

# Runtimes for Concurrency and Distribution

Gabriel Rovesti

Academic Year 2024-2025

April 20, 2025

# Contents

<b>1</b>	<b>The Notion of Run-time Support</b>	<b>6</b>
1.1	Introduction to Programming Abstractions . . . . .	6
1.2	Procedures as Abstractions . . . . .	6
1.3	The Abstraction Gap . . . . .	6
1.4	Levels of Abstraction . . . . .	6
1.5	Key Concepts . . . . .	7
<b>2</b>	<b>Multiprogramming</b>	<b>8</b>
2.1	The Run-time Environment . . . . .	8
2.2	Components of the Run-time Environment . . . . .	8
2.2.1	Operating System . . . . .	8
2.2.2	Language-Specific Run-time Support . . . . .	8
2.3	The Program Abstraction . . . . .	9
2.4	Multiprogramming Implementation . . . . .	9
2.4.1	Process Management . . . . .	9
2.4.2	Memory Management . . . . .	9
2.4.3	Resource Management . . . . .	9
2.5	Key Insights . . . . .	10
<b>3</b>	<b>Virtualization</b>	<b>11</b>
3.1	Introduction to Virtualization . . . . .	11
3.2	From Abstraction to Virtualization . . . . .	11
3.3	Types of Virtualization . . . . .	11
3.3.1	Platform Virtualization . . . . .	11
3.3.2	Resource Virtualization . . . . .	11
3.4	Hypervisors/Virtual Machine Monitors . . . . .	11
3.5	Benefits of Virtualization . . . . .	12
3.6	Virtualization Techniques . . . . .	12
3.6.1	Trap-and-Emulate . . . . .	12
3.6.2	Binary Translation . . . . .	12
3.7	Challenges in Virtualization . . . . .	12
3.8	Virtualization in Modern Platforms . . . . .	12
<b>4</b>	<b>Distribution and Scalability</b>	<b>13</b>
4.1	Introduction to Distribution . . . . .	13
4.2	Relationship Between Concurrency and Distribution . . . . .	13
4.3	Fundamental Concepts of Distribution . . . . .	13
4.3.1	Distributed System Definition . . . . .	13
4.3.2	Key Characteristics . . . . .	13
4.4	Scalability in Distributed Systems . . . . .	13
4.4.1	Dimensions of Scalability . . . . .	14

4.4.2	Scalability Challenges . . . . .	14
4.5	Distribution Models . . . . .	14
4.5.1	Client-Server Model . . . . .	14
4.5.2	Peer-to-Peer Model . . . . .	14
4.5.3	Hybrid Models . . . . .	14
4.6	Runtime Support for Distribution . . . . .	14
4.7	Key Insights . . . . .	15
<b>5</b>	<b>Models of Concurrency</b>	<b>16</b>
5.1	Design Challenges for Concurrency Models . . . . .	16
5.2	The Language Designer's Perspective . . . . .	16
5.2.1	Units of Concurrent Execution . . . . .	16
5.2.2	Syntactic Representation . . . . .	16
5.2.3	Execution Model . . . . .	17
5.3	Shared Memory vs. Message Passing . . . . .	17
5.3.1	Shared Memory Model . . . . .	17
5.3.2	Message Passing Model . . . . .	17
5.4	Specific Language Approaches to Concurrency . . . . .	17
5.4.1	Java's Concurrency Model . . . . .	17
5.4.2	Go's Concurrency Model . . . . .	17
5.4.3	Ada's Concurrency Model . . . . .	18
5.4.4	Erlang's Concurrency Model . . . . .	18
5.5	Evaluating Concurrency Models . . . . .	18
5.6	Key Insights . . . . .	18
<b>6</b>	<b>Communication Among Threads</b>	<b>19</b>
6.1	The Necessity of Communication . . . . .	19
6.2	Challenges in Inter-Thread Communication . . . . .	19
6.3	Communication Models . . . . .	19
6.3.1	Shared Memory Communication . . . . .	19
6.3.2	Message-Based Communication . . . . .	20
6.4	Synchronization Primitives . . . . .	20
6.4.1	Mutual Exclusion (Mutex) . . . . .	20
6.4.2	Semaphores . . . . .	20
6.4.3	Condition Variables . . . . .	20
6.4.4	Barriers . . . . .	20
6.5	Communication Patterns . . . . .	20
6.5.1	Producer-Consumer . . . . .	20
6.5.2	Readers-Writers . . . . .	21
6.5.3	Dining Philosophers . . . . .	21
6.6	Communication and Distribution . . . . .	21
6.7	Key Insights . . . . .	21
<b>7</b>	<b>Synchronous Communication</b>	<b>22</b>
7.1	Characteristics of Synchronous Communication . . . . .	22
7.2	Synchronous Communication Mechanisms . . . . .	22
7.2.1	Rendezvous . . . . .	22
7.2.2	Synchronous Message Passing . . . . .	22
7.2.3	Remote Procedure Call (RPC) . . . . .	22
7.3	Implementing Synchronous Communication . . . . .	23
7.4	Synchronous Communication in Programming Languages . . . . .	23
7.4.1	Ada's Rendezvous Mechanism . . . . .	23

7.4.2	Go's Channel-Based Communication . . . . .	23
7.5	Advantages of Synchronous Communication . . . . .	24
7.6	Limitations of Synchronous Communication . . . . .	24
7.7	Synchronous Communication in Distributed Systems . . . . .	24
7.8	Key Insights . . . . .	24
<b>8</b>	<b>Asynchronous Communication (The Monitor)</b>	<b>25</b>
8.1	Introduction to Asynchronous Communication . . . . .	25
8.2	The Monitor Concept . . . . .	25
8.2.1	Key Components of Monitors . . . . .	25
8.3	Monitor Implementation . . . . .	25
8.3.1	Basic Structure . . . . .	25
8.3.2	Condition Variables . . . . .	26
8.4	Monitor Variations . . . . .	26
8.4.1	Hoare Monitors . . . . .	26
8.4.2	Mesa Monitors . . . . .	26
8.5	Ada's Protected Objects . . . . .	27
8.6	Monitors vs. Other Synchronization Mechanisms . . . . .	27
8.7	Asynchronous Benefits and Challenges . . . . .	27
8.7.1	Benefits . . . . .	27
8.7.2	Challenges . . . . .	28
8.8	Key Insights . . . . .	28
<b>9</b>	<b>The Multiple Facets of Synchronization</b>	<b>29</b>
9.1	Dimensions of Synchronization . . . . .	29
9.1.1	Purpose of Synchronization . . . . .	29
9.1.2	Mechanism Characteristics . . . . .	29
9.2	Ada's Protected Objects: A Comprehensive Approach . . . . .	29
9.2.1	Key Capabilities . . . . .	29
9.2.2	Protected Object Types . . . . .	30
9.3	Advanced Synchronization Patterns . . . . .	30
9.3.1	Read-Write Locks . . . . .	30
9.3.2	Priority Inheritance . . . . .	30
9.3.3	Non-Blocking Synchronization . . . . .	30
9.4	Synchronization Problems and Solutions . . . . .	30
9.4.1	Deadlock . . . . .	30
9.4.2	Starvation . . . . .	30
9.4.3	Priority Inversion . . . . .	31
9.5	Case Study: The Mars Pathfinder Incident . . . . .	31
9.6	Synchronization in Distributed Contexts . . . . .	31
9.7	Key Insights . . . . .	31
<b>10</b>	<b>Back to Distribution</b>	<b>32</b>
10.1	Revisiting Distribution Fundamentals . . . . .	32
10.2	The Role of Runtimes in Distributed Systems . . . . .	32
10.3	Distributed System Architecture . . . . .	32
10.3.1	Logical Architecture . . . . .	32
10.3.2	Physical Architecture . . . . .	32
10.4	Distribution Transparency . . . . .	33
10.5	Challenges in Distributed Systems . . . . .	33
10.5.1	The CAP Theorem . . . . .	33
10.5.2	The FLP Impossibility Result . . . . .	33

