MYE017 Distributed Systems

Kostas Magoutis
magoutis@cse.uoi.gr
http://www.cse.uoi.gr/~magoutis

Course administration

- Meeting Tue 9-12pm at I2
- Office hours by apt (Office B.34)

Course administration

- Web site : http://www.cse.uoi.gr/~magoutis/MYE017
- Προτεινόμενα συγγράμματα:
  - Κατανεμημένα συστήματα αρχές και υποδείγματα ($1^{η}$ Έκδοση) Tanenbaum, van Steen (TvS)
  - Distributed Systems: Concepts and Design, Coulouris, Dollimore, Kindberg, Blair (CDK)
  - Κατανεμημένα Συστήματα με Java - 3η έκδοση, Ι.Κ. Κάβουρας, Ι.Ζ. Μήλης, Γ.Β. Ξηλωμένος, Α.Α. Ρουκουνάκη
  - Εισαγωγή στον Αντικειμενοστραφή Προγραμματισμό με Python (Κεφ. 11), Μαγκούτης, Νικολάου (MN)
- Grade breakdown
  - Final exam 75%, assignments 25%

Definition of a distributed system

- Tanenbaum/van Steen: “A distributed system is a collection of independent computers that appears to its users as a single coherent system”

- Lamport: “A distributed system is one in which the failure of a computer you didn't even know existed can render your own computer unusable”

Organization of a distributed system

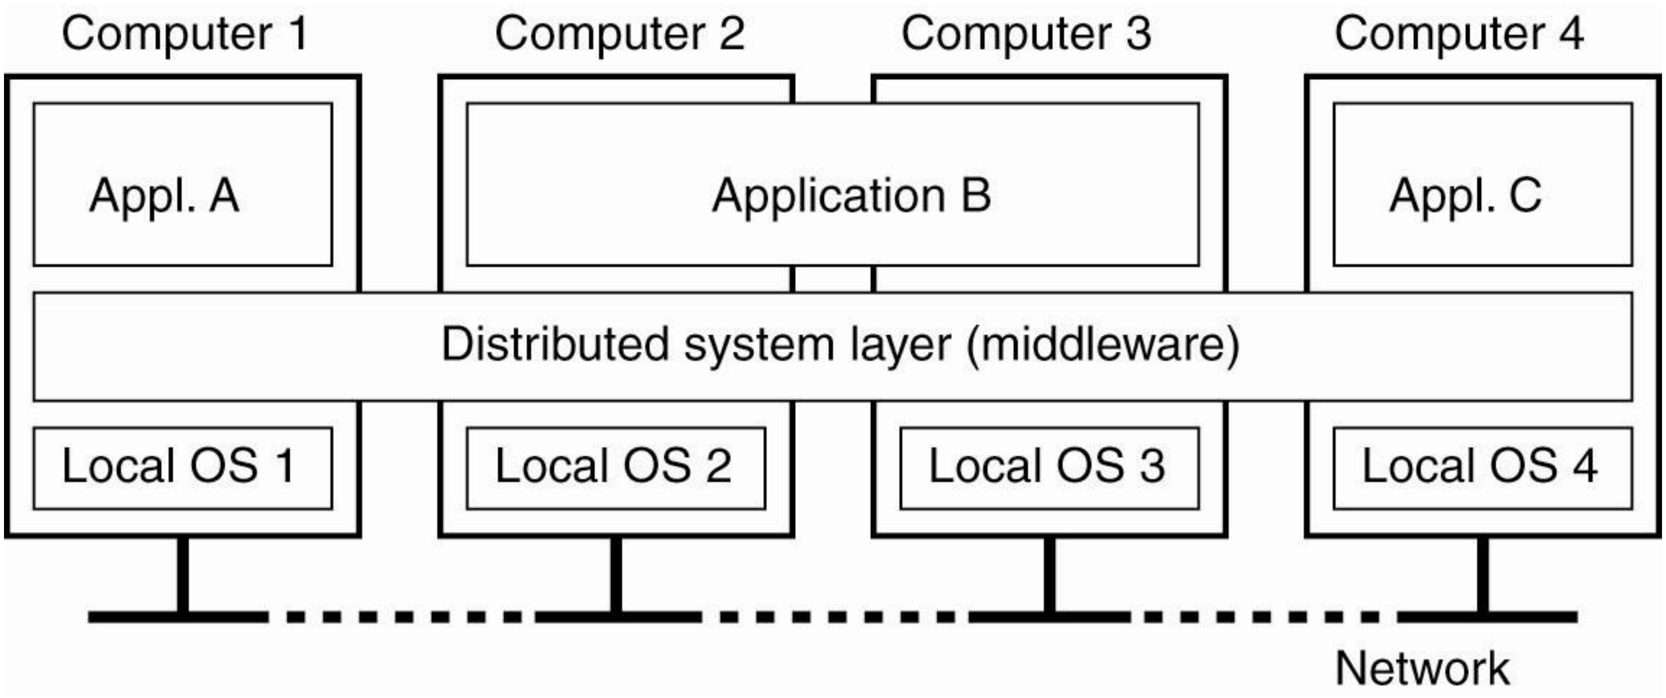

A distributed system organized as middleware. The middleware layer extends over multiple machines, and offers each application the same interface.

