

PhD

3rd
CYCLE

FCUP
UA
UM
2025

U.PORTO

Fast Probabilistic Consensus on a Set

Rafael Cizeski Nitchai

FC

U.PORTO
FACULDADE DE CIÊNCIAS
UNIVERSIDADE DO PORTO

universidade
de aveiro



Universidade do Minho

U.PORTO
FACULDADE DE CIÊNCIAS
UNIVERSIDADE DO PORTO

universidade
de aveiro



Universidade do Minho

Fast Probabilistic Consensus on a Set

Rafael Cizeski Nitchai

Doctoral Program in Applied Mathematics
Faculty of Sciences of the University of Porto, University of Aveiro and
University of Minho
2025

Fast Probabilistic Consensus on a Set

Rafael Cizeski Nitchai

Thesis carried out as part of the Doctoral Program in Applied
Mathematics

Department of Mathematics
2025

Supervisor

Serguei Popov, Senior researcher, Centre of Mathematics of the
University of Porto

Co-supervisor

Jorge Milhazes de Freitas, Full Professor, Faculty of Sciences of
the University of Porto

Acknowledgments

First and foremost, I would like to express my deepest gratitude to my supervisors, Prof. Serguei Popov and Prof. Jorge Freitas, and contributor, Prof. Sebastian Muller, for their invaluable guidance, unwavering support, and insightful feedback throughout the course of my research. Their expertise and encouragement have been instrumental in shaping the direction and completion of this thesis.

A special thanks goes to Prof. Alberto Saa, FAPESP (Project 2013/09357-9), and CMUP for their generosity in lending the computational resources that allowed my simulations to be run. Their support was crucial to the successful completion of the numerical aspects of my research.

I would like to extend my heartfelt thanks to my colleague and friend, Dr. Lucas Amorim, who have been a constant source of encouragement, advice, and camaraderie. His willingness to discuss ideas and provide feedback has made this journey more enriching and enjoyable.

A special mention goes to my family for their endless love, patience, and belief in me. To my aunt and uncle, Dra. Leyza Dorini and Dr. Fabio Dorini, thank you for your unwavering support and understanding during the most challenging times.

I would like to acknowledge the funding and resources provided by Fundação para Ciência e Tecnologia and CMUP, without which this research would not have been possible. Your support has been crucial in allowing me to pursue and realize my academic goals.

Finally, I want to extend my deepest appreciation to Dra. Olivia Saa and Eliel. Their love, companionship, and unwavering support have been the foundation that kept me going through every challenge and success along the way. You are, without a doubt, the most important part of this journey, and I couldn't have done it without you by my side.

Thank you all for being part of this journey.

Abstract

The fast probabilistic consensus (FPC) is a leaderless voting consensus protocol that allows a set of nodes to agree on a value of a single bit. FPC is robust and efficient in Byzantine infrastructures and presents a low communicational complexity. In this paper, we introduce a modification of the Fast Probabilistic Consensus protocol (FPC) capable of achieving consensus on a maximal independent set of a graph —hence named Fast Probabilistic Consensus on a Set (FPCS)— that still preserves the robustness, effectiveness, and low communicational complexity of FPC.

This thesis shows that FPCS effectively resolves the problem (with high probability) of achieving consensus on a maximal independent set of a graph of conflicts (i.e. a maximal set of nonconflicting transactions), even when a significant proportion of nodes is malicious. These nodes intend to delay the consensus or even completely break it (meaning that nodes would arrive at different conclusions about the maximal independent set).

Our study refers to a specific implementation of cryptocurrencies, but the results hold for more general majority models.

Keywords: Distributed systems, Consensus protocols, Byzantine infrastructures

Resumo

O Fast Probabilistic Consensus (FPC) é um protocolo de consenso por votação sem líder que permite a um conjunto de nós concordar sobre o valor de um único bit. O FPC é robusto e eficiente em infraestruturas bizantinas e apresenta uma baixa complexidade comunicacional. Neste trabalho, introduzimos uma modificação do protocolo FPC capaz de alcançar consenso em um conjunto independente máximo de um grafo — daí o nome Fast Probabilistic Consensus on a Set (FPCS) — que ainda preserva a robustez, eficácia e baixa complexidade comunicacional do FPC.

Esta tese demonstra que o FPCS resolve de forma eficaz (com alta probabilidade) o problema de alcançar consenso em um conjunto independente máximo de um grafo de conflitos (ou seja, um conjunto máximo de transações não conflitantes), mesmo quando uma proporção significativa de nós é maliciosa. Esses nós têm a intenção de atrasar o consenso ou até mesmo quebrá-lo completamente (o que significa que os nós chegariam a diferentes conclusões sobre o conjunto independente máximo).

Nosso estudo se refere a uma implementação específica de criptomoedas, mas os resultados são aplicáveis a modelos de maioria mais gerais.

Palavras-chave: Sistemas distribuídos, Protocolos de consenso, Infraestruturas bizantinas

Contents

List of Figures	xi
1 Introduction	1
1.1 How this thesis is structured	2
2 State of the Art	7
2.1 SIFT: Design and Analysis of a Fault-Tolerant Computer for Air- craft Control	7
2.2 Reaching Agreement in the Presence of Faults	8
2.3 The Byzantine Generals Problem	8
2.4 The Consensus Problem in Unreliable Distributed Systems	9
2.5 Impossibility of Distributed Consensus with One Faulty Process	10
2.6 Another Advantage of Free Choice: Completely Asynchronous Agreement Protocols	11
2.7 The correctness proof of Ben-Or's randomized consensus algorithm	12
2.8 FPC-BI: Fast Probabilistic Consensus within Byzantine Infrastructures	12
3 Fast Probability Consensus on a Set (FPCS)	15
3.1 Notation	15
3.2 Protocol	19
4 Solving n-spends with FPCS	23
4.1 A lower bound for k	30
4.2 What if nodes' visions are slightly different?	31
4.2.1 A simple model for the discovering process	34
4.3 Numerical analysis	36
4.3.1 Analysing the effects of the proportion of malicious nodes	37
4.3.2 Circumventing failures by increasing the number of queries	37
4.3.3 The importance of randomness	38
5 Generalization for arbitrary graphs	41
5.1 Star Graphs	41
5.2 General case	43
6 Conclusion	53

Bibliography	55
Appendix A Appendix	59
A.1 Proofs of Chapter 4	59
A.1.1 Lemma 4.0.4	59
A.2 Proofs of Chapter 5	60
A.2.1 Lemma 5.2.5	60
A.2.2 Lemma 5.2.7	61

List of Figures

3.1	Diagram representing the general structure of two transactions.	16
3.2	Graph representation of the example described in the text. The edges between transactions indicate that there is a common input.	16
3.3	Graph representation of a set with 6 transactions where all of them are in conflict with each other.	17
3.4	The intervals of control $\mathcal{I}_{q,t}^{(u)}$ for a transaction u which has a proportion of likes among honest nodes $p_t^{(u)} > 1 - \mu$ (in blue), and $\mathcal{I}_{q,t}^{(w)}$ for a transaction w such that $p_t^{(w)} < \mu$ (in red). Notice that these intervals are separated from the support $[\beta, 1 - \beta]$ of X_t	19
3.5	Representation of step (vi) of the protocol. The graph represents the node's vision at time t . The numbers represent the order according to the hash function. In the first subfigure, the transactions in yellow represent the set B ; in the second subfigure, transactions in green represent B' and transactions in red represent $N(B', A_t^{(n)})$; in the last subfigure, transactions in blue represent B'' , the chosen maximal independent set of the round.	21
4.1	The intervals of control for a transaction u which has a proportion of likes among honest nodes $p_t^{(u)} > 1 - \mu$ (in blue), and for a transaction v such that $p_t^{(v)} < \mu$ (in red). Notice that these intervals are separated from the support $[\beta, 1 - \beta]$ of X_t	24
4.2	The interval of control (in blue) of a transaction $u \in \mathbb{T}_t$ with $p_t^{(u)} = 1/2$ is always separated from the boundary of $[\beta, 1 - \beta]$ by a distance $h_c := (1 - q - 2\beta)/2$	25
4.3	In Case 1, every time X_t falls on the right of $\mathcal{I}_{q,t}^{(v_t)}$ (in blue), an honest node is likely to assign 0 to the auxiliary opinion of every transaction.	26
4.4	In Case 2, every time X_t falls on the left of $\mathcal{I}_{q,t}^{(v_t)}$ (in blue), an honest node might also like a transaction w if the upper boundary of $\mathcal{I}_{q,t}^{(w)}$ (in red) is also on the right of X_t	27
4.5	Effects of q on the performance of the protocol. Logarithmic scale.	37
4.6	Effects of k on the performance of the protocol. Logarithmic scale.	38
4.7	Finalization and agreement rate for different values of q and β	39

4.8	Agreement failures (in percentage) and agreement rate of a double spend $\mathbb{T} = \{u, v\}$ where malicious nodes reply with the same opinion to a query from every node.	40
5.1	The interval of control (in blue) of a transaction $u \in \mathbb{T}_t$ with $p_t^{(u)} = 1/(2(1-q))$ and of a transaction $v \in N(u)$. Notice that the sum of the intervals covers the whole support of X_t	42

Chapter 1

Introduction

Distributed consensus is a fundamental problem in the field of computer science, particularly within distributed systems. It involves multiple agents, or nodes, in a network agreeing on a single data value or decision, despite the possibility of faults or adversarial behavior by some of these nodes. The challenge of achieving consensus in such environments has profound implications for ensuring the reliability, security, and performance of distributed applications. As distributed systems have become increasingly prominent in today's technology landscape, the study of consensus mechanisms has grown in importance, driven by the demand for resilient, scalable, and secure infrastructures.

The significance of distributed consensus is most evident in the context of Distributed Ledger Technologies (DLTs), such as blockchains. DLTs rely on consensus algorithms to achieve agreement on the state of a ledger, which is often shared among a large number of decentralized participants. These technologies have revolutionized the way data is stored, validated, and secured, offering new paradigms for financial transactions, supply chain management, digital identity, and beyond. In DLTs, consensus mechanisms are not just a technical requirement; they are the cornerstone that enables trustless systems to operate without a central authority.

One of the primary use cases for DLTs is in cryptocurrencies, where consensus algorithms like Bitcoin's Longest Chain Rule [43] or Ethereum's Casper protocol [13] ensure that all participants in the network agree on the sequence of transactions in the ledger. This consensus is crucial to prevent double-spending, where the same digital asset could be used more than once, and to maintain the integrity of the financial system in a decentralized manner.

Beyond cryptocurrencies, DLTs are finding widespread application in areas such as supply chain management, where consensus ensures the accuracy and transparency of product tracking data across multiple stakeholders. For example, each step of a product's journey from manufacturer to consumer can be recorded on a blockchain, with consensus algorithms guaranteeing that the data remains tamper-proof and universally accepted among all parties involved. This capability is vital for industries that require traceability and compliance, such as pharmaceuticals and food production.

DLTs also have a significant impact on digital identity management, where consensus mechanisms help create secure, verifiable identities that are not reliant on centralized

databases. By leveraging consensus, these systems can prevent identity fraud and enable secure, privacy-preserving authentication across a variety of platforms and services. This approach is increasingly important in the digital economy, where the need for secure and interoperable identity solutions continues to grow.

Another critical use case is in decentralized finance (DeFi), which utilizes smart contracts on DLTs to create financial instruments without intermediaries. In this context, consensus algorithms play a vital role in validating transactions and executing agreements in a manner that is transparent, auditable, and immune to manipulation by any single participant. The reliability of these financial services hinges on the robustness of the underlying consensus protocols.

The importance of distributed consensus in these scenarios cannot be overstated. Achieving consensus in a distributed environment is challenging due to issues like network latency, asynchrony, and the presence of potentially malicious actors. Traditional approaches to consensus, such as those used in centralized systems, do not scale well in decentralized settings where trust must be established without relying on a single point of control. This has led to the development of innovative probabilistic and Byzantine fault-tolerant algorithms that are specifically designed to address these challenges in the context of DLTs.

The study and development of consensus algorithms are essential to the future of not only blockchain-based applications but also a wide range of distributed systems that demand resilience and robustness against failures and adversarial conditions. As the landscape of digital technologies continues to evolve, the role of distributed consensus will only grow in importance, driving innovations that enhance the security, transparency, and efficiency of decentralized platforms.

In this thesis, we present a comprehensive introduction and analysis of the Fast Probabilistic Consensus on a Set (FPCS) protocol. Our study focuses on the intricacies of this protocol and its application within different graph structures. We delve into the theoretical aspects of FPCS, providing detailed estimates for the probability of successfully achieving consensus on a maximal independent set in the context of complete graphs. Additionally, we extend our analysis to explore the protocol's behavior under a different security threshold when applied to arbitrary graphs, highlighting the conditions and scenarios under which consensus can be reliably reached.

1.1 How this thesis is structured

Chapter 2 in this thesis provides an overview of fundamental research on the problem of distributed consensus, starting with the fault-tolerant computer system SIFT. Developed for critical aircraft control applications, SIFT employs simple majority voting and introduces concepts such as fault tolerance, redundancy, and flexible synchronization. The chapter then moves on to more theoretical works, analyzing the seminal paper by Lamport, Shostak, and Pease on achieving agreement in the presence of faults. This paper establishes the crucial condition that the number of processors must be at least three times the number of faulty processors plus one to ensure consensus and introduces the Byzantine Generals Problem.

This problem, a metaphor for the challenges of consensus in systems with potentially malicious components, highlights the difficulties in guaranteeing agreement when generals (or processes) can send contradictory information. Algorithms using oral and signed messages to reach consensus are presented, with the latter being capable of achieving consensus regardless of the number of generals, as long as forgery is detectable.

The chapter further explores the challenges of consensus in unreliable distributed systems through Fischer's work. The paper emphasizes the importance of consensus problems for ensuring system reliability and examines various approaches, including voting. The impossibility of consensus in fully asynchronous systems with even a single process failure is discussed, highlighting the limitations of deterministic methods. To overcome these limitations, the chapter introduces a probabilistic solution proposed by Ben-Or. Ben-Or introduces asynchronous protocols that use randomness to reach consensus, bypassing previous impossibility results. Despite their potential inefficiency in terms of communication rounds, these protocols demonstrate the viability of probabilistic methods for achieving reliable consensus in failure-prone systems. The chapter then provides an analysis of Aguilera and Toueg's proof of correctness for Ben-Or's randomized consensus algorithm. This analysis shows that the algorithm remains correct even if less than half of the processes fail and addresses potential issues with using a "global coin" to speed up the algorithm, highlighting the complexity of designing randomized consensus protocols in adversarial environments.

The following chapter introduces FPCS, an advanced voting-based consensus protocol designed as an evolution of the FPC-BI protocol and the main focus of the thesis. While FPC-BI addresses the problem of reaching consensus on the value of a single bit, FPCS extends it to reach consensus on an Maximal Independent Set (MIS for short) of a graph of conflicting transactions. This extension is crucial for UTXO-based DLTs, where transactions sharing the same output are considered in conflict. Conflicts in these systems may arise due to malicious behavior or, more commonly, due to the asynchronous nature of transaction propagation on the network.

The chapter begins by establishing the notation and definitions used in the FPCS protocol. It introduces the concept of a conflict graph, where nodes represent transactions, and edges represent conflicts between transactions. The concept of cryptographic hash functions (CHF) are defined to establish a total order in the set of transactions. It is assumed that the CHF is pseudo-random and exhibits the diffusion property, meaning that any small change in the input results in a new random output. The chapter also introduces the notion of "liking" a transaction, which indicates a node's preference for one transaction over its conflicting transactions. The FPCS protocol operates in rounds, assuming the existence of a "global coin", a public sequence of random numbers, to guide decision-making.

The FPCS protocol is then described in detail. In each round, nodes randomly query other nodes about their known transactions and which ones they "like." Nodes update their "views" — the set of known transactions — and form their "opinions" — indicating which transactions they "like" — based on the gathered information. Nodes use the global coin and CHF to determine which transactions to include in their opinions, ensuring that the set

of transactions they “like” always forms an MIS. The goal of the FPCS protocol is to reach consensus among nodes on an MIS within the set of transactions, even in the presence of malicious nodes that may not adhere to the protocol.

The next chapter demonstrates the application of FPCS to solve the n -spend problem, a common scenario in UTXO-based DLTs where all transactions in a conflict set spend at least one common UTXO. The chapter focuses on complete conflict graphs, where each pair of transactions is in conflict, representing the topology of an n -spend. A theorem is presented, stating that, with high probability, FPCS will reach consensus on a single MIS in the conflict graph, provided that the number of queries (k) is sufficiently large and the proportion of malicious nodes (q) is below a certain threshold. The proof of the theorem relies on several lemmas that provide probabilistic estimates for maintaining and achieving a “pre-consensus” state, where a significant majority of nodes “like” the same set of transactions. The chapter also derives a lower bound for the value of k required to ensure the correctness of the protocol.

Through numerical analysis, the chapter further explores the influence of the proportion of malicious nodes (q) and the number of queries (k) on the protocol’s performance. Simulations, modeling malicious nodes as being controlled by a single omniscient adversary, show that as q increases, the protocol’s efficiency decreases, leading to more rounds needed to reach consensus. However, increasing k can mitigate this issue, reducing termination failures and allowing more nodes to reach consensus in fewer rounds. The chapter also investigates the role of randomness in the protocol’s robustness, showing that removing randomness from specific components can lead to consensus failures. Two scenarios are analyzed: a double-spending scenario where the support of the global coin is reduced until it becomes deterministic and a scenario where the random ordering of transactions is removed. Simulations show that, in both cases, the protocol may fail to reach consensus.

Chapter 5 extends the analysis of FPCS to arbitrary conflict graphs, considering more complex attack scenarios where malicious nodes can manipulate the set of transactions itself, not just their voting preferences. The chapter begins by exploring a limiting case using star graphs where the standard security threshold ($q < \beta < 1/3$) is insufficient to guarantee consensus. This case demonstrates that malicious nodes can exploit the conflict graph’s structure to influence the protocol’s outcome, even if they control a relatively small proportion of nodes. To address this issue, the chapter proposes an adjusted security threshold ($q < \beta < 1/4$), ensuring that the support of the global coin is large enough to counter the influence of malicious nodes.

Under this new security threshold, the chapter presents a theorem for the general case, stating that, with high probability, FPCS will reach consensus on a single MIS in an arbitrary conflict graph, provided that k is sufficiently large. The proof of the theorem relies on several lemmas that analyze the conditions and probabilities of achieving “pre-consensus” in different scenarios, depending on the size and structure of the conflict graph. The chapter shows that, in each round, nodes are likely to converge to a “pre-consensus” on a significant subset of transactions on which they are still undecided, eventually leading to full consensus on a single MIS.