10.5.3 Other Challenges . . . . .	33
10.6 Distributed Runtime Components . . . . .	34
10.6.1 Communication Middleware . . . . .	34
10.6.2 Service Discovery . . . . .	34
10.6.3 Orchestration and Management . . . . .	34
10.7 Key Insights . . . . .	34
<b>11 Distributed Inter-Process Communication</b>	<b>35</b>
11.1 The Challenge of Remote Communication . . . . .	35
11.2 Network Protocol Fundamentals . . . . .	35
11.2.1 Protocol Layers . . . . .	35
11.2.2 Connection-Oriented vs. Connectionless . . . . .	35
11.3 Remote Procedure Call (RPC) . . . . .	35
11.3.1 RPC Concept . . . . .	35
11.3.2 RPC Components . . . . .	36
11.3.3 RPC Operation . . . . .	36
11.3.4 RPC Challenges . . . . .	36
11.4 Message-Oriented Middleware (MOM) . . . . .	36
11.4.1 Characteristics . . . . .	36
11.4.2 MOM Components . . . . .	36
11.4.3 Delivery Semantics . . . . .	37
11.5 Web Services and RESTful Communication . . . . .	37
11.5.1 Web Services . . . . .	37
11.5.2 REST (Representational State Transfer) . . . . .	37
11.6 Emerging Communication Patterns . . . . .	37
11.6.1 GraphQL . . . . .	37
11.6.2 gRPC . . . . .	37
11.6.3 WebSockets . . . . .	37
11.7 Serialization and Data Formats . . . . .	38
11.7.1 Data Serialization . . . . .	38
11.7.2 Common Data Formats . . . . .	38
11.8 Key Insights . . . . .	38
<b>12 Distributed Concurrency</b>	<b>39</b>
12.1 The Convergence of Distribution and Concurrency . . . . .	39
12.2 Server-Side Scalability . . . . .	39
12.2.1 Scalability Dimensions . . . . .	39
12.2.2 Server Architectures for Scalability . . . . .	39
12.3 Concurrency Control in Distributed Systems . . . . .	40
12.3.1 Optimistic Concurrency Control . . . . .	40
12.3.2 Pessimistic Concurrency Control . . . . .	40
12.3.3 Distributed Transactions . . . . .	40
12.4 Balancing Scalability and Consistency . . . . .	41
12.4.1 Consistency Models . . . . .	41
12.4.2 Consistency vs. Performance Trade-offs . . . . .	41
12.5 Distributed Concurrency Patterns . . . . .	41
12.5.1 Sharding/Partitioning . . . . .	41
12.5.2 Command Query Responsibility Segregation (CQRS) . . . . .	41
12.5.3 Event Sourcing . . . . .	41
12.6 Distributed Concurrency Technologies . . . . .	42
12.6.1 Actor Model . . . . .	42
12.6.2 Conflict-Free Replicated Data Types (CRDTs) . . . . .	42

12.7 Key Insights . . . . .	42
<b>13 Distributed Synchronization</b>	<b>43</b>
13.1 The Challenge of Distributed Agreement . . . . .	43
13.2 Clock Synchronization . . . . .	43
13.2.1 Physical Clock Synchronization . . . . .	43
13.2.2 Logical Clocks . . . . .	43
13.3 Distributed Consensus . . . . .	44
13.3.1 The Consensus Problem . . . . .	44
13.3.2 Consensus Algorithms . . . . .	44
13.4 State Machine Replication . . . . .	44
13.5 Distributed Mutual Exclusion . . . . .	45
13.5.1 Token-Based Algorithms . . . . .	45
13.5.2 Permission-Based Algorithms . . . . .	45
13.5.3 Quorum-Based Algorithms . . . . .	45
13.6 Distributed Deadlock Detection . . . . .	45
13.6.1 Centralized Detection . . . . .	45
13.6.2 Distributed Detection . . . . .	45
13.7 Distributed Termination Detection . . . . .	45
13.8 Leader Election . . . . .	46
13.9 Key Insights . . . . .	46
<b>14 Enter the Cloud</b>	<b>47</b>
14.1 The Evolution of Cloud Computing . . . . .	47
14.1.1 Historical Context . . . . .	47
14.1.2 Visionaries and Early Adopters . . . . .	47
14.2 Cloud Computing Characteristics . . . . .	47
14.2.1 Essential Characteristics (NIST Definition) . . . . .	47
14.2.2 Service Models . . . . .	48
14.2.3 Deployment Models . . . . .	48
14.3 Cloud Infrastructure Architecture . . . . .	48
14.3.1 Compute Virtualization . . . . .	48
14.3.2 Storage Systems . . . . .	48
14.3.3 Networking . . . . .	48
14.4 Cloud Design Patterns . . . . .	49
14.4.1 Scalability Patterns . . . . .	49
14.4.2 Resilience Patterns . . . . .	49
14.4.3 Data Management Patterns . . . . .	49
14.5 Cloud-Native Application Architecture . . . . .	49
14.5.1 Microservices . . . . .	49
14.5.2 Containers and Orchestration . . . . .	49
14.5.3 DevOps and Continuous Delivery . . . . .	50
14.6 Challenges in Cloud Computing . . . . .	50
14.6.1 Technical Challenges . . . . .	50
14.6.2 Business Challenges . . . . .	50
14.7 Future Trends . . . . .	50
14.8 Key Insights . . . . .	50

# Chapter 1

## The Notion of Run-time Support

### 1.1 Introduction to Programming Abstractions

Programming languages provide abstractions that simplify the development process, allowing programmers to express complex operations through high-level constructs rather than low-level machine instructions. These abstractions are fundamental to modern software development but require mechanisms to bridge the gap between the programming model and the underlying hardware.

### 1.2 Procedures as Abstractions

The concept of a **procedure** (or function) is a fundamental abstraction present in virtually all modern programming languages. However, processors do not natively support this abstraction:

- Procedures allow code to be organized into reusable units
- They enable parameter passing and return values
- They support local variables with limited scope
- They implement control flow mechanisms (call and return)

### 1.3 The Abstraction Gap

There exists a fundamental gap between:

- What programming languages offer as "native" features
- What processors directly support in hardware

This gap must be filled by software that executes on the processor, underneath the program itself. This software operates at a lower level of abstraction than the program but higher than the bare processor.

### 1.4 Levels of Abstraction

The execution environment can be understood as a series of abstraction layers:

1. Application program (highest)
2. Programming language abstractions

3. Run-time support for these abstractions
4. Operating system services
5. Hardware capabilities (lowest)

## 1.5 Key Concepts

### Key Concepts

- **Programming abstraction:** High-level constructs that simplify programming but require translation for execution
- **Abstraction gap:** The difference between what programming languages offer and what processors directly support
- **Run-time support:** Software mechanisms that implement language abstractions on top of processor capabilities



# Chapter 2

## Multiprogramming

### 2.1 The Run-time Environment

Building on the previous chapter, we can now formally define the machinery that sits between program execution and the processor:

#### Definition

The **run-time environment** is the set of software mechanisms that enable program execution by implementing abstractions not directly supported by the underlying hardware.

### 2.2 Components of the Run-time Environment

The run-time environment consists of two complementary parts:

#### 2.2.1 Operating System

The operating system's primary purpose is to:

- Allow simultaneous execution of multiple programs
- Ensure isolation between programs (preventing interference)
- Manage shared resources (processor time, memory, I/O)
- Provide abstractions like processes, files, and I/O channels

#### 2.2.2 Language-Specific Run-time Support

Each programming language may require specific abstractions that are not provided by the operating system:

- Memory management (e.g., garbage collection in Java or C#)
- Thread management (for languages with built-in concurrency)
- Exception handling mechanisms
- Support for object-oriented features (inheritance, polymorphism)

## 2.3 The Program Abstraction

A central abstraction implemented by the operating system is the notion of a "program" itself:

- Programs are represented as processes with their own:
  - Address space (virtual memory)
  - Execution context (registers, stack, program counter)
  - Resource allocations (files, sockets, etc.)
- The operating system creates the illusion that each program runs in isolation
- In reality, the processor is shared among multiple programs through context switching

## 2.4 Multiprogramming Implementation

The implementation of multiprogramming involves several key mechanisms:

### 2.4.1 Process Management

- Process creation and termination
- Process scheduling algorithms
- Process state transitions (ready, running, blocked)
- Context switching between processes

### 2.4.2 Memory Management

- Virtual memory implementation
- Address translation (virtual to physical)
- Memory protection between processes
- Memory allocation and deallocation

### 2.4.3 Resource Management

- Allocation of processor time
- Management of I/O devices
- File system implementation
- Communication channels

## 2.5 Key Insights

### Key Insights

- The run-time environment bridges the gap between programming abstractions and hardware capabilities
- Multiprogramming creates the illusion of simultaneous program execution on sequential hardware
- Isolation between programs is a fundamental requirement for reliable computing
- The operating system implements the basic abstraction of a "program" as a process

## Chapter 3

# Virtualization

### 3.1 Introduction to Virtualization

Virtualization represents an extension of the abstraction concept discussed in previous chapters. It allows for creating multiple virtual instances of execution platforms (such as operating systems) on a single physical machine.

