

# Distributed Systems

- ▶ Lecturer: Justin Pearson [justin@it.uu.se](mailto:justin@it.uu.se).
- ▶ Recommended text book:  
*“Distributed Systems Concepts and Design”,  
Coulouris, Dollimore and Kindberg. Addison and  
Wesley.*
- ▶ Course webpage to follow. But follow the link from:  
[user.it.uu.se/%7ejustin](http://user.it.uu.se/%7ejustin)

## Other good books

- ▶ “*Distributed Systems an algorithmic approach*”, Sikumar Ghosh (Very good).
- ▶ “*Principles of concurrent and distributed programming*” M. Ben-Ari.
- ▶ “*Distributed Operating Systems and Algorithms*”, Randy Chow and Theodore Johnson.

# Course Plan

- ▶ Introduction.
- ▶ Remote Objects and Remote Procedure calls.
- ▶ Clocks, Clock Synchronisation, Logical time and global states.
- ▶ Coordination and Agreement and Transaction
- ▶ Transactions
- ▶ Distributed Snapshots and State collection.
- ▶ Two and three phase commit protocols.
- ▶ Shared Memory

# Course plan for the DVP students

- ▶ Distributed Transactions
- ▶ Introduction to Peer to Peer networks
- ▶ Possibly the CAP Theorem.

# What is a distributed System?

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

So multiple processors sharing the same memory bus would not be counted as a distributed system.

# Problems to be addressed

- ▶ Concurrency.
- ▶ No global Clock.
- ▶ Independent Failures.

# Concurrency

- ▶ Different Computers on a network. Each computer can be doing work at the same time.
- ▶ What happens if two computers want to access a resource at the same time?
- ▶ Network delays are not constant make synchronisation difficult.

# No Global Clock

- ▶ When programs need to cooperate they coordinate their actions by exchanging messages.
- ▶ Close coordination often depends on a shared idea of the time at which events occur. Need notions of before or after.
- ▶ There are limits to the accuracy with which components in a network can synchronise their clocks. There is no global notion of the correct time.
- ▶ Direct consequence of the fact that messages are sent through a network and network delays are not constant (or even bounded).



# Independent failures

- ▶ All computer systems can fail it is good design to build robustness in.
- ▶ Distributed systems fail in new ways.
- ▶ Network faults might result in components being isolated but not stopping. What happens if your waiting for an acknowledge message and you never get it? How do you know if the remote system got your message?

# Examples of Distributed Systems

- ▶ The Internet.
- ▶ Intranets. (Internal Internet inside a company)
- ▶ Grid Computing.
- ▶ P2P programming (Peer to Peer). Can be used for other things besides file sharing.
- ▶ Mobile and ubiquitous computing.

# Resource Sharing

- ▶ What is a resource?
  - ▶ Hardware. Shared printer. Shared processor
  - ▶ Pieces of running software, distributed objects, remote procedures.
  - ▶ Data.
- ▶ Can resources be replicated?
- ▶ Where does data live? Is the database in one place or distributed? File sharing networks.

# Challenges

- ▶ Heterogeneity. (Everybody is different).
- ▶ Security
- ▶ Scalability
- ▶ Failure Handling
- ▶ Concurrency
- ▶ Transparency

# Heterogeneity

- ▶ Different networks
- ▶ Different hardware
- ▶ Different operating systems
- ▶ Different programming languages
- ▶ Different implementations.

How do we solve all this?

# Middleware

- ▶ A software layer that provides programming abstractions as well as masking the heterogeneity of the underlying system.
- ▶ Examples include:
  - ▶ CORBA ( Çorba Turkish for soup)
  - ▶ Java RMI, Jini
- ▶ How do you deal with mobile code in a heterogeneous environment?
  - ▶ One possible solution is to use a virtual machine, Java VM or .Net.

# Security

- ▶ On its own network traffic is not secure. Any message could in principle be seen by anybody. Packet sniffers.
- ▶ Encryption can solve some problems.

More problems include:

- ▶ Denial of service attacks.
- ▶ How do you make mobile code secure.

# Scalability

- ▶ Short story, things get bigger all the time. More users, more computers, more data, more ...
- ▶ How do you make systems scalable?
- ▶ It is not as simple as just adding more hardware.
- ▶ One machine services 20 users, does not mean that two machines will service 40 users.
- ▶ Design problem.



# Failure Handling

- ▶ *Detecting Failures*: Some failures can be detected. But in a distributed system it becomes hard.
- ▶ *Failure Masking*: Retransmission, backups (data and servers)
- ▶ *Tolerating failures*: When should you not care?
- ▶ *Recovery*: Not always easy in a distributed system. What is the global notion of state?

# Concurrency

- ▶ Managing shared resources.
- ▶ For an object to be safe in a concurrent environment its operations must be synchronised in such a way that data remains consistent (bank account example). This is harder than simply using semaphores in non-distributed systems.

# Transparency

- ▶ Access transparency enables local and remote resources to be accessed using identical operations.
- ▶ Location transparency. Allows access without knowing object locations.
- ▶ Concurrency transparency
- ▶ Replication transparency: enables multiple instances to be active to increase reliability and performance.
- ▶ Failure transparency
- ▶ Mobility transparency.

# Architectural Models

- ▶ An Architectural model of a distributed system is concerned with the placements of its parts and the relationship between them.
- ▶ We are going to look the most popular design for a distributed system. The client server model.

Insert figure 2.2. from the book here.

# Client Server Model

- ▶ Process acts as a client sends request to a server.
- ▶ Examples:
  - ▶ Webserver. Client (Web browser) sends a request for a webpage the webserver then returns the requested server.
  - ▶ A SQL server, client processes send request for data or requests to modify data.
- ▶ Servers can become clients. For example the webserver might simply be a web proxy and pass requests on to other servers.

# Variations of the client-server model

- ▶ Mobile code and Mobile Agents
- ▶ Thin clients

# Mobile Code and Mobile Agents

- ▶ Web applets.
  - ▶ A client request results in downloading of applet code and then the client interacts with the applet.
- ▶ Mobile Agents
  - ▶ A mobile agent is a running program (code and data) that travels from one computer to another in a network performing some task.



# Thin Clients

- ▶ A thin client refers to software/hardware that supports a window-based user interface local to the user while executing application programs on a remote computer.
- ▶ Examples include:
  - ▶ Sun Rays
  - ▶ remote X clients.
  - ▶ VNC
  - ▶ Microsoft remote desktop.

# Summary and Conclusion

- ▶ Distributed Systems are everywhere.
- ▶ Distributed systems have their own design problems and issues.
- ▶ Middleware supplies abstractions to allow distributed systems to be designed. Focus of this course: What abstractions are necessary to a distributed system.
- ▶ Client-server architecture is a common way of designing distributed systems.