A STUDY OF DISTRIBUTED CONSENSUS IN UNRELIABLE NETWORKS

JANE A. RESEARCHER
B.Sc. Computer Science

A THESIS SUBMITTED FOR THE DEGREE OF DOCTOR OF PHILOSOPHY

DEPARTMENT OF COMPUTER SCIENCE
UNIVERSITY OF SOMEWHERE

2026

Declaration

I hereby declare that this thesis is my original work and it has been written by me in its entirety. I have duly acknowledged all the sources of information which have been used in the thesis.

This thesis has also not been submitted for any degree in any university previously.

Jane A. Researcher
June 2026

Acknowledgments

I would like to thank my supervisor and committee for their guidance throughout this project, and my colleagues in the research group for many useful discussions.

Abstract

Distributed systems must often reach agreement despite unreliable networks, message loss, and partial failures of participating nodes. This thesis studies consensus protocols under weakened network assumptions and proposes an algorithm that degrades gracefully as message loss increases, rather than failing outright once a fixed fault threshold is exceeded. We give a formal model of the network and failure assumptions, prove safety and liveness properties of the proposed protocol, and evaluate it against existing approaches under simulated partition and loss conditions. The results suggest that relaxing strict synchrony assumptions in exchange for probabilistic termination guarantees offers a favourable trade-off for systems operating over wide-area, lossy links.

Contents

List of Figures vi
List of Tables vii
List of Algorithms viii

1 Introduction **1**
1.1 Motivation 1
1.2 Contributions 1

2 Related Work **3**
2.1 Synchronous and Partially Synchronous Models 3
2.2 Gaps Addressed by This Thesis 3

References **4**

List of Figures

List of Tables

1.1 Notation used throughout this thesis. 2

List of Algorithms

Chapter 1

Introduction

Distributed consensus is a foundational problem in distributed computing: a set of nodes must agree on a single value even when some nodes may fail or messages may be lost or delayed. Classical results such as the FLP impossibility theorem show that deterministic consensus cannot be guaranteed in a fully asynchronous system with even a single faulty process [FLP85]. Practical systems work around this by making partial synchrony assumptions, but these assumptions are frequently violated in real wide-area deployments.

1.1 Motivation

Modern distributed applications increasingly run across wide-area networks with variable latency and non-trivial packet loss. Protocols designed and tuned for data-centre networks, where loss is rare and latency is predictable, often behave poorly once deployed across such links [OO14].

1.2 Contributions

This thesis makes the following contributions:

- A formal model of consensus under bounded but non-zero message loss.
- A protocol whose expected termination time degrades smoothly with loss rate, rather than failing once a fixed threshold is crossed.
- An experimental evaluation comparing the protocol against Raft-style baselines [OO14] under simulated partition and loss.

Table 1.1: Notation used throughout this thesis.

Symbol	Meaning
n	number of nodes
f	number of tolerated faulty nodes
p	message loss probability

Chapter 2

Related Work

Consensus protocols have a long history in distributed systems research, from early work on the Byzantine Generals problem [LSP82] through leader-based protocols designed for practical adoption [OO14].

2.1 Synchronous and Partially Synchronous Models

Most deployed consensus protocols assume partial synchrony: message delays are bounded, but the bound is unknown to participants [FLP85]. This assumption underlies both Paxos-family protocols and Raft.

2.2 Gaps Addressed by This Thesis

Prior evaluations of consensus protocols under loss typically treat loss as a binary partition event, rather than a continuous parameter. A later evaluation chapter (omitted from this starter) would test the proposed protocol across a continuous range of loss rates.

References

[FLP85] Michael J. Fischer, Nancy A. Lynch, and Michael S. Paterson. Impossibility of distributed consensus with one faulty process. *Journal of the ACM*, 32(2):374–382, 1985.

[LSP82] Leslie Lamport, Robert Shostak, and Marshall Pease. The byzantine generals problem. *ACM Transactions on Programming Languages and Systems*, 4(3):382–401, 1982.

[OO14] Diego Ongaro and John Ousterhout. In search of an understandable consensus algorithm. In *2014 USENIX Annual Technical Conference*, pages 305–319, 2014.