	206	18.0M	18.0M	6.6M	1	2
39	Sono	119.747K	21.338K	200.0M	128.0M	49.1M	2	2
40	ION	65.24K	0	0	24.2M	18.3M	1	2
41	PureVidz	36.697K	55	0	125.3M	125.3M	1	1
42	Rubies	33.559K	20	0	10.4M	10.4M	1	1
43	Mojo Coin	32.365K	32	0	12.3M	12.3M	1	1
44	PayCoin	24.596K	4	0	12.0M	12.0M	1	1
45	AllSafe	16.748K	0	15.0M	10.6M	9.1M	2	2
46	PostCoin	13.944K	27	0	15.9M	15.9M	1	1
47	Donu	10.409K	0	0	6.5M	5.1M	1	1
48	Draft Coin	8.611K	468	0	18.7M	8.7M	2	3
49	Cabbage	6.151K	0	0	10.5M	10.5M	1	1
50	Iconic	1.216K	36	0	592.894K	592.894K	1	1
51	Gridcoin	0	36.842K	0	459.4M	0	2	3
52	Ttcoin	0	29.527K	0	3.9B	0	7	8
53	Enecuum	0	23.4K	350.0M	288.0M	0	2	3
54	Navcoin	0	16.339K	0	76.6M	0	3	3
55	Nxt	0	5.822K	1.0B	999.0M	0	3	3
56	Crown	0	640	42.0M	31.8M	0	2	2

Continued on next page

Table G.1: Market and activities values as of 27 July 2023 [35] (Continued)

No	**Crypto**	**Market**	**Volume (24H)**	**Max Supply**	**Total Supply**	**Circ. Supply**	**Exchs**	**Pairs**
57	MintCoin	0	0	0	24.9B	0	1	1
58	Rubycoin	0	0	0	27.6M	0	1	1
59	Avatar Coin	0	0	0	0	0	1	1
60	Aces	0	0	0	0	0	1	1

Appendix H

Codebase Repository Activity

The following graphic illustrates the distribution of contributions to Algorand's codebase since its inception in 2019.

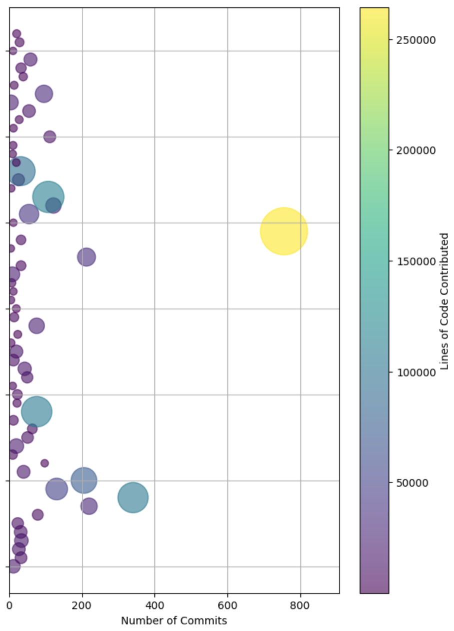

Figure H.1: Distribution of developer contributions to Algorand's codebase.

Within the Algorand codebase, there is evidence of a markedly dynamic participation: out of 105 contributors, 88 have made significant contributions, with a particularly prominent individual emerging as the primary contributor. Conversely, Bitcoin's developmental architecture appears more centralized [40, 100]. Notwithstanding Algorand's vibrant contributor landscape, the data suggests a centralization around a singular entity as the principal contributor to its codebase.

Appendix I

Relay Nodes Geographic Distribution

I.1 IP and Geo-location Details

The distribution of the relay nodes responsible for bootstrapping the Algorand network. The table focuses on providing information regarding the precise logical addresses and managing organisations associated with these nodes.

Table I.1: Algorand relay nodes across the world

IP	Continent	Country	Type	Organization
130.245.173.83	America	United States	University	Stony Brook
103.171.44.105	Asia	India	Cloud Provider	Cloudtechtiq
5.78.75.165	America	United States	Data Center	Hetzner
169.150.202.141	Asia	Israel	CDN	DataCamp
169.150.222.206	Asia	Hong Kong	CDN	DataCamp
185.24.9.80	America	Canada	CDN	DataCamp
74.118.139.188	Europe	Netherlands	Cloud Provider	TeraSwitch
69.65.31.5	America	United States	Cloud Provider	GigeNet
139.162.92.170	Asia	Japan	Cloud Provider	Akamai Technologies
190.103.179.60	America	Mexico	Cloud Provider	Sondatech
193.205.184.250	Europe	Italy	University	Salerno
208.91.104.51	America	Canada	Cloud Provider	TeraSwitch

Continued on next page

Table I.1: Algorand relay nodes across the world (Continued)