### 3.2 From Abstraction to Virtualization

- **Abstraction** hides implementation details, presenting a simplified interface
- **Virtualization** creates complete simulated environments that behave like real systems
- Virtualization creates the illusion of dedicated hardware resources

### 3.3 Types of Virtualization

#### 3.3.1 Platform Virtualization

- **Full virtualization:** Complete simulation of hardware to run unmodified operating systems
- **Paravirtualization:** Operating system is aware of virtualization and cooperates
- **Hardware-assisted virtualization:** CPU features that aid virtualization (e.g., Intel VT-x, AMD-V)

#### 3.3.2 Resource Virtualization

- **Memory virtualization:** Virtual memory space for each virtual machine
- **I/O virtualization:** Access to virtualized devices
- **Network virtualization:** Virtual network interfaces and connections

### 3.4 Hypervisors/Virtual Machine Monitors

The hypervisor is the key component that enables virtualization:

- **Type 1 (bare-metal):** Runs directly on hardware (e.g., VMware ESXi, Xen)
- **Type 2 (hosted):** Runs as an application on a host OS (e.g., VirtualBox, VMware Workstation)

### 3.5 Benefits of Virtualization

- **Consolidation:** Running multiple virtual machines on a single physical server
- **Isolation:** Containing failures within virtual machine boundaries
- **Flexibility:** Easy creation, migration, and management of virtual machines
- **Resource utilization:** Efficient use of computing resources

### 3.6 Virtualization Techniques

#### 3.6.1 Trap-and-Emulate

- VM executes non-privileged instructions directly
- Privileged instructions trigger traps to the hypervisor
- Hypervisor emulates the behavior of privileged instructions

#### 3.6.2 Binary Translation

- Translates privileged instructions on-the-fly
- Translated code executes safely without requiring traps
- Translation results can be cached for performance

### 3.7 Challenges in Virtualization

- **Performance overhead:** Virtualization introduces execution overhead
- **Resource contention:** Multiple VMs competing for physical resources
- **I/O virtualization:** Efficiently handling device access
- **Memory management:** Balancing memory allocation among VMs

### 3.8 Virtualization in Modern Platforms

#### Key Insights

- Virtualization is now commonplace in modern execution platforms
- It forms the foundation of cloud computing infrastructure
- Containers represent a lightweight form of virtualization
- Modern processors include hardware features specifically designed to support virtualization

## Chapter 4

# Distribution and Scalability

### 4.1 Introduction to Distribution

Distribution extends the concept of concurrent execution beyond the boundaries of a single computer system, allowing components to operate across multiple networked nodes while maintaining a coherent system behavior.

### 4.2 Relationship Between Concurrency and Distribution

It's essential to understand the deep connection between concurrency and distribution:

- Distribution inherently requires concurrency (multiple components executing simultaneously)
- Concurrency serves as a foundational building block for distributed systems
- Distribution amplifies concurrency challenges and introduces new ones

### 4.3 Fundamental Concepts of Distribution

#### 4.3.1 Distributed System Definition

##### Definition

A **distributed system** is a collection of independent computers that appears to its users as a single coherent system.

#### 4.3.2 Key Characteristics

- **Multiple autonomous components:** Independent nodes with their own processors and memory
- **Communication via message passing:** No shared memory between nodes
- **Lack of global clock:** Challenging to establish a universal time reference
- **Independent failures:** Components can fail independently

### 4.4 Scalability in Distributed Systems

Scalability refers to a system's ability to handle growing amounts of work or to be enlarged to accommodate that growth.

#### 4.4.1 Dimensions of Scalability

- **Size scalability:** Adding more users and resources
- **Geographic scalability:** Users and resources distributed over wider areas
- **Administrative scalability:** Manageable even with many independent administrative organizations

#### 4.4.2 Scalability Challenges

- **Performance:** Maintaining acceptable performance as system grows
- **Availability:** Ensuring system remains available despite component failures
- **Consistency:** Maintaining data consistency across distributed components

### 4.5 Distribution Models

#### 4.5.1 Client-Server Model

- Separation of responsibilities between service providers (servers) and consumers (clients)
- Centralized management but potential bottlenecks and single points of failure

#### 4.5.2 Peer-to-Peer Model

- Nodes act as both clients and servers
- No central coordination, enhancing robustness and scalability
- Challenges in consistency and security

#### 4.5.3 Hybrid Models

- Combining aspects of client-server and peer-to-peer
- Examples include hierarchical systems and edge computing

### 4.6 Runtime Support for Distribution

- **Networking infrastructure:** Protocols, addressing, routing
- **Remote communication mechanisms:** Remote procedure calls, message queues
- **Service discovery:** Finding and connecting to services in the network
- **Failure detection and recovery:** Identifying and handling node failures

## 4.7 Key Insights

### Key Insights

- Distribution and concurrency are deeply intertwined concepts
- Distributed systems introduce challenges beyond those of concurrent systems
- Scalability requires careful design considerations at multiple levels
- Distribution requires specialized runtime support mechanisms



## Chapter 5

# Models of Concurrency

### 5.1 Design Challenges for Concurrency Models

Creating a model of concurrent execution for a programming language presents unique challenges:

- Designing concurrent abstractions that run on inherently sequential processors
- Providing clear semantics and syntax for programmers
- Balancing expressiveness with safety and performance
- Ensuring predictable and understandable behavior

### 5.2 The Language Designer's Perspective

When designing concurrency features for a programming language, several fundamental decisions must be made:

#### 5.2.1 Units of Concurrent Execution

- **Threads:** Lightweight sequential execution streams sharing memory
- **Processes:** Independent execution units with separate address spaces
- **Tasks:** Higher-level abstraction often implemented on top of threads
- **Actors:** Independent entities communicating via message passing

#### 5.2.2 Syntactic Representation

How concurrency is expressed in the language:

- Explicit thread creation (e.g., `new Thread()` in Java)
- Language keywords (e.g., `async/await`, `go` in Go)
- Library functions vs. built-in language constructs
- Declaration-based approaches (e.g., Ada tasks)

### 5.2.3 Execution Model

Rules governing concurrent execution:

- Preemptive vs. cooperative scheduling
- Deterministic vs. non-deterministic execution
- Fair vs. priority-based scheduling
- Execution order guarantees (or lack thereof)

## 5.3 Shared Memory vs. Message Passing

Two fundamental paradigms for concurrency:

### 5.3.1 Shared Memory Model

- Concurrent threads access the same memory locations
- Requires synchronization mechanisms to prevent race conditions
- Examples: Java threads, POSIX threads, C# threads
- Challenges: data races, deadlocks, atomicity violations

### 5.3.2 Message Passing Model

- Concurrent entities communicate by explicit message exchange
- No shared state between concurrent entities
- Examples: Erlang processes, Go channels, actor model
- Challenges: message ordering, delivery guarantees, performance

## 5.4 Specific Language Approaches to Concurrency

### 5.4.1 Java's Concurrency Model

- Thread-based concurrency with shared memory
- Built-in synchronization primitives (`synchronized`, `wait/notify`)
- Java Concurrency Utilities (since Java 5) with higher-level abstractions
- Executor framework for thread pool management