Data centers

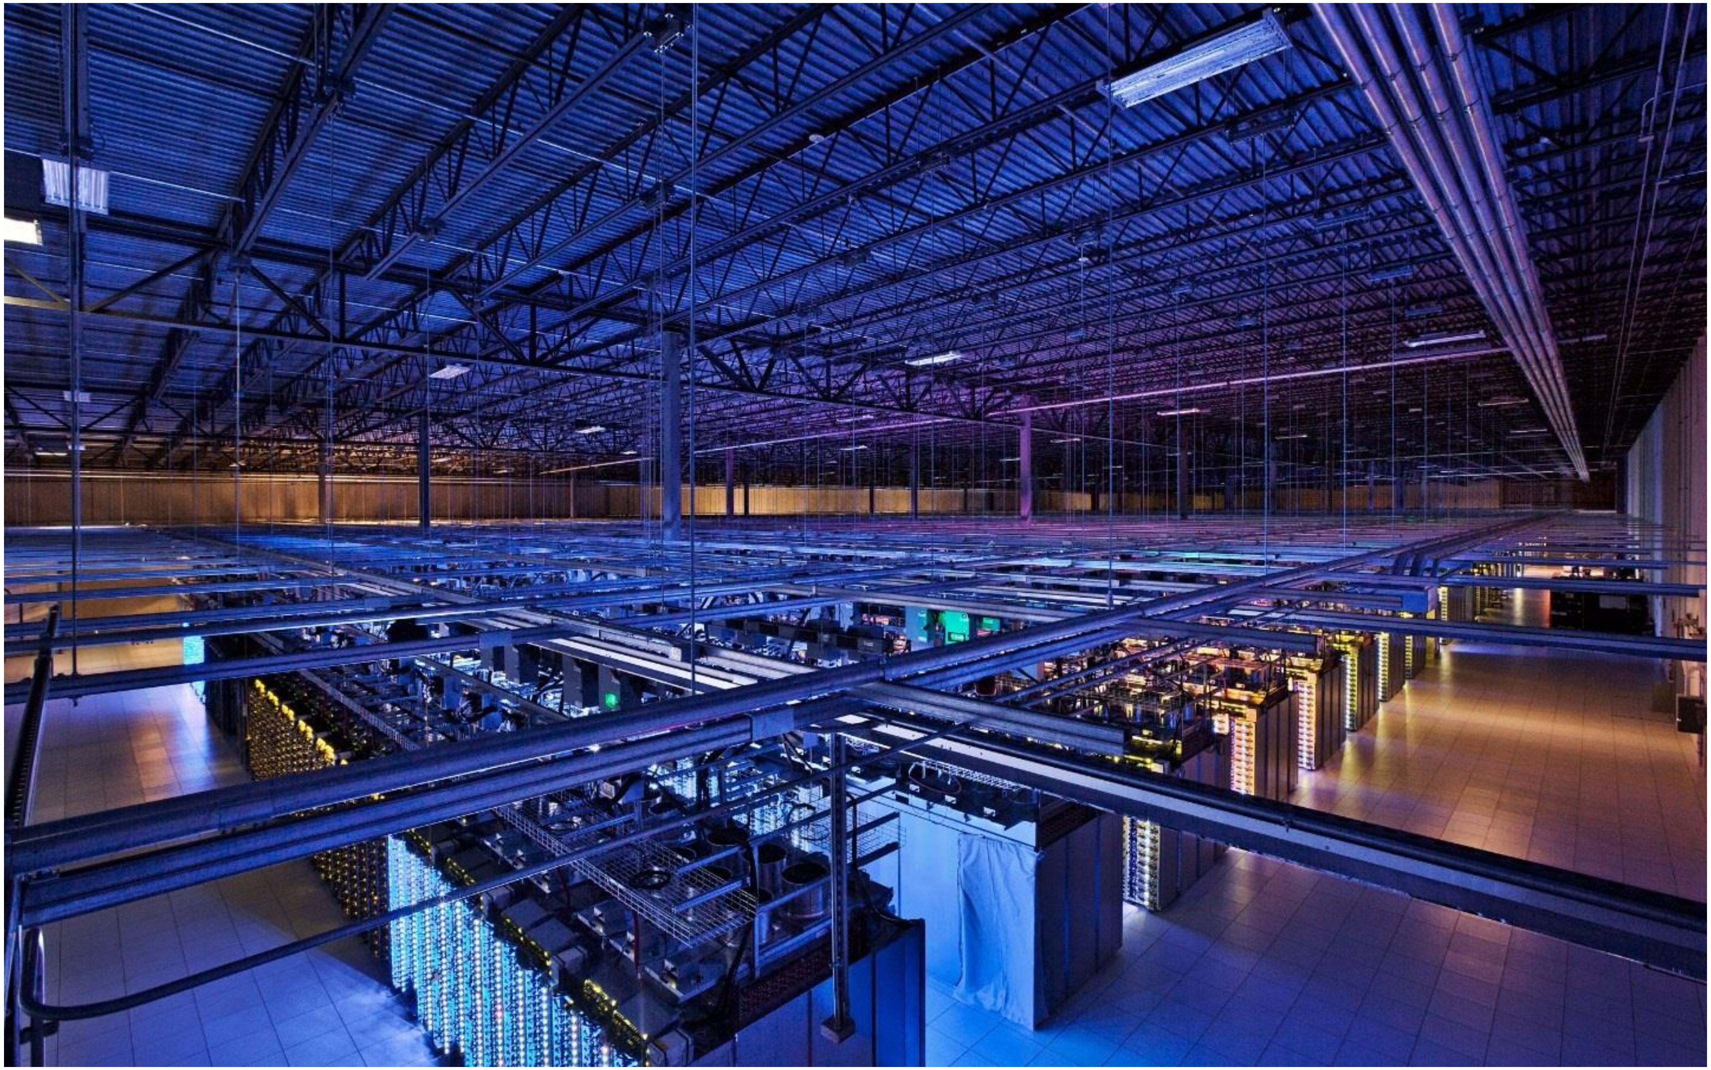

Networks

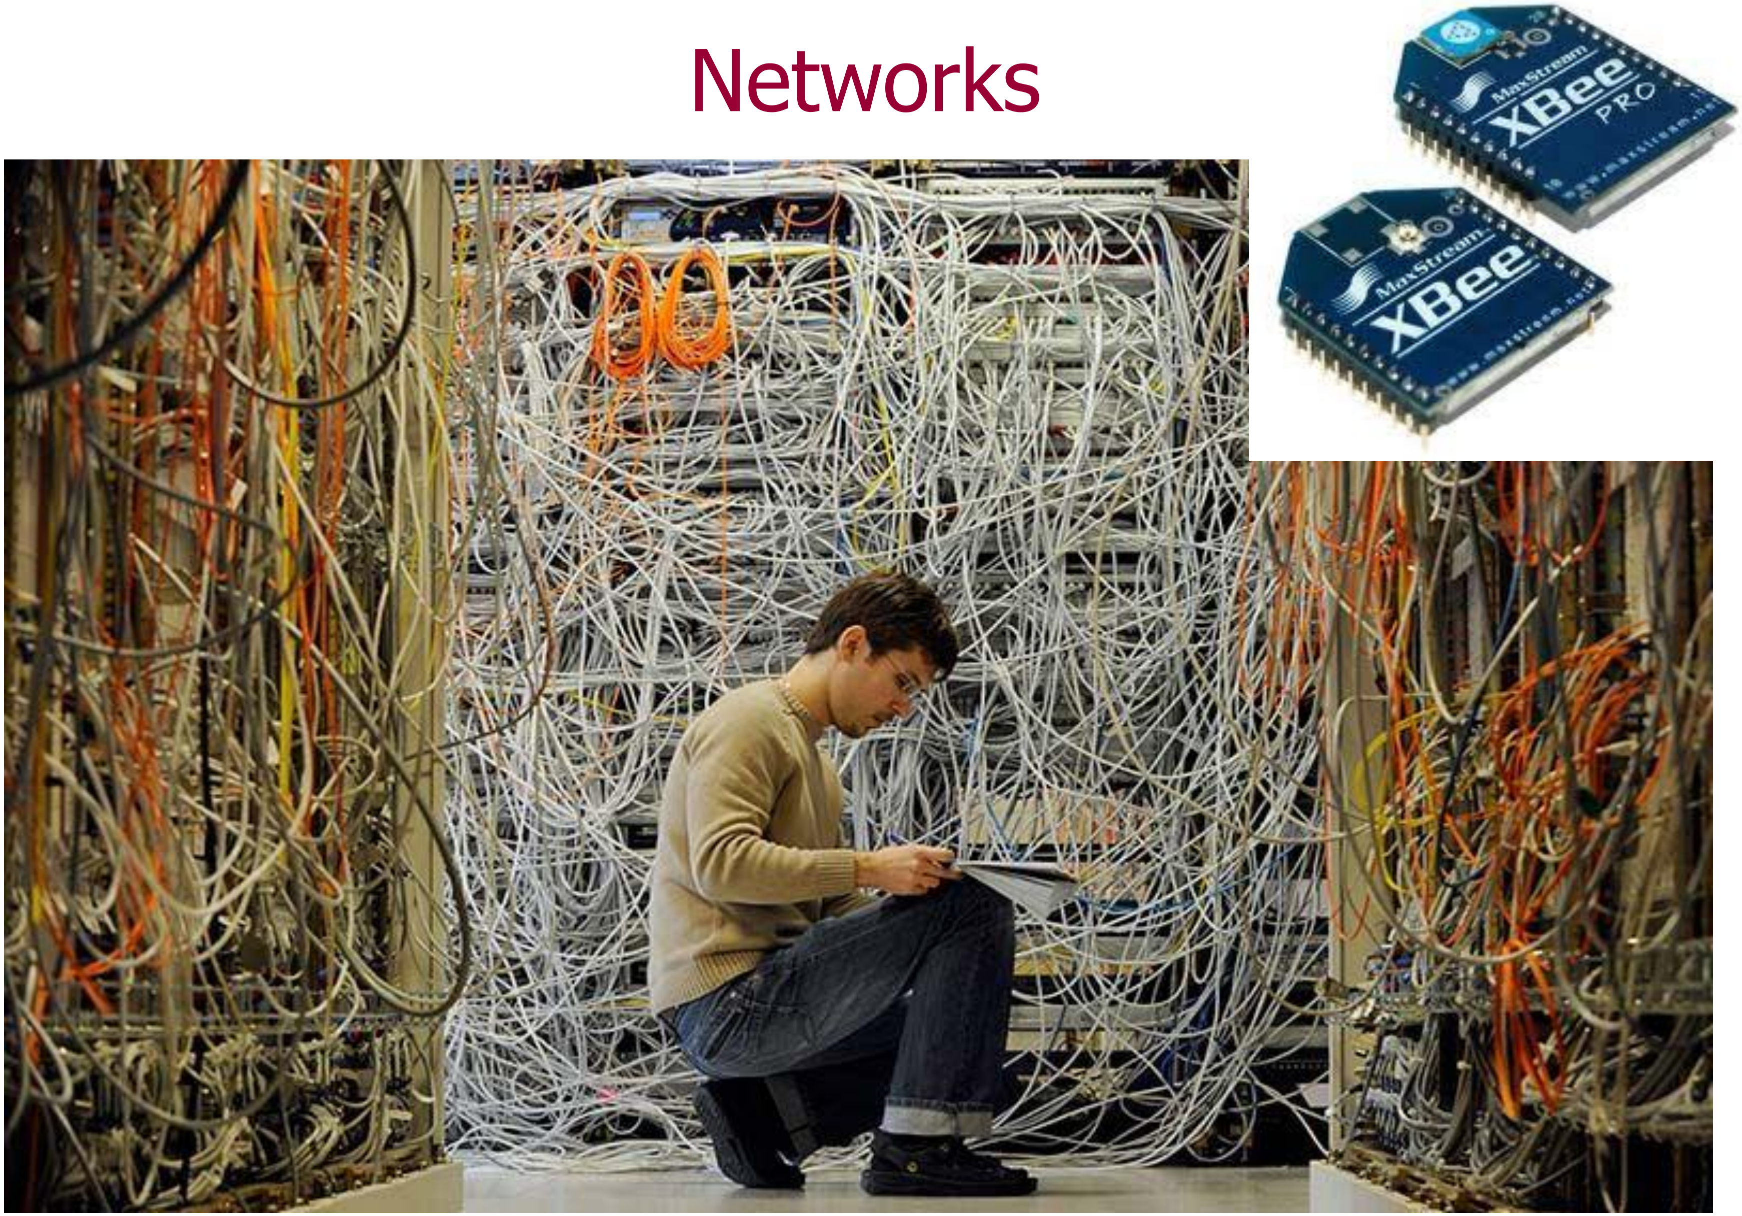

Connectivity