In conclusion, this thesis significantly contributes to the field of distributed consensus by providing a comprehensive analysis of FPCS, a new probabilistic consensus protocol. We demonstrate the effectiveness of FPCS in addressing practical challenges in UTXO-based DLTs, such as n -spends, and generalize the results to arbitrary conflict graphs, covering a broader range of attack scenarios. The theoretical analysis, combined with numerical simulations, provides strong evidence of FPCS's robustness, efficiency, and applicability in failure-prone distributed environments.

Chapter 2

State of the Art

2.1 SIFT: Design and Analysis of a Fault-Tolerant Computer for Air- craft Control

Written by Wensley et. al. [57], this paper discusses the development and design of a highly reliable, fault-tolerant computer system called SIFT (Software Implemented Fault Tolerance). SIFT is specifically intended for critical aircraft control applications where system failures can have severe consequences. The primary objective of SIFT is to create a computer system that can tolerate faults and continue operating reliably, even in the presence of hardware failures.

SIFT achieves fault tolerance through a software-centric approach, which includes mechanisms for error detection, analysis, and system reconfiguration. The architecture of the SIFT system consists of multiple processing units connected via a specially designed redundant bus system. These processing units are standard off-the-shelf minicomputers, while standard microcomputers handle the input and output operations. By executing tasks redundantly across different processing units, SIFT can compare results and use a voting mechanism to ensure correctness. SIFT uses a majority voting approach, assuming that if a majority of units produce the same output, it is likely correct. This method allows the system to tolerate failures in individual units without affecting the overall operation. However, while majority voting is effective against simple faults, its limitations in handling more complex Byzantine faults were recognized in subsequent research.

A significant feature of the SIFT system is its approach to synchronization. Unlike traditional systems that require tight synchronization, SIFT allows for independent execution by separate processors with only loose synchronization. This is achieved through fault-tolerant synchronization methods, which reduce the need for precise timing alignment between processors.

To prevent the spread of faults, SIFT employs fault isolation techniques. This means that if a fault occurs in one unit, it does not propagate to other units, thus containing the problem. Furthermore, SIFT can reconfigure itself by reallocating tasks away from faulty units to maintain system integrity and ensure continued operation. The reliability of the SIFT system is rigorously analyzed using Markov models, which help predict system behavior

and confirm that it meets the stringent reliability requirements essential for flight control computers.

To ensure the correctness and reliability of both the hardware and software components, SIFT utilizes formal verification methods. These methods provide a mathematical proof of correctness, ensuring that the system operates as intended under all expected conditions.

In summary, SIFT is designed to be a robust and highly reliable computer system for aircraft control applications. By leveraging software-based fault tolerance, redundancy, and flexible synchronization, SIFT aims to provide an ultrareliable solution that meets the critical demands of modern aviation safety and efficiency.

It is important to highlight that while SIFT is mainly associated with Wensley, two of the co-authors are Leslie Lamport and Robert Shostak, known for their significant contributions to the field of distributed consensus.

2.2 Reaching Agreement in the Presence of Faults

The realization that simple majority voting—like the one employed on SIFT—are insufficient in certain scenarios, led Lamport, Shostak, and Pease to address in this paper [46] the challenge of achieving consensus among a set of isolated processors, some of which may be faulty, communicating only through two-party messages. The goal is to develop an algorithm that allows nonfaulty processors to agree on consistent values, even if some processors lie or withhold information. The authors show that this problem is solvable only if the number of processors n is at least three times the number of faulty processors m plus one, i.e., $n \geq 3m + 1$. They also demonstrate that under a weaker assumption, where faulty processors can refuse to pass on information but cannot falsify it, agreement can be achieved regardless of the number of processors, using cryptographic methods to ensure message authenticity.

The paper begins by explaining the need for such algorithms in fault-tolerant systems, such as synchronizing clocks or agreeing on sensor data. In cases where a simple majority cannot be trusted due to faulty processors, the authors propose more complex voting schemes involving multiple rounds of information exchange. They illustrate their concepts with scenarios involving different numbers of processors and faults, leading to the development of interactive consistency algorithms that can reliably achieve consensus. The study highlights the importance of achieving consistent agreement in distributed systems and sets the foundation for further research into more efficient and adaptable algorithms.

2.3 The Byzantine Generals Problem

In their subsequent paper [34] in this topic, Lamport, Shostak, and Pease explores again the challenge of achieving consensus in a distributed computing system where some components may behave maliciously or unreliably, a scenario they illustrate with the metaphor of Byzantine generals trying to agree on a common battle plan. Each general can communicate

only through messengers, and one or more generals may be traitors who send conflicting information to prevent the loyal generals from reaching an agreement.

The primary objective is to design an algorithm that ensures all loyal generals reach the same decision (either attack or retreat), even in the presence of traitors. The authors establish two critical conditions: firstly, that all loyal generals must agree on the same plan of action, and secondly, that if the commanding general is loyal, all loyal lieutenants must follow his order. The study reveals that if communication is limited to oral messages, the problem is solvable only if more than two-thirds of the generals are loyal. This means that with three generals, a single traitor can prevent agreement, but with four generals, one traitor can be managed.

The authors present several algorithms to address the problem, depending on the system's assumptions. One solution, based on oral messages, requires each general to send and receive messages to and from every other general in multiple rounds, using majority voting to decide on the action. Another, more robust solution involves signed messages, ensuring that once a general receives a message, its authenticity and content cannot be tampered with, making it possible to reach consensus regardless of the number of generals, as long as the system can detect forgery and ensure message delivery.

The paper further discusses how these algorithms can be applied to real-world systems to enhance reliability. For instance, in computer systems requiring high reliability, such as flight control systems or distributed databases, using these consensus protocols can prevent faulty components from leading to system-wide failures. However, achieving this level of reliability comes at a high cost in terms of time and communication overhead, as these algorithms require multiple rounds of message exchanges and robust mechanisms for detecting and managing faults.

In conclusion, the Byzantine Generals Problem highlights the complexities of achieving fault-tolerant consensus in distributed systems and provides foundational algorithms that have shaped the development of reliable computing and fault-tolerant protocols.

2.4 The Consensus Problem in Unreliable Distributed Systems

In this paper Fischer [25] explores the fundamental challenges of achieving consensus among processes in distributed systems, where some processes may be unreliable or faulty. The problem is crucial for fault-tolerant distributed computing because it ensures that non-faulty processes can still reach a consensus, which is necessary for the system to function reliably despite failures.

Fischer reviews various agreement problems, which involve ensuring that all processes in a system agree on a particular piece of data. This agreement is vital in many contexts, such as in distributed databases where transactions must be consistently committed or aborted, in replicated file systems, and in systems like flight control, where modules must agree on critical decisions like whether to abort a landing.

A basic approach to achieving consensus involves voting and adopting the majority value. However, in the presence of faulty processes, this method can fail, as a single faulty process

can send conflicting votes, leading reliable processes to disagree. To address this, researchers have developed more complex protocols that involve multiple rounds of communication to filter out faulty behavior and ensure consistency among the non-faulty processes.

Fischer discusses different models of computation and types of faults, including crash faults, where a process simply stops functioning, and Byzantine faults, where a process behaves arbitrarily and possibly maliciously. He explains that solving these problems depends heavily on the assumptions made about the system, such as whether failures are detectable and whether processes run synchronously or asynchronously. In synchronous systems, processes can rely on timing to detect faults, but in asynchronous systems, where no timing guarantees exist, detecting faults becomes much more challenging.

The paper outlines various theoretical results, including the conditions under which consensus can be achieved. For example, it is shown that without authentication, a solution to the Byzantine Generals Problem is possible only if less than one-third of the processes are faulty. Fischer also presents results showing that in a fully asynchronous system, achieving consensus is impossible if even a single process can fail, highlighting the difficulties of fault tolerance in distributed systems.

Overall, Fischer's survey provides a comprehensive overview of the consensus problem in distributed systems, highlighting both the challenges and the theoretical solutions developed over the years to address these issues in unreliable environments.

2.5 Impossibility of Distributed Consensus with One Faulty Process

As one of the most influential in the field of Distributed Consensus, this paper [26] by Michael J. Fischer, Nancy A. Lynch, and Michael S. Paterson addresses the problem of achieving consensus in asynchronous distributed systems with the possibility of process failures. The central challenge is for all nonfaulty processes to agree on a binary value, even when some processes may fail or behave unreliably.

The paper proves a significant result: in a completely asynchronous system, it is impossible to guarantee consensus if even a single process can fail. This is true even under minimal failure conditions where the only issue is the sudden stopping (crash) of one process. The proof does not consider more complex failures like Byzantine faults, where a process might behave arbitrarily or maliciously, but focuses on the basic case where messages are always delivered correctly and in order.

The proof relies on the asynchronous nature of the system, where no assumptions are made about the relative speeds of processes or the timing of message delivery. Without synchronized clocks or the ability to detect failures reliably, any protocol aiming to achieve consensus might encounter situations where it cannot progress towards a decision, leading to nontermination. The result shows that any such protocol has a "window of vulnerability," during which a single fault can cause the system to remain indecisive indefinitely.

To demonstrate a practical application of these concepts, the paper also presents a protocol that successfully achieves consensus in cases where no processes fail during execution

and a majority of processes are initially active. This protocol operates in two stages: first, processes construct a graph of communication links, and then they decide based on the values communicated through this graph. The protocol's correctness is guaranteed as long as a majority of the processes are nonfaulty and remain active throughout its execution.

The findings emphasize the need for more refined models of distributed computing that account for practical constraints, such as partial synchrony or probabilistic guarantees of termination, to circumvent the impossibility result and develop robust consensus protocols. These insights have shaped subsequent research in distributed computing, leading to new approaches and algorithms that consider partial synchrony or probabilistic consensus to cope with the inherent limitations of asynchronous systems.

2.6 Another Advantage of Free Choice: Completely Asynchronous Agreement Protocols

Michael Ben-Or presents in this paper [9] a probabilistic solution to the consensus problem in asynchronous distributed systems, where processes may fail or behave unreliably. The study builds on previous work by Fischer, Lynch, and Paterson. Ben-Or introduces a protocol that overcomes this limitation by using randomness, ensuring that a decision will be reached with high probability, provided a majority of the processes remain operational.

In the proposed protocol, processes exchange information in rounds. Each process starts with a binary input value and communicates through messages. If a process receives enough consistent information from other processes, it can make a decision. If no consensus is reached in a round, processes randomly choose their next actions, which leads to a situation where, eventually, all nonfaulty processes will agree with high probability. This randomness prevents the adversary from indefinitely delaying consensus.

Ben-Or's protocol also addresses the Byzantine Agreement problem, where faulty processes can behave arbitrarily or even maliciously. By using a similar probabilistic approach, the protocol can reach consensus as long as the number of faulty processes is less than one-fifth of the total number of processes. The key innovation is the use of randomness, which allows the protocol to succeed even in the presence of Byzantine faults, a scenario where deterministic protocols would typically fail.

Although the proposed protocols are not the most efficient—potentially requiring a large number of communication rounds—their ability to reach consensus in an unpredictable and faulty environment highlights the strength of probabilistic methods. The study shows that, in contrast to deterministic approaches that are limited by the impossibility results, probabilistic protocols can offer a viable path to achieving reliable consensus in distributed systems.

2.7 The correctness proof of Ben-Or’s randomized consensus algorithm

This paper by Marcos K. Aguilera and Sam Toueg [2] provides a detailed analysis and proof of Ben-Or’s randomized consensus algorithm, which is used in asynchronous distributed systems where processes may fail by crashing. The algorithm is significant because it was the first to achieve consensus with high probability, circumventing earlier impossibility results by using randomization. Despite its simplicity, a full correctness proof that considers a strong adversary—one that can observe the state of processes and messages—had not been provided until this work.

The authors demonstrate that Ben-Or’s algorithm remains correct under the assumption that fewer than half of the processes can fail (i.e., $f < n/2$, where n is the total number of processes, and f is the number of faulty ones). They show that consensus can be achieved even if a strong adversary tries to manipulate the system. The paper also discusses a potential issue with using a “global coin” to speed up the algorithm. While global coins are often used to synchronize random decisions across processes, the authors found that in certain scenarios, such as when $n/3 \leq f < n/2$, using a global coin can actually prevent the algorithm from terminating, highlighting the complexity and subtlety of designing randomized consensus protocols in adversarial environments. The work concludes that while Ben-Or’s algorithm is robust under certain conditions, care must be taken in its implementation, particularly regarding the use of shared randomness.

2.8 FPC-BI: Fast Probabilistic Consensus within Byzantine Infrastructures

Proposed by Serguei Popov and William J. Buchanan [50], this paper introduces a novel consensus protocol designed for use in environments where some nodes may be controlled by an adversary. This protocol, called Fast Probabilistic Consensus within Byzantine Infrastructures (FPC-BI), is leaderless and maintains low communication complexity, making it suitable for large-scale, decentralized systems like cryptocurrencies.

FPC-BI operates under the assumption that a portion of the nodes may act maliciously, aiming to delay or disrupt the consensus process. Despite this, the protocol achieves consensus with high probability by leveraging random numbers, functioning as global coins, either from a trusted source or generated via decentralized methods, to guide decision-making. This randomness helps overcome adversarial influence, ensuring that honest nodes can reach a consistent decision on the value of a bit, which could relate to the validity of a transaction. A noteworthy aspect of FPC-BI is its resilience; it does not require perfect consensus on the values of these global coins—an approximate agreement is sufficient to drive the decision-making process.

The protocol’s design ensures that consensus is achieved efficiently, even when a significant number of nodes are Byzantine, by adjusting parameters to optimize the likelihood of

fast convergence to a single consensus state. The use of probabilistic methods and decentralized randomness provides a robust framework that enhances the security and reliability of distributed ledger systems, making FPC-BI a valuable tool for maintaining integrity in environments prone to faults and attacks.

Chapter 3

Fast Probability Consensus on a Set (FPCS)

The Fast Probabilistic Consensus on a Set (FPCS) is an advanced voting consensus protocol designed as an evolution of the FPC-BI. Both protocols share several key characteristics: they operate in rounds, are probabilistic in nature, achieve termination, integrity, and agreement with high probability, and assume the existence of a global coin.

While FPC-BI addresses the classical problem of achieving consensus on the value of a bit, FPCS extends this to achieving consensus on a Maximal Independent Set (MIS) of a graph of conflicting transactions. This extension is crucial for managing more complex scenarios, particularly in UTXO-based Distributed Ledger Technologies (DLTs)—like Bitcoin, Cardano, Kaspas, and IOTA—where transactions sharing the same output are considered to be in conflict.

Conflicts in these systems can arise from malicious behavior but are more commonly the result of faulty node behavior, and concurrency issues in UTXO-based smart contracts, and transaction duplicates in DAG-based DLTs with low block times. Despite these varied origins, our results remain unaffected as we consider a proportion q of nodes to be Byzantine, meaning their behavior is arbitrary, encompassing both malicious and faulty nodes.

Consider a scenario with three transactions: u (consuming UTXOs u_1 and u_2), v (consuming u_1), and w (consuming u_2). Then, u conflicts with both v and w , while v does not conflict with w . To update their ledgers, nodes must choose between accepting the set $\{u\}$ or the set $\{v, w\}$ as legitimate. Each round of the FPCS protocol generates a set of transactions to be liked by the nodes, which may change significantly in the initial rounds. However, the protocol is designed to stabilize after a few rounds, with nodes consistently liking similar transactions, eventually triggering a stop criterion and rendering the consensus final.

3.1 Notation

Consider a set of N nodes denoted by $\mathcal{N} = \{1, \dots, N\}$ and a set of conflicting transactions $\mathbb{T}_t = \{u_1, \dots, u_{T_t}\}$, which we call the *conflict set*. As the time dependency suggests, we assume that the conflict set can change over time. The exact dynamics of how new transactions

emerge will not be explored and we consider it can happen arbitrarily, nevertheless new transactions must always be *liked*¹ by at least one node. In particular, we assume that nodes themselves can issue new conflicting transactions, which can potentially be seen as a vector of attack. For our purposes, we assume that a transaction is composed of a unique transaction identifier (“Id”, for short), a set of inputs – often referred to as UTXOs (unspent transaction outputs) – and a set of outputs.

We also assume that outputs are unequivocally associated to a transaction that generated it, meaning that, by construction, a single output will be generated by exactly one transaction. Additionally, a transaction will never consume the same UTXO more than once. Let $\mathcal{U}_y$ be the set of outputs consumed by a certain transaction $y \in \mathbb{T}_t$, and $\mathcal{V}_y$ be the set of transactions that generated $\mathcal{U}_y$. We denote that the *past cone* $\mathcal{P}_y$ of a transaction y as $\mathcal{P}_y = y \cup \mathcal{V}_y \cup \{\mathcal{P}_x\}_{x \in \mathcal{V}_y}$.




Figure 3.1 Diagram representing the general structure of two transactions.

We say two transactions $x, y \in \mathbb{T}_t$ are in conflict if there is a UTXO being consumed by more than one transactions in the union of their past cones have a non-empty intersection (or, in other words, if accepting both transactions would imply there is a double spending somewhere in the ledger) and denote this by $x \leftrightarrow y$. If x and y are not in conflict, we write $x \nleftrightarrow y$. A transaction x conflicts with a set $B \subset \mathbb{T}$ if it conflicts with every element of B and this is represented by $x \leftrightarrow B$. It is natural to represent the set $\mathbb{T}_t$ and its conflicts as a graph $G = (\mathbb{T}_t, E_t)$, where given $x, y \in \mathbb{T}_t$ an edge $(x, y) \in E_t$ denotes that $x \leftrightarrow y$.

The conflict graph $\mathbb{T}_t$ is not necessarily complete², but can have a more complicated structure. For example, imagine a set of four transactions $\{w, x, y, z\} \in \mathbb{T}_t$ and two UTXOs, u_1, u_2 . Suppose that w consumes u_1 , x consumes u_1 and u_2 and both y and z consume u_2 ; then, $w \leftrightarrow x$ and $x \leftrightarrow \{y, z\}$ but, at the same time, w is not in conflict with y or z (see Fig. 3.2).




Figure 3.2 Graph representation of the example described in the text. The edges between transactions indicate that there is a common input.

¹The precise definition of a liked transaction is given below

²A graph $G = (V, E)$ is *complete* if, for any two vertices in V , there is an edge in E directly connecting them.

Considering the graph representation of conflicting transactions, it is useful to introduce a few additional definitions. Given a graph $G = (V, E)$, we say a subset $S \subset V$ is an *independent set* if, for any two elements of S , there is no edge in E connecting them. Furthermore, a *maximal independent set* (MIS for short) is an independent set that is not a proper subset of any other independent set. In the previous example, the singletons, $\{w, y\}$ and $\{w, z\}$ are independent sets, while $\{x\}$, $\{w, y\}$ and $\{w, z\}$ are maximal independent sets. Complete graphs are a particular case that we will study closely, it represents the scenario where all transactions in $\mathbb{T}_t$ have at least one input in common with every other transaction. This case is also known in cryptocurrency literature as a *n-spend*.




