




中国科学院深圳先进技术研究院
SHENZHEN INSTITUTE OF ADVANCED TECHNOLOGY
CHINESE ACADEMY OF SCIENCES

Distributed System Models



Dr. Minxian Xu
Associate Professor
Research Center for Cloud Computing
Shenzhen Institute of Advanced Technology, CAS
<http://www.minxianxu.info/dcp>

Most concepts are
drawn from Chapter 2
© Pearson Education

千磨万击还坚劲，任尔东西南北风。
——（清）郑板桥


Review



- Q1: When do we need multi-threading?

Review

- Allow multiple operations performed at the same time to achieve concurrency at the application level
- Improve application performance
- Improve resource utilization rate



■ Application scenarios

- Implementation of reactive user interfaces.
- Handling concurrent requests and serving multiple clients
- Non-blocking I/O operations.
- Asynchronous behaviour.
- Timer and alarms implementation

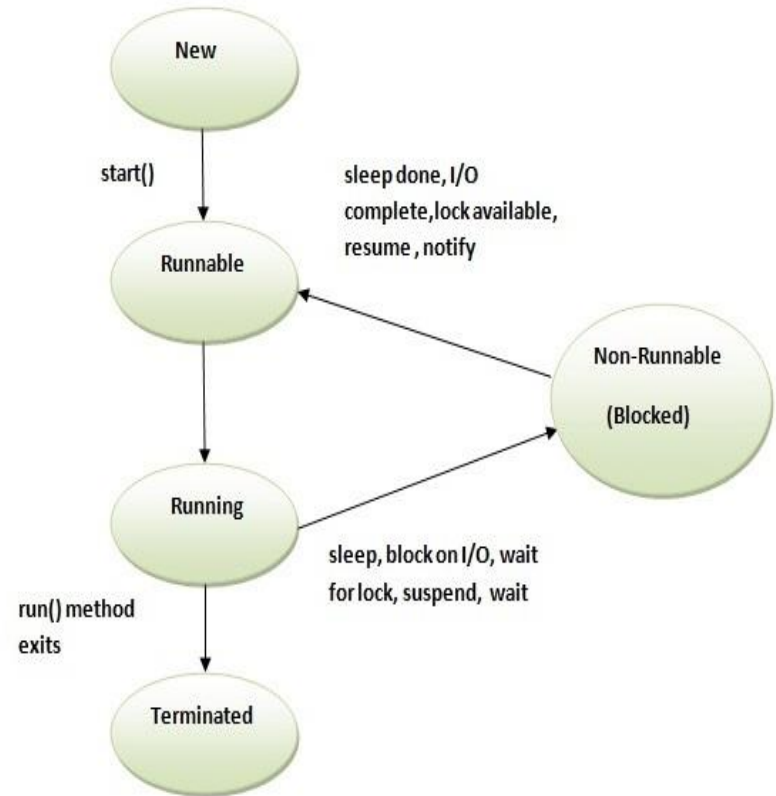
Review



- Q2: What is a thread and life cycle of a thread?

Review

- A Thread is a piece of code that runs in concurrent with other threads.
- Each thread is a statically ordered sequence of instructions.
- Threads are used to express concurrency on both single and multiprocessors machines.
- Threads share the same address space and therefore can share both data and code



Review



- Q3: What is synchronous access to shared resources and how can we achieve it?

Review

- If one thread tries to read the data and other thread tries to update the same data, it leads to inconsistent state.
- This can be prevented by synchronising access to the data.
- Use “synchronized” method:

```
public synchronized void update()  
{  
    ...  
}
```


Review



- Q4: Compare the worker pool multi-threading architecture with the thread-per-request architecture.

■ **Worker pool architecture**

- The server creates a fixed number of threads called a worker pool. As requests arrive at the server, they are put into a queue by the I/O thread and from there assigned to the next available worker thread.
- Server creates worker pool → request comes in, put into a queue → assigned to an available worker thread.

■ **Thread-per-request architecture**

- Thread created for each request, when the request is finished, the thread is deallocated.

Presentation Outline

- Introduction
- Physical Models:
 - Three Generations of DS: Early, Internet-Scale, Contemporary
- Architectural Models
 - Software Layers
 - System Architectures
 - Client-Server
 - Clients and a Single Server, Multiple Servers, Proxy Servers with Caches, Peer Model
 - Alternative Client-Sever models driven by:
 - Mobile code, mobile agents, network computers, thin clients, mobile devices, and spontaneous networking
 - Design Challenges/Requirements
- Fundamental Models – formal description
 - Interaction, failure, and security models.
- Summary

- Distributed systems should be designed to function **correctly** in **ALL circumstances/scenarios**.
- Distributed system models helps in...
 - ..classifying and understanding different implementations
 - ..identifying their weaknesses and their strengths
 - ..crafting new systems outs of pre-validated building blocks
- We will study distributed system models from different perspectives
 - Structure, organization, and placement of components
 - Interactions
 - Fundamental properties of systems

Characterization

- The structure and the organization of systems and the relationship among their components should be designed with the following goals in mind:
 - To cover the widest possible range of circumstances.
 - To cope with possible difficulties and threats.
 - To meet the current and possibly the future demands.
 - Architectural models provide both:
 - a pragmatic starting point
 - a conceptual view
- to address these challenges.

In terms of implementation models and basic blocks

In terms of logical view of the system, interaction flow, and components

Characterization: Challenges

(Difficulties and Threats)



- **Widely varying models of use**
 - High variation of workload, partial disconnection of components, or poor connection.
- **Wide range of system environments**
 - Heterogeneous hardware, operating systems, network, and performance.
- **Internal problems**
 - Non synchronized clocks, conflicting updates, various hardware and software failures.
- **External threats**
 - Attacks on data integrity, secrecy, and denial of service.



Characterization: Dealing with Challenges



- **Widely varying models of use**
 - The structure and the organization of systems allow for distribution of workloads, redundant services, and high availability.
- **Wide range of system environments**
 - A flexible and modular structure allows for implementing different solutions for different hardware, OS, and networks.
- **Internal problems**
 - The relationship between components and the patterns of interaction can resolve concurrency issues, while structure and organization of component can support failover mechanisms.
- **External threats**
 - Security has to be built into the infrastructure and it is fundamental for shaping the relationship between components.