IP	**Continent**	**Country**	**Type**	**Organization**
95.217.168.85	Europe	Finland	Data Center	Hetzner
65.109.101.235	Europe	Finland	Data Center	Hetzner
45.77.38.175	Asia	Singapore	Cloud Provider	The Constant Company
204.16.242.186	America	United States	Cloud Provider	TeraSwitch
172.105.45.161	Asia	India	Cloud Provider	Akamai Technologies
169.150.242.16	Europe	Croatia	CDN	DataCamp
110.43.96.214	Asia	China	Cloud Provider	Beijing Kingsoft
198.244.212.72	Europe	United Kingdom	Cloud Provider	OVH
67.209.54.77	Asia	Singapore	Cloud Provider	TeraSwitch
169.150.228.12	America	Colombia	CDN	DataCamp
112.80.39.155	Asia	China	Data Center	China Unicom
67.209.54.88	Asia	Singapore	Cloud Provider	TeraSwitch
136.244.116.215	Europe	France	Cloud Provider	The Constant Company
128.1.59.226	Asia	United Arab Emirates	Cloud Provider	Zenlayer
167.235.110.199	Europe	Germany	Data Center	Hetzner
79.172.193.81	Europe	Hungary	Data Center	Deninet
151.100.181.25	Europe	Italy	University	Sapienza
45.134.141.81	America	Brazil	CDN	DataCamp
172.105.44.124	Asia	India	Cloud Provider	Akamai Technologies
74.118.142.181	America	United States	Cloud Provider	TeraSwitch

Continued on next page

Table I.1: Algorand relay nodes across the world (Continued)

IP	Continent	Country	Type	Organization
149.28.36.5	America	United States	Cloud Provider	The Constant Company
185.156.44.72	Europe	Romania	CDN	DataCamp
49.12.17.217	Europe	Germany	Data Center	Hetzner
3.106.131.131	Australia	Australia	Cloud Provider	Amazon
136.243.69.89	Europe	Germany	Data Center	Hetzner
38.154.253.194	Europe	United Kingdom	Data Center	24 Shells
141.95.126.156	Europe	Germany	Cloud Provider	OVH
132.67.252.201	Asia	Israel	University	Tel Aviv
169.150.252.66	Europe	Greece	CDN	DataCamp
208.91.104.52	America	Canada	Cloud Provider	TeraSwitch
200.25.81.103	America	Mexico	Data Center	Edgeuno
67.209.55.43	Asia	Hong Kong	Cloud Provider	TeraSwitch
139.162.95.37	Asia	Japan	Cloud Provider	Akamai Technologies
138.199.14.117	Europe	France	CDN	DataCamp
128.32.157.58	America	United States	University	Berkeley
139.84.143.235	Asia	India	Cloud Provider	The Constant Company
169.150.221.193	America	United States	CDN	DataCamp
37.59.22.30	Europe	France	Cloud Provider	OVH
84.17.55.163	Europe	Poland	CDN	DataCamp
67.209.55.54	Asia	Hong Kong	Cloud Provider	TeraSwitch
143.244.58.98	Europe	Czechia	CDN	DataCamp
5.161.197.23	America	United States	Data Center	Hetzner

Continued on next page

Table I.1: Algorand relay nodes across the world (Continued)

IP	Continent	Country	Type	Organization
208.76.221.37	Europe	Spain	Cloud Provider	The Constant Company
169.150.224.230	America	United States	CDN	DataCamp
155.138.254.45	America	United States	Cloud Provider	The Constant Company
185.37.151.122	Asia	Israel	Data Center	Interhost
35.216.83.233	Asia	South Korea	Cloud Provider	Google
149.28.246.89	America	United States	Cloud Provider	The Constant Company
129.97.74.19	America	Canada	University	Waterloo
104.238.188.119	Europe	France	Cloud Provider	Choopa
200.25.81.100	America	Mexico	Data Center	Edgeuno
23.229.78.130	America	United States	Data Center	24 Shells
162.19.234.131	Europe	Germany	Cloud Provider	OVH
45.77.190.182	America	United States	Cloud Provider	The Constant Company
119.252.189.15	Australia	Australia	Data Center	Zone Networks
167.235.107.245	Europe	Germany	Data Center	Hetzner
74.118.139.61	Europe	Netherlands	Cloud Provider	TeraSwitch
74.118.136.215	Europe	Netherlands	Cloud Provider	TeraSwitch
13.246.12.50	Africa	South Africa	Cloud Provider	Amazon
169.150.246.93	Africa	South Africa	CDN	DataCamp
149.28.127.155	America	United States	Cloud Provider	The Constant Company
67.209.54.111	Asia	Singapore	Cloud Provider	TeraSwitch
138.199.41.58	America	United States	CDN	DataCamp
141.98.217.84	Europe	Ireland	Cloud Provider	TeraSwitch