Figure 3.3 Graph representation of a set with 6 transactions where all of them are in conflict with each other.

Depending on the network throughput, the set $\mathbb{T}_t$ can be very large and it is convenient to find a global (i.e., known to all nodes) way to totally order it. A popular way to do this is through a *cryptographic hash function* (CHF for short). Following the definition introduced by Merkle [36], a CHF is a function F such that:

1. F can be applied to any argument of any size and produces a fixed-size output.
2. Given F and x , it is easy to compute $F(x)$.
3. Given F , it is computationally infeasible to find any pair x, x' such that $x \neq x'$ and $F(x) = F(x')$. We call this the *no-collision condition*.

For the purpose of the paper, we will also assume that the CHF is a pseudo-random function. Pseudo-random means that the outcome of the function is deterministic, but it “appears” (passes a wide class of statistical tests) to be uniformly distributed in the interval $[0, 1]$. In other words, our hash function satisfies the property that any random perturbation in the input results in a uniformly distributed independent new output. We refer to this as the *diffusion property*. Let us note that the hash function allows us to define an order on some arbitrary data x, y : one can say that $x < y$ if $\text{hash}(x) < \text{hash}(y)$.

Considering discrete time $t = 0, 1, 2, \dots$ (we refer to it as the round t), we define by $A_t^{(n)}$ the set of transactions known by the node n at time t and call it the *node’s vision*. Furthermore, we assume that the node never forgets transactions, i.e., we assume that $A_0^{(n)} \subset A_1^{(n)} \subset \dots \subset \mathbb{T}_t$ for any $n \in \mathcal{N}$.

We say a node *likes* a transaction if it prefers it to its conflicting transactions. Moreover, we define *node n ’s opinion* at round t as the collection $O_t^{(n)} = \{\theta_t^{(n,x)}; x \in A_t^{(n)}\}$, where $\theta_t^{(n,x)}$

assumes the value:

$$\theta_t^{(n,x)} = \begin{cases} 1, & \text{if node } n \text{ likes transaction } x \text{ at time } t, \\ 0, & \text{otherwise.} \end{cases}$$

For a set $W \subset \mathbb{T}_t$, we say a node n likes W if it likes every transaction in W , or explicitly, $\theta_t^{(n,W)} = 1$ if $\theta_t^{(n,w)} = 1$ for all $w \in W$.

We also assume there exists a public sequence of random numbers $X_t \sim U[\beta, 1 - \beta]$, which is either provided by a trusted source or generated by the nodes themselves using some decentralized random number generating protocol. This approach is referred to as a *global coin* in many works on Byzantine consensus, for example, in [2, 14, 15, 27]. We assume all random numbers and messages between the nodes are delivered on time in every round.

Our objective is to formulate a protocol that facilitates consensus among the nodes $\mathcal{N}$ regarding a MIS within $\mathbb{T}_t$. A designated proportion q of the nodes, referred to as *malicious*, may opt not to adhere to our protocol, thereby choosing to impede or disrupt the consensus process. For a constant $c \in [0, 1]$, we say the protocol is *resistant* up to a threshold c if, for any $q < c$, consensus can be achieved with high probability.

Without loss of generality, we assume that the first $(1 - q)N$ nodes are *honest* (those that are not malicious) and define the proportion of likes among honest nodes of a transaction $u \in \mathbb{T}_t$ as

$$p_t^{(u)} := \frac{1}{(1 - q)N} \sum_{j=1}^{(1-q)N} \theta_t^{(j,u)}. \quad (3.1)$$

For a set $U \in \mathbb{T}_t$, the proportion of likes $p_t^{(U)}$ is defined as the proportion of honest nodes that like the whole set U .

We define the *Interval of Control* of the malicious nodes over a transaction v at round t as

$$\mathcal{I}_{q,t}^{(v)} := [(1 - q)p_t^{(v)}, (1 - q)p_t^{(v)} + q].$$

The lower/upper boundary of this interval is precisely the overall proportion of likes (i.e., considering both honest and malicious opinions) that the transaction has when all malicious nodes dislike/like it.

Our results hinge on the observation that once a significant majority of honest nodes align on a specific transaction or set of transactions, it becomes difficult for malicious nodes to reverse this opinion. To precisely delineate the threshold for a significant majority, we introduce the abbreviation

$$\mu := \frac{\beta - q}{2(1 - q)}.$$

Fig. 3.4 illustrates the concept of Intervals of Control and the relations between β and q . In particular, we will assume that $q < \beta$, or in other words, that μ is positive.




Figure 3.4 The intervals of control $\mathcal{I}_{q,t}^{(u)}$ for a transaction u which has a proportion of likes among honest nodes $p_t^{(u)} > 1 - \mu$ (in blue), and $\mathcal{I}_{q,t}^{(w)}$ for a transaction w such that $p_t^{(w)} < \mu$ (in red). Notice that these intervals are separated from the support $[\beta, 1 - \beta]$ of X_t .

With this in mind, we say that at time t , the system is in *pre-consensus* about a transaction $u \in \mathbb{T}_t$ if either $p_t^{(u)} > 1 - \mu$ or there exists $v \in N(u)$ ³ such that $p_t^{(v)} > 1 - \mu$. We denote this event (i.e., when either of the conditions above is true) by $\mathcal{PC}_t(u)$. For a set $U \in \mathbb{T}_t$ the system is in pre-consensus about U if $\mathcal{PC}_t(u)$ is true for every $u \in U$ and we represent this event by $\mathcal{PC}_t(U)$. If $\mathcal{PC}_t(u)$ is true for all $u \in \mathbb{T}_t$, we simply say the system is in pre-consensus and denote this event by $\mathcal{PC}_t$. Note that, under $\mathcal{PC}_t$, each node will have its preferred MIS, and the intersection could be very large, but in order to prove our theorem, we need the stronger condition that a large proportion of honest nodes like the exact same MIS. Hence, if $\mathcal{PC}_t$ is true and on top of that $p_t^{(U)} > 1 - \mu$ for some maximal independent set $U \subset \mathbb{T}_t$ we say the system is in *hard pre-consensus* and denote this event by $\mathcal{PC}_t^*$.

3.2 Protocol

For $t = 0$, the initial visions, $A_0^{(n)}$, and opinions, $O_0^{(n)}$, can be arbitrary, as long as the set of liked transactions for each node forms a maximal independent set of the conflict graph induced by $A_0^{(n)}$. Then, the following protocol should be executed iteratively once for every round $t \geq 1$ until the stop criterion is met:

1. At the beginning of every round t , each node n stores any new transaction it became aware of between rounds $t - 1$ and t , call this set $N_t^{(n)}$.
2. The node proceeds to randomly query (with uniform distribution) k nodes about their known transactions and which ones of them are liked.
3. The node then also stores any shared-by-others transactions it was not aware of, call this set $S_t^{(n)}$.
4. The node's vision at round t is then given by $A_t^{(n)} = A_{t-1}^{(n)} \cup S_t^{(n)} \cup N_t^{(n)}$

³Here we use the definition of a neighborhood from graph theory. Explicitly, for sets $U, V \subset \mathbb{T}_t$ the neighborhood of U with respect to V is defined as $N(U, V) := \{v \in V; \text{ such that there exists } u \in U \text{ for which } u \leftrightarrow v\}$. For simplicity we abbreviate $N(U) := N(U, \mathbb{T}_t)$

5. The node stores the collection $\{\eta_t^{(n,x)}; x \in A_t^{(n)}\}$, where $\eta_t^{(n,x)}$ corresponds to the number of 1-opinions the node n received from the queries in round t with respect to the transaction x .
6. The node receives the random value $X_t \sim U[\beta, 1 - \beta]$.
7. The node defines an auxiliary collection of opinions $\{\theta'(x); x \in A_t^{(n)}\}$, that will not be shared and will last only until the end of the round (hence we omit the dependence on n and t), using the following rule:

$$\theta'(x) = \begin{cases} 1, & \text{if } \eta_t^{(n,x)} / k > X_t \\ 0, & \text{otherwise.} \end{cases}$$

8. Let $B := \{x \in A_t^{(n)}; \theta'(x) = 1\}$. The node must find a way to assign 1 only to the opinions of a maximal independent subset of $A_t^{(n)}$. To do so, it iteratively removes from B the transaction $x \in B$ with the largest hash(Id_x, X_t) (this means the hash of Id_x concatenated with the random number X_t) until it obtains an independent set. Note that using the “largest hash” is not crucial, as any deterministic rule leading to unpredictable results is sufficient. Explicitly, it performs the following algorithm:

Algorithm 1 $\text{elim}(U, X_t)$

- 1: $W = U$
 - 2: **while** W is not an independent set **do**
 - 3: Compute
 - 4: $W = W \setminus \{y\}$
 - 5: **end while**
 - 6: **return** W
-

$$y = \underset{x \in W: \exists z \in W: z \leftrightarrow x}{\text{argmax}} \text{ hash}(\text{Id}_x, X_t),$$

Consider $B' := \text{elim}(B, X_t)$. While this set is independent by construction, it may not be maximal. Then, starting with B' , the node includes iteratively the non-conflicting transaction with the smallest hash(x, X_t) until a maximal independent set is obtained. Explicitly, the node executes the following:

Let $B'' := \text{compl}(B', A_t^{(n)}, X_t)$. Finally, the node assigns value 1 to the opinion $\theta_{t+1}^{(n,x)}$ of every transaction $x \in B''$ and zero to the others.

If the node’s opinion about a transaction does not change for ℓ rounds, then it is considered final and will not be further modified in the subsequent rounds.

Algorithm 2 $\text{compl}(U, V, X_t)$

-
- 1: $W = U$
 - 2: **while** W is not a maximal independent set **do**
 - 3: Compute

$$y = \underset{x \in V \setminus N(W, V)}{\operatorname{argmin}} \text{hash}(x, X_t)$$
 - 4: $W = W \cup \{y\}$
 - 5: **end while**
 - 6: **return** W
-




Figure 3.5 Representation of step (vi) of the protocol. The graph represents the node's vision at time t . The numbers represent the order according to the hash function. In the first subfigure, the transactions in yellow represent the set B ; in the second subfigure, transactions in green represent B' and transactions in red represent $N(B', A_t^{(n)})$; in the last subfigure, transactions in blue represent B'' , the chosen maximal independent set of the round.

Chapter 4

Solving n -spends with FPCS

Complete graphs are fundamental in graph theory, where each pair of distinct vertices is connected by a unique edge, forming the most densely connected type of graph. In the context of this work, this means that every transaction spend at least one UTXO in common, and, because of that, it is popularly known in cryptocurrency context as n -spends. The topology of this type of graphs allows us to assume very weak conditions regarding their size. In particular, we consider that $|\mathbb{T}_t| = f(t)$ for any nondecreasing function f taking the naturals into themselves. Another rather trivial characteristic is that in the complete case, pre-consensus in a single transaction is equivalent to hard pre-consensus. Indeed, if $p_t^{(u)} > 1 - \mu$ for some $u \in \mathbb{T}_t$, then for any $v \in \mathbb{T}_t \setminus \{u\}$, we have $u \in N(v)$ and then, by definition, $\mathcal{PC}_t(v)$ is also true.

At first, we consider that every node is aware of all the conflicting transactions at every round, i.e., $A_t^{(n)} = \mathbb{T}_t$ for every $t \in \mathbb{N}$ and all $n \in \mathcal{N}$.

Defining $\mathcal{C}$ as the event where honest nodes achieve consensus, our main result states that with high probability (depending on k and N), for any distribution of initial opinions, a pre-consensus state is achieved and maintained for enough consecutive rounds so that event $\mathcal{C}$ happens in the round $\mathcal{R} \leq 2\ell$.

Theorem 4.0.1. *For sufficiently large k , given an arbitrary distribution of the initial opinions of the honest nodes it holds that*

$$\mathbb{P}[\mathcal{C} \cap \{\mathcal{R} < 2\ell\}] \geq \left[1 - (1 - W_c(N, k))^{\ell-1}\right] \left[1 - \exp\left\{-\frac{k}{2}(\beta - q)^2\right\}\right]^{(1-q)N\ell}$$

where

$$W_c(N, k) = \frac{h_c}{2(1 - 2\beta)N} \left(1 - \exp\{-2N(1 - q)\varphi_3^2\}\right)$$

$$\varphi_3 := \mu - \exp\left\{-\frac{kh_c^2}{2}\right\}, \tag{4.1}$$

$$h_c := (1 - q - 2\beta)/2. \tag{4.2}$$

In order to prove the theorem, we need first to prove some additional lemmas. The first one estimates the probability of a pre-consensus state being maintained after one iteration of our algorithm. This result hinges on the fact that for transactions $u, v \in \mathbb{T}_t$ such that $p_t^{(v)} > 1 - \mu$ and $p_t^{(v)} < \mu$, there will be a distance larger than $(\beta - q)/2$ between the intervals of control of these transactions and the support of X_t (see Fig. 4.1).



Figure 4.1 The intervals of control for a transaction u which has a proportion of likes among honest nodes $p_t^{(u)} > 1 - \mu$ (in blue), and for a transaction v such that $p_t^{(v)} < \mu$ (in red). Notice that these intervals are separated from the support $[\beta, 1 - \beta]$ of X_t .

Lemma 4.0.1. *Suppose that, at round t , the system is in a state where $p_t^{(v_0)} > 1 - \mu$ for some $v_0 \in \mathbb{T}_t$. Then, for a sufficiently large k the probability of hard pre-consensus at round $t + 1$ (or in other words, that $\mathcal{PC}_{t+1}^*$ is true) satisfies*

$$\mathbb{P}[\mathcal{PC}_{t+1}^*] = \mathbb{P}[\mathcal{PC}_{t+1}(v_0)] \geq 1 - \exp\{-2N(1 - q)\varphi_1^2\}.$$

Proof. Note that if $p_t^{(v_0)} > 1 - \mu$ the lower boundary of $\mathcal{I}_{q,t}^{(v_0)}$ is greater than $1 - (\beta + q)/2$. For a node n we consider the event $G^{(1)} := \{\eta_t^{(n,v_0)} k^{-1} > 1 - \beta\}$ ¹ that a proportion larger than $1 - \beta$ of the queried nodes like v_0 . Hence,

$$\mathbb{P}[G^{(1)} \mid p_t^{(v_0)} > 1 - \mu] \geq \mathbb{P}[k^{-1}S_k \geq 1 - \beta],$$

where $S_k \sim \mathcal{B}(k, 1 - (\beta + q)/2)$. This latter probability can be estimated using the Hoeffding inequality

$$\begin{aligned} \mathbb{P}[k^{-1}S_k \geq 1 - \beta] &\geq 1 - \exp\left\{-2k\left[1 - \beta - \left(1 - \frac{\beta + q}{2}\right)\right]^2\right\} \\ &= 1 - \exp\left\{-\frac{k}{2}(\beta - q)^2\right\}. \end{aligned}$$

Under the event $G^{(1)}$, every other transaction is liked by a proportion of the queried nodes smaller than β and, consequently, every other auxiliary opinion is equal to zero. Then, the probability that an honest node likes transaction v_0 at round $t + 1$ is estimated by

$$\begin{aligned} \mathbb{P}[\theta_{t+1}^{(n,v_0)} = 1 \mid p_t^{(v_0)} > 1 - \mu] &\geq \mathbb{P}[\theta'(v_0) = 1, \theta'(v) = 0 \quad \forall v \neq v_0 \mid p_t^{(v_0)} > 1 - \mu] \\ &\geq \mathbb{P}[G^{(1)} \mid p_t^{(v_0)} > 1 - \mu] \geq 1 - \exp\left\{-\frac{k}{2}(\beta - q)^2\right\}. \end{aligned} \quad (4.3)$$

¹The superscript on events G s denote Lemma 1, 2, 3, etc, while subscripts denote different events within the same Lemma

Now, the probability that a proportion of at least $1 - \mu$ of the honest nodes likes transaction v_0 at round $t + 1$ can be estimated by the probability of a random variable $S_{N(1-q)} \sim \mathcal{B}(N(1-q), 1 - \exp\{-(k/2)(\beta - q)^2\})$ being larger than $(1 - \mu)$. Considering that k can be taken sufficiently large² so that the probability of success of this binomial is larger than $1 - \mu$, using again the Hoeffding inequality we find that

$$\begin{aligned} \mathbb{P}[\mathcal{PC}_{t+1}(v_0)] &\geq \mathbb{P}[p_{t+1}^{(v_0)} > 1 - \mu] \geq \mathbb{P}\left[\frac{1}{N(1-q)} S_{N(1-q)} > 1 - \mu\right] \\ &> 1 - \exp\{-2N(1-q)\varphi_1^2\}. \end{aligned}$$

□

We have shown that if $p_t^{(v_0)} > 1 - \mu$ and for k are sufficiently large, then $p_{t+1}^{(v_0)} > 1 - \mu$ will happen high probability. Our next objective is to prove that a pre-consensus state can, in fact, be achieved relatively fast, given any distribution of honest nodes' likes. For this purpose, we distinguish between the two following cases; Case 1 represents the situation where the proportion of likes assigned to the favorite transaction is smaller or equal to $1/2$, whereas Case 2 represents the complementary case.

First notice that, since intervals of control have length q , an interval that is centered around $1/2$ will be completely contained in the support of the random variable X_t as long as $q < 1 - 2\beta$. Considering this condition together with $q < \beta$ introduced before and maximizing³ for q gives us the restriction $q < \beta < 1/3$ which is the security threshold considered on the remaining of this chapter.

Now given a transaction u such that $p_t^{(u)} = 1/2$, its interval of control $\mathcal{I}_{q,t}^{(v)} = [(1 - q)/2, (1 + q)/2]$ is not only contained in the support of X_t but is separated from the boundary of this domain by a distance $h_c = (1 - q - 2\beta)/2$ (see Fig. 4.2). Consequently, in Case 1/Case 2, the distance between the interval of control of the favorite transaction and the lower/upper boundary of $[\beta, 1 - \beta]$ is always larger than h_c .




Figure 4.2 The interval of control (in blue) of a transaction $u \in \mathbb{T}_t$ with $p_t^{(u)} = 1/2$ is always separated from the boundary of $[\beta, 1 - \beta]$ by a distance $h_c := (1 - q - 2\beta)/2$.

Denote by v_t the transaction with the highest proportion⁴ of honest likes at round t ; in other words, $p_t^{(v_t)} \geq p_t^{(w)}$ for any $w \in \mathbb{T}_t$. We start with Case 1. The main idea, in this case, is that with a probability $h/(2 - 4\beta)$, the random number X_t belongs to the interval $(1 - \beta - h/2, 1 - \beta)$ that is separated from $\mathcal{I}_{q,t}^{(v_t)}$ by a distance larger than $h/2$, which means that an honest node will likely assign 0 to the auxiliary opinion of v_t (and every other

²This will be a recurrent argument; we investigate how large k must be in Section 4.1.

