

Introduction to Distributed Systems

IN5020/9020
Lecturer: Amir Taherkordi

August 21, 2023

UiO  University of Oslo

Outline

- **Definition** of a distributed system
- **Goals** of a distributed system
- **Implications** of distributed systems
- **Pitfalls** in developing distributed systems
- **Types** of distributed systems

From Single Computers to Distributed Systems

- 1945-1985: Computers
 - Large and expensive
 - Operated independently
- 1985-now: two advances in technology
 - Powerful microprocessors
 - High-speed computer networks
- Result: **Distributed Systems**

Putting together **computing** systems composed of a large number of **computers connected** by a high-speed network

What Is a Distributed System?

■ Definition

Operational perspective:

A distributed system is one in which hardware or software **components**, located at **networked** computers, **communicate and coordinate** their actions only by **passing messages**.

[Coulouris]

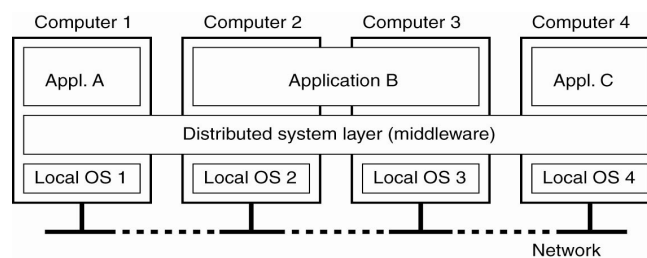
User perspective:

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

[Tanenbaum]

What Is a Distributed System?

- A DS is organized as middleware
 - a software layer placed between users and applications, and operating systems and networks
 - extending over multiple machines
 - offering each application the same interface



IN5020, IFI/UiO

5

Examples of Distributed Systems

- Web search
 - Indexing the entire contents of the Web
- Large scale retailers
 - Storing and accessing transactions
- Massively multiplayer online games
 - Very large number of users sharing a virtual world
- Financial trading
 - Real time access and processing of a wide range of information sources
 - Delivery of items of interest in a timely manner
 - Blockchain
- More examples?

IN5020, IFI/UiO

6

What is not a distributed system?

- Coffee shop example
- Distribution brings complexity
 - Important to know why you are distributing



Outline

- Definition of a distributed system
- Goals of a distributed system
- Implications of distributed systems
- Pitfalls in developing distributed systems
- Types of distributed systems

Goals of Distributed Systems

- Resource sharing
- Distribution transparency
- Openness
- Scalability
- Fault tolerance
- Allowing heterogeneity (interoperability)

Resource Sharing

- Making resources **accessible**:
 - accessing remote resources
 - sharing them in a controlled and efficient way
- Examples: printers, storages, files, etc.
- **Resource managers** control access, offer a scheme for naming, and control concurrency
- A **resource sharing model** describes how
 - resources are made available
 - resources can be used
 - service provider and user interact with each other

Models for Resource Sharing

- **Client-server** resource model (often **resource-based**)
 - Server processes act as resource managers, and offer services (collection of procedures)
 - Client processes send requests to servers
 - HTTP implements a client-server resource model
- **Object-based** resource model
 - Any **shared resource** is modeled as an **object**
 - Object provides access to its operations with a **message based interface**
 - Object-based middleware (CORBA, Java RMI) defines object-based resource model

Distribution Transparency

- An important goal of a DS:
 - hiding the fact that its **processes** and **resources** are physically distributed across multiple computers
- Definition:
 - A distributed system that is able to present itself to its users and applications as if it were only a single computer system is said to be **transparent**.
- What kind of transparency?