Models at a Glance

Physical, Architectural, and
Fundamental Models

Physical Models

- A representation of the underlying H/W elements of a DS that abstracts away specific details of the computer/networking technologies.
- Baseline physical model – a small set of nodes.
- Three Generations of DSs (Distributed Systems):
 - Early DSs [70-80s]: LAN-based, 10-100 nodes
 - Internet-scale DSs [early 90-2010]: Clusters, Grids, P2P (with autonomous nodes)
 - Contemporary DSs: dynamic nodes in **Mobile Systems** that offer location-aware services, **Clouds** with resource pools offering services on pay-as-you-go basis, and **Internet of Things (IoT)** (seamless interaction between physical and cyber world for smart * applications such as Smart Health and Smart Cities) Daily visits of ChatGPT is **70 million** at Dec 2023, requiring **512,000** A100 HGX servers if deployed as Google Search

- An Architectural model of a distributed system is concerned with the placement of its parts and relationship between them. Examples:
 - Client-Server (CS) and Peer Process models.
 - CS can be modified by:
 - The partitioning of data/replication at cooperative servers
 - The caching of data by proxy servers or clients
 - The use of mobile code and mobile agents
 - The requirements to add or remove mobile devices.

Fundamental Models

- Fundamental Models are concerned with a **formal description** of the properties that are **common in all** of the architectural models
- Models addressing time synchronization, message delays, failures, security issues are addressed as:
 - **Interaction Model** – deals with performance and the difficulty of setting of time limits in a distributed system.
 - **Failure Model** – specification of the faults that can be exhibited by processes
 - **Security Model** – discusses possible threats to processes and communication channels.

Presentation Outline



- Introduction
- Architectural Models
 - Software Layers
 - System Architectures
 - Client-Server
 - Clients and a Single Server, Multiple Servers, Proxy Servers with Caches, Peer Model
 - Alternative Client-Server models driven by:
 - Mobile code, mobile agents, network computers, thin clients, mobile devices and spontaneous networking
 - Design Challenges/Requirements
- Fundamental Models – formal description
 - Interaction, Failure, and Security models.
- Summary

Architectural Models – Intro [1]

- The architecture of a system is its structure in terms of separately specified components.
 - Its goal is to meet present and likely future demands.
 - Major concerns are making the system reliable, manageable, adaptable, and cost-effective.
- Architectural Model:
 - **Simplifies and abstracts** the functions of individual components
 - The **placement of the components** across a network of computers – define patterns for the distribution of data and workloads
 - The **interrelationship** between the components – ie., functional roles and the patterns of communication between them.

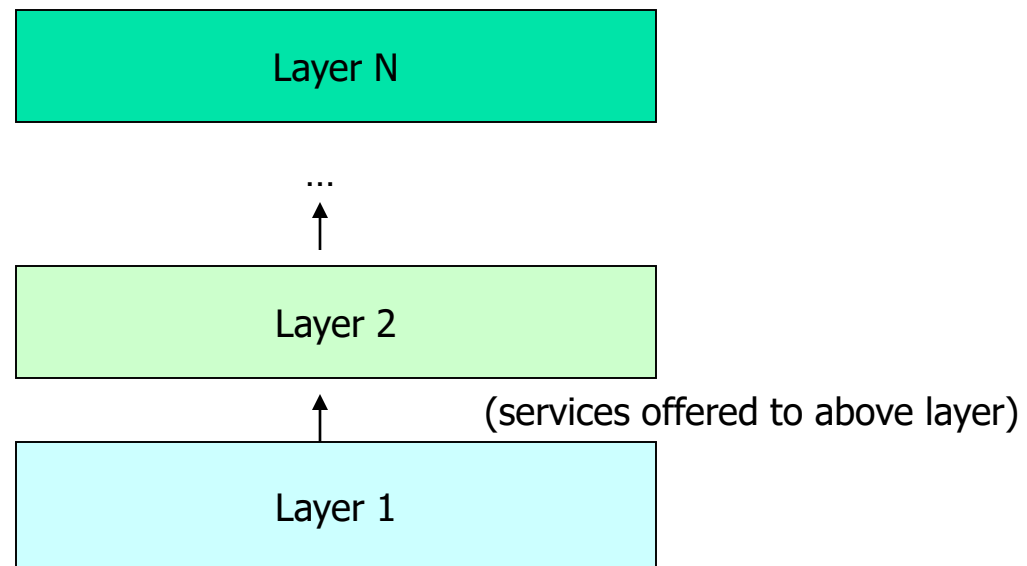
Architectural Models – Intro [2]

- Architectural Model - simplifies and abstracts the functions of individual components:
 - An initial simplification is achieved by classifying processes as:
 - Server processes
 - Client processes
 - Peer processes
 - Cooperate and communicate in a symmetric manner to perform a task.