³We want our protocol to be resistant to the largest q possible.

⁴We assume that this maximum is unique, otherwise we could choose the one with largest hash.

transaction, since they all have a smaller proportion of likes) and then, following our protocol, choose the transaction with the smallest hash as the favorite (see Fig. 4.3).




Figure 4.3 In Case 1, every time X_t falls on the right of $\mathcal{I}_{q,t}^{(v_t)}$ (in blue), an honest node is likely to assign 0 to the auxiliary opinion of every transaction.

Lemma 4.0.2. Let $p_t^{(v_t)} \leq 1/2$ be the largest proportion of honest likes among all transactions at round t . Then, for a sufficiently large k , there exists a transaction $u \in \mathbb{T}_t$ (not necessarily v_t) such that the probability that a proportion of at least $1 - \mu$ of honest nodes likes u at round $t + 1$ satisfies

$$\mathbb{P} \left[p_{t+1}^{(u)} > 1 - \mu \right] \geq \frac{h}{2(1-2\beta)} \left(1 - \exp \{ -2N(1-q)\varphi_2^2 \} \right),$$

where

$$\varphi_2 = \mu - N \exp \left\{ -\frac{kh_c^2}{2} \right\}$$

Proof. Let $G^{(2)}$ be the event that the random threshold X_t lies in the interval $(1 - \beta - h/2, 1 - \beta)$ and $u = \operatorname{argmin}_{v \in \mathbb{T}_t} \operatorname{hash}(v, X_t)$. Since the upper boundary of $\mathcal{I}_{q,t}^{(v_t)}$ is smaller than $1 - \beta - h/2$, we consider $S_k \sim \mathcal{B}(k, 1 - \beta - h)$. Now, we use Hoeffding's inequality to bound the probability that an honest node assigns 0 to the auxiliary opinion of v_t :

$$\begin{aligned} \mathbb{P} \left[\theta'(v_t) = 0 \mid p_t^{(v_t)} \leq 1/2, G^{(2)} \right] &\geq \mathbb{P} \left[k^{-1} S_k < X_t \right] \\ &\geq \mathbb{P} [k^{-1} S_k < 1 - \beta - h_c/2] \\ &\geq 1 - \exp \{ -2k \{ 1 - \beta - h_c/2 - (1 - \beta - h_c) \}^2 \} \\ &\geq 1 - \exp \left\{ -\frac{kh_c^2}{2} \right\}. \end{aligned} \tag{4.4}$$

While this estimate holds for all transactions, in some cases it is excessively conservative. Notice, for example, that in general, the auxiliary opinion for a transaction v can take value 1 only if at least one of the N nodes likes v , while all others will be disliked with a probability 1. Consequently, the probability of an honest node assigning $\theta'(v) = 0$ for every transaction and then, following the protocol, liking transaction u at round $t + 1$ satisfies:

$$\mathbb{P} \left[\theta_{t+1}^{(n,u)} = 1 \mid p_t^{(v_t)} \leq 1/2, G^{(2)} \right] \geq \left(1 - \exp \left\{ -\frac{kh_c^2}{2} \right\} \right)^N \geq 1 - N \exp \left\{ -\frac{kh_c^2}{2} \right\}, \tag{4.5}$$

where this last term will be positive for a sufficiently large k . Considering $S_{N(1-q)} \sim \mathcal{B}(N(1-q), 1 - N \exp \{ -kh_c^2/2 \})$ then, analogously to the previous lemma, assuming k is sufficiently

large, it follows that the probability of a proportion of at least $1 - \mu$ of the honest nodes liking transaction u at round $t + 1$ satisfies

$$\begin{aligned}
& \mathbb{P} \left[p_{t+1}^{(u)} > 1 - \mu \mid p_t^{(v_t)} \leq 1/2 \right] \\
& \geq \mathbb{P}[G^{(2)}] \mathbb{P} \left[p_{t+1}^{(u)} > 1 - \mu \mid p_t^{(v_t)} \leq 1/2, G^{(2)} \right] \\
& \geq \frac{h_c}{2(1-2\beta)} \mathbb{P} \left[\frac{1}{N(1-q)} S_{N(1-q)} > 1 - \mu \right] \\
& \geq \frac{h_c}{2(1-2\beta)} \left(1 - \exp \left\{ -2N(1-q) \left(\left[1 - N \exp \left\{ -\frac{kh_c^2}{2} \right\} \right] - 1 + \mu \right)^2 \right\} \right) \\
& = \frac{h_c}{2(1-2\beta)} \left(1 - \exp \{ -2N(1-q)\varphi_2^2 \} \right).
\end{aligned}$$

□

We now consider Case 2, i.e., if $p_t^{(v_t)} > 1/2$. In this situation, a random threshold smaller than $p_t^{(v_t)}(1-q)$ (the lower boundary of $\mathcal{I}_{q,t}^{(v_t)}$, the blue interval in Figure 4.4) will not only favor v_t , but can also favor another transaction w , which overall proportion can be increased to $p_t^{(w)}(1-q) + q$ (the upper boundary of $\mathcal{I}_{q,t}^{(w)}$, the red interval in Figure 4.4). This is not necessarily a problem, but in order to guarantee that a large proportion of honest nodes like the same transaction we rely on the diffusion property of the hash function.




Figure 4.4 In Case 2, every time X_t falls on the left of $\mathcal{I}_{q,t}^{(v_t)}$ (in blue), an honest node might also like a transaction w if the upper boundary of $\mathcal{I}_{q,t}^{(w)}$ (in red) is also on the right of X_t .

Lemma 4.0.3. Let $p_t^{(v_t)} > 1/2$ be the largest proportion of honest likes among all transactions at round t . Then, for a sufficiently large k the probability that at least a proportion $1 - \mu$ of honest nodes likes the transaction v_t at round $t + 1$ satisfies

$$\mathbb{P}[p_{t+1}^{(v_t)} > 1 - \mu] \geq \frac{k\beta h}{2(1-2\beta)(N+k\beta)} \left(1 - \exp \{ -2N(1-q)\varphi_3^2 \} \right),$$

where

$$\varphi_3 := \mu - \exp \left\{ -\frac{kh_c^2}{2} \right\}.$$

Proof. Let G_3 be the event that the random threshold X_t belongs to the interval $(\beta, \beta + h/2)$. Now, the lower boundary of $\mathcal{I}_{q,t}^{(v_t)}$ is greater than $\beta + h/2$. Let $S_k \sim \mathcal{B}(k, \beta + h)$, then, the

probability that an honest node likes transaction v_t conditioned on $G_1^{(3)}$ is estimated again with Hoeffding's inequality by

$$\begin{aligned} \mathbb{P}[\theta'(v_t) = 1 \mid p_t^{(v_t)} > 1/2, G_1^{(3)}] &\geq \mathbb{P}[k^{-1}S_k \geq X_t] \\ &\geq \mathbb{P}[k^{-1}S_k \geq \beta + h_c/2] \\ &\geq 1 - \exp\left\{-\frac{kh_c^2}{2}\right\}. \end{aligned} \quad (4.6)$$

As stated in the previous lemma, the number of auxiliary opinions that are 1 is bounded by N . Instead of studying every possible combination of likes and dislikes of v_t and its potential competitors, we consider the event $G_2^{(3)}$ that v_t , at round t , has the smallest hash⁵ among them. Then, by the design of the protocol, if an honest node assigns 1 to the auxiliary opinion of v_t , it will be chosen as the node's favorite independently of its auxiliary opinion on the competitors. By the diffusion property of the hash function, the event $G_2^{(3)}$ occurs roughly with probability $1/N$. This means that, for an honest node n , it holds that

$$\mathbb{P}[\theta_{t+1}^{(n,v_t)} = 1 \mid p_t^{(v_t)} > 1/2, G_1^{(3)}, G_2^{(3)}] \geq \mathbb{P}[\theta'(v_t) = 1 \mid p_t^{(v_t)}, G_1^{(3)}]. \quad (4.7)$$

It follows that for a sufficiently large k , the probability of a proportion $1 - \mu$ of honest nodes liking transaction v_t at round $t + 1$ is estimated by

$$\begin{aligned} \mathbb{P}[p_{t+1}^{(v_t)} > 1 - \mu \mid p_t^{(v_t)} > 1/2] &\geq \mathbb{P}[G_1^{(3)}, G_2^{(3)}] \mathbb{P}[p_{t+1}^{(v_t)} > 1 - \mu \mid p_t^{(v_t)} > 1/2, G_1^{(3)}, G_2^{(3)}] \\ &\geq \frac{h_c}{2(1-2\beta)} \frac{1}{N} \mathbb{P}\left[\frac{1}{N(1-q)} S_{N(1-q)} > 1 - \mu\right] \\ &\geq \frac{h_c}{2(1-2\beta)} \frac{1}{N} \left(1 - \exp\left\{-2N(1-q) \left(\mu - \exp\left\{-\frac{kh_c^2}{2}\right\}\right)^2\right\}\right) \\ &= \frac{h_c}{2(1-2\beta)N} \left(1 - \exp\{-2N(1-q)\varphi_3^2\}\right). \end{aligned}$$

□

To ensure the results remain as clean as possible, it is prudent to determine the minimum between the probabilities of the events associated with Lemmas 4.0.2 and 4.0.3. The following result addresses this need and the proof is provided in the Appendix⁶.

Lemma 4.0.4. *For any positive constant c and sufficiently large k , it holds that*

$$\frac{1}{N} (1 - \exp\{-c\varphi_3^2\}) \leq 1 - \exp\{-c\varphi_2^2\} \quad (4.8)$$

⁵Rigorously, that $v_t = \operatorname{argmin}_{u \in \mathbb{T}} \operatorname{hash}(\operatorname{Id}_u, X_t)$

⁶We will follow this approach whenever a proof is purely mathematical.

Armed with these results, we now proceed to prove the theorem stated at the beginning of this chapter.

Proof of theorem 3.1 First, we define the random variables

$$\Psi := \min\{r \geq 1 : p_r^{(v_i)} > 1 - \mu\},$$

that is the first round in which the system is in the pre-consensus state, and

$$\hat{\tau}_n := \min\{r \geq \Psi + \ell : \theta_r^{(n, v_i)} = \dots = \theta_{r-\ell+1}^{(n, v_i)} = 1\},$$

that is the first time in which v_i is liked by node n for ℓ subsequent rounds after round Ψ . We then define the events

$$D_t^{(1)} = \{\Psi = t\}, \quad D_{[t,s]}^{(1)} = \{\Psi \in [t,s]\},$$

$$D^{(2)} = \{\hat{\tau}_n = \Psi + \ell; \forall n = 1, \dots, (1-q)N\}.$$

Notice that if $\Psi < \ell$, then

$$\mathcal{C} \cap \{\mathcal{R} \leq t + \ell\} \supset D_{[1,t]}^{(1)} \cap D^{(2)}.$$

This way

$$\mathbb{P}[\mathcal{C} \cap \{\mathcal{R} < 2\ell\}] \geq \mathbb{P}[D_{[1,t]}^{(1)} \cap D^{(2)}] = \mathbb{P}\left[\left\{\bigcup_{i=1}^{\ell-1} D_i^{(1)}\right\} \cap D^{(2)}\right],$$

since the events are disjoint, it follows that

$$\begin{aligned} \mathbb{P}[\mathcal{C} \cap \{\mathcal{R} < 2\ell\}] &\geq \mathbb{P}\left[\bigcup_{i=1}^{\ell-1} D_i^{(1)} \cap D^{(2)}\right] = \sum_{i=1}^{\ell-1} \mathbb{P}[D_i^{(1)} \cap D^{(2)}] \\ &= \sum_{i=1}^{\ell-1} \mathbb{P}[D^{(2)} | D_i^{(1)}] \mathbb{P}[D_i^{(1)}] = \sum_{i=1}^{\ell-1} \mathbb{P}[D^{(2)}] \mathbb{P}[D_i^{(1)}] \\ &= \mathbb{P}[D^{(2)}] \mathbb{P}[D_{[1,\ell-1]}^{(1)}] \end{aligned}$$

To calculate $\mathbb{P}[D^{(2)}]$ we write

$$\begin{aligned} \mathbb{P}[D^{(2)}] &= \mathbb{P}\left[\{\hat{\tau}_n = \Psi + \ell; \forall n = 1, \dots, (1-q)N\}\right] \\ &= \mathbb{P}\left[\left\{\theta_{\Psi+1}^{(n, v_i)} = \dots = \theta_{\Psi+\ell}^{(n, v_i)} = 1; \forall n = 1, \dots, (1-q)N\right\}\right] \\ &= \mathbb{P}\left[\bigcap_{i=1}^{\ell} \bigcap_{n=1}^{(1-q)N} \left\{\theta_{\Psi+i}^{(n, v_i)} = 1\right\}\right], \end{aligned}$$

To simplify the notation, define $D_i^{(3)} = \bigcap_{n=1}^{(1-q)N} \left\{ \theta_{\Psi+i}^{(n, v_i)} = 1 \right\}$, then

$$\mathbb{P}[D^{(2)}] = \mathbb{P} \left[\bigcap_{i=1}^{\ell} D_i^{(3)} \right] = \mathbb{P} \left[D_{\ell}^{(3)} | D_1^{(3)} \dots D_{\ell-1}^{(3)} \right] \dots \mathbb{P} \left[D_2^{(3)} | D_1^{(3)} \right] \mathbb{P} \left[D_1^{(3)} \right]. \quad (4.9)$$

But notice that if $D_i^{(3)}$ is true, then $\mathcal{PC}_{i+1}$ is also true, and using a similar bound to the one used in Lemma 4.0.1 we find that

$$\mathbb{P}[D^{(2)}] \geq \left[1 - \exp \left\{ -\frac{k}{2} (\beta - q)^2 \right\} \right]^{(1-q)N\ell}.$$

Now before calculating $\mathbb{P} \left[D_{[1, \ell-1]}^{(1)} \right]$, we define

$$W(n, k) := \min \left\{ \frac{h_c}{2(1-2\beta)} \left(1 - \exp \{ -2N(1-q)\varphi_2^2 \} \right), \right. \\ \left. \frac{h_c}{2(1-2\beta)N} \left(1 - \exp \{ -2N(1-q)\varphi_3^2 \} \right) \right\}$$

(these are the two bounds found in Lemmas 4.0.2 and 4.0.3. It follows directly from Lemma 4.0.4 that for a sufficiently large k , it holds that

$$W_c(n, k) = \frac{h_c}{2(1-2\beta)N} \left(1 - \exp \{ -2N(1-q)\varphi_3^2 \} \right).$$

Then

$$\mathbb{P} \left[D_{[1, \ell-1]}^{(1)} \right] = \sum_{i=1}^{\ell} \mathbb{P}[\Psi = i] \geq 1 - (1 - W_c)^{\ell-1}.$$

We finally conclude that

$$\mathbb{P}[\mathcal{C} \cap \{\mathcal{R} < 2\ell\}] \geq \left[1 - (1 - W_c(n, k))^{\ell-1} \right] \left[1 - \exp \left\{ -\frac{k}{2} (\beta - q)^2 \right\} \right]^{(1-q)N\ell}.$$

□

4.1 A lower bound for k

During the previous Section, we repeatedly used the argument that k was sufficiently large in our reasoning. Here we find bounds for it.

In Lemma 4.0.2 we assumed k was large enough so that the probability of success of $S_{N(1-q)} \sim \mathcal{B}(N(1-q), 1 - N \exp \{ -kh_c^2/2 \})$ to be larger than $1 - \mu$, this guarantees that

Hoeffding can be applied. Explicitly, we want

$$1 - N \exp \left\{ -\frac{kh_c^2}{2} \right\} > 1 - \mu$$

which is equivalent to the following condition to hold for k

$$k > \frac{2(\log N - \log \mu)}{h_c^2}. \quad (4.10)$$

Analogously, in Lemma 4.0.3 we required that k was sufficiently large so that

$$1 - \exp \left\{ -\frac{kh_c^2}{2} \right\} > 1 - \mu$$

which is equivalent to

$$k > -\frac{2 \log \mu}{h_c^2}.$$

Finally, in Lemma 4.0.4 we required k to be sufficiently large so that φ_2 and φ_3 to be larger than 0. This leads us to identical conditions on k .

We conclude then, that in order for our results to work, the stronger condition (4.10) must hold.

4.2 What if nodes' visions are slightly different?

For the remainder of this chapter, we no longer consider that every node is aware of the same set of transactions (i.e., the whole conflict set) at every round. Instead, we assume that the node's vision can be slightly different, as long as they maintain a sufficiently large – with respect to $\mathbb{T}_t$ – intersection; or explicitly we consider that there is some small $\delta > 0$ such that

$$\frac{\left| \bigcap_{n \in \mathcal{N}} A_t^{(n)} \right|}{|\mathbb{T}_t|} > 1 - \delta. \quad (4.11)$$

Since, by the design of our protocol, before calculating the auxiliary opinions a node includes in its vision any new transaction (which was issued or shared by another node), the estimates (4.4) and (4.6) are still valid. Furthermore, since in (4.7) the probability of an honest node n liking v_t is estimated by the event of the node assigning 1 to their auxiliary opinions—which means that these transactions were indicated by a queried node and consequently are included in $A_t^{(n)}$ —then Lemmas 4.0.1 and 4.0.3 are not affected by the condition (4.11). On the other hand, since in (4.5) we are assuming that a node assigned $\theta'(v) = 0$ for every transaction v , there is the possibility that some v are still unknown to this node and consequently the transaction with the smallest hash in the node's vision may not be the transaction with the smallest hash in the whole conflict set. This would invalidate the

final bound in Lemma 4.0.2 and to fix it, we rely again on the diffusion property of the hash function.