### 5.4.2 Go's Concurrency Model

- Goroutines: lightweight concurrent functions
- Channels for communication between goroutines
- Select statement for multiplexing channel operations
- "Share memory by communicating, don't communicate by sharing memory" philosophy

### 5.4.3 Ada's Concurrency Model

- Built-in task type for concurrent execution
- Rendezvous mechanism for synchronous communication
- Protected objects for safe shared data access
- Strong compile-time checks for concurrency safety

### 5.4.4 Erlang's Concurrency Model

- Lightweight processes with no shared state
- Pure message passing for inter-process communication
- "Let it crash" philosophy with supervision hierarchies
- Designed for fault-tolerant, distributed systems

## 5.5 Evaluating Concurrency Models

Criteria for assessing concurrency models:

- **Safety:** Prevention of data races and other concurrency hazards
- **Liveness:** Ensuring progress and avoiding deadlocks
- **Composability:** Ability to combine concurrent components
- **Scalability:** Performance with increasing concurrency
- **Understandability:** Ease of reasoning about concurrent behavior

## 5.6 Key Insights

### Key Insights

- Language designers must make fundamental choices about concurrency models
- Different languages offer distinct approaches to concurrency based on their design principles
- The choice of concurrency model has profound implications for program structure and behavior
- No single concurrency model is optimal for all use cases

## Chapter 6

# Communication Among Threads

### 6.1 The Necessity of Communication

Concurrent programs are inherently collaborative, requiring mechanisms for threads to:

- Exchange data and results
- Coordinate activities and execution order
- Signal events and state changes
- Share resources safely

### 6.2 Challenges in Inter-Thread Communication

Enabling communication between threads introduces several complex challenges:

- **Race conditions:** Non-deterministic behavior due to timing of operations
- **Data inconsistency:** Partial updates leading to invalid states
- **Deadlocks:** Circular wait conditions causing system standstill
- **Livelocks:** Threads continuously change state without progressing
- **Priority inversion:** Lower-priority threads preventing higher-priority ones from running

### 6.3 Communication Models

Different approaches to enabling thread communication:

#### 6.3.1 Shared Memory Communication

- Threads access common memory locations for data exchange
- Requires explicit synchronization mechanisms
- Generally offers lower latency for local communication
- Examples: Global variables, heap objects in Java/C++

### 6.3.2 Message-Based Communication

- Threads exchange discrete messages containing data
- Can be implemented with or without shared memory
- Promotes clearer separation of concerns
- Examples: Channels in Go, mailboxes in actor systems

## 6.4 Synchronization Primitives

Basic building blocks for coordinating thread communication:

### 6.4.1 Mutual Exclusion (Mutex)

- Ensures only one thread can access a critical section at a time
- Prevents concurrent access to shared resources
- Examples: `synchronized` in Java, `mutex` in C++

### 6.4.2 Semaphores

- Count-based synchronization mechanism
- Can be used for signaling and resource management
- Supports both mutual exclusion and condition synchronization

### 6.4.3 Condition Variables

- Allow threads to wait for specific conditions to be met
- Used in conjunction with mutual exclusion
- Examples: `wait/notify` in Java, `condition_variable` in C++

### 6.4.4 Barriers

- Synchronize multiple threads at specific execution points
- Ensure all threads reach a certain stage before any proceed further

## 6.5 Communication Patterns

### 6.5.1 Producer-Consumer

- One or more threads produce data items
- One or more threads consume these items
- Typically implemented using a bounded buffer
- Synchronization ensures buffer integrity and proper signaling

### 6.5.2 Readers-Writers

- Multiple readers can access data simultaneously
- Writers need exclusive access
- Various policies for balancing reader vs. writer priorities

### 6.5.3 Dining Philosophers

- Classic problem illustrating deadlock and resource allocation
- Demonstrates challenges in concurrent resource acquisition

## 6.6 Communication and Distribution

The connection between thread communication and distributed systems:

- Local communication models may or may not scale to distributed settings
- Shared memory doesn't naturally extend to distributed environments
- Message passing aligns well with distributed communication requirements
- Distribution adds latency, reliability, and ordering challenges

## 6.7 Key Insights

### Key Insights

- Communication is essential for meaningful concurrent collaboration
- Thread communication introduces complex synchronization challenges
- Different communication models have distinct strengths and weaknesses
- The choice of communication model affects scalability to distributed settings

## Chapter 7

# Synchronous Communication

### 7.1 Characteristics of Synchronous Communication

Synchronous communication between threads has several defining characteristics:

- **Temporal coupling:** Sender and receiver must coordinate in time
- **Blocking behavior:** Sender typically waits until communication completes
- **Immediate feedback:** Results or acknowledgments are received directly
- **Simplified reasoning:** Sequential-like programming model

### 7.2 Synchronous Communication Mechanisms

#### 7.2.1 Rendezvous

- Threads synchronize at a specific point in their execution
- Both sender and receiver block until the communication occurs
- Prominent in Ada's task communication model
- Enables direct exchange of data with strong synchronization guarantees

#### 7.2.2 Synchronous Message Passing

- Sender blocks until the message is received
- Provides guaranteed delivery and immediate feedback
- Examples: Synchronous operations on channels in Go, CSP model

#### 7.2.3 Remote Procedure Call (RPC)

- Client invokes a procedure that executes in another thread or process
- Client blocks until the procedure completes and returns results
- Makes distributed communication appear like local procedure calls
- Forms the basis for many distributed system interactions

## 7.3 Implementing Synchronous Communication

Key components required for synchronous communication:

- **Synchronization primitives:** Mutexes, condition variables, semaphores
- **Thread blocking mechanisms:** Suspending and resuming threads
- **Message buffering:** Handling temporary timing differences
- **Context management:** Preserving state during blocking operations

## 7.4 Synchronous Communication in Programming Languages

### 7.4.1 Ada's Rendezvous Mechanism

```
-- Server task (entry provider)
task Server is
    entry Request(Data : in Item_Type; Result : out Result_Type);
end Server;

task body Server is
begin
    loop
        accept Request(Data : in Item_Type; Result : out Result_Type) do
            -- Process the request and set Result
            Result := Process(Data);
        end Request;
    end loop;
end Server;

-- Client task
task body Client is
    Data : Item_Type;
    Result : Result_Type;
begin
    -- Prepare data
    Data := Prepare_Data;

    -- Make synchronous call - blocks until complete
    Server.Request(Data, Result);

    -- Use result
    Use_Result(Result);
end Client;
```

### 7.4.2 Go's Channel-Based Communication

```
// Synchronous communication with unbuffered channels
func server(requests chan int, responses chan int) {
    for {
        // Receive request (blocks until a value is sent)
        req := <-requests

        // Process request
        result := process(req)

        // Send response (blocks until received)
        responses <- result
    }
}
```



```
}  
  
func client(requests chan int, responses chan int) {  
    // Send request (blocks until received)  
    requests <- prepareData()  
  
    // Receive response (blocks until a value is sent)  
    result := <-responses  
  
    // Use result  
    useResult(result)  
}
```

## 7.5 Advantages of Synchronous Communication

- **Simplicity:** Programming model similar to sequential code
- **Determinism:** Clear execution order and data flow
- **Immediate feedback:** Errors are reported directly to the caller
- **Backpressure:** Natural flow control when systems are overloaded

## 7.6 Limitations of Synchronous Communication

- **Reduced concurrency:** Blocking limits parallel execution
- **Potential deadlocks:** Circular dependencies can cause system standstill
- **Vulnerability to failures:** Caller is affected by receiver failures
- **Performance concerns:** Blocking can waste computational resources
- **Limited scalability:** May not perform well in highly distributed systems

## 7.7 Synchronous Communication in Distributed Systems

- Implemented through technologies like RPC, RESTful APIs, gRPC
- Adds challenges of network latency and failures
- Often requires timeout mechanisms to handle non-responsive peers
- May use network protocols that provide reliable delivery (e.g., TCP)

## 7.8 Key Insights

### Key Insights

- Synchronous communication requires coordination between sender and receiver
- It simplifies reasoning about program behavior but limits concurrency
- Multiple implementation strategies exist with different trade-offs
- Synchronous communication can be challenging to scale in distributed systems

## Chapter 8

# Asynchronous Communication (The Monitor)

### 8.1 Introduction to Asynchronous Communication

In contrast to synchronous communication, asynchronous communication decouples the timing of sender and receiver operations:

- Sender and receiver do not need to be active simultaneously
- Communication operations typically don't block the sender
- Intermediate storage holds messages until the receiver is ready

### 8.2 The Monitor Concept

#### Definition

A **monitor** is a synchronization construct that encapsulates shared data with procedures that provide mutually exclusive access to that data, along with condition variables for thread coordination.