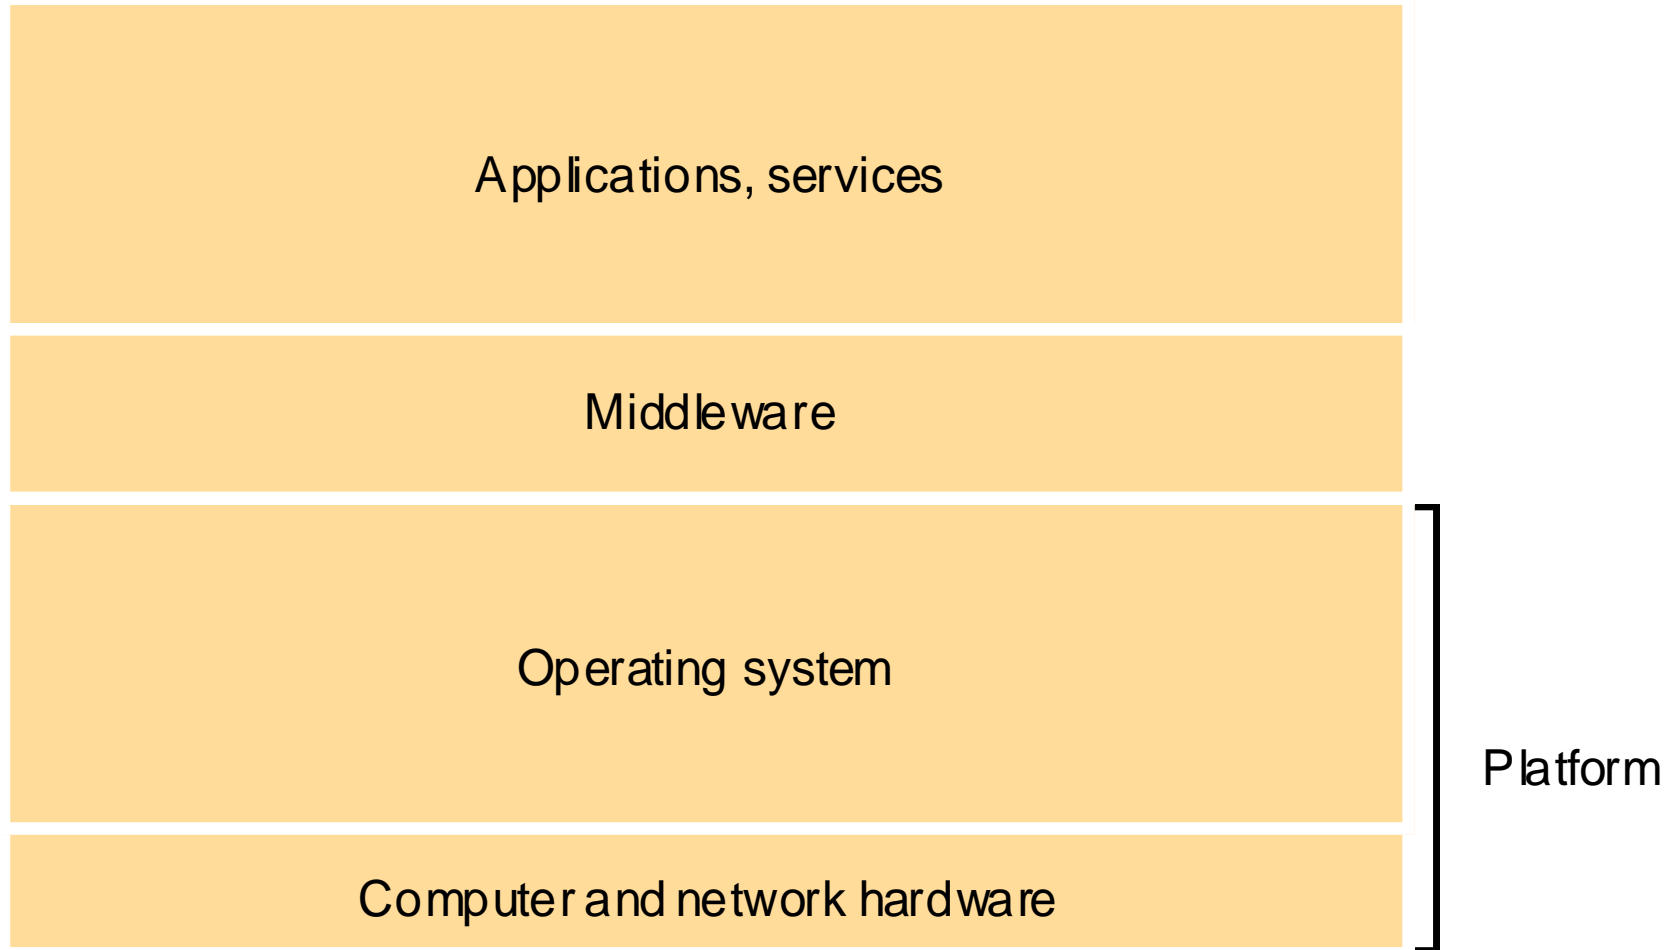


Software Architecture and Layers

- The term *software architecture* referred:
 - Originally to the structure of software as *layers* or modules in a single computer.
 - More recently in terms of *services* offered and requested between processes in the same or different computers.
- Breaking up the complexity of systems by designing them through layers and services
 - Layer: a group of related functional components
 - Service: functionality provided to the next layer.



Software and hardware service layers in distributed systems




Platform

- The lowest hardware and software layers are often referred to as a platform for distributed systems and applications.
- These low-level layers provide services to the layers above them, which are implemented independently in each computer.
- Major Examples
 - Intel x86/Windows
 - Intel x86/Linux
 - Intel x86/Solaris
 - PowerPC/MacOS
 - iPhone/iOS
 - Samsung Galaxy/Android/HarmonyOS

Middleware

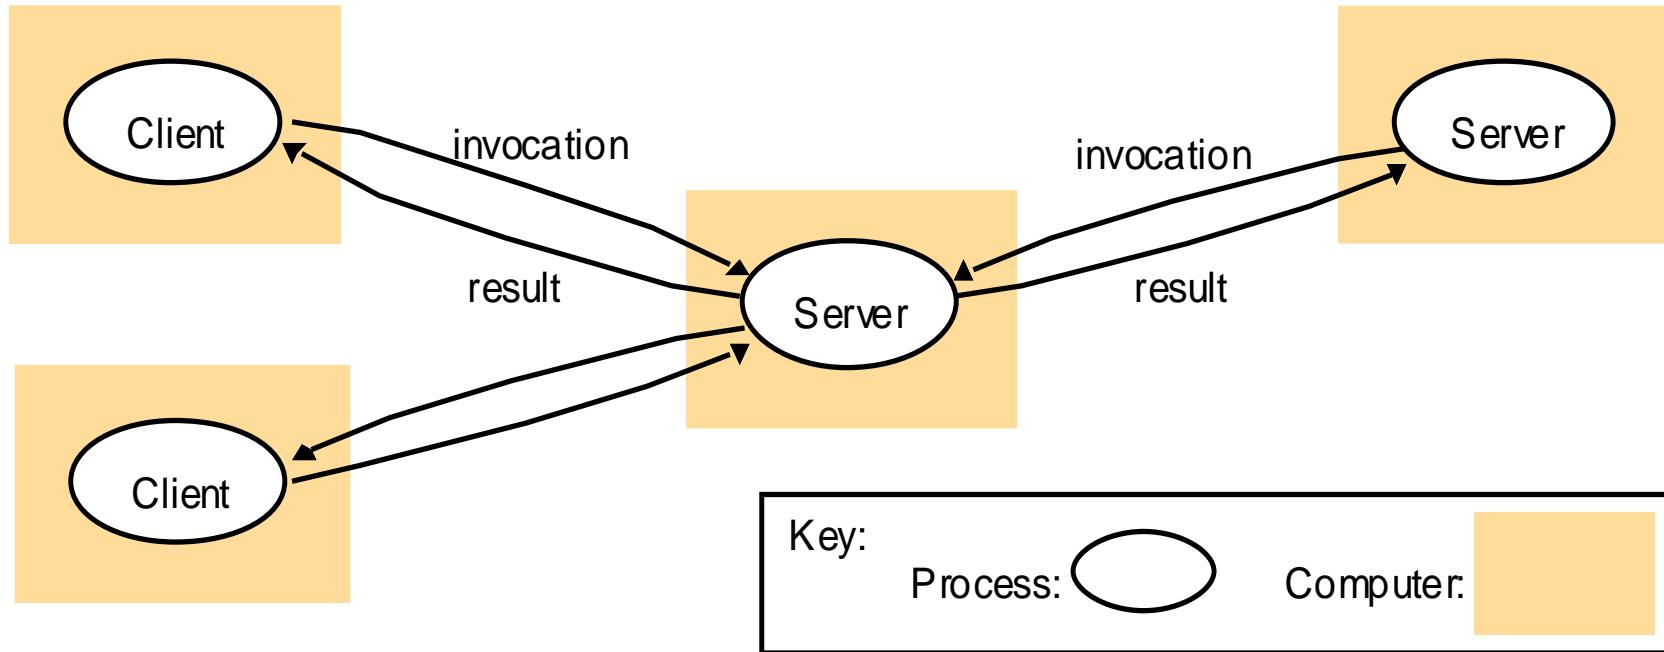
- A layer of software whose purpose is to **mask heterogeneity** present in distributed systems and to provide a convenient programming model to application developers.
- Major Examples:
 - Sun RPC (Remote Procedure Call)
 - OMG CORBA (Common Object Request Broker Architecture)
 - Microsoft D-COM (Distributed Components Object Model)
 - Sun Java RMI (Remote Method Invocation)
 - Modern Middleware Examples:
 - Manjrasoft Aneka– for Cloud computing
 - IBM WebSphere
 - Microsoft .NET
 - Sun J2EE
 - Google AppEngine
 - Microsoft Azure