Lemma 4.2.1. *Considering condition (4.11), let $p_t^{(v_t)} \leq 1/2$ be the largest proportion of likes among all transactions at round t ; then, for a sufficiently large k , the probability that a proportion of at least $1 - \mu$ of honest nodes likes a transaction $u \in \mathbb{T}_t$ (not necessarily v_t) at round $t + 1$ satisfies*

$$\mathbb{P}[p_{t+1}^{(u)} > 1 - \mu \mid p_t^{(v_t)} \leq 1/2] \geq \frac{h_c}{2(1-2\beta)} \left(1 - \exp\{-2N(1-q)\hat{\phi}_2^2\}\right),$$

where

$$\hat{\phi}_2 := \mu - 1 + (1 - \delta) \left(1 - N \exp\left\{-\frac{kh_c^2}{2}\right\}\right).$$

Proof. Let G_1 be the event where the random threshold X_t is in the interval $(1 - \beta - h/2, 1 - \beta)$ and $u = \operatorname{argmin}_{v \in \mathbb{T}_t} \operatorname{hash}(v, X_t)$. The estimate (4.4) is not affected by condition (4.11), i.e., it still holds that

$$\mathbb{P}[\theta'(v_0) = 1 \mid p_t^{(v_0)} > 1 - \mu] \geq 1 - \exp\left\{-\frac{k}{2}(\beta - q)^2\right\}.$$

Now consider the event $G_2 = \{u \in \bigcap_{n \in \mathcal{N}} A_t^{(n)}\}$. By the properties of the hash function, it is immediate that $\mathbb{P}[G_2] = 1 - \delta$. Then, the probability of an honest node liking transaction u is bounded by

$$\begin{aligned} \mathbb{P}[\theta_{t+1}^{(n,u)} = 1 \mid p_t^{(v_t)}, G_1] &\geq (1 - \delta) \mathbb{P}[\theta_{t+1}^{(n,u)} = 1 \mid p_t^{(v_t)}, G_1, G_2] \\ &\geq (1 - \delta) \left(1 - N \exp\left\{-\frac{kh_c^2}{2}\right\}\right). \end{aligned}$$

The remaining proof follows analogously to the proof of Lemma 4.0.2. $\square$

As we have shown, the difference between φ_2 and $\hat{\varphi}_2$ is marginal, and consequently, it is natural to assume that Lemma 4.0.4 will hold by replacing φ_2 with $\hat{\varphi}_2$ at least for small values of δ .

Corollary 4.2.2. *Under condition (4.11), for any positive constant c , sufficiently large k , and $\delta < \mu$ we have*

$$\frac{1}{N} (1 - \exp\{-c\varphi_3^2\}) \leq 1 - \exp\{-c\hat{\varphi}_2^2\}. \quad (4.12)$$

Proof. We only have to prove that $\varphi_3 > \hat{\varphi}_2$ and that $\hat{\varphi}_2 > 0$, for a sufficiently large k . The first inequality is trivial. The second one will be satisfied if, and only if,

$$k > \frac{2 \log\left(\frac{(1-\delta)N}{\mu-\delta}\right)}{h_c^2} = \frac{2 \left(\log N + \log \frac{1-\delta}{\mu-\delta}\right)}{h_c^2}. \quad (4.13)$$

Notice that this last condition only holds for $\delta < \mu$. Now, all logical steps taken in subsection A.1.1, up to inequality (A.4) will hold by simply replacing φ_2 with $\hat{\varphi}_2$. It remains to show that

$$\varphi_3^2 - \hat{\varphi}_2^2 < (1 - \mu)^2. \quad (4.14)$$

For the remainder of this proof, abbreviate

$$x(k) := \exp \left\{ -\frac{kh_c^2}{2} \right\},$$

and notice that

$$\begin{aligned} \hat{\varphi}_2 &= \mu - 1(1 - \delta)(1 - Nx) \\ &= \mu - Nx + \delta Nx \\ &= \varphi_2 + \delta Nx. \end{aligned}$$

We now expand the left-hand side of (4.14)

$$\begin{aligned} \varphi_3^2 - \hat{\varphi}_2^2 &= \varphi_3^2 - [\varphi_2^2 - 2\varphi_2\delta Nx + \delta^2 N^2 x^2] \\ &= (\varphi_3^2 - \varphi_2^2) + 2\varphi_2\delta Nx - \delta^2 N^2 x^2. \end{aligned} \quad (4.15)$$

In the Appendix A.1.1 and Section 4.1, we already proved that for $k > 2(\log N - \log \mu)/h_c^2$, the term between parentheses in (4.14) is smaller than $(1 - \mu)^2$. Since (4.13) is a stronger condition on k , this implies that (4.13) will be satisfied if

$$\delta Nx(2\varphi_2 - \delta Nx) < 0,$$

or replacing the value of φ_2 and isolating δ , if

$$\delta < \frac{2\mu - 2Nx}{Nx}.$$

This completes the proof. Nevertheless, we compare the two conditions on δ found. In particular, we want the values of k for which

$$\mu < \frac{2\mu - 2Nx}{Nx}.$$

Replacing the value of $x(k)$ back into the right-hand side of the last inequality, we find that

$$\frac{2\mu - 2Nx}{Nx} = 2 \left(\frac{\mu - N \exp \left\{ -\frac{kh_c^2}{2} \right\}}{N \exp \left\{ -\frac{kh_c^2}{2} \right\}} \right) = \frac{2\mu}{N} \exp \left\{ \frac{kh_c^2}{2} \right\} - 2.$$

Comparing this result with μ and solving for k yields the following condition

$$k > \frac{2 \left(\log N + \log \frac{\mu + 2}{2\mu} \right)}{h_c^2}$$

which is a weaker condition than (4.13) for $\delta > 1/78$ (and stronger otherwise). □

Finally, as a corollary, we have the more general version of Theorem 4.0.1.

Corollary 4.2.3. *Under condition (4.11), given an arbitrary distribution of the initial opinions of the honest nodes it holds that*

$$\mathbb{P}[\mathcal{C} \cap \{\mathcal{R} < 2\ell\}] \geq \left[1 - (1 - W_c(n, k))^{\ell-1} \right] \left[1 - \exp \left\{ -\frac{k}{2}(\beta - q)^2 \right\} \right]^{(1-q)N\ell},$$

as long as either

$$k > \frac{2 \left(\log N + \log \frac{\mu + 2}{2\mu} \right)}{h_c^2}$$

and $\delta < 1/78$; or

$$k > \frac{2 \left(\log N + \log \frac{1 - \delta}{\mu - \delta} \right)}{h_c^2}$$

and $1/78 < \delta < \mu$.

4.2.1 A simple model for the discovering process

The previous corollary relies on the condition (4.11) that a sufficiently large portion of the conflict set is known by all nodes, and one can ask whether that is reasonable. To answer that, in this subsection, we explore a simple model for the arrival and discovery process for new transactions that can be used as a proxy for a real-life scenario. This model is independent of the rest of this work and the hypothesis considered here should not be carried any further.

We consider that at each round t a new set of transactions N_t is issued. For each transaction $u \in N_t$ and each node $n \in \mathcal{N}$ independently, suppose that the probability of n receiving u is given by r , or equivalently, following the notation introduced in the definition of our protocol,

$$\mathbb{P}[u \in N_t^{(n)}] = r.$$

Furthermore, we also consider that the network is sufficiently efficient so that a node can not be unaware of a transaction for two consecutive rounds. Explicitly, we consider that at round t , for any $n \in \mathcal{N}$ it holds that

$$\mathbb{T}_{t-1} \subset A_t^{(n)}. \tag{4.16}$$

In addition to receiving a new transaction u from the network, an honest node can include u in its vision if it queries a second node that is already aware of the transaction. Let $\mathcal{K}_t^{(n)}$ be the set of queried nodes by n at round t , then it follows that for any $u \in N_t$

$$\mathbb{P}[u \notin A_t^{(n)}] = \mathbb{P}[u \notin N_t^{(n)}] \mathbb{P}[u \notin N_t^{(j)} \text{ for all } j \in \mathcal{K}_t^{(n)}] = (1-r)^{k+1}, \quad (4.17)$$

or in other words, the probability of a node n finishing round t without knowing transaction u is equal to the probability of n not receiving u through the network times the probability of n querying only nodes that also have not received u .

Now, considering (4.16), we can rewrite the left hand side of condition (4.11) in the following way

$$\begin{aligned} \frac{\left| \bigcap_{n \in \mathcal{N}} A_t^{(n)} \right|}{|\mathbb{T}_t|} &= \frac{\left| \bigcap_{n \in \mathcal{N}} [(A_t^{(n)} \setminus \mathbb{T}_{t-1}) \cup \mathbb{T}_{t-1}] \right|}{|\mathbb{T}_t|} = \frac{\left| \left(\bigcap_{n \in \mathcal{N}} A_t^{(n)} \setminus \mathbb{T}_{t-1} \right) \cup \mathbb{T}_{t-1} \right|}{|\mathbb{T}_t|} \\ &= \frac{\left| \bigcap_{n \in \mathcal{N}} A_t^{(n)} \setminus \mathbb{T}_{t-1} \right| + |\mathbb{T}_{t-1}|}{|\mathbb{T}_t|} = \frac{\left| \bigcap_{n \in \mathcal{N}} A_t^{(n)} \setminus \mathbb{T}_{t-1} \right| + |\mathbb{T}_{t-1}|}{|N_t \cup \mathbb{T}_{t-1}|} \\ &= \frac{\left| \bigcap_{n \in \mathcal{N}} A_t^{(n)} \setminus \mathbb{T}_{t-1} \right| + |\mathbb{T}_{t-1}|}{|N_t| + |\mathbb{T}_{t-1}|} \end{aligned} \quad (4.18)$$

Again considering (4.16) it is natural to expect that condition (4.11) will be satisfied if $|N_t|$ is sufficiently small when compared to $|\mathbb{T}_{t-1}|$. In fact if

$$|N_t| < \frac{\delta}{1-\delta} |\mathbb{T}_{t-1}|$$

then

$$\frac{\left| \bigcap_{n \in \mathcal{N}} A_t^{(n)} \right|}{|\mathbb{T}_t|} = \frac{\left| \bigcap_{n \in \mathcal{N}} A_t^{(n)} \setminus \mathbb{T}_{t-1} \right| + |\mathbb{T}_{t-1}|}{|N_t| + |\mathbb{T}_{t-1}|} \geq \frac{\left| \bigcap_{n \in \mathcal{N}} A_t^{(n)} \setminus \mathbb{T}_{t-1} \right| + |\mathbb{T}_{t-1}|}{\delta(1-\delta)^{-1} |\mathbb{T}_{t-1}| + |\mathbb{T}_{t-1}|} \geq 1 - \delta. \quad (4.19)$$

While this is a good first approach, with equation (4.17) we can easily estimate the probability of the condition (4.11) given a set of new transactions of arbitrary size. By looking at equation (4.18) it is clear that in order to (4.11) to hold, it is sufficient that

$$\frac{\left| \bigcap_{n \in \mathcal{N}} A_t^{(n)} \setminus \mathbb{T}_{t-1} \right|}{|N_t|} \geq 1 - \delta,$$

or, in other words, that at most $\delta|N_t|$ new transactions are not discovered by at least one node at round t . Now if we abbreviate the event $H_t^{(n)} = \{u \notin A_t^{(n)}\}$, then the probability that

at least one node does not know u at the end of round t can be estimated using equation (4.17) and the union bound

$$\mathbb{P}[\bigcup_{n \in \mathcal{N}} H_t^{(n)}] \leq \sum_{n \in \mathcal{N}} \mathbb{P}[H_t^{(n)}] = (1 - q)|\mathcal{N}|(1 - r)^{k+1}$$

4.3 Numerical analysis

In this section, we delve into three of the most critical aspects of our protocol, exploring their influence on performance and resilience. First, we thoroughly examine how varying the proportion of malicious nodes, q , affects the overall performance of the system. Next, we shift our focus to the parameter k , investigating how incrementally increasing k can lead to significant improvements in outcomes, even in challenging or critical scenarios where performance is most at risk. Finally we explore the role of β —and consequently the random variable X_t —in the protocol.

To provide a comprehensive understanding, we conducted a series of simulations for each scenario, totaling 10000 simulations per configuration. This robust dataset allows us to explore the effects of the parameters in depth and with statistical reliability. The parameter settings used in these simulations are detailed in the table below, which outlines the key elements of our experimental setup, unless stated otherwise. Moreover, we impose a hard-cap on the number of rounds before consensus, i.e., if nodes are still undecided after 100 rounds, the simulation is considered as a *termination failure*.

Parameter		Value
N	Number of nodes	1000
T	Size of the conflict set	1000
q	Proportion of malicious nodes	0.25
β	Support of the random variable X_t	0.301
ℓ	Number of subsequent rounds before consensus	5
$p_t^{(v_t)}$	Largest proportion of likes between honest nodes	0.45

In these simulations, malicious nodes are modeled as though they are under the control of a single, all-knowing adversary—an omniscient entity that possesses complete awareness of the system’s state at every round. This adversary can strategically manipulate the malicious nodes with full knowledge of the protocol’s dynamics, simulating a worst-case scenario where the system is facing a highly intelligent and coordinated opponent.

Moreover, the strategy adopted is similar to the strategy introduced by Caposelle, Mueller and Penzkofer in [16]. Explicitly, the adversary waits until all honest nodes have exchanged opinions with one another. Then, the adversary attempts to split the honest nodes into two equally sized groups with opposing views, aiming to maximize the variance of their η -values.

To achieve this, the adversary must have continuous access to the η -values of the honest nodes. It then responds to undecided nodes in a way that increases the variance while keeping the median of the η -values near 0.5.




Figure 4.5 Effects of q on the performance of the protocol. Logarithmic scale.

This detailed analysis will shed light on how the interplay between q and k contributes to the robustness and efficiency of the protocol, offering valuable insights for optimizing performance under a wide range of conditions.

4.3.1 Analysing the effects of the proportion of malicious nodes

The histograms in Fig. 4.5 show different distributions for the number of rounds taken by the protocol as the parameter q varies from 0.10 to 0.30 (incremented by 0.05). For lower values of q (e.g., $q = 0.10$), the distribution is more concentrated around fewer rounds, with a sharp peak early on and rapid decay. As q increases, the distribution spreads out, becoming wider, with a more even spread over a greater number of rounds. For higher q values like 0.25 and 0.30, we see a significant number of termination failures (indicated by the bars on the far right). Nevertheless, we highlight the fact that every simulation that did not finalized in a termination failure was successful in achieving consensus.

As both our theoretical analysis and intuitive reasoning suggest, the efficiency of the protocol decreases as the value of q increases. This means that as the proportion of malicious nodes rises, the protocol encounters greater difficulties in reaching a consensus within a reasonable timeframe.

Although the malicious entity was unable to completely disrupt the consensus process—meaning it could not successfully cause the honest nodes to reach final, differing opinions—it was still able to significantly prolong the time required for consensus. Specifically, the adversary's actions led to a substantial delay in the number of rounds needed for the honest nodes to terminate the process.

4.3.2 Circumventing failures by increasing the number of queries

To investigate the effects of the number of queries, we simulate the critical scenario represented on the last histogram of Fig. 4.5, but for different values of k .

The histogram set bellow illustrates the distribution of the number of rounds for k ranging from 50 to 200. In each of the subplots, there is a noticeable spike at exactly 100 rounds,




Figure 4.6 Effects of k on the performance of the protocol. Logarithmic scale.

which we identified as a termination failure. This indicates that, for certain configurations, some nodes are unable to reach consensus, even after 100 rounds. The height of these spikes diminishes as k increases, which suggests that a larger k value reduces the likelihood of nodes experiencing termination failure.

For $k = 50$, a significant number of nodes fail to finalize within 100 rounds, as evidenced by the large spike at the 100-round mark. This suggests that smaller k values are more prone to failure. As k increases to 75, 100, 150, and finally 200, the spike at 100 rounds becomes less prominent. In fact, at $k = 200$, the number of nodes that fail to terminate within 100 rounds is considerably lower compared to $k = 50$. This trend suggests that increasing k contributes to the stability of the protocol by reducing termination failure and enabling more nodes to reach consensus in fewer rounds.

4.3.3 The importance of randomness

The proposed protocol is designed to depend on the inherent randomness of a global coin, represented by the random variable X_t , in two of its key components. This randomness plays a crucial role in ensuring the system's robustness and its ability to reach consensus among honest nodes, even in the presence of adversarial influence.

To demonstrate the significance of this randomness, in the following subsection, we explore what happens when it is removed from specific parts of the protocol. By doing so, we can observe how the protocol behaves in edge cases—scenarios where consensus is more difficult to achieve or completely fails. These edge cases serve to illustrate how essential the inclusion of randomness is to the protocol's ability to prevent consensus from being broken. This analysis will highlight the vulnerabilities that emerge when the protocol is deprived of this crucial random component.

FPCS without a random threshold

To better understand the impact of the random threshold on the protocol's performance, we conducted a series of simulations, varying the parameters β and q , focusing on a straightforward case of double-spending (i.e., a conflict set with only two transactions). As β




Figure 4.7 Finalization and agreement rate for different values of q and β .

approaches $1/2$, the support of the random variable X_t narrows, eventually reaching a point where it becomes fixed at $1/2$. In this situation, nodes will adopt the opinion of the simple majority of their queries, occasionally relying on the *elim* and *compl* mechanisms to solidify their decisions.

In Fig. 4.7, the left panel illustrates the *finalization rate*, showing how many simulations reached consensus before round 100. The right panel depicts the *agreement rate*, defined as the ratio between the largest group of nodes agreeing on a particular transaction and the total number of honest nodes. Our results indicate an optimal value for β around $1/3$, which aligns with the theoretical value derived in the previous section.

We also note that these findings closely resemble those reported by Capossele, Mueller, and Penzkofer. This similarity is expected, as resolving a double-spending scenario can be viewed as a binary decision problem. In this case, our protocol (FPCS) behaves similarly to the FPC protocol, further validating the theoretical insights drawn from both approaches.

FPCS without random ordering

Now we remove the effect of the random variable X_t from the order of transactions. Explicitly, in step (vi) of the protocol, instead of considering $\text{hash}(\text{Id}_x, X_t)$, in its calculations, the node should consider simply $\text{hash}(\text{Id}_x)$.

In this case, the goal of a malicious node is to use the fact that the order is known a priori to force honest nodes to finalize with different opinions. For that, consider the case where the conflict set is fixed (hence we omit the time dependency) and constituted only by transactions $u \in \mathbb{T}$ such that $p_t^{(u)} = 1/[2(1-q)]$ and $v \in \mathbb{T}$ such that $p_t^{(v)} = 1 - p_t^{(u)} = (1-2q)/[2(1-q)]$.

The strategy the malicious nodes will adopt is to reply with the same opinion to a query from every node, i.e. when queried by a node n that likes u or v at round t , the malicious will reply that it likes u or v respectively. The reason why this is a viable strategy will be explained in the next chapter.

Fig. 4.8 shows the results of 10000 simulations of this scenario for different values of q . On the left, we see the percentage of these simulations that resulted in an *agreement*




Figure 4.8 Agreement failures (in percentage) and agreement rate of a double spend $\mathbb{T} = \{u, v\}$ where malicious nodes reply with the same opinion to a query from every node.

failure – i.e. iterations that finalized without all nodes in consensus – and on the right, the boxplots for the *agreement rates* – i.e. the proportional size of the largest group that finalized in consensus with respect to the total number of honest nodes – of these failures. We see that, even though the attacker achieved an agreement failure in only 4% of the simulations in the worst scenario (when $q = 0.25$), these successful attacks had the first quartile of the agreement rate just above 0.7 and with extreme cases touching 0.5. In other words, it is not often that this strategy results in an agreement failure, but when it does, it causes a considerable split in the opinion of honest nodes.

