APPROXIMATE CONSENSUS UNDER NETWORK PARTITION

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF COMPUTER SCIENCE
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Jane Doe
August 2026

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.

Acknowledgements

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

Contents

Abstract **iv**

Acknowledgements **v**

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

List of Tables

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

List of Figures

2.1 Median latency versus simulated partition rate. 3

Chapter 1

Introduction

1.1 Motivation

Large-scale distributed systems increasingly rely on approximate consensus protocols to balance latency and consistency [Lamport, 1998]. 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 [Lamport, 1998] and Raft [Ongaro and Ousterhout, 2014] 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 [Ongaro and Ousterhout, 2014].

Figure 2.1: Median latency versus simulated partition rate.

Bibliography

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

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