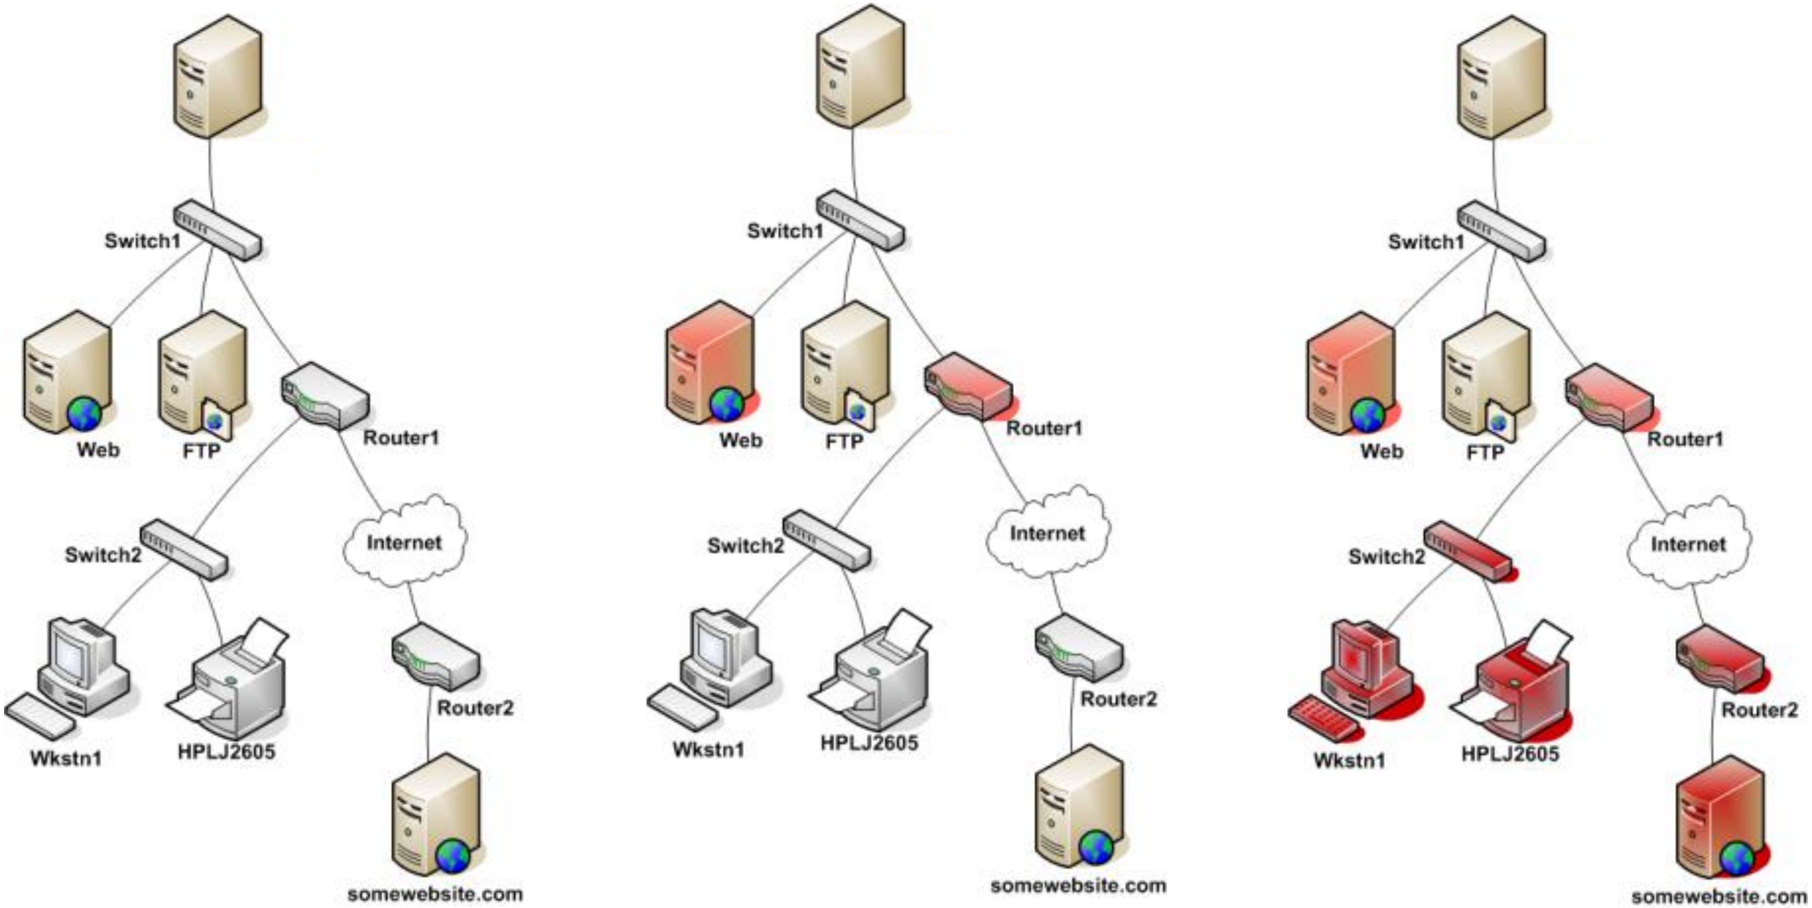

Internet

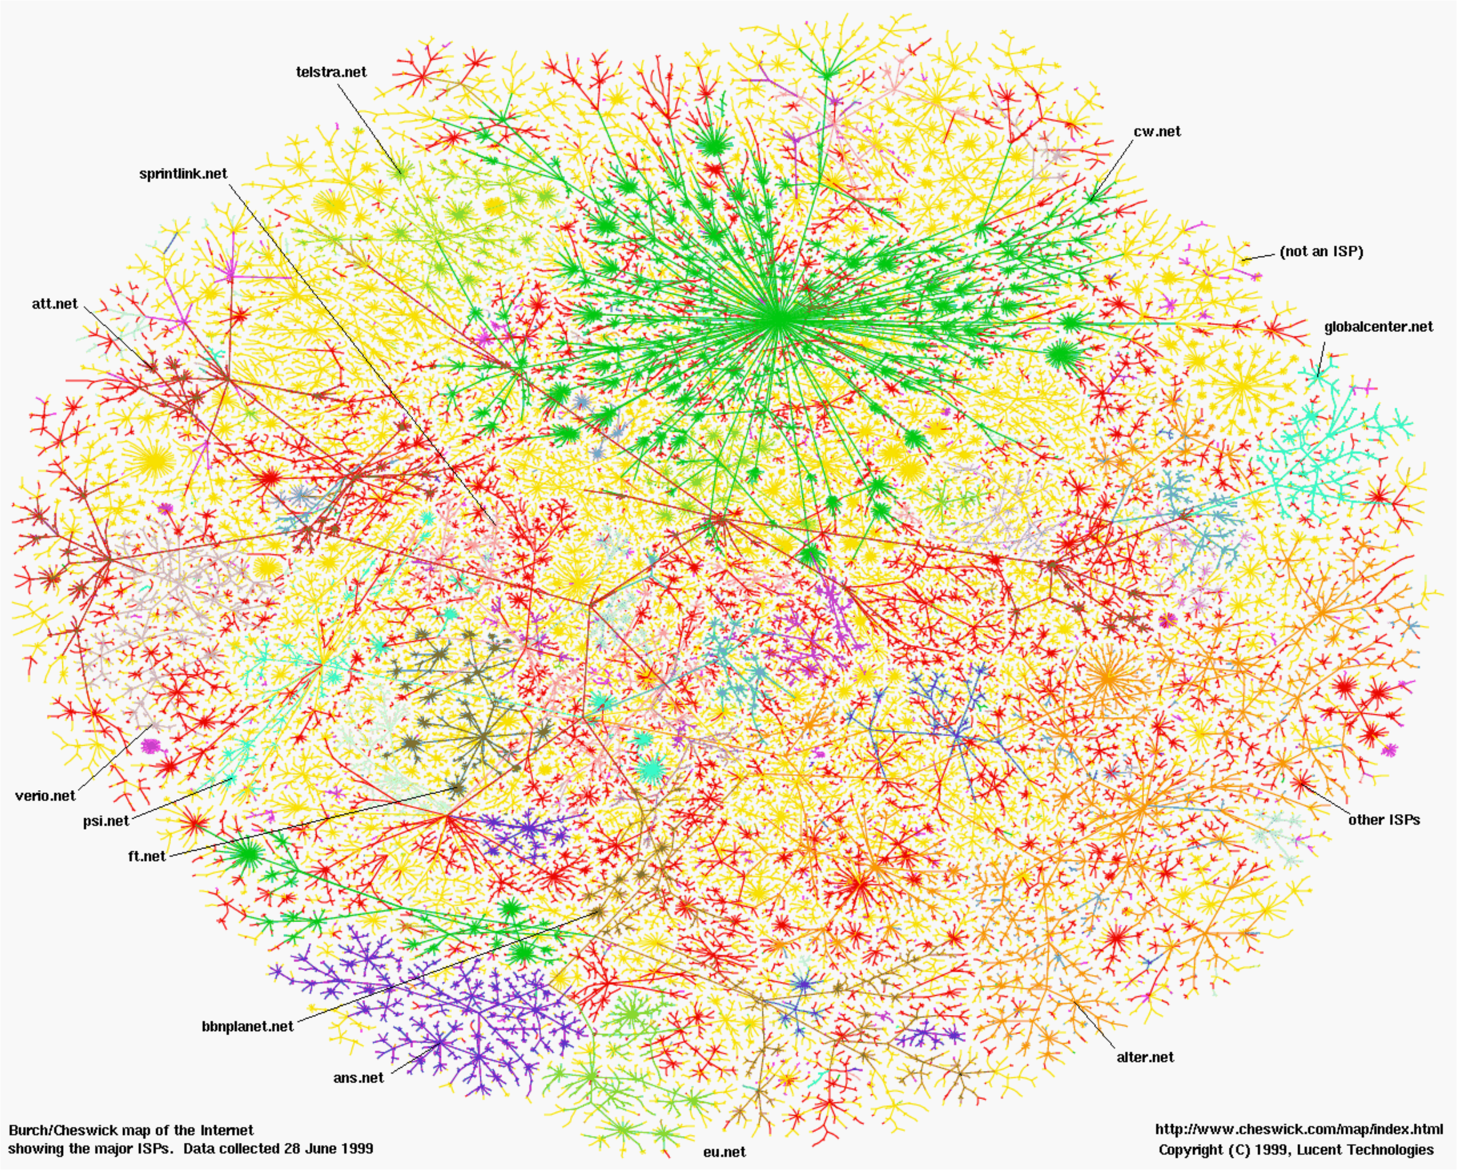
telstra.net
cw.net
sprintlink.net
(not an ISP)
att.net
globalcenter.net
verio.net
psi.net
ft.net
other ISPs
bbnplanet.net
ans.net
alter.net
eu.net
Burch/Cheswick map of the Internet
showing the major ISPs. Data collected 28 June 1999
http://www.cheswick.com/map/index.html
Copyright (C) 1999, Lucent Technologies