Chapter 5

Generalization for arbitrary graphs

While complete graphs represent a typical form of attack, they do not constitute the most sophisticated strategy. In more intricate situations, malicious nodes extend their influence beyond mere voting, manipulating $\mathbb{T}$ itself—for instance, by introducing new conflicting transactions—to advance their objectives. Indeed, this section commences with an exploration of an edge case in which the threshold $q < \beta < 1/3$ proves insufficient to guarantee consensus.

In this chapter, we aim to generalize our results for arbitrary conflict graphs, which introduces certain constraints. We now assume that the conflict set is not only known by all nodes but also fixed, meaning no new transactions are issued. Consequently, we omit the subscript and refer to the conflict set simply as $\mathbb{T}$.

5.1 Star Graphs

A S_j star graph is a complete bipartite graph constituted by one internal vertex connected to a set of j external vertices, called leaves. Of course, the two only possible maximal independent sets in this graph are the set of leaves and the singleton of the interior vertex.

Assume that $\mathbb{T} = S_j$ for some integer $j \geq 2$ and designate u as the interior vertex. Consider also that $k = N$ or, in other words, that nodes will query every other node every round. Assume malicious nodes will adopt the following strategy: when queried by a node that likes u (resp. $N(u)$), the malicious will reply it also likes u (resp. $N(u)$). Moreover, consider that $p_t^{(u)} = 1/[2(1 - q)]$ and define as $\hat{p}_t^{(x)}$ and $\tilde{p}_t^{(x)}$ the overall proportion of likes (i.e. including the opinions of malicious nodes) that a transaction $x \in \mathbb{T}$ has at round t from the perspective of a node that likes u and $N(u)$ respectively.

Now consider the case $1/6 < q < \beta < 1/3$ where malicious nodes control a significant, though not critical, proportion of nodes. Then it can easily be verified that

$$\begin{aligned}\tilde{p}_t^{(N(u))} &= p_t^{(N(u))}(1 - q) < \beta; \\ \tilde{p}_t^{(u)} &= p_t^{(u)}(1 - q) + q > 1 - \beta; \\ \hat{p}_t^{N(u)} &= p_t^{(N(u))}(1 - q) + q = p_t^{(u)}(1 - q) = \hat{p}_t^{(u)} = 1/2.\end{aligned}$$

From the first two relations, we observe that regardless of the outcome of X_t , nodes that liked u at the beginning of the round will see no reason to change their mind since $\tilde{p}_t^{(u)} > 1 - \beta > X_t > \beta > \tilde{p}_t^{(N(u))}$. On the other hand, nodes that originally liked $N(u)$ will encounter a tie $\hat{p}_t^{(u)} = \hat{p}_t^{(N(u))} = 1/2$ and then for every $x \in \mathbb{T}$ either assign $\theta'(x) = 1$ if $X_t < 1/2$, or assign $\theta'(x) = 0$ if $X_t > 1/2$ (see Fig. 5.1). In both cases, step 7) of our protocol will pick between $\{u\}$ and $N(u)$, the set that contains the transaction with the smallest $\text{hash}(\text{Id}_x, X_t)$. Due to the uniformity property of the hash function, the smallest hash will be in $N(u)$ with probability $j/(j+1)$. This implies that malicious nodes can, with high probability, bypass the random component of the protocol and compel nodes that liked $N(u)$ initially to persist in liking $N(u)$. If this situation persists for ℓ rounds, opinions become final, and consensus is broken.




Figure 5.1 The interval of control (in blue) of a transaction $u \in \mathbb{T}_t$ with $p_t^{(u)} = 1/(2(1-q))$ and of a transaction $v \in N(u)$. Notice that the sum of the intervals covers the whole support of X_t .

This attack is only possible because the union of the intervals of control $\mathcal{I}_{q,t}^{(u)} \cup \mathcal{I}_{q,t}^{(N(u))}$ covers the whole support of X_t . An intuitive way to solve this is to decrease β (and consequently enlarge the support of X_t), but by doing that, since q must be smaller than β , we are also getting less resistant to byzantine actors.

To find a middle ground, we decrease β by a margin just enough to guarantee that the support of X_t is larger than $2q$ (two times the length of an interval of control). In other words, we want to maximize β subject to the constraints $2q < 1 - 2\beta$ and $q < \beta$. The result is the security threshold $q < \beta < 1/4$ and it is a standing assumption on all the following results.

A fundamental property of the system under this security threshold is that, under certain outcomes of X_t , nodes are not only likely to approve any transaction $u \in \mathbb{T}$ that has a sufficiently large proportion of likes but also, at the same time, disapprove any transaction in $N(u)$.

To find exactly how large this proportion of likes has to be, notice that if $u \in \mathbb{T}$ has $p_t^{(u)} = (1 - \beta - q)/(1 - q)$ then the upper boundary of $\mathcal{I}_{q,t}^{(u)}$ is $1 - \beta$ (which is equal to the upper boundary of the support of X_t); on the other hand, if $p_t^{(u)} = 1/[2(1 - q)]$, then the

lower boundary of $\mathcal{I}_{q,t}^{(u)}$ is $1/2$. We define p^* as the middle point between these two values:

$$p^* := \frac{1}{2} \left[\frac{1 - \beta - q}{1 - q} + \frac{1}{2(1 - q)} \right] = \frac{1}{2(1 - q)} + \frac{1 - 2\beta - 2q}{4(1 - q)}.$$

This way if $p_t^{(u)} < p^*$ or $p_t^{(u)} \geq p^*$ there will be a gap of at least size $h := (1/2 - \beta - q)/2$ between $\mathcal{I}_{q,t}^{(u)}$ and $1 - \beta$ and $1/2$, respectively. For clarity in presenting the results, it is advantageous to define $h_* := \min\{h, \beta - q\}$.

It is convenient to use p^* and $1 - \mu$ to partition the set of transactions. Thus, we denote

$$\mathcal{U}_t^- := \{u \in \mathbb{T} \text{ such that } p_t^{(u)} < p^*\}; \quad (5.1)$$

$$\mathcal{U}_t := \{u \in \mathbb{T} \text{ such that } p^* \geq p_t^{(u)} < 1 - \mu\}; \quad (5.2)$$

$$\mathcal{U}_t^+ := \{u \in \mathbb{T} \text{ such that } p_t^{(u)} \geq 1 - \mu\}. \quad (5.3)$$

Note that, since $p^* > 1/2 > \mu$, there are no conflicts between transactions in $\mathcal{U}_t$ and $\mathcal{U}_t^+$, otherwise a proportion at least $p^* - \mu > 0$ would have to like both transactions in a conflict. Moreover, the set $\mathcal{U}_t^+$ possesses two additional properties delineated in the following Lemma.

Lemma 5.1.1. *The set $\mathcal{U}_t^+$ is independent. Furthermore, under $\mathcal{PC}_t$, it is a MIS.*

Proof. For independence, suppose the contrary, namely, that there exist $u_i, u_j \in \mathcal{U}_t^+$ such that $u_i \leftrightarrow u_j$. Consequently, a proportion of at least $1 - \mu - 1/2 > 0$ nodes must like both u_i and u_j , which cannot be true.

Now assume $\mathcal{PC}_t$ is true and that $\mathcal{U}_t^+$ is not maximal, i.e., there is $v \in \mathbb{T} \setminus \mathcal{U}_t^+$ such that $v \leftrightarrow u$ for every $u \in \mathcal{U}_t^+$. But then $p_t^{(v)} < 1 - \mu$ since it is not in $\mathcal{U}_t^+$. There is also no $w \in N(u)$ such that $p_t^{(w)} > 1 - \mu$ since the only transactions that have a proportion larger than $1 - \mu$ are in $\mathcal{U}_t^+$ and by hypothesis we assumed they were not in conflict with v . In this case, $\mathcal{PC}_t$ must be false. $\square$

5.2 General case

Our main result relies on two key factors: firstly, in Lemma 5.2.1 we observe that pre-consensus on a set of transactions is maintained (with high probability) from one round to the next; and secondly, Lemmas 5.2.3, 5.2.4 and 5.2.6 show that on each round and with positive probability, nodes will converge to a pre-consensus about a significant (with respect to size) subset of the transactions they remain undecided about.

Lemma 5.2.1. *Suppose that $\mathcal{PC}_t(U)$ is true for a set $U \subset \mathbb{T}$, then for sufficiently large k and N , defining*

$$\psi_1 := \mu - T \exp \left\{ -\frac{k}{2} (\beta - q)^2 \right\}$$

it holds that

$$\mathbb{P}[\mathcal{PC}_{t+1}(U)] \geq 1 - \exp\{-2N(1 - q)\psi_1^2\}.$$

Proof. First, notice that for an honest node $n \in \mathcal{N}$ and a transaction $u \in U$, there is a gap of size $(\beta + q)/2$ between $\mathcal{I}_{q,t}^{(u)}$ and the support of X_t . Now consider the event $H_1^{(1)} := \{\eta_t^{(n,u)} k^{-1} > 1 - \beta\}$ (we use the superscript to denote Lemma 1, 2, etc.) that a proportion larger than $1 - \beta$ of the queried nodes like u . Hence

$$\mathbb{P}[H_1^{(1)} | \mathcal{PC}_t(U)] \geq \mathbb{P}[k^{-1} S_k \geq 1 - \beta],$$

where $S_k \sim \mathcal{B}(k, 1 - (\beta + q)/2)$. This later probability can be estimated using Hoeffding inequality and we find

$$\mathbb{P}[H_1^{(1)} | \mathcal{PC}_t(U)] \geq 1 - \exp\left\{-\frac{k}{2}(\beta - q)^2\right\}.$$

By symmetry, for $v \in N(U)$, considering the event $H_2^{(1)} := \{\eta_t^{(n,v)} k^{-1} < \beta\}$, it also holds that

$$\mathbb{P}[H_2^{(1)} | \mathcal{PC}_t(U)] \geq 1 - \exp\left\{-\frac{k}{2}(\beta - q)^2\right\}.$$

Then, if we denote by $H_3^{(1)}$ the event where an honest node assigns $\theta'(u) = 1$ to all transactions $u \in U$ and $\theta'(v) = 0$ to all transactions $v \in N(U)$, it follows that

$$\mathbb{P}[H_3^{(1)} | \mathcal{PC}_t(U)] \geq 1 - T \exp\left\{-\frac{k}{2}(\beta - q)^2\right\}. \quad (5.4)$$

Finally, the probability of $\mathcal{PC}_{t+1}(U)$ can be estimated by the probability of a random variable $S_{N(1-q)} \sim \mathcal{B}(N(1-q), 1 - T \exp(-(k/2)(\beta - q)^2))$ being larger than $N(1-q)(1 - \mu)$. Considering that N is sufficiently large so that the probability of success of this binomial is larger than $1 - \mu$, using again Hoeffding inequality, we find that

$$\mathbb{P}[\mathcal{PC}_{t+1}(U)] > 1 - \exp\{-2N(1-q)\psi_1^2\}.$$

□

It is worth noticing that neither the previous proof nor the estimate itself depends on the size of U ; instead, we relied on the size of the entire conflict set in the estimate (5.4). Consequently, this result remains valid even when $U = \mathbb{T}$ and, in this case, we also get the following corollary.

Corollary 5.2.2. *Suppose $\mathcal{PC}_t$ is true, then for a sufficiently large k and N , it holds that*

$$\mathbb{P}[\mathcal{PC}_{t+1}^*] \geq 1 - \exp\{-2N(1-q)\psi_1^2\}$$

Proof. Taking $U = \mathbb{T}$ in Lemma 5.2.1, $H_3^{(2)}$ becomes the event where an honest node assigns 1 to the auxiliary opinions of all transactions in $\mathcal{U}_t^+$ and 0 to the rest. As seen in Lemma 5.1.1, under $\mathcal{PC}_t$, $\mathcal{U}_t^+$ is a MIS, hence algorithms *elim* and *compl* will have no effect on the

auxiliar opinions, and the liked MIS will be $\mathcal{U}_t^+$ itself. The remaining of the proof follows analogously to Lemma 5.2.1. $\square$

Now, to show that on each round nodes will come to an agreement about a significant portion of the transactions they are still undecided about, we will distinguish between three different cases for the size of $\mathcal{U}_t \cup N(\mathcal{U}_t)$. For simplicity of notation, given a set $U \subset \mathbb{T}$, we define its closure as $\bar{U} := U \cup N(U)$.

The first and simplest case is when $|\bar{\mathcal{U}}_t| = 0$ or, in other words, when there are no clear favorite between the transactions. In this scenario, given a favorable outcome of X_t , nodes will assign 0 to the auxiliary opinion of all transactions with high probability, and then, using the algorithm *compl*, choose the same MIS to like.

Lemma 5.2.3. *If $|\bar{\mathcal{U}}_t| = 0$, then for sufficiently large k and N and defining*

$$\psi_2 := \mu - T \exp \left\{ -\frac{kh_*^2}{2} \right\}$$

it holds that

$$\mathbb{P}[\mathcal{P}\mathcal{C}_{t+1}^*] \geq \frac{h}{2(1-2\beta)} (1 - \exp\{-2N(1-q)\psi_2^2\}).$$

Proof. Let $H_1^{(3)}$ be the event where $X_t \in [1 - \beta - h/2, 1 - \beta]$. Then, for any $u \in \mathcal{U}_t^-$

$$\begin{aligned} \mathbb{P}[\theta'(u) = 1 | H_1^{(3)}] &\leq \mathbb{P} \left[S_k \geq \left(1 - \beta - \frac{h}{2} \right) k \right] \\ &\leq \exp \left\{ -\frac{kh^2}{2} \right\}, \end{aligned}$$

where $S_k \sim \mathcal{B}(k, p^*(1-q) + q)$. On the other hand, for every $v \in \mathcal{U}_t^+$ it holds that

$$\begin{aligned} \mathbb{P}[\theta'(v) = 0 | H_1^{(3)}] &\leq \mathbb{P} [S_k \leq (1 - \beta) k] \\ &\leq \exp \left\{ -\frac{k(\beta - q)^2}{2} \right\}, \end{aligned} \tag{5.5}$$

where in this case $S_k \sim \mathcal{B}(k, (1 - \mu)(1 - q))$. Now let W be the output of $\text{compl}(\mathcal{U}_t^+, \mathbb{T}, X_t)$ and define $H_2^{(3)}$ as the event where $\theta'(u) = 0$ for all $u \in \mathcal{U}_t^-$ and $\theta'(v) = 1$ for all $v \in \mathcal{U}_t^+$. Using Hoeffding again, we find

$$\begin{aligned} \mathbb{P}[\theta_{t+1}^{(n,W)} = 1 | H_1^{(3)}] &\geq \mathbb{P}[H_2^{(3)} | H_1^{(3)}] \\ &\geq 1 - T \exp \left\{ -\frac{kh_*^2}{2} \right\}. \end{aligned}$$

Then, for sufficiently large k and conditioning on $H_1^{(3)}$, it follows that

$$\begin{aligned}\mathbb{P}[\mathcal{PC}_{t+1}^*] &\geq \mathbb{P}[H_1^{(3)}] \mathbb{P}[p_{t+1}^{(W)} > 1 - \mu \mid H_1^{(3)}] \\ &\geq \frac{h}{2(1-\beta)} \mathbb{P}[S_{N(1-q)} > N(1-\mu)(1-q)] \\ &\geq \frac{h}{2(1-\beta)} (1 - \exp\{-2N(1-q)\psi_2^2\}),\end{aligned}$$

where $S_{N(1-q)} \sim \mathcal{B}(N(1-q), 1 - T \exp\{-kh_*^2/2\})$. $\square$

The second case, is when $|\tilde{\mathcal{U}}_t|$ contains a significantly large portion of transactions. Specifically, by significantly large, we mean it is at least the size of the square root of the number of transactions nodes are still undecided about. It is convenient to abbreviate

$$\gamma_t := \left\lceil \sqrt{|\mathbb{T} \setminus \mathcal{U}_t^+|} \right\rceil.$$

Lemma 5.2.4. *If $|\tilde{\mathcal{U}}_t| > \gamma_t$, then for sufficiently large k*

$$\mathbb{P}[\mathcal{PC}_{t+1}(\tilde{\mathcal{U}}_t)] \geq \frac{h}{2(1-2\beta)} (1 - \exp\{-2N(1-q)\psi_2^2\}).$$

Proof. Let $H^{(4)}$ be the event $X_t \in [(1-h)/2, (1+h)/2]$. Then for a transaction $u \in \mathcal{U}_t$

$$\mathbb{P}[\theta'(u) = 0 \mid H^{(4)}] \leq \mathbb{P}[S_k \leq k(1+h)/2] \leq \exp\left\{-\frac{kh^2}{2}\right\},$$

where $S_k \sim \mathcal{B}(k, p^*(1-q))$. On the other hand, for $v \in N(\mathcal{U}_t)$

$$\mathbb{P}[\theta'(v) = 1 \mid H^{(4)}] \leq \mathbb{P}[S_k \geq k(1-h)/2] \leq \exp\left\{-\frac{kh^2}{2}\right\},$$

where in this case $S_k \sim \mathcal{B}(k, (1-p^*)(1-q) + q)$. Moreover, the estimate (5.5) is still valid for any transaction in $\mathcal{U}_t^+$. Then, for an honest node n

$$\mathbb{P}[\theta_{t+1}^{(n, \mathcal{U}_t)} = 1, \theta_{t+1}^{(n, N(\mathcal{U}_t))} = 0 \mid H^{(4)}] \geq 1 - T \exp\{-kh_*^2/2\},$$

and the remaining of the proof follows analogously to the previous Lemma. $\square$

The result for the final case, when $0 < |\tilde{\mathcal{U}}_t| < \gamma_t$, will depend not only on a favorable outcome of X_t but also on a favorable ordering (with respect to the order induced by the *hash* function and X_t) of certain transactions, hence we will rely on the diffusion property again. We need an additional Lemma to find a bound for the probability of this favorable order. The proof can be found in Appendix A.2.1.

Lemma 5.2.5. *Let $Y_1, Y_2, \dots, Y_{M-m}$ and $X_1, X_2, \dots, X_m$, with $m < M$, be samples of i.i.d. random variables following an uniform distribution in the interval $[0, 1]$. Then for a integer $1 \leq c \leq M - m$,*

it holds that

$$\mathbb{P}[Y_{(c)} < X_{(1)}] > \left(\frac{M - m - c + 1}{M - m + 1} \right)^m,$$

where $Y_{(c)}$ and $X_{(1)}$ correspond respectively to the c -th and 1-st order statistic of the samples.