#### 8.2.1 Key Components of Monitors

- **Encapsulated shared state:** Private data accessible only through monitor procedures
- **Mutual exclusion:** Only one thread can execute within the monitor at a time
- **Condition variables:** Allow threads to wait for specific conditions
- **Entry/exit protocol:** Handles acquisition and release of the monitor lock

### 8.3 Monitor Implementation

#### 8.3.1 Basic Structure

```
// Java-like pseudocode for a monitor
class BoundedBuffer {
    private Object[] buffer;
    private int count = 0, in = 0, out = 0;
    private int size;
```

```
public BoundedBuffer(int size) {
    this.size = size;
    buffer = new Object[size];
}

// Monitor operations provide mutual exclusion
public synchronized void put(Object item) throws InterruptedException {
    // Wait until space is available
    while (count == size)
        wait(); // Release monitor and wait

    // Add item to buffer
    buffer[in] = item;
    in = (in + 1) % size;
    count++;

    // Notify waiting consumers
    notify();
}

public synchronized Object get() throws InterruptedException {
    // Wait until an item is available
    while (count == 0)
        wait(); // Release monitor and wait

    // Remove item from buffer
    Object item = buffer[out];
    buffer[out] = null;
    out = (out + 1) % size;
    count--;

    // Notify waiting producers
    notify();
    return item;
}
}
```

### 8.3.2 Condition Variables

- Allow threads to wait for specific conditions
- Associated with the monitor and automatically release monitor lock
- Operations: wait, signal/notify, signalAll/notifyAll
- Different signaling semantics: signal-and-continue vs. signal-and-exit

## 8.4 Monitor Variations

### 8.4.1 Hoare Monitors

- When a thread signals a condition, it immediately transfers control to a waiting thread
- Signaler resumes only after the waiting thread exits the monitor or waits again
- Provides stronger guarantees about conditions when threads are awakened

### 8.4.2 Mesa Monitors

- When a thread signals a condition, waiting threads are moved to the ready queue

- Signaler continues execution; awakened threads compete for monitor entry
- Requires re-checking conditions in a loop (the "wake and re-check" pattern)
- Most common implementation (Java, C#, and many other languages)

## 8.5 Ada's Protected Objects

Ada's interpretation of the monitor concept:

```
protected type Resource is
    -- Entry operations (may involve queuing)
    entry Acquire;

    -- Protected procedures (modify the state)
    procedure Release;

    -- Protected functions (read-only access)
    function Is_Available return Boolean;
private
    Available : Boolean := True;
end Resource;

protected body Resource is
    entry Acquire when Available is
    begin
        Available := False;
    end Acquire;

    procedure Release is
    begin
        Available := True;
    end Release;

    function Is_Available return Boolean is
    begin
        return Available;
    end Is_Available;
end Resource;
```

## 8.6 Monitors vs. Other Synchronization Mechanisms

- **Monitors vs. Semaphores:** Monitors provide higher-level abstraction with encapsulation
- **Monitors vs. Message Passing:** Monitors use shared memory rather than explicit messages
- **Monitors vs. Locks:** Monitors combine data protection with condition synchronization

## 8.7 Asynchronous Benefits and Challenges

### 8.7.1 Benefits

- **Increased concurrency:** Threads spend less time blocked
- **Improved responsiveness:** Systems remain interactive during operations
- **Better resource utilization:** Computation can continue during I/O or waiting

- **Enhanced scalability:** Better suited for distributed environments

### 8.7.2 Challenges

- **Complex programming model:** More difficult to reason about execution flow
- **Error handling:** Errors may be discovered far from their cause
- **Callback hell:** Nested callbacks can lead to unreadable code
- **State management:** Maintaining context across asynchronous operations

## 8.8 Key Insights

### Key Insights

- Monitors provide structured access to shared resources with built-in synchronization
- Asynchronous communication through monitors decouples sender and receiver timing
- Different monitor implementations offer various guarantees and semantics
- Monitor-based communication balances safety with performance in concurrent systems

## Chapter 9

# The Multiple Facets of Synchronization

### 9.1 Dimensions of Synchronization

Synchronization in concurrent systems serves multiple purposes and can be categorized along several dimensions:

#### 9.1.1 Purpose of Synchronization

- **Mutual exclusion:** Preventing simultaneous access to shared resources
- **Condition synchronization:** Coordinating threads based on specific conditions
- **Ordering:** Ensuring operations occur in a specific sequence
- **Barrier synchronization:** Making threads wait at specific points until all arrive

#### 9.1.2 Mechanism Characteristics

- **Blocking vs. non-blocking:** Whether threads wait by yielding the processor
- **Pessimistic vs. optimistic:** Preventing conflicts vs. detecting and resolving them
- **Fine-grained vs. coarse-grained:** Scope of resources protected
- **Fair vs. unfair:** Policies for selecting which waiting thread proceeds

### 9.2 Ada's Protected Objects: A Comprehensive Approach

Ada's protected objects provide a particularly rich synchronization model:

#### 9.2.1 Key Capabilities

- **Integrated mutual exclusion:** All operations are automatically protected
- **Differentiated access modes:** Read-only vs. read-write operations
- **Condition-based entry:** Built-in guards on entry operations
- **Priority inheritance:** Prevents priority inversion problems

## 9.2.2 Protected Object Types

```
protected type Resource_Manager(Max_Resources : Positive) is
  -- Entry operations (may queue)
  entry Allocate(Amount : Positive; ID : out Resource_ID);
  entry Deallocate(ID : Resource_ID);

  -- Protected procedures (modify state, mutually exclusive)
  procedure Initialize;

  -- Protected functions (read-only, allow concurrent readers)
  function Available_Resources return Natural;
  function Is_Allocated(ID : Resource_ID) return Boolean;
private
  Available : Natural := Max_Resources;
  -- Other implementation details
end Resource_Manager;
```

## 9.3 Advanced Synchronization Patterns

### 9.3.1 Read-Write Locks

- Allow multiple concurrent readers but exclusive writers
- Improve throughput for read-heavy workloads
- Various policies for reader vs. writer preference

### 9.3.2 Priority Inheritance

- Addresses the problem of priority inversion
- Lower-priority thread inherits priority of higher-priority waiting thread
- Ensures critical sections are executed at appropriate priority

### 9.3.3 Non-Blocking Synchronization

- Atomic operations and compare-and-swap primitives
- Lock-free and wait-free algorithms
- Eliminates deadlocks and reduces contention

## 9.4 Synchronization Problems and Solutions

### 9.4.1 Deadlock

- **Definition:** Circular waiting where each thread holds resources needed by others
- **Prevention:** Resource ordering, single resource allocation, deadlock detection
- **Recovery:** Thread termination, resource preemption

### 9.4.2 Starvation

- **Definition:** Thread never gets resources it needs despite being eligible
- **Prevention:** Fair scheduling, aging mechanisms, resource reservation

### 9.4.3 Priority Inversion

- **Definition:** Lower-priority thread blocks higher-priority thread indirectly
- **Solutions:** Priority inheritance, priority ceiling protocol

## 9.5 Case Study: The Mars Pathfinder Incident

### Historical Note

The Mars Pathfinder mission in 1997 experienced system resets caused by a priority inversion problem. A low-priority task held a mutex needed by a high-priority task, while a medium-priority task prevented the low-priority task from completing. This real-world example illustrates the critical importance of proper synchronization in concurrent systems.