- The most evident aspect of DS design is the division of responsibilities between system components (applications, servers, and other processes) and the placement of the components on computers in the network.
- It has major implication for:
 - Performance, reliability, and security of the resulting system.

Client-Server Basic Model:

Clients invoke individual servers




- Client processes interact with individual server processes in a separate computer in order to access data or resource. The server in turn may use services of other servers.
- Example:
 - A Web Server is often a client of file server.
 - Browser → search engine → crawlers → other web servers.

Client-Server Architecture Types


(Tier arch compliments layer architecture)




- Two-tier model (classic)



- Three-tier (when the server, becomes a client)




- Multi-tier (cascade model)



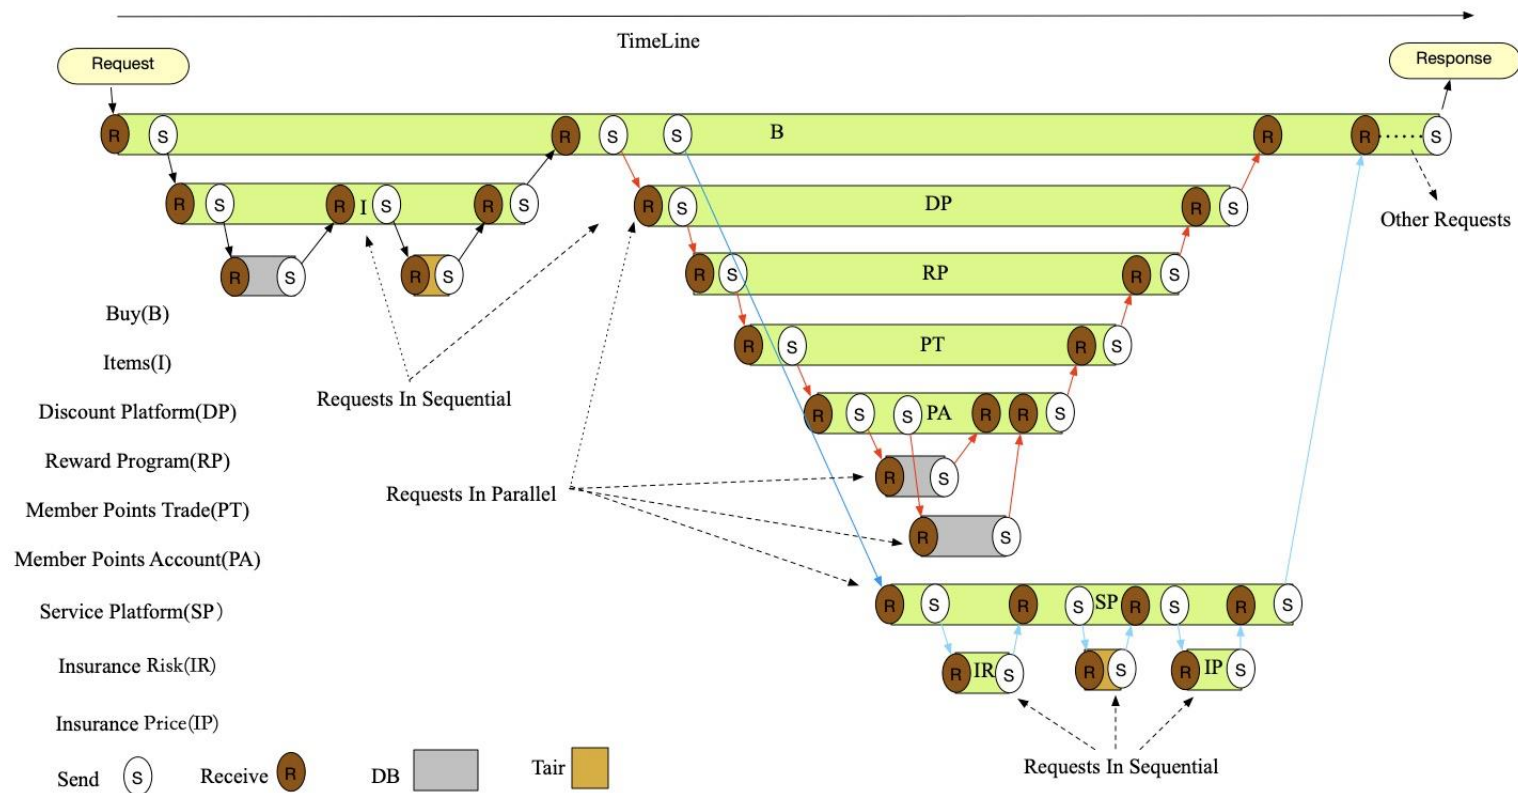
Clients and Servers

- General interaction between a client and a server.




Clients and Servers

- Complex interaction between a client and a server (Alibaba scenario).




Minxian Xu, et al., Practice of Alibaba Cloud on Elastic Resource Provisioning for Large-scale Microservices Cluster, *Software: Practice and Experience (SPE)*, Vol. 54, Issue 1, pp. 39-57, 2024.