Lemma 5.2.6. *If $0 < |\bar{\mathcal{U}}_t| < \gamma_t$, then for sufficient large k and some random independent set $W \in \mathbb{T}$ such that $|\bar{W}| \geq \gamma_t$, it holds that*

$$\mathbb{P}[\mathcal{PC}_{t+1}(\bar{W})] \geq \frac{h}{324(1-2\beta)}(1 - \exp\{-2N(1-q)\psi_2^2\}).$$

Proof. Let $H_1^{(7)}$ be the event $X_t \in [1 - \beta - h/2, 1 - \beta]$. For any $v \in \mathcal{U}_t^-$ we have that

$$\mathbb{P}[\theta'(v) = 1 | H_1^{(7)}] \leq \mathbb{P}[S_k \geq (1 - \beta - h/2)k] \leq \exp\left\{-\frac{kh^2}{2}\right\},$$

where $S_k \sim \mathcal{B}(k, p^*(1-q) + q)$. It follows that

$$\mathbb{P}[\theta'(v) = 0 \text{ for all } v \in \mathcal{U}_t^- | H_1^{(7)}] \geq 1 - T \exp\left\{-\frac{kh^2}{2}\right\}.$$

Now define V as the set of the first¹ γ_t transactions in $\mathbb{T} \setminus \{\bar{\mathcal{U}}_t^+ \cup \bar{\mathcal{U}}_t\}$. Consider the event $H_2^{(6)}$ where V comes before the first transaction in $\bar{\mathcal{U}}_t$. From Lemma 5.2.5, we find that

$$\mathbb{P}[H_2^{(7)}] > \left(\frac{M_t - |\bar{\mathcal{U}}_t| - \gamma_t + 1}{M_t - |\bar{\mathcal{U}}_t| + 1} \right)^{|\bar{\mathcal{U}}_t|} \quad (5.6)$$

$$\geq \left(\frac{M_t - 2\gamma_t + 1}{M_t - \gamma_t + 1} \right)^{\gamma_t}, \quad (5.7)$$

where $M_t := |T \setminus \bar{\mathcal{U}}_t^+|$. But since $\gamma_t = \lceil \sqrt{M_t} \rceil$, then $(\gamma_t - 1)^2 < M_t \leq \gamma_t^2$ and it follows from (5.6) that

$$\mathbb{P}[H_2^{(7)}] > \left(\frac{(\gamma_t - 1)^2 - 2\gamma_t + 1}{(\gamma_t - 1)^2 - \gamma_t + 1} \right)^{\gamma_t}.$$

The function above, for all $\gamma_t \geq 4$ is always larger or equal than $\left(\frac{1}{3}\right)^4 = \frac{1}{81}$. Now let's study the cases where the bound above is not applicable.

- When $\gamma_t = 1$ in this case, $\bar{\mathcal{U}}_t = \emptyset$.
- When $\gamma_t = 2$ in this case, $|\bar{\mathcal{U}}_t| \leq 1$ and $M_t \in [3, 4]$. From Lemma 5.2.5, we find that, in this case, $\mathbb{P}[H_2^{(7)}] > \frac{1}{3}$.

¹According to the order induced by X_t and the hash function.

- When $\gamma_t = 3$ in this case, $|\vec{\mathcal{U}}_t| \leq 2$ and $M_t \in [5, 9]$. From Lemma 5.2.5, we find that, in this case, $\mathbb{P}[H_2^{(7)}] > \left(\frac{1}{4}\right)^2$.

Combining those cases, we find that, in general, $\mathbb{P}[H_2^{(7)}] > \left(\frac{1}{3}\right)^4 = \frac{1}{81}$. Define W as the output of $\text{compl}(\emptyset, V, X_t)$. Notice that no transaction in $(\mathcal{U}_t \cup N(\mathcal{U}_t))^C$ is in conflict with $\mathcal{U}_t$. Moreover, for an honest node n , if $\theta'(v) = 0$ for all $v \in N(\mathcal{U}_t)$, then W must be in the MIS liked by n , independently of the node's auxiliary opinion about transactions in $\mathcal{U}_t$. This means that

$$\begin{aligned} \mathbb{P}[\theta_{t+1}^{(n,W)} = 1 | H_1^{(7)}, H_2^{(7)}] \\ &= \mathbb{P}[\theta'(v) = 0 \text{ for all } v \in \mathcal{U}_t^C | H_1^{(7)}, H_2^{(7)}] \\ &\geq 1 - T \exp \left\{ -\frac{kh^2}{2} \right\} \end{aligned}$$

Then, for a sufficiently large k , using the same reasoning as in previous Lemmas and conditioning on $H_1^{(7)}$ and $H_2^{(7)}$

$$\begin{aligned} \mathbb{P}[\mathcal{PC}_{t+1}(W)] &\geq \mathbb{P}[H_1^{(7)}] \mathbb{P}[H_2^{(7)}] (1 - \exp - 2N(1-q)\psi_2^2) \\ &\geq \frac{h}{324(1-2\beta)} (1 - \exp\{-2N(1-q)\psi_2^2\}). \end{aligned}$$

□

From Lemmas 5.2.3, 5.2.4 and 5.2.6 we conclude that nodes will agree on a significant portion of transactions they are still undecided about with probability bounded by the minimum of the estimates given by the Lemmas. Or explicitly, at least with probability $W(k, N)$, defined as

$$W(k, N) := \frac{h}{324(1-2\beta)} (1 - \exp\{-2N(1-q)\psi_2^2\}).$$

Given $W(k, N)$ and a set of transactions, in the next Lemma, we model this dynamics of agreeing on a significant portion of transactions per round to find an estimate for the probability of getting to pre-consensus before a specific round. The proof can be found in Appendix A.2.2.

Lemma 5.2.7. *Given a set $V = \{v_1, \dots, v_T\}$ and a function $l : V \times \mathbb{N} \rightarrow \{0, 1\}$ define $V_t^i := \{v \in V : l(v, t) = i\}$ for $i = 0, 1$. Suppose l satisfies:*

1. *Let B_t be the event where $|V_t^1| - |V_{t-1}^1| \geq \sqrt{|V_{t-1}^0|}$. For any $t > 0$, $\mathbb{P}[B_t] \geq W$.*
2. *The events B_t are independent for all $t \in [0, \infty)$.*

Let S be the random variable representing the first time such that $V_S^0 = 0$ and call $t_0 = 2\sqrt{T}$. Then for $t \geq 0$

$$\mathbb{P}[S \leq t_0 + t] \geq 1 - \left(1 - \frac{W^{t_0}}{t_0}\right)^t.$$

Defining $\mathcal{C}$ as the event where honest nodes achieve consensus, our main result states that with high probability (depending on k), for any distribution of initial opinions, a pre-consensus state is achieved and maintained for enough consecutive rounds so that event $\mathcal{C}$ happens in the round $\mathcal{R} \leq 2\ell$.

Theorem 5.2.8. *For sufficiently large k , given an arbitrary distribution of the initial opinions of the honest nodes it holds that*

$$\begin{aligned} \mathbb{P}[\mathcal{C} \cap \{\mathcal{R} < 2\ell\}] &\geq \left(1 - \exp\{-2N(1-q)\psi_1^2\}\right) \\ &\quad \left(1 - (1-W)^{t_0+t-1}\right) \\ &\quad \left[1 - T \exp\left\{-\frac{k}{2}(\beta-q)^2\right\}\right]^{(1-q)N\ell} \end{aligned}$$

Proof. The proof is very similar to the complete case. For the sake of completeness, we repeat the arguments.

First, we define the random variables

$$\Psi := \min\{r \geq 1 : \mathcal{PC}^* \text{ is true}\},$$

that is the first round in which the system is in the hard pre-consensus state, and

$$\hat{\tau}_n := \min\{r \geq \Psi + \ell : \theta_r^{(n,U)} = \dots = \theta_{r-\ell+1}^{(n,U)} = 1\},$$

that is the first time in which a maximal independent set U is liked by node n for ℓ subsequent rounds after round Ψ . We then define the events

$$\begin{aligned} D_t^{(1)} &= \{\Psi = t\}, \quad D_{[t,s]}^{(1)} = \{\Psi \in [t,s]\}, \\ D^{(2)} &= \{\hat{\tau}_n = \Psi + \ell; \forall n = 1, \dots, (1-q)N\}. \end{aligned}$$

Notice that if $t < \ell$, then

$$\mathcal{C} \cap \{\mathcal{R} \leq t + \ell\} \supset D_{[1,t]}^{(1)} \cap D^{(2)} \supset D_{[2,t]}^{(1)} \cap D^{(2)}.$$

This way

$$\mathbb{P}[\mathcal{C} \cap \{\mathcal{R} < 2\ell\}] \geq \mathbb{P}\left[D_{[2,\ell-1]}^{(1)} \cap D^{(2)}\right] = \mathbb{P}\left[\left\{\bigcup_{i=2}^{\ell-1} D_i^{(1)}\right\} \cap D^{(2)}\right],$$

since the events $D_i^{(1)}$ are disjoint and the system is time-invariant, it follows that

$$\begin{aligned} \mathbb{P}[\mathcal{C} \cap \{\mathcal{R} < 2\ell\}] &\geq \mathbb{P}\left[\bigcup_{i=2}^{\ell-1} D_i^{(1)} \cap D^{(2)}\right] = \sum_{i=2}^{\ell-1} \mathbb{P}[D_i^{(1)} \cap D^{(2)}] \\ &= \sum_{i=2}^{\ell-1} \mathbb{P}[D^{(2)} | D_i^{(1)}] \mathbb{P}[D_i^{(1)}] = \sum_{i=2}^{\ell-1} \mathbb{P}[D^{(2)}] \mathbb{P}[D_i^{(1)}] \\ &= \mathbb{P}[D^{(2)}] \mathbb{P}[D_{[1, \ell-1]}^{(1)}] \end{aligned}$$

To calculate $\mathbb{P}[D^{(2)}]$ we write

$$\begin{aligned} \mathbb{P}[D^{(2)}] &= \mathbb{P}\left[\{\hat{\tau}_n = \Psi + \ell; \forall n = 1, \dots, (1-q)N\}\right] \\ &= \mathbb{P}\left[\{\theta_{\Psi+1}^{(n,U)} = \dots = \theta_{\Psi+\ell}^{(n,U)} = 1; \forall n = 1, \dots, (1-q)N\}\right] \\ &= \mathbb{P}\left[\bigcap_{i=1}^{\ell} \bigcap_{n=1}^{(1-q)N} \{\theta_{\Psi+i}^{(n,U)} = 1\}\right], \end{aligned}$$

To simplify the notation, define $D_i^{(3)} = \bigcap_{n=1}^{(1-q)N} \{\theta_{\Psi+i}^{(n,U)} = 1\}$, then

$$\mathbb{P}[D^{(2)}] = \mathbb{P}\left[\bigcap_{i=1}^{\ell} D_i^{(3)}\right] = \mathbb{P}[D_{\ell}^{(3)} | D_1^{(3)} \dots D_{\ell-1}^{(3)}] \dots \mathbb{P}[D_2^{(3)} | D_1^{(3)}] \mathbb{P}[D_1^{(3)}]. \quad (5.8)$$

But notice that if $D_i^{(3)}$ is true, then $\mathcal{PC}_{i+1}^*$ is also true, and using a similar bound to the one used in Lemma 5.2.1 we find that

$$\mathbb{P}[D^{(2)}] \geq \left[1 - T \exp\left\{-\frac{k}{2}(\beta - q)^2\right\}\right]^{(1-q)N\ell}.$$

Now, define the random variable

$$\hat{\Psi} := \min\{r \geq 1 : \mathcal{PC} \text{ is true}\},$$

that is the first round in which the system is in the pre-consensus state, and the event

$$\hat{D}_t^{(1)} = \{\hat{\Psi} = t\}, \quad \hat{D}_{[t,s]}^{(1)} = \{\hat{\Psi} \in [t, s]\},$$

Then

$$\begin{aligned} \mathbb{P}[D_{[2, \ell-1]}^{(1)}] &= \sum_{i=2}^{\ell} \mathbb{P}[\Psi = i] \geq \sum_{i=2}^{\ell} \mathbb{P}[\{\Psi = \hat{\Psi} + 1\} \cap \{\hat{\Psi} = i - 1\}] \\ &= \sum_{i=2}^{\ell} \mathbb{P}[\Psi = \hat{\Psi} + 1 | \hat{\Psi} = i - 1] \mathbb{P}[\hat{\Psi} = i - 1]. \end{aligned}$$

Again, since the system is time-invariant

$$\begin{aligned}\mathbb{P}\left[D_{[2,\ell-1]}^{(1)}\right] &\geq \sum_{i=2}^{\ell} \mathbb{P}[\Psi = \hat{\Psi} + 1] \mathbb{P}[\hat{\Psi} = i - 1] = \mathbb{P}[\Psi = \hat{\Psi} + 1] \sum_{i=2}^{\ell} \mathbb{P}[\hat{\Psi} = i - 1] \\ &= \mathbb{P}[\Psi = \hat{\Psi} + 1] \hat{D}_{[1,\ell-1]}^{(1)}.\end{aligned}$$

By Corollary 5.2.2 and Lemma 5.2.7,

$$\mathbb{P}\left[D_{[2,\ell-1]}^{(1)}\right] \geq \left(1 - \exp\{-2N(1-q)\psi_1^2\}\right) \left(1 - \left(1 - \frac{W^{t_0}}{t_0}\right)^t\right).$$

We finally conclude that

$$\begin{aligned}\mathbb{P}[\mathcal{C} \cap \{\mathcal{R} < 2\ell\}] &\geq \left(1 - \exp\{-2N(1-q)\psi_1^2\}\right) \\ &\quad \left(1 - \left(1 - \frac{W^{t_0}}{t_0}\right)^t\right) \\ &\quad \left[1 - T \exp\left\{-\frac{k}{2}(\beta - q)^2\right\}\right]^{(1-q)N\ell}\end{aligned}$$

□

Chapter 6

Conclusion

This paper presents the Fast Probabilistic Consensus on a Set (FPCS), a new probabilistic consensus protocol designed to address the challenges of distributed consensus in failure-prone systems, particularly in the context of UTXO-based Distributed Ledger Technologies. FPCS stands out for its ability to achieve consensus on a Maximum Independent Set within a graph of conflicting transactions, a critical issue in UTXO-based DLTs where transactions sharing the same output are considered to be in conflict.

Building on fundamental work in the field of distributed consensus, such as Ben'Or's and, more recently, FPC-BI, FPCS leverages the power of probabilistic methods to overcome the limitations of deterministic approaches, particularly in asynchronous systems where the impossibility of consensus with even a single process failure has been demonstrated. The use of randomness, through a "global coin" and a cryptographic hash function, enables FPCS to reach consensus with high probability, even in the presence of Byzantine nodes that may behave arbitrarily or maliciously.

The effectiveness of FPCS is demonstrated through rigorous theoretical analysis and numerical simulations, initially focusing on the n -spend problem, a common scenario in UTXO-based DLTs where all transactions in a conflict set spend at least one common UTXO. The paper establishes a security threshold ($q < \beta < 1/3$) under which FPCS ensures consensus with high probability, provided that the number of queries k is sufficiently large. Furthermore, the paper explores the impact of the number of queries k and the proportion of malicious nodes q on the protocol's performance, highlighting the importance of k in mitigating the effect of increasing q on the protocol's efficiency.

Beyond complete graphs, the paper generalizes the results to arbitrary conflict graphs, capturing more complex attack scenarios where the graph's structure itself can be manipulated by malicious nodes. An adjusted security threshold ($q < \beta < 1/4$) is introduced to guarantee consensus in such scenarios, and the paper proves that, with high probability, FPCS will reach consensus on a single MIS in an arbitrary conflict graph, provided that k is sufficiently large.

In summary, FPCS offers a promising solution to the distributed consensus problem, particularly in scenarios where robustness against Byzantine nodes and the ability to handle complex sets of conflicting transactions are crucial. Its probabilistic nature allows it to cir-

circumvent the impossibility results inherent to deterministic approaches, and its effectiveness is validated through theoretical analysis and simulations.

Bibliography

- [1] Abraham, I., Gueta, G., Malkhi, D., Alvisi, L., Kotla, R., and Martin, J.-P. (2017). Revisiting fast practical byzantine fault tolerance.
- [2] Aguilera, M. and Toueg, S. (2012). Correctness proof of ben-or’s randomized consensus algorithm. *Distributed Computing*, 25.
- [3] Alighanbari, M. and How, J. (2006). An unbiased kalman consensus algorithm. volume 5, page 6 pp.
- [4] Amin, M. and Draief, M. (2015). Global majority consensus by local majority polling on graphs of a given degree sequence. *Discrete Applied Mathematics*, 180.
- [5] Aumann, Y. and Lindell, Y. (2007). Security against covert adversaries: Efficient protocols for realistic adversaries. *Journal of Cryptology*, 23:281–343.
- [6] Barborak, M. and Malek, M. (1993). The consensus problem in fault-tolerant computing. *ACM Comput. Surv.*, 25:171–220.
- [7] Becchetti, L., Clementi, A., Natale, E., Pasquale, F., and Trevisan, L. (2015). Stabilizing consensus with many opinions.
- [8] Belotti, M., Bozic, N., Pujolle, G., and Secci, S. (2019). A vademecum on blockchain technologies: When, which and how. *IEEE Communications Surveys and Tutorials*, PP:1–1.
- [9] Ben-Or, M. (1983). Another advantage of free choice: Completely asynchronous agreement protocols (extended abstract). pages 27–30.
- [10] Benjamini, I., Chan, S.-O., O’Donnell, R., Tamuz, O., and Tan, L.-Y. (2014). Convergence, unanimity and disagreement in majority dynamics on unimodular graphs and random graphs. *Stochastic Processes and their Applications*, 126.
- [11] Bertrand, N., Gramoli, V., Konnov, I., Lazic, M., Tholoniati, P., and Widder, J. (2022). Brief announcement: Holistic verification of blockchain consensus. In *Proceedings of the 2022 ACM Symposium on Principles of Distributed Computing, PODC’22*, page 424–426, New York, NY, USA. Association for Computing Machinery.
- [12] Bracha, G. (1987). Asynchronous byzantine agreement protocols. *Inf. Comput.*, 75:130–143.
- [13] Buterin, V. and Griffith, V. (2017). Casper the friendly finality gadget.
- [14] Cachin, C., Kursawe, K., and Shoup, V. (2000). Random oracles in constantinople: Practical asynchronous byzantine agreement using cryptography. *Journal of Cryptology*, 18.
- [15] Canetti, R. and Rabin, T. (1993). Fast asynchronous byzantine agreement with optimal resilience. In *Proceedings of the 25th Annual ACM Symposium on Theory of Computing*.