## 9.6 Synchronization in Distributed Contexts

How synchronization extends to distributed systems:

- **Challenges:** Lack of shared memory, network latency, partial failures
- **Distributed locks:** Consensus protocols, distributed mutexes
- **Time synchronization:** Logical clocks, vector clocks
- **Consistency models:** Strong vs. eventual consistency

## 9.7 Key Insights

### Key Insights

- Synchronization encompasses multiple facets beyond simple mutual exclusion
- Advanced synchronization mechanisms balance safety with performance
- Well-designed synchronization prevents problems like deadlock and starvation
- Proper synchronization is crucial for both concurrent and distributed systems



# Chapter 10

## Back to Distribution

### 10.1 Revisiting Distribution Fundamentals

Having explored concurrency in depth, we can now return to distribution with a richer understanding:

- **Distribution** is fundamentally built on concurrent execution
- **Concurrency mechanisms** provide the foundation for distributed coordination
- **Distribution** extends concurrency challenges across network boundaries

### 10.2 The Role of Runtimes in Distributed Systems

- **Layered architecture**: Multiple levels of runtime support
- **Increasing complexity**: From local concurrency to distributed coordination
- **Transparency**: Hiding distribution details from application logic
- **Failure handling**: Managing partial failures across distributed components

### 10.3 Distributed System Architecture

#### 10.3.1 Logical Architecture

- **Component model**: How functionality is divided into distributable units
- **Communication patterns**: Request-response, pub-sub, streaming
- **State management**: Centralized vs. distributed, replication strategies
- **Consistency models**: Strong, eventual, causal consistency

#### 10.3.2 Physical Architecture

- **Deployment models**: On-premises, cloud, hybrid, edge
- **Network topology**: Hierarchical, mesh, peer-to-peer
- **Resource allocation**: Static vs. dynamic, auto-scaling

## 10.4 Distribution Transparency

Types of transparency that distributed runtimes aim to provide:

- **Access transparency:** Local and remote resources accessed with the same operations
- **Location transparency:** Services can be used without knowing their physical location
- **Concurrency transparency:** Multiple users can share resources without interference
- **Replication transparency:** Multiple copies maintained without user knowledge
- **Failure transparency:** Failures are hidden or automatically recovered from
- **Migration transparency:** Resources can move without affecting operations
- **Performance transparency:** System can be reconfigured for performance without application changes
- **Scaling transparency:** System can expand in scale without affecting structure

## 10.5 Challenges in Distributed Systems

### 10.5.1 The CAP Theorem

#### Definition

The CAP theorem states that a distributed system cannot simultaneously provide all three of the following guarantees:

- **Consistency:** All nodes see the same data at the same time
- **Availability:** Every request receives a response
- **Partition tolerance:** The system continues to operate despite network failures

In practice, partition tolerance is necessary for distributed systems, forcing a trade-off between consistency and availability.

### 10.5.2 The FLP Impossibility Result

- Proves that no distributed consensus algorithm can guarantee termination in an asynchronous system with even one faulty process
- Has profound implications for distributed algorithms
- Leads to probabilistic approaches to consensus

### 10.5.3 Other Challenges

- **Network latency:** Time delay in communication affecting performance
- **Partial failures:** Some components fail while others continue operating
- **Clock synchronization:** Difficulty establishing a consistent time view
- **Security:** Increased attack surface in distributed environments

## 10.6 Distributed Runtime Components

### 10.6.1 Communication Middleware

- **Remote Procedure Call (RPC) frameworks:** gRPC, Java RMI
- **Message queues:** RabbitMQ, Apache Kafka
- **Pub/sub systems:** MQTT, NATS
- **API gateways:** Managing service access and protocols

### 10.6.2 Service Discovery

- **Registry-based:** Services register with a central registry
- **Client-side discovery:** Clients query registry and choose services
- **Server-side discovery:** Load balancer routes requests to services
- **DNS-based:** Using DNS for service lookup

### 10.6.3 Orchestration and Management

- **Container orchestration:** Kubernetes, Docker Swarm
- **Service mesh:** Istio, Linkerd
- **Monitoring and telemetry:** Prometheus, OpenTelemetry
- **Configuration management:** Centralized vs. distributed

## 10.7 Key Insights

### Key Insights

- Distribution builds on concurrency but introduces additional complexity
- Distributed runtimes provide multiple layers of abstraction and transparency
- Fundamental theoretical limits constrain what distributed systems can achieve
- Modern distributed systems rely on sophisticated runtime support for reliability and performance

# Chapter 11

## Distributed Inter-Process Communication

### 11.1 The Challenge of Remote Communication

Remote communication between distributed components differs fundamentally from local communication:

- Components exist in different address spaces
- Communication crosses network boundaries
- Network introduces latency, bandwidth constraints, and failure modes
- No shared memory for direct data access

### 11.2 Network Protocol Fundamentals

#### 11.2.1 Protocol Layers

- **Physical layer:** Hardware transmission of bits
- **Data link layer:** Reliable transmission between adjacent nodes
- **Network layer:** Routing across networks (IP)
- **Transport layer:** End-to-end communication (TCP, UDP)
- **Application layer:** Application-specific protocols (HTTP, SMTP, etc.)

#### 11.2.2 Connection-Oriented vs. Connectionless

- **Connection-oriented (TCP):** Establishes session, ensures reliability, orders messages
- **Connectionless (UDP):** No session, best-effort delivery, lower overhead

### 11.3 Remote Procedure Call (RPC)

#### 11.3.1 RPC Concept

- Makes remote service invocation appear like local procedure calls
- Hides the complexity of network communication
- Provides location transparency to applications

### 11.3.2 RPC Components

- **Client stub:** Marshals parameters, initiates remote call
- **Server skeleton:** Unmarshals parameters, invokes service, returns results
- **Interface Definition Language (IDL):** Specifies service interfaces
- **Name server:** Helps clients locate servers

### 11.3.3 RPC Operation

1. Client calls client stub (appears as normal procedure call)
2. Client stub marshals parameters into message
3. Client runtime sends message to server
4. Server runtime receives message, passes to server skeleton
5. Server skeleton unmarshals parameters, calls service implementation
6. Service executes, returns results to skeleton
7. Skeleton marshals results, returns to client
8. Client stub unmarshals results, returns to client

### 11.3.4 RPC Challenges

- **Parameter passing:** Handling complex data types, references
- **Error handling:** Network failures, server crashes
- **Idempotence:** Ensuring operations can be safely retried
- **Performance:** Overhead of marshaling, network latency

## 11.4 Message-Oriented Middleware (MOM)

### 11.4.1 Characteristics

- **Asynchronous communication model**
- **Message queues** decouple senders from receivers
- Supports various interaction patterns (point-to-point, pub/sub)
- Provides message persistence and delivery guarantees

### 11.4.2 MOM Components

- **Message queues:** Store messages until consumed
- **Message brokers:** Route messages between producers and consumers
- **Topics/exchanges:** Enable publish-subscribe patterns
- **Message acknowledgments:** Ensure delivery confirmation

### 11.4.3 Delivery Semantics

- **At-most-once:** Message may be lost but never delivered twice
- **At-least-once:** Message is guaranteed delivery but may be duplicated
- **Exactly-once:** Message is delivered once and only once (most challenging)

## 11.5 Web Services and RESTful Communication

### 11.5.1 Web Services

- **SOAP:** XML-based protocol with formal interface definitions
- **WSDL:** Web Services Description Language for interface specification
- **UDDI:** Universal Description, Discovery, and Integration for service registry

### 11.5.2 REST (Representational State Transfer)

- Architectural style using standard HTTP methods
- Resources identified by URIs
- Stateless communication model
- Multiple representation formats (JSON, XML, etc.)
- Hypermedia as the engine of application state (HATEOAS)

## 11.6 Emerging Communication Patterns

### 11.6.1 GraphQL

- Query language for APIs
- Clients specify exactly what data they need
- Single endpoint for multiple resources
- Reduces over-fetching and under-fetching of data

### 11.6.2 gRPC

- High-performance RPC framework
- Uses Protocol Buffers for interface definition and serialization
- Supports streaming (unary, server, client, and bidirectional)
- Built on HTTP/2 for multiplexing and header compression

