

About the book

The field of distributed systems is becoming increasingly important to the theory and practice of Computer Science. The progress in distributed systems over the past few decades has led to a vast repertoire of techniques, models, systems, and programming frameworks. The influence of distributed computing can be seen in many traditional and emerging areas of Computer Science such as programming languages, algorithms, operating systems, file systems, storage systems, group communication, Internet-of-Things (IoT), Blockchain systems, and so on.

The past decade has seen tremendous advances in the development and deployment of several systems and technologies aimed at large scale systems such as the Map-Reduce programming framework, DynamoDB key-value stores, and consensus protocols such as Paxos and RAFT. These developments and their foundational underpinnings have sadly not been lucidly described in textbook style narratives. In this backdrop, there is a strong need for a textbook on distributed systems that covers the fundamental aspects and design principles that have evolved over the last three decades.

This book is divided into three parts.

The first part covers foundational concepts: the notion of time, mutual exclusion, distributed algorithms, consensus protocols and distributed hash tables.

The second part covers the application of distributed systems: to services file systems, database systems, distributed machine learning and distributed ledgers.

The final part covers emerging topics and applications of distributed systems: fog and edge computing, security, and performance evaluation.



The field of distributed systems now spans across multiple areas of Computer Science. This book is suitable for final year undergraduate and postgraduate students (2 semester course).

The entire book and related instructional material such as lecture slides, videos and other supplementary material will be made available for free download through our website. The July 2024 release will include all the chapters in Part 1 and some chapters in Part 2. We plan to release the remaining chapters of Part 2 and Part 3 by December 2024.

For any feedback, comments, or error reports, please email the authors at thedistsysbook@gmail.com. We believe that the book will serve the community in a great way and we shall thus appreciate the feedback of the community. This will help us improve the book.

Contents

Part I

Foundational Concepts



- Time in Distributed Systems
- Distributed Data Storage
- Mutual Exclusion Algorithms
- Distributed Algorithms
- Consensus and Agreement
- Consistency

Part II

Applications & Systems



- Distributed File Systems
- Distributed Databases and Transaction Processing
- Blockchains
- Distributed Machine Learning

Part III

Management



- Middleware
- Fog and Edge Computing
- Distributed Programming Frameworks
- Security in Distributed Systems
- Performance Evaluation and Benchmarking

Case studies