- [16] Caposelle, A., Müller, S., and Penzkofer, A. (2021). Robustness and efficiency of voting consensus protocols within byzantine infrastructures. *Blockchain: Research and Applications*, 2(1):100007.
- [17] Castro, M. and Liskov, B. (1999). Practical byzantine fault tolerance. In *Proceedings of the Third Symposium on Operating Systems Design and Implementation*, OSDI '99, page 173–186, USA. USENIX Association.
- [18] Cooper, C., Elsässer, R., and Radzik, T. (2014). The power of two choices in distributed voting.
- [19] Cox, J. T., Durrett, R., and Perkins, E. A. (2000). Rescaled voter models converge to super-brownian motion. *Annals of Probability*.
- [20] Crain, T., Gramoli, V., Larrea, M., and Raynal, M. (2018). Dbft: Efficient leaderless byzantine consensus and its application to blockchains. pages 1–8.
- [21] Cruise, J. and Ganesh, A. (2013). Probabilistic consensus via polling and majority rules. *Queueing Systems*, 78.
- [22] Doerr, B., Goldberg, L., Minder, L., Sauerwald, T., and Scheideler, C. (2011). Stabilizing consensus with the power of two choices. *Annual ACM Symposium on Parallelism in Algorithms and Architectures*.
- [23] Elsässer, R., Friedetzky, T., Kaaser, D., Mallmann-Trenn, F., and Trinker, H. (2017). Brief announcement: Rapid asynchronous plurality consensus. pages 363–365.
- [24] Feldman, P. and Micali, S. (1997). An optimal probabilistic protocol for synchronous byzantine agreement. *SIAM J. Comput.*, 26:873–933.
- [25] Fischer, M. (2000). The consensus problem in unreliable distributed systems (a brief survey).
- [26] Fischer, M., Lynch, N., and Paterson, M. (1985). Impossibility of distributed consensus with one faulty process. *J. ACM*, 32:374–382.
- [27] Friedman, R., Mostéfaoui, A., and Raynal, M. (2004). Simple and efficient oracle-based consensus protocols for asynchronous byzantine systems. pages 228–237.
- [28] Gao, S., Zhan, B., Wu, Z., and Zhang, L. (2024). Verifying randomized consensus protocols with common coins. In *2024 54th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, volume abs/1712.01367, page 403–415. IEEE.
- [29] Gutierrez, A., Müller, S., and Šebek, S. (2023). On asymptotic fairness in voting with greedy sampling. *Advances in Applied Probability*, 55:1–34.
- [30] Gärtner, B. and Zehmakan, A. (2018). *Majority Model on Random Regular Graphs*, pages 572–583.
- [31] Kar, S. and Moura, J. (2007). Distributed average consensus in sensor networks with random link failures and communication channel noise. pages 676 – 680.
- [32] Konnov, I., Lazic, M., Veith, H., and Widder, J. (2017). A short counterexample property for safety and liveness verification of fault-tolerant distributed algorithms. *ACM SIGPLAN Notices*, 52:719–734.
- [33] Kwiatkowska, M. and Norman, G. (2002). Verifying randomized byzantine agreement. In Peled, D. A. and Vardi, M. Y., editors, *Formal Techniques for Networked and Distributed Systems — FORTE 2002*, pages 194–209, Berlin, Heidelberg. Springer Berlin Heidelberg.

- [34] Lamport, L., Shostak, R., and Pease, M. (2002). The byzantine generals problem. *ACM Trans. Program. Lang. Syst.*, 4.
- [35] Lin, B.-Y., Dziubałtowska, D., Macek, P., Penzkofer, A., and Müller, S. (2023). Robustness of the tangle 2.0 consensus. In Hyytiä, E. and Kavitha, V., editors, *Performance Evaluation Methodologies and Tools*, pages 259–276, Cham. Springer Nature Switzerland.
- [36] Merkle, R. (1989). One way hash functions and des. pages 428–446.
- [37] Miller, A., Xia, Y., Croman, K., Shi, E., and Song, D. (2016). The honey badger of bft protocols. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, CCS '16*, page 31–42, New York, NY, USA. Association for Computing Machinery.
- [38] Mossel, E., Neeman, J., and Tamuz, O. (2012). Majority dynamics and aggregation of information in social networks. *Autonomous Agents and Multi-Agent Systems*, 28.
- [39] Muselli, M. (2000). New improved bounds for reliability of consecutive-k-out-of-n:f systems. *Journal of Applied Probability*, 37.
- [40] Müller, S., Penzkofer, A., Camargo, D., and Saa, O. (2021a). *On Fairness in Voting Consensus Protocols*, pages 927–939.
- [41] Müller, S., Penzkofer, A., Kuśmierz, B., Camargo, D., and Buchanan, W. (2021b). *Fast Probabilistic Consensus with Weighted Votes*, pages 360–378.
- [42] Müller, S., Penzkofer, A., Polyanskii, N., Theis, J., Sanders, W., and Moog, H. (2022). Tangle 2.0 leaderless nakamoto consensus on the heaviest dag. *IEEE Access*, 10:105807–105842.
- [43] Nakamoto, S. (2009). Bitcoin: A peer-to-peer electronic cash system. *Cryptography Mailing list at <https://metzdowd.com>*.
- [44] Nitchai, R., Popov, S., and Müller, S. (2023). FPCS: Solving n-spends on a utxo-based dlt. *Blockchain and Applications, 5th International Congress. BLOCKCHAIN 2023. Lecture Notes in Networks and Systems*.
- [45] Nitchai, R., Popov, S., Müller, S., and Saa, O. (2024). Security threshold for fpcs on star graphs.
- [46] Pease, M., Shostak, R., and Lamport, L. (1979). Reaching agreement in the presence of faults. *Journal of the ACM (JACM)*, 27:228–234.
- [47] Penzkofer, A., Müller, S., and Capossele, A. (2019). Robustness and efficiency of leaderless probabilistic consensus protocols within byzantine infrastructures.
- [48] Popov, S. (2015). The tangle.
- [49] Popov, S. (2017). On a decentralized trustless pseudo-random number generation algorithm. *Journal of Mathematical Cryptology*, 11.
- [50] Popov, S. and Buchanan, W. (2020). FPC-BI: Fast probabilistic consensus within byzantine infrastructures. *Journal of Parallel and Distributed Computing*, 147:77–86.
- [51] Popov, S. and Müller, S. (2021). Voting-based probabilistic consensus and their applications in distributed ledgers. *Annals of Telecommunications*, 77.
- [52] Popov, S., Saa, O., and Finardi, P. (2019). Equilibria in the tangle. *Computers and Industrial Engineering*, 136:160–172.

- [53] Rabin, M. (1983). Randomized byzantine generals. pages 403–409.
- [54] Rocket, T. (2018). Snowflake to avalanche : A novel metastable consensus protocol family for cryptocurrencies.
- [55] Rocket, T., Yin, M., Sekniqi, K., Van Renesse, R., and Sirer, E. (2019). Scalable and probabilistic leaderless bft consensus through metastability.
- [56] Tanaka-Yamawaki, M., Kitamikado, S., and Fukuda, T. (1996). Consensus formation and the cellular automata. *Robotics and Autonomous Systems*, 19:15–22.
- [57] Wensley, J., Lamport, L., Goldberg, J., Green, M., Levitt, K., Shostak, R., and Weinstock, C. (1978). Sift: Design and analysis of a fault-tolerant computer for aircraft control. *Proceedings of the IEEE*, 66:1240 – 1255.
- [58] Zhang, G., Pan, F., Mao, Y., Tijanic, S., Dang’ana, M., Motepalli, S., Zhang, S., and Jacobsen, H.-A. (2024). Reaching consensus in the byzantine empire: A comprehensive review of bft consensus algorithms. *ACM Comput. Surv.*, 56(5).
- [59] Zhang, Y., Setty, S., Chen, Q., Zhou, L., and Alvisi, L. (2020). Byzantine ordered consensus without byzantine oligarchy. In *Proceedings of the 14th USENIX Conference on Operating Systems Design and Implementation, OSDI’20*, USA. USENIX Association.

Appendix A

Appendix

A.1 Proofs of Chapter 4

A.1.1 Lemma 4.0.4

For any positive constant c and $k > 2(\log N - \log \mu)/h_c^2$, it holds that

$$\frac{1}{N}(1 - \exp\{-c\varphi_3^2\}) \leq 1 - \exp\{-c\varphi_2^2\} \quad (\text{A.1})$$

Proof. We will rewrite this inequality in terms of $\varphi_3^2 - \varphi_2^2$. Since $\varphi_3 > \varphi_2$, it is clear that $\varphi_3^2 - \varphi_2^2 \geq 0$ if both φ_3 and φ_2 are nonnegative. Now $\varphi_3 \geq 0$ if, and only if,

$$k > -\frac{2\log \mu}{h_c^2},$$

while $\varphi_2 \geq 0$ if, and only if,

$$k > \frac{2(\log N - \log \mu)}{h_c^2} = -\frac{2\log(\mu/N)}{h_c^2}. \quad (\text{A.2})$$

Then, assuming the stronger condition (A.2) holds, we can rewrite the left-hand-side of (A.1) as

$$\begin{aligned} \frac{1}{N}(1 - \exp\{-c\varphi_3^2\}) &= \frac{1 - \exp\{-c(\varphi_3^2 - \varphi_2^2)\}}{N} \\ &\quad + \frac{\exp\{-c(\varphi_3^2 - \varphi_2^2)\}}{N}(1 - \exp\{-c\varphi_2^2\}), \end{aligned}$$

which will be smaller or equal than $1 - \exp\{-c\varphi_2^2\}$ if, and only if,

$$\frac{1 - \exp\{-c(\varphi_3^2 - \varphi_2^2)\}}{1 - \exp\{-c\varphi_2^2\}} + \exp\{-c(\varphi_3^2 - \varphi_2^2)\} \leq N. \quad (\text{A.3})$$

Since $N \geq 2$, then (A.3) will hold as long as

$$\frac{1 - \exp\{-c(\varphi_3^2 - \varphi_2^2)\}}{1 - \exp\{-c\varphi_2^2\}} < 1$$

which is equivalent to the following condition to hold

$$\varphi_3^2 - \varphi_2^2 < \varphi_2^2.$$

Since φ_2 is larger than $\mu - 1$, this last condition will hold as long as

$$\varphi_3^2 - \varphi_2^2 < (1 - \mu)^2. \quad (\text{A.4})$$

Before expanding the left hand-side of (A.4), we abbreviate

$$x := \exp\left\{-\frac{kh_c^2}{2}\right\},$$

then (A.4) is rewritten as

$$(\mu - x)^2 - (\mu - Nx)^2 < (1 - \mu)^2, \quad (\text{A.5})$$

but since $(\mu - x)^2 > (\mu - Nx)^2$, and that we took k sufficiently large so that $(\mu - Nx)$ is non-negative, then (A.5) will hold if

$$(\mu - x)^2 < (1 - \mu)^2. \quad (\text{A.6})$$

But from condition (A.2) it follows that

$$x = \exp\left\{-\frac{kh_c^2}{2}\right\} < \exp\log\left(\frac{\mu}{m_2}\right) < \mu,$$

and then, (A.6) will hold as long as

$$\mu^2 < (1 - \mu)^2, \quad (\text{A.7})$$

which holds for $\mu < 1/2$. Since μ is bounded from above by $1/6$, the proof is complete. $\square$

A.2 Proofs of Chapter 5

A.2.1 Lemma 5.2.5

Let $Y_1, Y_2, \dots, Y_{M-m}$ and $X_1, X_2, \dots, X_m$, with $m < M$, be samples of i.i.d. random variables following an uniform distribution in the interval $[0, 1]$. Then for a integer $1 \leq c \leq M - m$, it

holds that

$$\mathbb{P}[Y_{(c)} < X_{(1)}] > \left(\frac{M - m - c + 1}{M - m + 1} \right)^m,$$

where $Y_{(c)}$ and $X_{(1)}$ correspond respectively to the c -th and 1-st order statistic of the samples.

Proof. Conditioning on $X_{(1)}$ we get

$$\begin{aligned} \mathbb{P}[Y_{(c)} < X_{(1)}] &= \int_0^1 \mathbb{P}[Y_{(c)} < X_{(1)} \mid X_{(1)} = z] f_{X_{(1)}}(z) dz \\ &= m \sum_{j=c}^{M-m} \binom{M-m}{j} \int_0^1 z^j (1-z)^{M-j-1} dz \\ &= m \sum_{j=c}^{M-m} \binom{M-m}{j} \frac{\Gamma(j+1)\Gamma(M-j)}{\Gamma(M+1)} \\ &= \frac{(M-c)!}{M!} \frac{(M-m)!}{(M-m-c)!} > \left(\frac{M-m-c+1}{M-m+1} \right)^m. \end{aligned}$$

□

A.2.2 Lemma 5.2.7

Given a set $V = \{v_1, \dots, v_T\}$ and a function $l: V \times \mathbb{N} \rightarrow \{0, 1\}$ define $V_t^i := \{v \in V : l(v, t) = i\}$ for $i = 0, 1$. Suppose l satisfies:

1. Let B_t be the event where $|V_t^1| - |V_{t-1}^1| \geq \sqrt{|V_{t-1}^0|}$. For any $t > 0$, $\mathbb{P}[B_t] \geq W$.
2. The events B_t are independent for all $t \in [0, \infty)$.

Let S be the random variable representing the first time such that $V_S^0 = 0$ and call $t_0 = 2\sqrt{T}$. Then for $t \geq 0$

$$\mathbb{P}[S \leq t_0 + t] \geq 1 - \left(1 - \frac{W^{t_0}}{t_0} \right)^t$$

Proof. Let D_t be the event where $\sqrt{|V_{t+1}^0|} \leq \sqrt{|V_{t-1}^0|} - 1$. We begin by proving that $B_t B_{t+1} \subset D_t$. Assuming B_t , there is an $\epsilon \in [0, |V_{t-1}^0| - \sqrt{|V_{t-1}^0|}]$ such that

$$|V_t^1| = |V_{t-1}^1| + \sqrt{|V_{t-1}^0|} + \epsilon. \tag{A.8}$$

ϵ cannot be larger than $|V_{t-1}^0| - \sqrt{|V_{t-1}^0|}$, otherwise

$$|V_t^1| > |V_{t-1}^1| + |V_{t-1}^0| = T.$$

Then, assuming B_{t+1} :

$$\begin{aligned} |V_{t+1}^1| - |V_{t-1}^1| &= |V_{t+1}^1| - |V_t^1| + |V_t^1| - |V_{t-1}^1| \\ &\geq \sqrt{|V_t^0|} + \sqrt{|V_{t-1}^0|} + \epsilon \\ &= \sqrt{n - |V_t^1|} + \sqrt{|V_{t-1}^0|} + \epsilon. \end{aligned}$$

By (A.8), we get that $|V_{t+1}^1| - |V_{t-1}^1|$ is bounded below by

$$\begin{aligned} &\sqrt{n - |V_{t-1}^1|} - \sqrt{|V_{t-1}^0|} - \epsilon + \sqrt{|V_{t-1}^0|} + \epsilon \\ &= \sqrt{|V_{t-1}^0|} - \sqrt{|V_{t-1}^0|} - \epsilon + \sqrt{|V_{t-1}^0|} + \epsilon \\ &:= g(\epsilon). \end{aligned}$$

The function g is concave on ϵ , so its minimum lies at the boundaries for ϵ . We have

$$\begin{cases} g(0) = \sqrt{|V_{t-1}^0|} - \sqrt{|V_{t-1}^0|} + \sqrt{|V_{t-1}^0|} \\ g(|V_{t-1}^0| - \sqrt{|V_{t-1}^0|}) = |V_{t-1}^0|. \end{cases}$$

Then, whenever $|V_{t-1}^0| \geq 3$, its minimum is achieved when $\epsilon = 0$, i.e.:

$$|V_{t+1}^1| - |V_{t-1}^1| \geq \sqrt{|V_{t-1}^0|} - \sqrt{|V_{t-1}^0|} + \sqrt{|V_{t-1}^0|}$$

However, since $|V_{t-1}^0| \geq 3$,

$$\begin{aligned} |V_{t+1}^1| - |V_{t-1}^1| &\geq 2\sqrt{|V_{t-1}^0|} - 1 \\ &\implies 2\sqrt{|V_{t-1}^0|} - 1 \leq |V_{t-1}^0| - |V_{t+1}^0| \\ &\implies \sqrt{|V_{t+1}^0|} \leq \sqrt{|V_{t-1}^0|} - 1. \end{aligned}$$

On the other hand, if $|V_{t-1}^0| = 2$, by B_t :

$$|V_{t-1}^0| - |V_t^0| \geq \sqrt{2} \implies |V_t^0| \leq 2 - \sqrt{2} \implies |V_t^0| = 0.$$

Then, by B_{t+1} , $|V_{t+1}^0| = 0$, and D_t also holds. If $|V_{t-1}^0| = 1$, by B_t :

$$|V_{t-1}^0| - |V_t^0| \geq 1 \implies |V_t^0| \leq 0.$$

So D_t also holds. Now, since $B_t B_{t+1} \subset D_t$:

$$\mathbb{P}[D_t] \geq \mathbb{P}[B_t B_{t+1}] = \mathbb{P}[B_t] \mathbb{P}[B_{t+1}] \geq W^2.$$

Now, if $B_t B_{t+1}$ implies that $\sqrt{|V_{t-1}^0|} - \sqrt{|V_{t+1}^0|} \geq 1$,

$$\bigcap_{t \leq i \leq t+2\delta} B_i \implies \sqrt{|V_{t-1}^0|} - \sqrt{|V_{t+2\delta-1}^0|} \geq \delta,$$

and

$$\bigcap_{t \leq i \leq t+2\sqrt{T}} B_i \implies \sqrt{|V_{t-1}^0|} - \sqrt{|V_{t+2\sqrt{T}-1}^0|} \geq \sqrt{T}.$$

But since $\sqrt{|V_{t-1}^0|}$ is always smaller or equal than $\sqrt{T}$, then $|V_{t+2\sqrt{T}-1}^0| = 0$. Then, if B_t is true for $t_0 = 2\sqrt{T}$ subsequent rounds, V^0 after those rounds will be the empty set.

To find a bound for this event, divide the interval $[0, t_0 + t]$ in sub-intervals of length t_0 . Then, the probability of not having t_0 consecutive successes in the interval of length $t_0 + t$ is always smaller than the probability of not having t_0 consecutive successes in each one of the sub-intervals. Or explicitly

$$\begin{aligned} & \mathbb{P}[\text{less than } t_0 \text{ consecutive successes in } t_0 + t] \\ & < \mathbb{P}[\text{less than } t_0 \text{ consecutive successes in every interval of length } t_0] \\ & = (1 - W^{t_0})^{\lfloor (t_0+t)/t_0 \rfloor}. \end{aligned}$$

This means that

$$\begin{aligned} & \mathbb{P}[S \leq t_0 + t] \\ & \geq \mathbb{P}[t_0 \text{ consecutive successes in } t_0 + t] \\ & > 1 - (1 - W^{t_0})^{\lfloor (t_0+t)/t_0 \rfloor} \\ & \geq 1 - (1 - W^{t_0})^{t/t_0} \end{aligned}$$

But since $(1 - W^{t_0})^{t/t_0} < (1 - W^{t_0}/t_0)^t$, it follows that

$$\mathbb{P}[S \leq t_0 + t] \geq 1 - \left(1 - \frac{W^k}{k}\right)^{n-k}$$

which completes the proof. □