A service provided by multiple servers

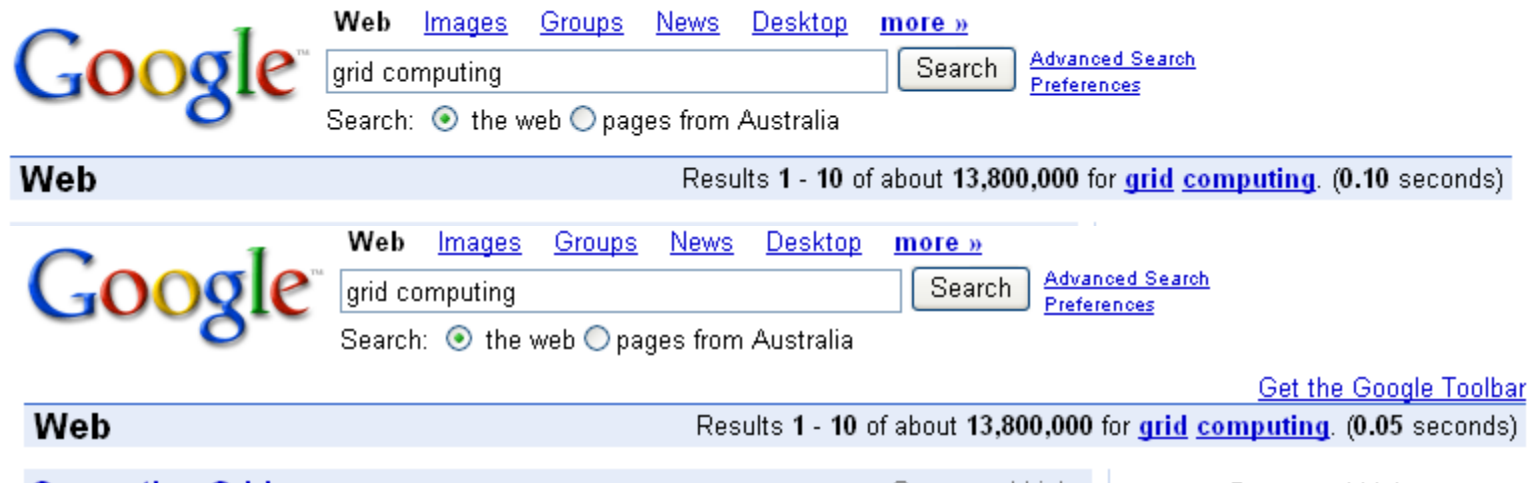


- Services may be implemented as several server processes in separate host computers.
- Example: Cluster based Web servers and apps such as Google, parallel databases Oracle

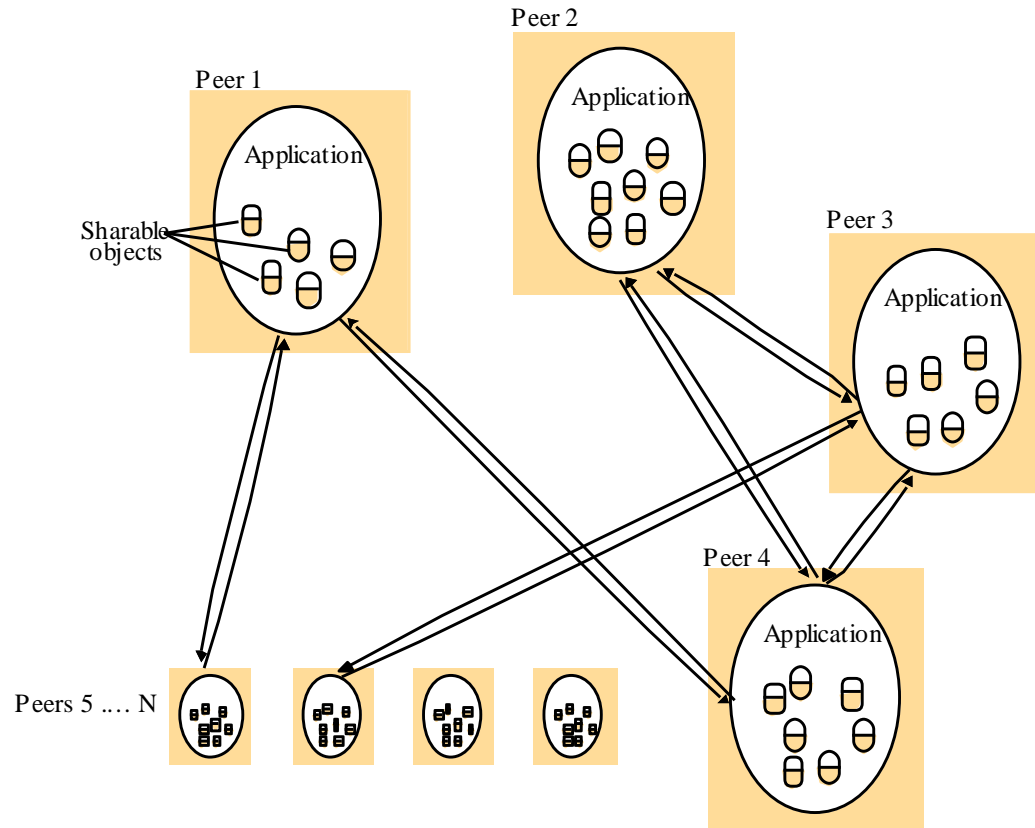
Proxy servers (replication transparency) and caches: Web proxy server



- A cache is a store of recently used data.




Peer Processes: A distributed application based on peer processes



- All of the processes play similar roles, interacting cooperatively as peers to perform distributed activities or computations without distinction between clients and servers. E.g., music sharing systems Napster, Gnutella, Kaza, BitTorrent.
- Distributed “**white board**” – users on several computers to view and interactively modify a picture between them.


P2P with a Centralized Index Server (e.g. Napster Architecture)



Variants of Client Sever Model: Mobile code and Web applets



a) client request results in the downloading of applet code



b) client interacts with the applet




- Applets downloaded to clients give good interactive response
- Mobile codes such as Applets are potential security threat, so the browser gives applets limited access to local resources (e.g. NO access to local/user file system).

Variants of Client Sever Model: Mobile Agents



- A running program (code and data) that travels from one computer to another in a network carrying out an autonomous task, usually on behalf of some other process
 - advantages: flexibility, savings in communications cost
 - virtual markets, software maintain on the computers within an organisation.
- Potential security threat to the resources in computers they visit. The environment receiving agent should decide which of the local resource to allow. (e.g., crawlers and web servers).
- Agents themselves can be vulnerable – they may not be able to complete task if they are refused access.
- **Example technology:**
 - **Java Agent Development Framework (JADE)**


Thin clients and compute servers



- **Network computer:** download OS and applications from the network and run on a desktop (solve up-gradation problem) at runtime.
- **Thin clients:** Windows-based UI on the user machine and application execution on a remote computer. E.g, X-11 system.

Mobile devices and spontaneous networking


[3rd Generation Distributed System]




- The world is increasingly populated by small and portable computing devices.
- W-LAN needs to handle constantly changing heterogeneous, roaming devices
- Need to provide discovery services: (1) *registration service* to enable servers to publish their services and (2) *lookup service* to allow clients to discover services that meet their requirements.