Disks

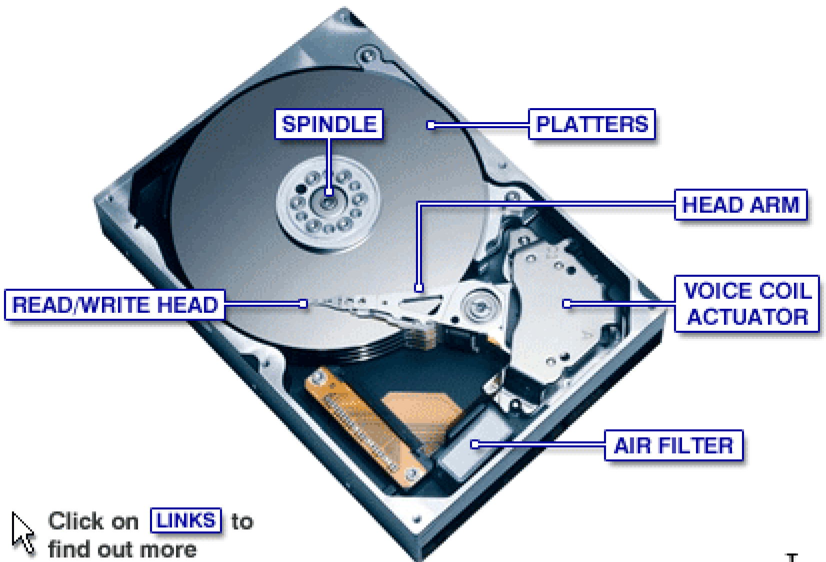
SPINDLE
PLATTERS
HEAD ARM
READ/WRITE HEAD
VOICE COIL ACTUATOR
AIR FILTER
Click on LINKS to find out more

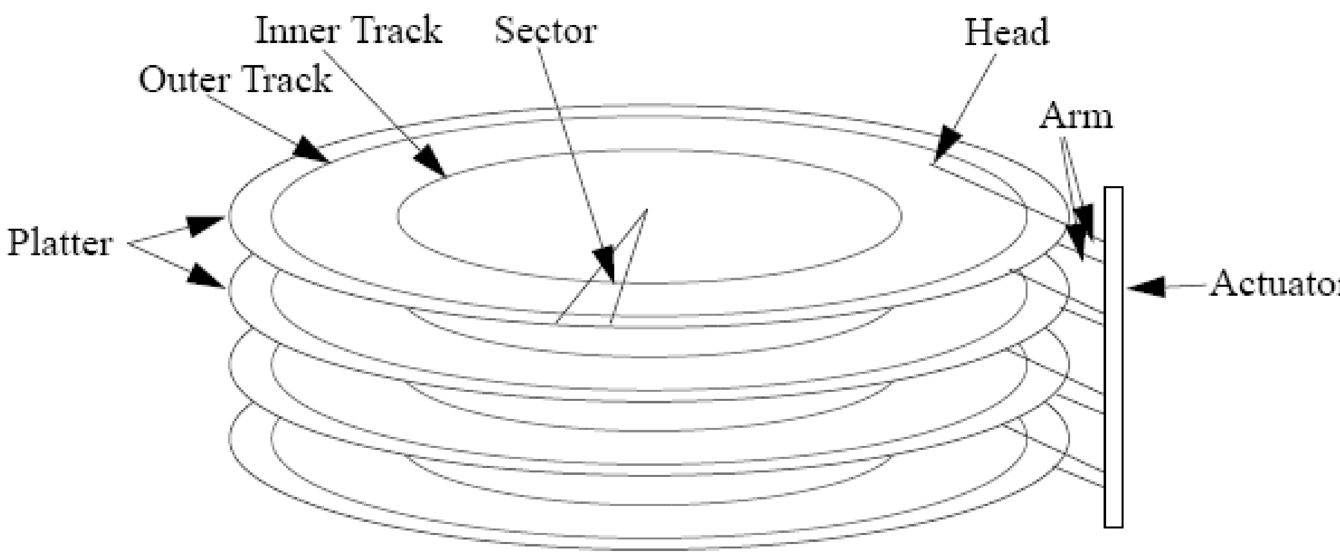
Inner Track
Sector
Head
Outer Track
Arm
Platter
Actuator