Forms of Transparency

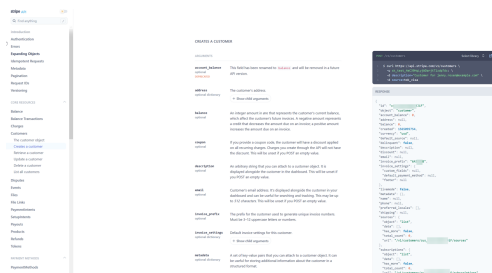
Transparency	Description
Access	Hide differences in data representation and how a resource is accessed
Location	Hide where a resource is located
Migration	Hide that a resource may move to another location
Relocation	Hide that a resource may be moved to another location while in use
Replication	Hide that a resource is replicated
Concurrency	Hide that a resource may be shared by several competitive users
Failure	Hide the failure and recovery of a resource

- Degree of transparency
 - Situations in which full transparency is not good – better to expose than to mask effects of distribution?
 - Trade-off between a high degree of transparency and performance

Openness

- Definition:

An **open** distributed system is a system that offers **services** according to **standard rules** that describe **syntax** and **semantics** of those services.
- E.g., in computer networks: format, content, and meaning of messages
- In DS: services specified through interfaces
 - Interface Definition Language (IDL): capturing syntax
 - Web Service Definition Language (WSDL): capturing syntax
 - Semantics? an informal way
- Extensibility: an open DS can be extended and improved incrementally
 - add or replace components



Scalability

- **Scalability** denotes the ability of a system to handle an increasing future load
- A **scalable** system: remains effective when there is a significant increase in the amount of resources (data) and number of users, e.g.
 - Internet: number of users and services has grown enormously
 - Google: scaled over the years to handle $O(100)$ billion queries a month, expected query time 0.2 sec.
- **Scalability in three dimensions:**
 - **in size:** Users and resources can easily be added
 - More servers, more databases, etc.
 - **geographically scalable:** Users and resources may lie far apart
 - Coffee shop in the US and Europe.
 - **administratively scalable:** The system spans many administrative organizations or persons.
 - A org. function is handled by another system. Less system admins

Scalability Problems

- Problems with **size** scalability:
 - Often caused by centralized solutions

Concept	Example
Centralized services	A single server for all users
Centralized data	A single on-line telephone book
Centralized algorithms	Doing routing based on complete information

- Problems with **geographical** scalability:
 - traditional synchronous communication in LAN
 - unreliable communications in WAN
- Problems with **administrative** scalability:
 - Conflicting policies, complex management, security problems

Scaling Techniques


- **Replication**
 - replicate resources (services, data) across the system
 - increases availability, helps to balance load
 - caching (special form of replication)
- **Sharding (Distribution)**
 - splitting a resource (such as data) into smaller parts, and spreading the parts across the system (cf DNS)
 - Consistent hashing
 - Uniform distribution in sharding.
- **Hiding communication latencies**
 - avoid waiting for responses to remote service requests (use asynchronous communication or reduce the number of remote requests)

Fault Tolerance

- Hardware, software and network fail!!
- DS must maintain **availability** even in cases where hardware/software/network have **low** reliability
- Failures in distributed systems are partial
 - makes error handling particularly difficult
- Many techniques for handling failures
 - Detecting failures (checksum)
 - Masking failures (retransmission in protocols)
 - Tolerating failures (as in web-browsers, 404)
 - Recovery from failures (roll back)
 - Redundancy (replicate servers in failure-independent ways)

Example: Google File System

- Goals: Fault tolerance



IN5020, IFI/UiO

19

Outline

- Definition of a distributed system
- Goals of a distributed system
- **Implications** of distributed systems
- **Pitfalls** in developing distributed systems
- Types of distributed systems

IN5020, IFI/UiO

20

Implications of Distributed Systems

- **Concurrency**
 - components execute in concurrent processes that read and update shared resources. Requires coordination, (race conditions)
- **No global clock**
 - makes coordination difficult (ordering of events)
- **Independent failure of components**
 - “partial failure” & incomplete information
- **Unreliable communication**
 - Loss of connection and messages. Message bit errors
- **Unsecure communication**
 - Possibility of unauthorised recording and modification of messages
- **Expensive communication**
 - Compared to independent processes on the same computer: communication between computers usually has **less bandwidth, longer latency**, and costs more

Pitfalls When Developing DS

- False assumptions made by first time developer:
 - The network is **reliable**.
 - The network is **secure**.
 - The network is **homogeneous**.
 - The topology does **not change**.
 - Latency is **zero**.
 - Bandwidth is **infinite**.
 - Transport cost is **zero**.
 - There is **one administrator**.