Distributed Deep Reinforcement Learning

- Deep reinforcement learning needs to address data inefficiency, exploration-exploitation trade-off, and multi-task learning
- Distributed modifications of DRL
 - Agents could be run on many machines simultaneously



GORILA Architecture



A3C Architecture

Summary - Models and Implications

- The use of CS (Client-Server) has impact on the software architecture followed:
 - Distribution of responsibilities
 - Synchronization mechanisms between client and server
 - Admissible types of requests/responses
- Basic CS model, responsibility is statically allocated.
 - File server is responsible for file, not for web pages.
- **Peer Process model**, responsibility is dynamically allocated:
 - In fully decentralized music file sharing system, search process may be delegated to different peers at runtime.

Design Requirements/Challenges of Distributed Systems



- **Performance Issues**

- Responsiveness
 - Support interactive clients
 - Use caching and replication
- Throughput
- Load balancing and timeliness

- **Quality of Service:**

- Reliability
- Security
- Adaptive performance.

- **Dependability issues:**

- Correctness, security, and fault tolerance
- Dependable applications continue to work in the presence of faults in hardware, software, and networks.

Presentation Outline

- Introduction
- Architectural Models
 - Software Layers
 - System Architectures
 - Client-Server
 - Clients and a Single Server, Multiple Servers, Proxy Servers with Caches, Peer Model
 - Alternative Client-Server models driven by:
 - Mobile code, mobile agents, network computers, thin clients, mobile devices and spontaneous networking
 - Design Challenges/Requirements
- Fundamental Models – formal description
 - Interaction, Failure, and Security models.
- Summary

Fundamental Models at Glance

- Fundamental Models are concerned with a **formal description** of the properties that are **common in all** of the architectural models
- All architectural models are composed of processes that communicate with each other by sending messages over a computer networks.
- Models addressing time synchronization, message delays, failures, security issues are addressed as:
 - **Interaction Model** – deals with performance and the difficulty of setting of time limits in a distributed system.
 - **Failure Model** – specification of the faults that can be exhibited by processes
 - **Security Model** – discusses possible threats to processes and communication channels.

Interaction Model

- Computation occurs within processes;
- The processes interact by passing messages, resulting in:
 - Communication (information flow)
 - Coordination (synchronization and ordering of activities) between processes.
- Two significant factors affecting interacting processes in a distributed system are:
 - Communication performance is often a limiting characteristic.
 - It is impossible to maintain a single global notion of time.

Performance of Communication Channel

- The communication channel in our model is realised in a variety of ways in DSs. E.g., by implementation of:
 - Streams
 - Simple message passing over a network.
- Communication over a computer network has performance characteristics:
 - Latency:
 - A delay between the start of a message's transmission from one process to the beginning of reception by another.
 - Bandwidth:
 - the total amount of information that can be transmitted over in a given time. Eg. **100 Mbps** (*megabits per second*)
 - Communication channels using the same network, have to share the available bandwidth.
 - Jitter
 - The variation in the time taken to deliver a series of messages. It is very relevant to multimedia data.

Interaction Model:

Computer clocks and timing events

- Each computer in a DS has its own internal clock, which can be used by local processes to obtain the value of the current time.
- Therefore, two processes running on different computers can associate timestamp with their events.
- However, even if two processes read their clocks at the same time, their local clocks may supply different time.
 - This is because computer clock drifts from perfect time and their drift rates differ from one another.
- Even if the clocks on all the computers in a DS are set to the same time initially, their clocks would eventually vary quite significantly unless corrections are applied.
 - There are several techniques to correct time on computer clocks. For example, computers may use radio receivers to get readings from GPS (Global Positioning System) with an accuracy about 1 microsecond.

Interaction Model:

Two variants of the interaction model

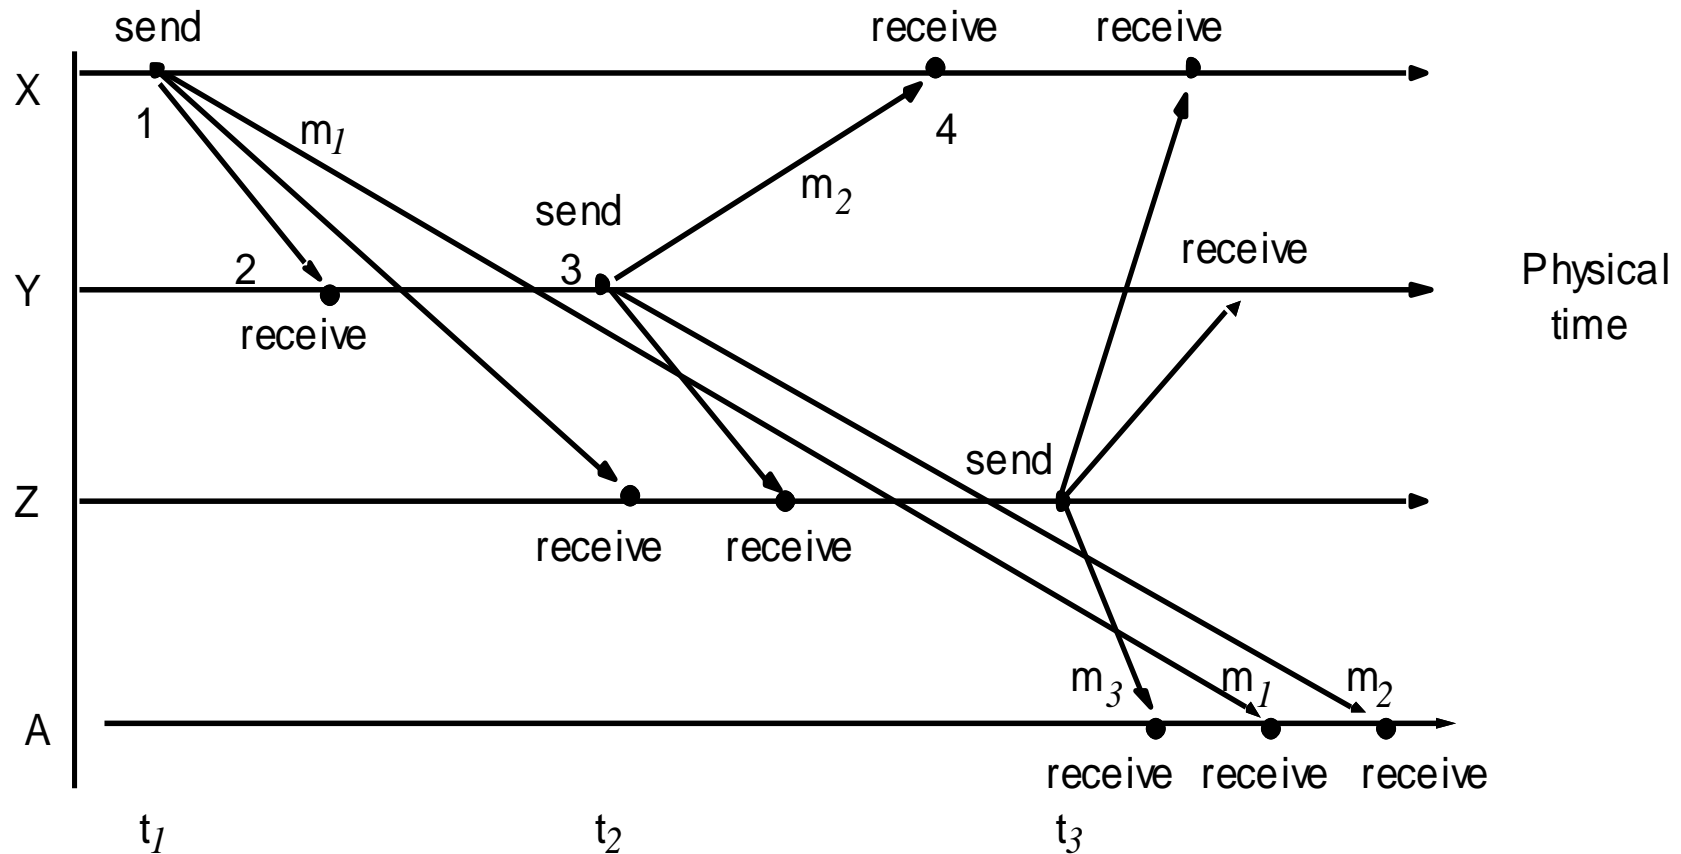
- In a DS it is hard to set time limits on the time taken for process execution, message delivery or clock drift.
- Synchronous DS – hard to achieve:
 - The time taken to execute a step of a process has known lower and upper bounds.
 - Each message transmitted over a channel is received within a known bounded time.
 - Each process has a local clock whose drift rate from real time has known bound.
- Asynchronous DS: There is NO bounds on:
 - Process execution speeds
 - Message transmission delays
 - Clock drift rates.