Continued on next page

Table I.1: Algorand relay nodes across the world (Continued)

IP	Continent	Country	Type	Organization
37.19.203.103	Europe	Bulgaria	CDN	DataCamp
5.161.200.141	America	United States	Data Center	Hetzner
65.20.96.33	Europe	Spain	Cloud Provider	The Constant Company
128.31.0.83	America	United States	University	MIT
3.140.75.230	America	United States	Cloud Provider	Amazon
198.244.229.79	Europe	United Kingdom	Cloud Provider	OVH
37.19.207.114	America	United States	CDN	DataCamp
69.160.65.232	America	United States	Data Center	Fibernet
74.118.142.78	America	United States	Cloud Provider	TeraSwitch
69.160.65.233	America	United States	Data Center	Fibernet
204.16.244.28	America	United States	Cloud Provider	TeraSwitch
204.16.244.94	America	United States	Cloud Provider	TeraSwitch
37.59.22.29	Europe	France	Cloud Provider	OVH
195.176.181.144	Europe	Switzerland	Cloud Provider	Zone Networks
146.59.81.201	Europe	Poland	Cloud Provider	OVH
74.118.143.38	Europe	Netherlands	Cloud Provider	TeraSwitch
45.179.88.15	America	Brazil	Cloud Provider	SWITCH.101
208.91.104.74	America	Canada	Cloud Provider	TeraSwitch
204.16.242.174	America	United States	Cloud Provider	TeraSwitch
74.118.142.175	America	United States	Cloud Provider	TeraSwitch
155.138.228.163	America	United States	Cloud Provider	The Constant Company
102.129.144.20	America	United States	Data Center	Hostzone
54.160.254.254	America	United States	Cloud Provider	Amazon
141.98.218.50	America	United States	Cloud Provider	TeraSwitch

Continued on next page

Table I.1: Algorand relay nodes across the world (Continued)

IP	**Continent**	**Country**	**Type**	**Organization**
195.12.59.106	Europe	Ukraine	Data Center	K-Link
146.59.81.200	Europe	Poland	Cloud Provider	OVH
79.172.193.82	Europe	Hungary	Data Center	Deninet
148.251.154.180	Europe	Germany	Data Center	Hetzner
141.98.217.71	Europe	Ireland	Cloud Provider	TeraSwitch
136.243.69.88	Europe	Germany	Data Center	Hetzner
155.138.238.199	America	United States	Cloud Provider	The Constant Company
67.209.55.53	Asia	Hong Kong	Cloud Provider	TeraSwitch
172.105.46.167	Asia	India	Cloud Provider	Akamai Technologies
45.63.88.202	America	United States	Cloud Provider	The Constant Company

I.2 Evaluation of Economic, Social, and Demographic Factors

The dispersion of Algorand Relay Nodes among the leading five nations - namely the United States (30%), Germany (7%), Canada (5%), India (5%), and France (5%) - intimates a calculated emphasis on countries characterized by significant economic influence and diverse demographic densities [13, 14, 43, 72]. The United States and Germany, recognized for their considerable Gross Domestic Product (GDP), accommodate a prominent quantity of nodes, thereby underscoring the confluence between economic might and advanced technological frameworks. In contrast, Canada exemplifies the capacity of nations with less populous demographics but commendable economic yields to serve as integral constituents in the network. The nodes presence in India, in light of its extensive populace, accentuates the obstacles and inequities in forging a technological imprint in areas of high population density. Nonetheless, the strategic choice of this location can be construed as advantageous, given India's burgeoning economic trajectory in the regional context.

From the perspective of physical security, the bulk of Algorand's relay nodes are strategically positioned in regions renowned for their relative physical stability [124]. For instance, Europe—hosting 36% of nodes (40 nodes) with prominent presences in countries such as Germany, Netherlands, and the United Kingdom—offers comparatively minimum potential disruption risks linked to natural disasters. Nonetheless, a subset of nodes is placed in regions recognized for their vulnerability to such occurrences, with a specific focus on nodes within Los Angeles, Tokyo and Hong Kong, susceptible to seismic activity and typhoons, respectively.

Legal compliance is another crucial aspect, where Algorand's network draws benefits from a substantial number of nodes being hosted in countries featuring supportive (27% with 30 nodes), or at the very least, non-restrictive (66% with 73 nodes), blockchain regulations [34, 98]. European countries such as France, Switzerland, and Estonia, and Asian countries like Singapore, Hong Kong, and Japan, are recognized for their progressive stance and clear regulatory frameworks towards blockchain technology. However, there exist nodes within countries with unclear or volatile regulatory perspectives. For instance, nodes in China (Asia) may encounter legal and regulatory obstacles due to the country's rigid stance on cryptocurrencies, a concern that similarly applies to countries like India with fluctuating legal positions.