Outline

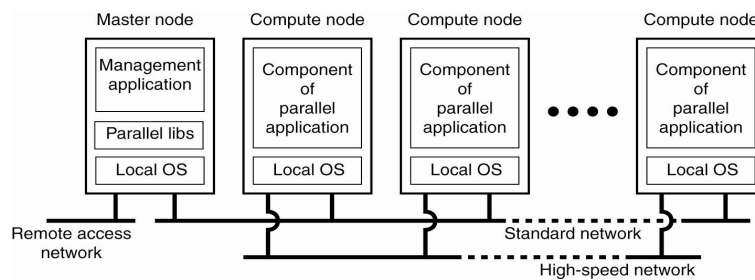
- Definition of a distributed system
- Goals of a distributed system
- Implications of distributed systems
- Pitfalls in developing distributed systems
- Types of distributed systems

Types of Distributed Systems

- Type1: Distributed **Computing** Systems
 - Used for high performance computing tasks
 - Cluster and Cloud computing systems
 - Grid computing systems
 - MapReduce, Hadoop, Storm, Ethereum
- Type2: Distributed **Information** Systems
 - Systems mainly for management and integration of business functions
 - Transaction processing systems
 - Enterprise application integration
 - NSQ, Stripe, Amazon
- Type3: Distributed **Pervasive** (or **Ubiquitous**) Systems
 - Mobile and embedded systems
 - Home systems
 - Sensor networks, IoT
 - Ikea Trådfri, Philips hue, Google NEST

Type1: Cluster Computing Systems


- Collection of similar PCs, closely connected, all run same OS



- A collection of **computing nodes** + **master node**
- Master runs **middleware**: parallel execution and management

Type1: Grid Computing Systems

- Federation of autonomous and heterogeneous computer systems (HW,OS,...), several adm. domains

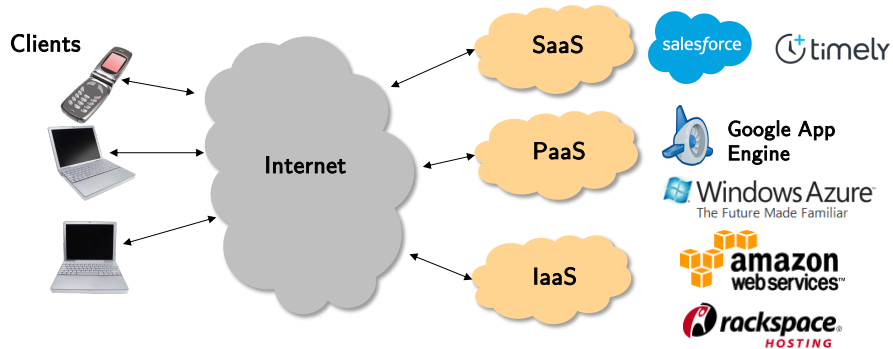


A layered architecture for grid computing systems

- Example: SETI@HOME

Type1: Cloud Computing


- View: distributed resources as a commodity or utility
 - Resources are provided by service suppliers and effectively rented rather than owned by the end user.
- The term **cloud computing** capture the vision of computing as a utility



- Cloud vs. Grid?

Type2: Enterprise Application Integration

- Allowing existing applications to directly exchange information using communication middleware



- Middleware as a communication facilitator in enterprise application integration, e.g.
 - RabbitMQ, Kafka, RocketMQ


Type3: Distributed Pervasive Systems

- Pervasive systems:

exploiting the increasing integration of **services and (small/tiny) computing devices** in our everyday **physical world**

- (Mobile) Devices: discover the environment (its services) and establish themselves in this environment as best as possible.
- Requirements for pervasive applications
 - Embrace contextual changes
 - Encourage ad-hoc and dynamic composition
 - Recognize sharing as the default

Type3: Example: Smart Home System



Summary

- Distributed systems:
 - components located in a network that communicates and coordinates their actions exclusively by sending messages.
- Goals like resource sharing, distribution transparency, openness, scalability, fault tolerance and heterogeneity can be satisfied by distributed systems
- Consequences of distributed systems
 - Independent failure of components
 - Unsecure communication
 - No global clock
- Many pitfalls when developing distributed systems
- Three main types of distributed systems