Interaction Model: Event Ordering



- In many DS applications we are interested in knowing whether an event occurred before, after, or concurrently with another event at other processes.
 - The execution of a system can be described in terms of events and their ordering despite the lack of accurate clocks.
- Consider a mailing list with:
users X, Y, Z, and A.

Real-time ordering of events



Inbox of User A looks like:

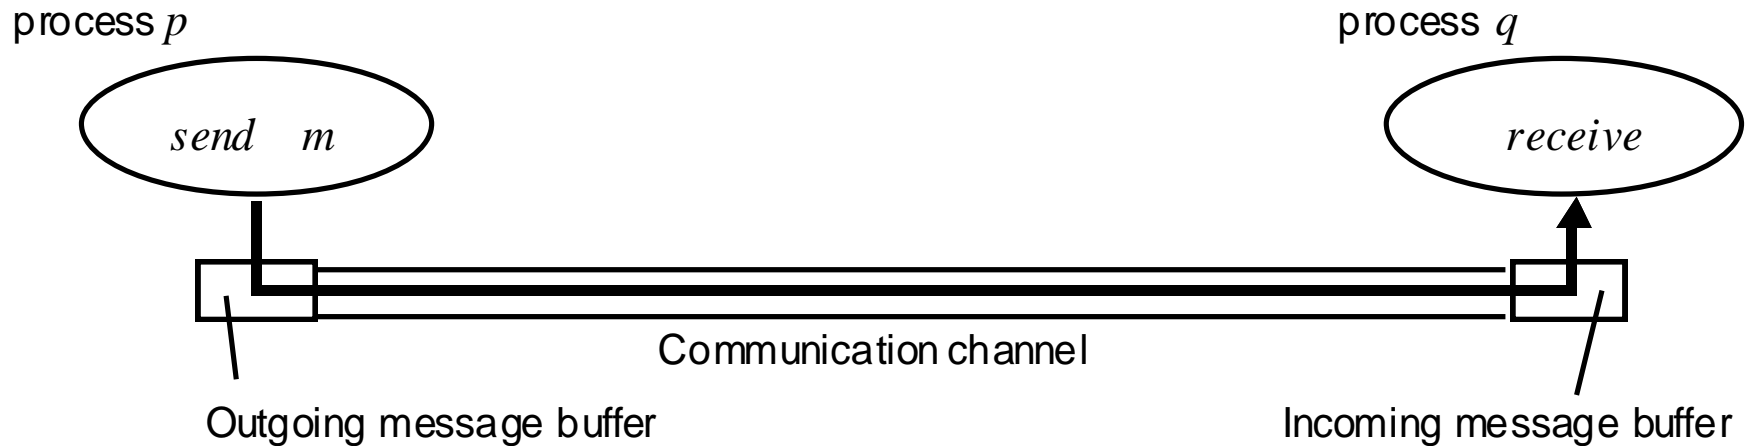
<i>Item</i>	<i>From</i>	<i>Subject</i>
23	Z	Re: Meeting
24	X	Meeting
26	Y	Re: Meeting

- Due to independent delivery in message delivery, message may be delivered in different order.
- If messages m1, m2, m3 carry their time t1, t2, t3, then they can be displayed to users accordingly to their time ordering.

Failure Model

- In a DS, both processes and communication channels may fail – i.e., they may depart from what is considered to be correct or desirable behavior.
- Types of failures:
 - Omission Failure
 - Arbitrary Failure
 - Timing Failure

Processes and channels



- Communication channel produces an omission failure if it does not transport a message from “*p*”’s outgoing message buffer to “*q*”’s incoming message buffer. This is known as “dropping messages” and is generally caused by a lack of buffer space at the receiver or at gateway or by a network transmission error.

Omission and arbitrary failures

<i>Class of failure</i>	<i>Affects</i>	<i>Description</i>
Fail-stop	Process	Process halts and remains halted. Other processes may detect this state.
Crash	Process	Process halts and remains halted. Other processes may not be able to detect this state.
Omission	Channel	A message inserted in an outgoing message buffer never arrives at the other end's incoming message buffer.
Send-omission	Process	A process completes a <i>send</i> , but the message is not put in its outgoing message buffer.
Receive-omission	Process	A message is put in a process's incoming message buffer, but that process does not receive it.
Arbitrary (Byzantine)	Process or channel	Process/channel exhibits arbitrary behaviour: it may send/transmit arbitrary messages at arbitrary times, commit omissions; a process may stop or take an incorrect step.

Timing failures



<i>Class of Failure</i>	<i>Affects</i>	<i>Description</i>
Clock	Process	Process's local clock exceeds the bounds on its rate of drift from real time.
Performance	Process	Process exceeds the bounds on the interval between two steps.
Performance	Channel	A message's transmission takes longer than the stated bound.

Masking Failures


- It is possible to construct reliable services from components that exhibit failures.
 - For example, multiple servers that hold replicas of data can continue to provide a service when one of them crashes.
- A knowledge of failure characteristics of a component can enable a new service to be designed to mask the failure of the components on which it depends:
 - Checksums are used to mask corrupted messages.

