Approximate Consensus Under Network Partition

Jane Doe

August 27, 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.

Contents

1 Introduction 5

1.1 Motivation 5
1.2 Contributions 5
1.3 Related Work 5

2 Evaluation 7

2.1 Experimental Setup 7
2.2 Latency Under Load 7
2.3 Discussion 7

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.

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.

Figure 2.1. Median latency versus simulated partition rate..

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].

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.