Transparency in a Distributed System

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 may be shared by several competitive users
Concurrency	Hide that a resource may be shared by several competitive users
Failure	Hide the failure and recovery of a resource
Persistence	Hide whether a (software) resource is in memory or on disk

Different forms of transparency in a distributed system.

Scalability Problems

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

Examples of scalability limitations.

Decentralized algorithms

- No machine has complete information about the system state
- Machines make decisions based only on local information
- Failure of one machine does not ruin the algorithm
- There is no implicit assumption that a global clock exists

Scaling Techniques (1)

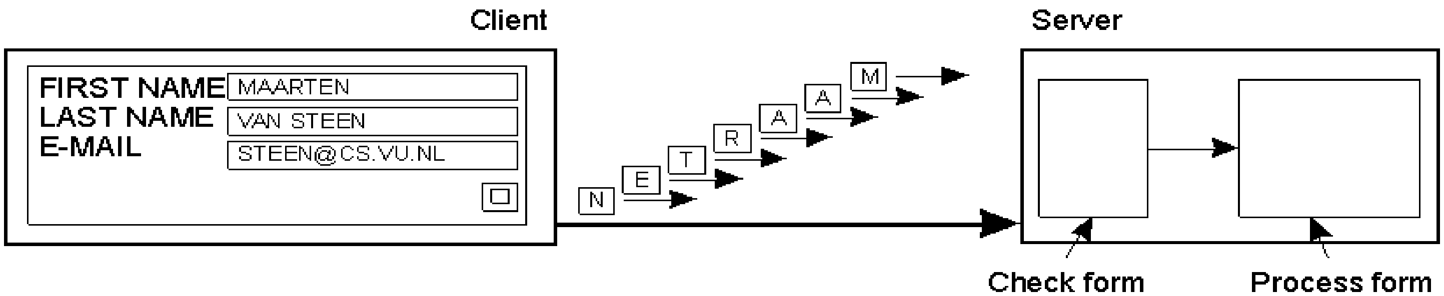

(a)

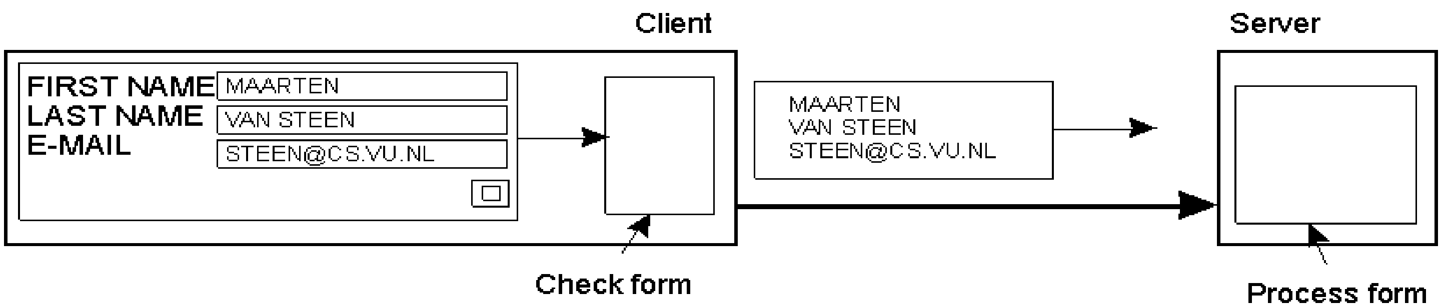

(b)

The difference between letting:

a) a server or

b) a client check forms as they are being filled

Scaling Techniques (2)