Security Model




- The security of a DS can be achieved by securing the processes and the channels used in their interactions and by protecting the objects that they encapsulate against unauthorized access.

Protecting Objects: Objects and principals




- Use “**access rights**” that define who is allowed to perform operation on a object.
- The server should verify the identity of the principal (user) behind each operation and checking that they have sufficient access rights to perform the requested operation on the particular object, rejecting those who do not.

The enemy



- To model security threats, we assume an enemy that is capable of sending any process or reading/copying message between a pair of processes
- Threats form a potential enemy: threats to processes, threats to communication channels, and denial of service.

Defeating security threats: Secure channels



- Encryption and authentication are used to build secure channels.
- Each of the processes knows the identity of the principal on whose behalf the other process is executing and can check their access rights before performing an operation.

Presentation Outline

- Introduction
- Architectural Models
 - Software Layers
 - System Architectures
 - Client-Server
 - Clients and a Single Server, Multiple Servers, Proxy Servers with Caches, Peer Model
 - Alternative Client-Server models driven by:
 - Mobile code, mobile agents, network computers, thin clients, mobile devices and spontaneous networking
 - Design Challenges/Requirements
- Fundamental Models – formal description
 - Interaction, Failure, and Security models.
- Summary

Summary



- Most DSs are arranged accordingly to one of a variety of architectural models:
 - Client-Server
 - Clients and a Single Server, Multiple Servers, Proxy Servers with Cache, Peer Model
 - Alternative Client-Server models driven by:
 - Mobile code, mobile agents, network computers, thin clients, mobile devices and spontaneous networking
- Fundamental Models – formal description
 - Interaction, failure, and security models.
- The concepts discussed in the module play an important role while architecting DS and apps.

- Demo: Building a Swing GUI

Install WindowBuilder


1. Open Eclipse
2. Click Help...
3. Select Install new software...
4. Access the following website:
<https://www.eclipse.org/windowbuilder/download.php>
<http://download.eclipse.org/windowbuilder/WB/integration/4.6/> (search windowbuilder in Eclipse Marketplace)
5. Copy the installation URL that matches your eclipse version to the “work with:” field.
6. Install WindowsBuilder
7. Restart Eclipse

Create a Project with WindowBuilder

1. Create a new Java Project
2. Then Right Click on the src folder and do:

New > Other > WindowBuilder > Swing
Designer > Application Window

```
1 import java.awt.EventQueue;
4
5 public class MainWindow {
6
7     private JFrame frame;
8
9     /**
10      * Launch the application.
11      */
12     public static void main(String[] args) {
13         EventQueue.invokeLater(new Runnable() {
14             public void run() {
15                 try {
16                     MainWindow window = new MainWindow();
17                     window.frame.setVisible(true);
18                 } catch (Exception e) {
19                     e.printStackTrace();
20                 }
21             }
22         });
23     }
```



1.9.0.201608250402

Layouts - A crash course

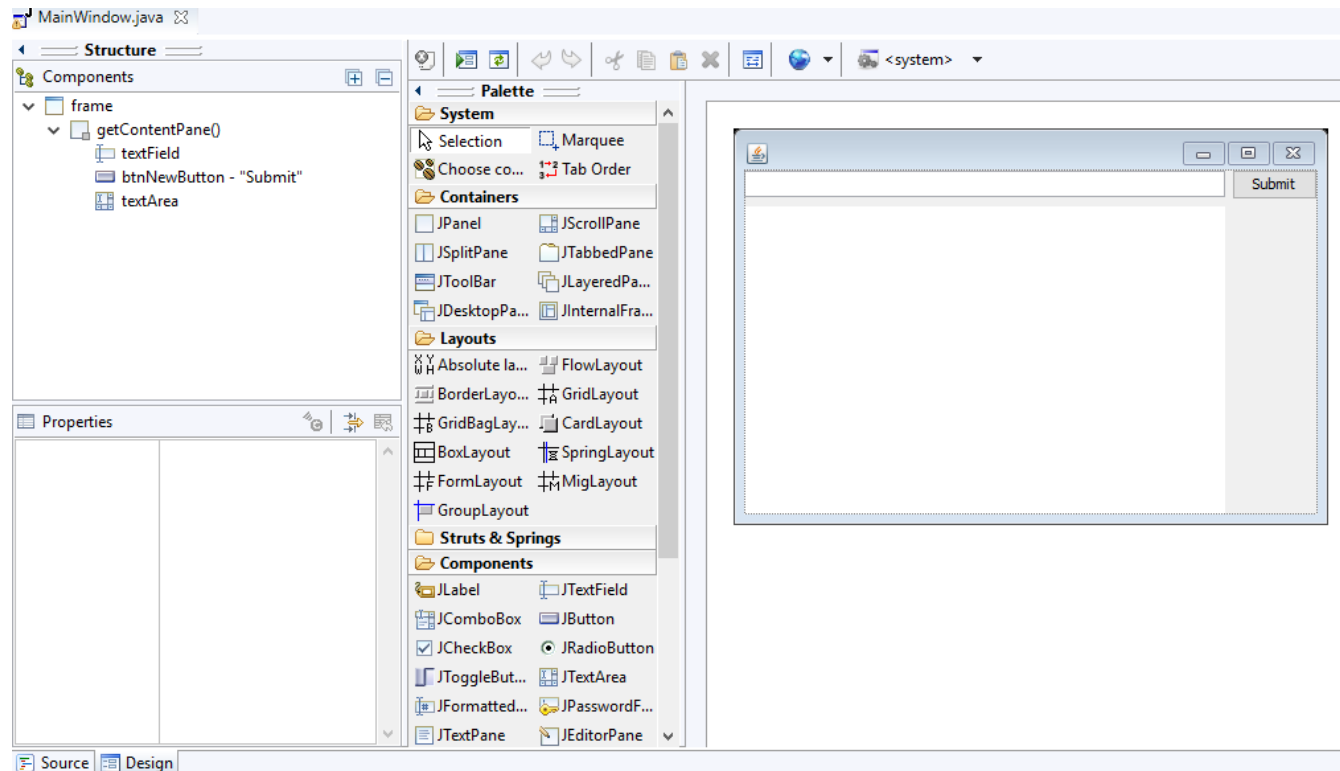
Absolute Layout : Does not scale when resized

Grid Layout: Items are rendered on a grid and will stack or unstack depending on the size of the window

GridBagLayout:

<https://docs.oracle.com/javase/tutorial/uiswing/layout/gridbag.html>

Add things to the GUI as desired



Event Handlers

Right click a component (i.e. a button) and select

Event Handler > action > actionPerformed

See line 85 - 89 of MainWindow.java