Approximate Consensus Under Network Partition

by

Jane Doe

B.Sc. Computer Science, Some University, 2018

A THESIS SUBMITTED IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF

Doctor of Philosophy

in

THE FACULTY OF GRADUATE AND POSTDOCTORAL STUDIES
(Computer Science)

The University of British Columbia
(Vancouver)

August 2026

© Jane Doe, 2026

The following individuals certify that they have read, and recommend to the Faculty of Graduate and Postdoctoral Studies for acceptance, the thesis entitled:

Approximate Consensus Under Network Partition

submitted by **Jane Doe** in partial fulfillment of the requirements for the degree of **Doctor of Philosophy** in **Computer Science**.

Examining Committee:

Advisor Name, Professor, Computer Science, UBC
Supervisor

Reader One, Professor, Computer Science, UBC
Supervisory Committee Member

Reader Two, Professor, Electrical Engineering, UBC
Additional Examiner

Additional Supervisory Committee Members:

Advisor Name, Professor, Computer Science, UBC
Supervisory Committee Member

Abstract

This thesis studies how approximate consensus protocols degrade under partial network partition, and proposes a protocol that trades a small latency overhead for improved throughput.

Lay Summary

Computer systems that run across many machines need to agree on a shared view of the world even when parts of the network occasionally fail to communicate. This thesis shows that a small, deliberate compromise in how quickly machines respond keeps the whole system working smoothly when the network is unreliable, instead of slowing to a crawl.

Table of Contents

Abstract . . . . . **iii**

Lay Summary . . . . . **iv**

Table of Contents . . . . . **v**

List of Tables . . . . . **vi**

List of Figures . . . . . **vii**

Acknowledgments . . . . . **viii**

1 Introduction **1**

1.1 Motivation 1

1.2 Contributions 1

1.3 Related Work 1

2 Evaluation **2**

2.1 Experimental Setup 2

2.2 Latency Under Load 2

2.3 Discussion 2

Bibliography **4**

List of Tables

Table 2.1 Throughput (requests/s) under simulated partition rates. 2

List of Figures

Figure 2.1 Median latency versus simulated partition rate. 3

Acknowledgments

The author thanks their advisor and lab colleagues for their support throughout this research.

Chapter 1

Introduction

1.1 Motivation

Large-scale distributed systems increasingly rely on approximate consensus protocols to balance latency and consistency [1]. This work studies how such protocols degrade under network partition.

1.2 Contributions

The main contributions are:

1. A formal model of partial-partition tolerance.
2. An empirical evaluation summarised in Table 2.1.
3. A reference implementation used to produce Figure 2.1.

1.3 Related Work

Classical consensus protocols such as Paxos [1] and Raft [2] assume a fully connected network. This thesis relaxes that assumption.

Chapter 2

Evaluation

2.1 Experimental Setup

Experiments were run on a five-node cluster. Table 2.1 summarises throughput under increasing partition rates.

Table 2.1: Throughput (requests/s) under simulated partition rates.

Partition rate	Baseline	Proposed
0%	12,400	12,100
5%	9,800	11,600
10%	6,200	10,900

2.2 Latency Under Load

Figure 2.1 shows median request latency as the partition rate increases.

2.3 Discussion

The proposed protocol maintains throughput more effectively than the baseline once the partition rate exceeds 5%, at the cost of a small constant latency overhead [2].

Figure 2.1: Median latency versus simulated partition rate.

Bibliography

[1] L. Lamport. The part-time parliament. *ACM Transactions on Computer Systems*, 16(2):133–169, 1998.

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