### 11.6.3 WebSockets

- Full-duplex communication channels over TCP
- Enables real-time, bidirectional communication
- Maintains persistent connection after initial handshake
- Lower overhead than repeated HTTP requests

## 11.7 Serialization and Data Formats

### 11.7.1 Data Serialization

- Process of converting in-memory structures to transmittable format
- Must handle different data types, object references, inheritance
- Performance considerations for encoding/decoding

### 11.7.2 Common Data Formats

- **JSON**: Human-readable, widely supported, schema-optional
- **XML**: Extensible, self-descriptive, schema support
- **Protocol Buffers**: Compact binary format, schema-required, efficient
- **Avro**: Schema-based serialization with dynamic typing
- **MessagePack**: Compact binary format similar to JSON

## 11.8 Key Insights

### Key Insights

- Distributed IPC requires specialized mechanisms to handle network constraints
- Different communication patterns serve different application needs
- The choice of serialization format impacts performance and interoperability
- Modern distributed systems often combine multiple communication paradigms

## Chapter 12

# Distributed Concurrency

### 12.1 The Convergence of Distribution and Concurrency

Distributed systems inherently require concurrency at multiple levels:

- **Inter-node concurrency:** Multiple nodes operating simultaneously
- **Intra-node concurrency:** Multiple threads/processes within each node
- **Request concurrency:** Handling multiple client requests simultaneously
- **Data concurrency:** Managing concurrent access to distributed data

### 12.2 Server-Side Scalability

#### 12.2.1 Scalability Dimensions

- **Vertical scaling (scale up):** Adding more resources to existing nodes
- **Horizontal scaling (scale out):** Adding more nodes to the system
- **Functional partitioning:** Dividing the system by functionality
- **Data partitioning:** Dividing data across multiple nodes

#### 12.2.2 Server Architectures for Scalability

##### Thread-per-Request Model

- Creates a new thread for each incoming request
- Simple programming model
- Limited scalability due to thread creation overhead
- Resource consumption grows with concurrent requests

##### Thread Pool Model

- Maintains a pool of worker threads
- Reuses threads for multiple requests
- Limits maximum resource consumption
- Queues requests when all threads are busy



**Event-Driven Model**

- Single-threaded event loop processes requests asynchronously
- Non-blocking I/O operations
- Callbacks or promises/futures for handling completion
- Efficient use of resources for I/O-bound workloads
- Examples: Node.js, Twisted, nginx

**Reactor Pattern**

- Demultiplexes and dispatches requests to appropriate handlers
- Uses efficient I/O polling mechanisms (select, epoll, kqueue)
- Single-threaded or multi-threaded variations

**Proactor Pattern**

- Initiates asynchronous operations and handles their completion
- Completion handlers process results when operations finish
- More complex than reactor but potentially more efficient

## 12.3 Concurrency Control in Distributed Systems

### 12.3.1 Optimistic Concurrency Control

- Assumes conflicts are rare
- Allows operations to proceed without locking
- Validates changes before committing
- Aborts and retries if conflicts are detected

### 12.3.2 Pessimistic Concurrency Control

- Assumes conflicts are likely
- Acquires locks before operations
- Prevents conflicts but reduces concurrency
- Requires distributed lock management

### 12.3.3 Distributed Transactions

- Ensure ACID properties across multiple nodes
- Two-Phase Commit (2PC) protocol:
  - Preparation phase: All participants vote on transaction
  - Commit/abort phase: Coordinator decides based on votes
- Challenges: Blocking, coordinator failures, performance impact

## 12.4 Balancing Scalability and Consistency

### 12.4.1 Consistency Models

- **Strong consistency:** All reads reflect the most recent write
- **Eventual consistency:** System will become consistent over time
- **Causal consistency:** Causally related operations seen in same order by all nodes
- **Session consistency:** Client's operations seen in correct order within session

### 12.4.2 Consistency vs. Performance Trade-offs

- Stronger consistency typically means lower performance
- Weaker consistency enables higher scalability and availability
- Application requirements should determine appropriate consistency level

## 12.5 Distributed Concurrency Patterns

### 12.5.1 Sharding/Partitioning

- Divides data across multiple nodes
- Reduces contention and increases throughput
- Enables parallel processing of independent data
- Challenges: Cross-shard operations, rebalancing, hot spots

### 12.5.2 Command Query Responsibility Segregation (CQRS)

- Separates read (query) operations from write (command) operations
- Can use different models and storage for reads vs. writes
- Allows independent scaling of read and write sides
- Often combined with event sourcing

### 12.5.3 Event Sourcing

- Stores changes as a sequence of events
- Current state derived by replaying events
- Provides complete audit trail and temporal queries
- Enables event-driven architecture and integration

## 12.6 Distributed Concurrency Technologies

### 12.6.1 Actor Model

- Encapsulates state within actors that communicate via messages
- Natural fit for distributed systems
- Location transparency (local and remote actors treated the same)
- Examples: Erlang/OTP, Akka

### 12.6.2 Conflict-Free Replicated Data Types (CRDTs)

- Data structures that can be replicated across multiple nodes
- Automatically resolve conflicts without coordination
- Support eventual consistency with strong convergence guarantees
- Examples: G-sets, LWW-registers, OR-sets

## 12.7 Key Insights

### Key Insights

- Distributed systems require effective concurrency management at multiple levels
- Server architectures must balance resource utilization with scalability
- The choice of concurrency control mechanism affects consistency and performance
- Modern distributed systems often use specialized patterns to manage concurrency

## Chapter 13

# Distributed Synchronization

### 13.1 The Challenge of Distributed Agreement

Achieving agreement among distributed nodes is fundamentally challenging due to:

- Unreliable networks with variable latency
- Independent node failures
- Lack of global clock or shared memory
- Partial network partitions

### 13.2 Clock Synchronization

#### 13.2.1 Physical Clock Synchronization

- Network Time Protocol (NTP)
- Precision Time Protocol (PTP)
- Challenges: Network delays, clock drift, leap seconds
- Limitations: Cannot achieve perfect synchronization

#### 13.2.2 Logical Clocks

- **Lamport clocks:** Scalar timestamps establishing a partial ordering
- **Vector clocks:** Vector timestamps capturing causal relationships
- **Matrix clocks:** Track knowledge about other nodes' knowledge
- **Hybrid logical clocks:** Combine physical and logical time

## 13.3 Distributed Consensus

### 13.3.1 The Consensus Problem

#### Definition

The distributed consensus problem requires agreement among distributed nodes on a single value or sequence of values, ensuring:

- **Agreement:** All correct nodes decide on the same value
- **Validity:** The value decided was proposed by some node
- **Termination:** All correct nodes eventually decide

### 13.3.2 Consensus Algorithms

#### Paxos

- Classic algorithm for distributed consensus
- Roles: proposers, acceptors, learners
- Two-phase protocol: prepare and accept
- Guarantees safety but not liveness under certain conditions
- Complex to understand and implement correctly

#### Raft

- Designed for understandability
- Leader-based approach with terms
- Three subproblems: leader election, log replication, safety
- Strong leader principle simplifies algorithm

#### Byzantine Fault Tolerance (BFT)

- Handles malicious (Byzantine) failures
- Requires more than two-thirds of nodes to be honest
- Higher communication overhead than non-Byzantine algorithms
- Examples: PBFT, Tendermint, HotStuff

## 13.4 State Machine Replication

- Implements replicated services using consensus on operation sequence
- Each node applies the same operations in the same order
- Deterministic operations ensure consistent state
- Provides fault tolerance and high availability

## 13.5 Distributed Mutual Exclusion

### 13.5.1 Token-Based Algorithms

- Single token circulates among nodes
- Node holding token can enter critical section
- Examples: Raymond's algorithm, Suzuki-Kasami algorithm

### 13.5.2 Permission-Based Algorithms

- Node requests permission from other nodes to enter critical section
- Examples: Ricart-Agrawala algorithm, Maekawa's algorithm