In regard to social dynamics, a majority of the nodes are established in countries marked by their political stability [15, 49]. Regions such as North America (specifically the United States and Canada, comprising 35% of the nodes), Europe (notably Germany and the Netherlands, accounting for 12%), and Australia (2%) are highlighted for their socio-political stability and pronounced level of technological adoption. Nonetheless, the existence of nodes in certain areas necessitates enhanced vigilance to maintain uninterrupted operations. Although they constitute a smaller fraction, nodes situated in South Africa, contributing to 2% of the total, could potentially be impacted by such socio-political conditions.

Examining the empirical economic aspect, operational expenditures associated with running nodes vary considerably. The Algorand network's nodes are situated in both high and low operational cost regions [14, 48]. Nodes situated in cities like New York (14%) and Los Angeles (9%), and in countries like Singapore (4%), might confront elevated operational costs due to high living expenses and business operation costs. While this potentially influence the sustainability, the broad distribution of nodes likely mitigates the overall network impact. In this context, certain parts of Asia, such as India (5%), offer lower operational costs relative to regions like North America and Europe.

Appendix J

Distribution of Relay Nodes and Their Corresponding Submarine Cable Infrastructure

The following table presents a detailed distribution of relay nodes across various countries. Additionally, it provides information regarding the submarine cables connected to each country, which serve as the foundational communication backbone.

High			Medium			Low			Very Low		
Country	N	C	Country	N	C	Country	N	C	Country	N	C
US	33	90	France	5	28	Greece	1	17	Israel	3	5
UK	3	59	Switzerland	1	27	Hong Kong	4	16	Hungary	2	5
Singapore	3	38	China	2	24	Brazil	2	16	Poland	3	2
Japan	2	34	Australia	2	23	Ireland	2	16	Czechia	1	2
Spain	2	33	India	5	22	Finland	2	12	Ukraine	1	2
Italy	2	33	Canada	5	20	Colombia	1	12	Bulgaria	1	2
			UAE	1	20	South Africa	2	11	Croatia	1	2
						South Korea	1	11	Romania	1	1
						Netherlands	4	10			
						Mexico	3	10			
						Germany	8	8			

Table J.1: Relay nodes (N) and connected submarine cables (C) to the country

Appendix K

Blockchain/Cryptocurrency Regulation and Laws

The subsequent table provides an exploration of the global regulatory frameworks and legislative provisions concerning blockchain technology and cryptocurrencies, focusing specifically on the countries hosting relay nodes.

Table K.1: Regulation of blockchain and cryptocurrency within countries hosting relay nodes.[34, 98]

Country	**Regulatory Framework**	**AML/CTF**	**Travel Rule**	**Stablecoins**
United States	▷	✓	✓	↻
United Kingdom	▷	✓	✓	↻
Australia	▷	✓	▷	▷
Austria	▷	✓	▷	▷
The Bahamas	✓	✓	✓	✓
Bahrain	✓	✓	✗	✗
Canada	▷	✓	✓	▷
Cayman Islands	✓	✓	✓	✓

Continued on next page

Table K.1: Regulation of blockchain and cryptocurrency within countries hosting relay nodes.[34, 98] (Continued)

Country	**Regulatory Framework**	**AML/CTF**	**Travel Rule**	**Stablecoins**
China	⏸	⏸	⏸	⏸
Denmark	▷	✓	✗	✗
Estonia	✓	✓	▷	✗
France	✓	✓	✓	✗
Germany	✓	✓	✓	✗
Gibraltar	✓	✓	✓	✓
Hong Kong	✓	✓	✗	▷
Hungary	▷	✓	✗	✗
India	✗	▷	✗	✗
Italy	▷	✓	▷	▷
Japan	✓	✓	✓	✓
Jordan	✗	✓	✗	✗
Kuwait	✗	✗	✗	✗
Luxembourg	▷	✓	▷	▷
Malaysia	✓	✓	✓	✗
Mauritius	✓	✓	✓	✓
New Zealand	▷	▷	✗	▷
Oman	✗	✗	✗	✗
Panama	▷	▷	✗	✗
Qatar	⏸	⏸	⏸	⏸

Continued on next page

Table K.1: Regulation of blockchain and cryptocurrency within countries hosting relay nodes.[34, 98] (Continued)

Country	**Regulatory Framework**	**AML/CTF**	**Travel Rule**	**Stablecoins**
Saudi Arabia	‖	‖	‖	‖
Singapore	✓	✓	✓	▷
South Africa	▷	✓	▷	▷
Switzerland	✓	✓	✓	✓
Taiwan	×	✓	✓	×
Turkey	×	▷	×	×
United Arab Emirates	✓	✓	✓	↻

Legends: ▷ *Initiated*, ✓ *Available*, ‖ *Not Initiated*, × *Prohibited*, ↻ *Finalizing*