

1 Introduction

Main Points

- What we are going to cover
- What distributed systems are
- What characteristics well-designed systems have

1.1 Course Outline

What are we going to do?

- 1 lectures on computer communications and networks and loosely coupled systems, revising material from the Multimedia Communications Systems course. The Postgraduate course will get an additional two hours in their seminar slot.
- 5 lectures on Remote Procedure Call, distributed objects and classic distributed systems such as NFS
- 3 lectures on integrating distributed systems through the web, content distribution networks and Peer to Peer computing.
- 5 lectures on closely coupled systems for distributed transactions
- 1 lecture on recent advances in networks and distributed systems

1.1.1 Exercises

Two exercises:

1. Building a distributed chat program based on RPC, using RMI (33% weighted)
2. Building a distributed hash table (subject to change) (66% weighted)

The exercises will be in Java, and where appropriate, will be building upon and extending provided skeleton code.

- The exercises will be peer-assessed.
- In the next exercise class after the handin you will be required to mark two of your classmates' assignments.
- Successfully marking two assignments will provide you with 5 marks for the course (2 assignments = 10 marks).
- At the end of the term, you will hand in your two assignments along with the assessments from your classmates.
- These will be formally given a course mark out of 30 marks.

1.1.2 Detailed Timetable

Week	Lecture	Lecture	Exercise Class	Assignment
1 (10/1)	Introduction	Fundamentals	No class	
2 (17/1)	OS Issues	RPC	Fundamentals	
3 (24/1)	Object Systems	Enterprise Computing and CORBA	Programming Exercise	
4 (31/1)	Security	Security	Security	Ass 1 due
5 (7/2)	Naming	Distributed File Systems	Marking	
6 (14/2)	P2P Networks	P2P Networks	P2P	
7 (21/2)	Replication	Replication	Programming Exercise	
8 (28/2)	Transactions	Concurrency Control	Exam sample	Ass 2 due
9 (7/3)	Distributed Transactions	Coping with Failure	Marking	
10 (14/3)	Pervasive Computing Management	No lecture	No class	Portfolio due

Table 1: Course Timetable 2004

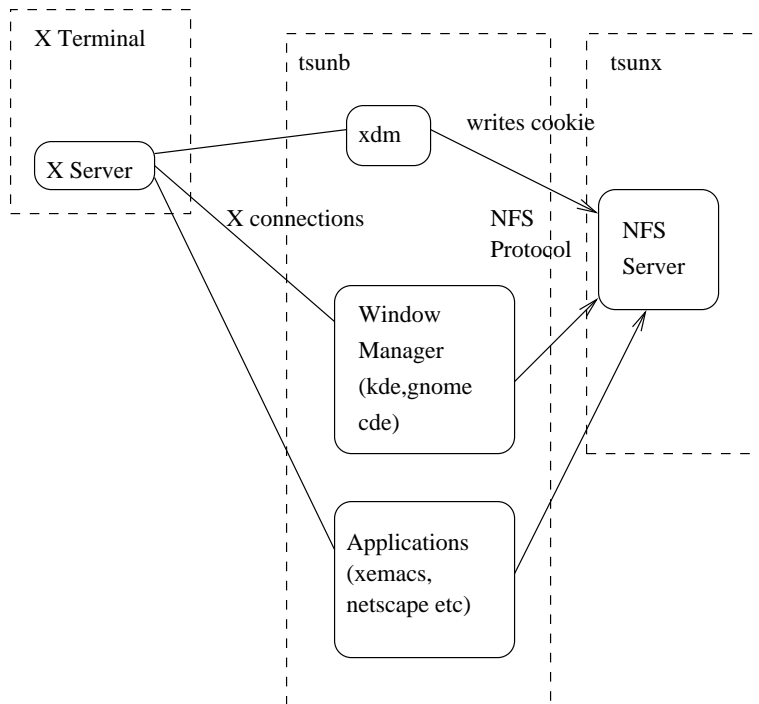
1.1.3 MSc Programme

- Details to be discussed in the MSc seminar
- Read from reading list
- Work on programming exercises

1.2 What's a Distributed System

A distributed system: physically separate computers working together

- Cheaper and easier to build lots of simple computers
- Easier to add power incrementally
- Machines may necessarily be remote, but system should work together
- Higher availability - one computer crashes, others carry on working
- Better reliability - store data in multiple locations
- More security - each piece easier to secure to right level.



1.2.1 The real world...

In real life, can get:

- Worse availability - every machine must be up. "A distributed system is one where some machine you've never heard of fails and prevents you from working"
- Worse reliability
- Worse security

Problem: Coordination is more difficult because multiple people involved, and communication is over network.

Your Task: What are the distributed interactions when you login at an x-terminal?

1.2.2 Interactions at an X terminal

Simplified interactions

1.3 Example Distributed Systems

Electronic Mail Mail delivered to remote mailbox. Requires global name space to identify users, transport mechanisms to get mail to mailbox

Distributed Information - WWW Remote information hidden below hypertext browser. Caching and other features operate transparently

Distributed File System Files stored on many machines, generally not machine you're working on. Files accessed transparently by OS knowing they're remote and doing remote operations on them such as read and write e.g. Network File System (NFS)

Trading Floor System Bids made, stocks sold, screens updated.

1.3.1 Network assumptions

1. The network is reliable
2. Latency is zero
3. Bandwidth is infinite
4. The network is secure
5. Topology doesn't change
6. There is one administrator
7. Transport cost is zero
8. The network is homogeneous

(Source: *The Eight Fallacies of Distributed Computing - Peter Deutsch*)

1.4 What do we want from a Distributed System?

1. Resource Sharing
2. Openness
3. Concurrency
4. Scalability
5. Fault Tolerance
6. Transparency

Your Task: order the importance of each of these features for the example systems in the previous slide.

1.5 Elements of a Distributed System

Another way to view a system...

- Communications system
- Messages
- Machines
- Processes on Machines
- Programs
- People

1.6 Conclusion

- Its difficult to design a good distributed system: there are a lot of problems in getting “good” characteristics, not the least of which is people
- Over the next ten weeks you will gain some insight into how to design a good distributed system.