### 13.5.3 Quorum-Based Algorithms

- Node must acquire locks from a quorum of nodes
- Overlapping quorums ensure mutual exclusion
- Examples: Majority quorum, tree quorum

## 13.6 Distributed Deadlock Detection

### 13.6.1 Centralized Detection

- Single coordinator collects wait-for information
- Periodically checks for cycles in global wait-for graph
- Single point of failure

### 13.6.2 Distributed Detection

- Nodes cooperate to detect deadlocks
- Edge-chasing algorithms propagate dependency information
- Challenges: Phantom deadlocks, concurrent detection

## 13.7 Distributed Termination Detection

- Determining when all distributed computation has completed
- Dijkstra-Scholten algorithm for diffusing computations
- Credit-recovery schemes
- Wave algorithms

## 13.8 Leader Election

- Selecting a unique coordinator among distributed nodes
- Ring-based algorithms (Chang-Roberts, Franklin)
- Bully algorithm
- Challenges: Network partitions, simultaneous failures

## 13.9 Key Insights

### Key Insights

- Distributed synchronization is fundamentally more challenging than local synchronization
- Consensus algorithms provide the foundation for reliable distributed systems
- Different synchronization problems require specialized distributed algorithms
- Trade-offs exist between consistency guarantees and performance

# Chapter 14

## Enter the Cloud

### 14.1 The Evolution of Cloud Computing

#### 14.1.1 Historical Context

- Mainframe era: Centralized computing with time-sharing
- Client-server era: Distributed computing with specialized servers
- Grid computing: Sharing resources across organizations
- Utility computing: Computing resources as metered services
- Cloud computing: On-demand, elastic resources with self-service capabilities

#### 14.1.2 Visionaries and Early Adopters

- John McCarthy's utility computing vision (1961)
- Amazon's infrastructure transformation (early 2000s)
- Amazon Web Services launch (2006)
- Google App Engine (2008)
- Microsoft Azure (2010)

### 14.2 Cloud Computing Characteristics

#### 14.2.1 Essential Characteristics (NIST Definition)

- **On-demand self-service:** Resources provisioned without human interaction
- **Broad network access:** Capabilities available over the network
- **Resource pooling:** Provider resources serve multiple consumers
- **Rapid elasticity:** Capabilities can be scaled quickly
- **Measured service:** Resource usage is monitored, controlled, and reported



### 14.2.2 Service Models

- **Infrastructure as a Service (IaaS):** Virtual machines, storage, networks
- **Platform as a Service (PaaS):** Application platforms, middleware, development tools
- **Software as a Service (SaaS):** Complete applications delivered over the network
- **Function as a Service (FaaS):** Serverless computing executing individual functions

### 14.2.3 Deployment Models

- **Public cloud:** Services available to general public
- **Private cloud:** Infrastructure operated for a single organization
- **Community cloud:** Infrastructure shared by several organizations
- **Hybrid cloud:** Composition of two or more distinct cloud infrastructures
- **Multi-cloud:** Using services from multiple cloud providers

## 14.3 Cloud Infrastructure Architecture

### 14.3.1 Compute Virtualization

- **Hypervisors:** KVM, Xen, VMware ESXi, Hyper-V
- **Containers:** Docker, containerd, CRI-O
- **Serverless:** Event-driven execution environments

### 14.3.2 Storage Systems

- **Object storage:** Amazon S3, Google Cloud Storage, Azure Blob Storage
- **Block storage:** Amazon EBS, Google Persistent Disk, Azure Disk Storage
- **File storage:** Amazon EFS, Google Filestore, Azure Files
- **Database services:** RDBMSs, NoSQL, NewSQL, time-series, graph

### 14.3.3 Networking

- **Virtual networks:** VPCs, subnets, security groups
- **Load balancing:** Application and network load balancers
- **Content delivery:** CDNs for global distribution
- **DNS services:** Route 53, Cloud DNS

## 14.4 Cloud Design Patterns

### 14.4.1 Scalability Patterns

- **Auto-scaling:** Automatically adjusting resources based on demand
- **Queue-based load leveling:** Smoothing workloads with message queues
- **Static content hosting:** Offloading static content to specialized storage
- **Sharding:** Horizontally partitioning data across multiple databases

### 14.4.2 Resilience Patterns

- **Circuit breaker:** Preventing cascading failures
- **Bulkhead:** Isolating failures to parts of the system
- **Retry with exponential backoff:** Handling transient failures
- **Chaos engineering:** Proactively testing system resilience

### 14.4.3 Data Management Patterns

- **CQRS:** Separating read and write operations
- **Event sourcing:** Storing state changes as event sequences
- **Materialized view:** Pre-computing query results
- **Polyglot persistence:** Using multiple data storage technologies

## 14.5 Cloud-Native Application Architecture

### 14.5.1 Microservices

- Small, autonomous services focused on specific business capabilities
- Independent deployment and scaling
- Technology diversity (polyglot programming)
- Challenges: Distributed system complexity, service coordination

### 14.5.2 Containers and Orchestration

- **Containers:** Lightweight, consistent runtime environments
- **Kubernetes:** Container orchestration platform
- **Service mesh:** Istio, Linkerd for service-to-service communication
- **Serverless:** Event-driven, auto-scaling compute without infrastructure management

### 14.5.3 DevOps and Continuous Delivery

- **Infrastructure as Code:** Terraform, CloudFormation
- **CI/CD pipelines:** Automated testing and deployment
- **Observability:** Metrics, logging, tracing
- **GitOps:** Git-based infrastructure and application deployment

## 14.6 Challenges in Cloud Computing

### 14.6.1 Technical Challenges

- **Data gravity:** Difficulty moving large datasets
- **Distributed system fallacies:** Network reliability, latency, etc.
- **Multi-cloud complexity:** Managing across cloud providers
- **Security and compliance:** Data protection, regulatory requirements

### 14.6.2 Business Challenges

- **Cost management:** Understanding and controlling cloud spending
- **Vendor lock-in:** Dependency on specific cloud provider services
- **Skills gap:** Finding expertise in cloud technologies
- **Organizational change:** Adapting processes for cloud-native operations

## 14.7 Future Trends

- **Edge computing:** Processing closer to data sources
- **Serverless evolution:** Expanding beyond functions to more complex workloads
- **AI/ML integration:** Embedding intelligence in cloud services
- **Multi-cloud and hybrid clouds:** Standardization and interoperability
- **Sustainable computing:** Energy-efficient and environmentally conscious operations

## 14.8 Key Insights

### Key Insights

- Cloud computing represents the culmination of distributed computing evolution
- Cloud platforms rely on sophisticated runtimes for virtualization and orchestration
- Cloud-native design requires rethinking application architecture for distributed environments
- Understanding concurrency and distribution fundamentals is crucial for effective cloud computing

# Bibliography

- [1] Tanenbaum, A. S., & Van Steen, M. (2017). *Distributed Systems: Principles and Paradigms (3rd Edition)*. CreateSpace Independent Publishing Platform.
- [2] Coulouris, G., Dollimore, J., Kindberg, T., & Blair, G. (2011). *Distributed Systems: Concepts and Design (5th Edition)*. Addison-Wesley.
- [3] Burns, B. (2018). *Designing Distributed Systems: Patterns and Paradigms for Scalable, Reliable Services*. O'Reilly Media.
- [4] Kleppmann, M. (2017). *Designing Data-Intensive Applications*. O'Reilly Media.
- [5] Ongaro, D., & Ousterhout, J. (2014). *In Search of an Understandable Consensus Algorithm*. In USENIX Annual Technical Conference (pp. 305-319).
- [6] Lamport, L. (1998). *The Part-Time Parliament*. ACM Transactions on Computer Systems (TOCS), 16(2), 133-169.
- [7] Van Roy, P., & Haridi, S. (2007). *Concepts, Techniques, and Models of Computer Programming*. MIT Press.
- [8] Fielding, R. T. (2000). *Architectural Styles and the Design of Network-based Software Architectures*. Doctoral dissertation, University of California, Irvine.
- [9] Mell, P., & Grance, T. (2011). *The NIST Definition of Cloud Computing*. National Institute of Standards and Technology.