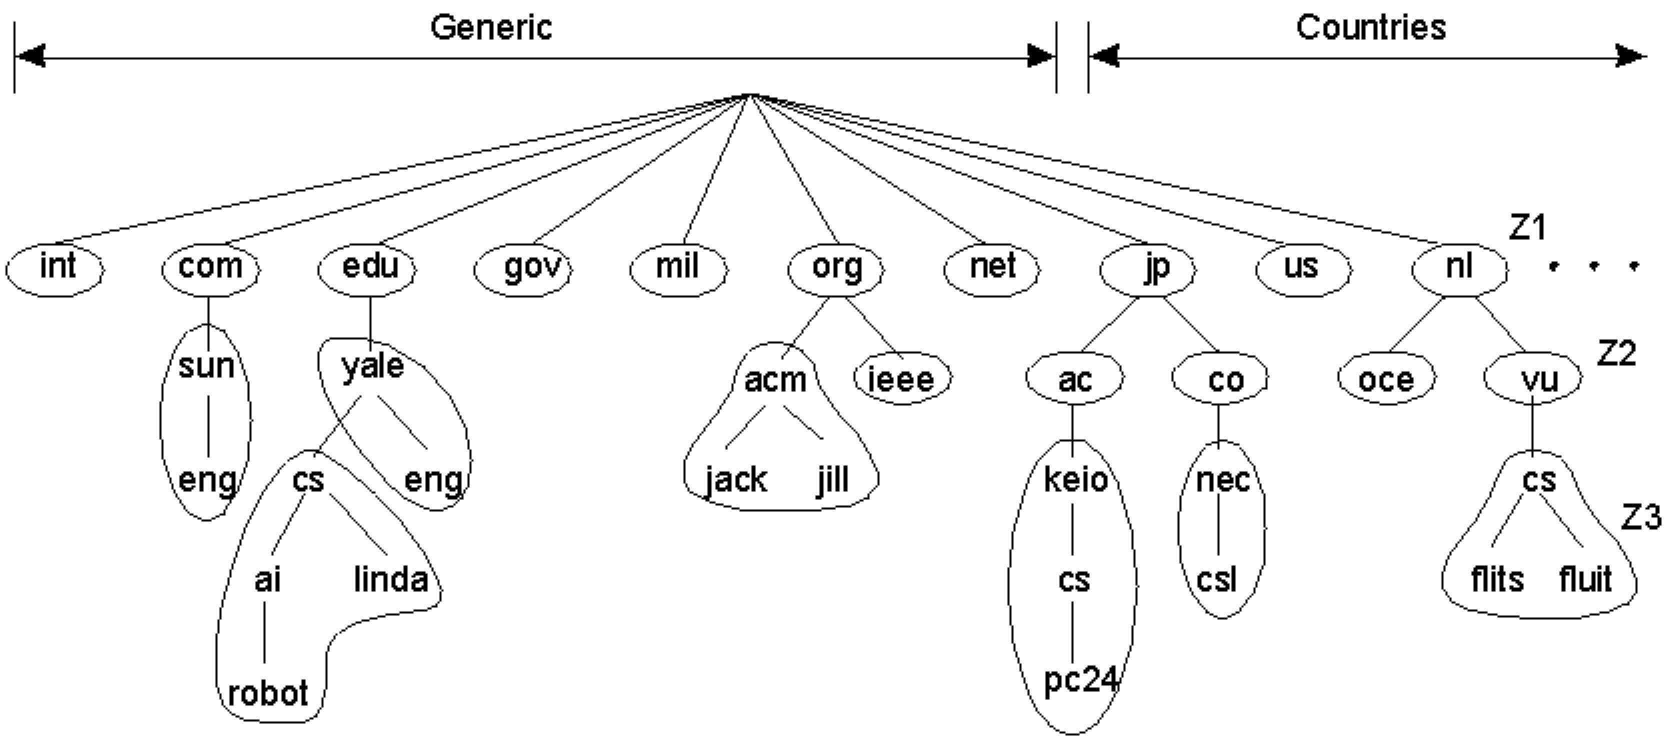

An example of dividing the DNS name space into zones.

Scaling Techniques (3)

- Asynchronous communication
- Replication
- Caching

Clients and Servers

- General interaction between a client and a server.

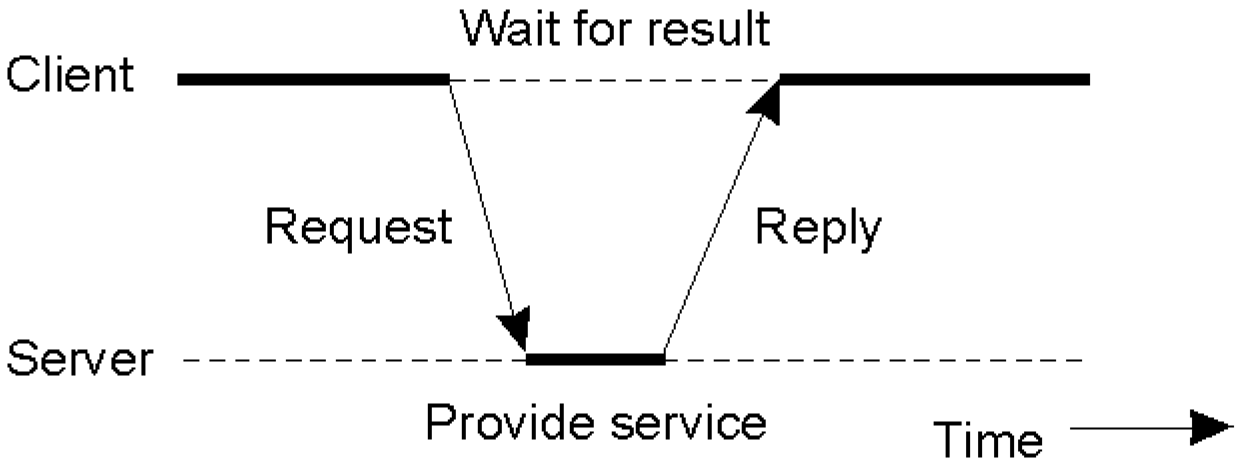

Multitiered Architectures (1)

- Alternative client-server organizations (a) – (e).

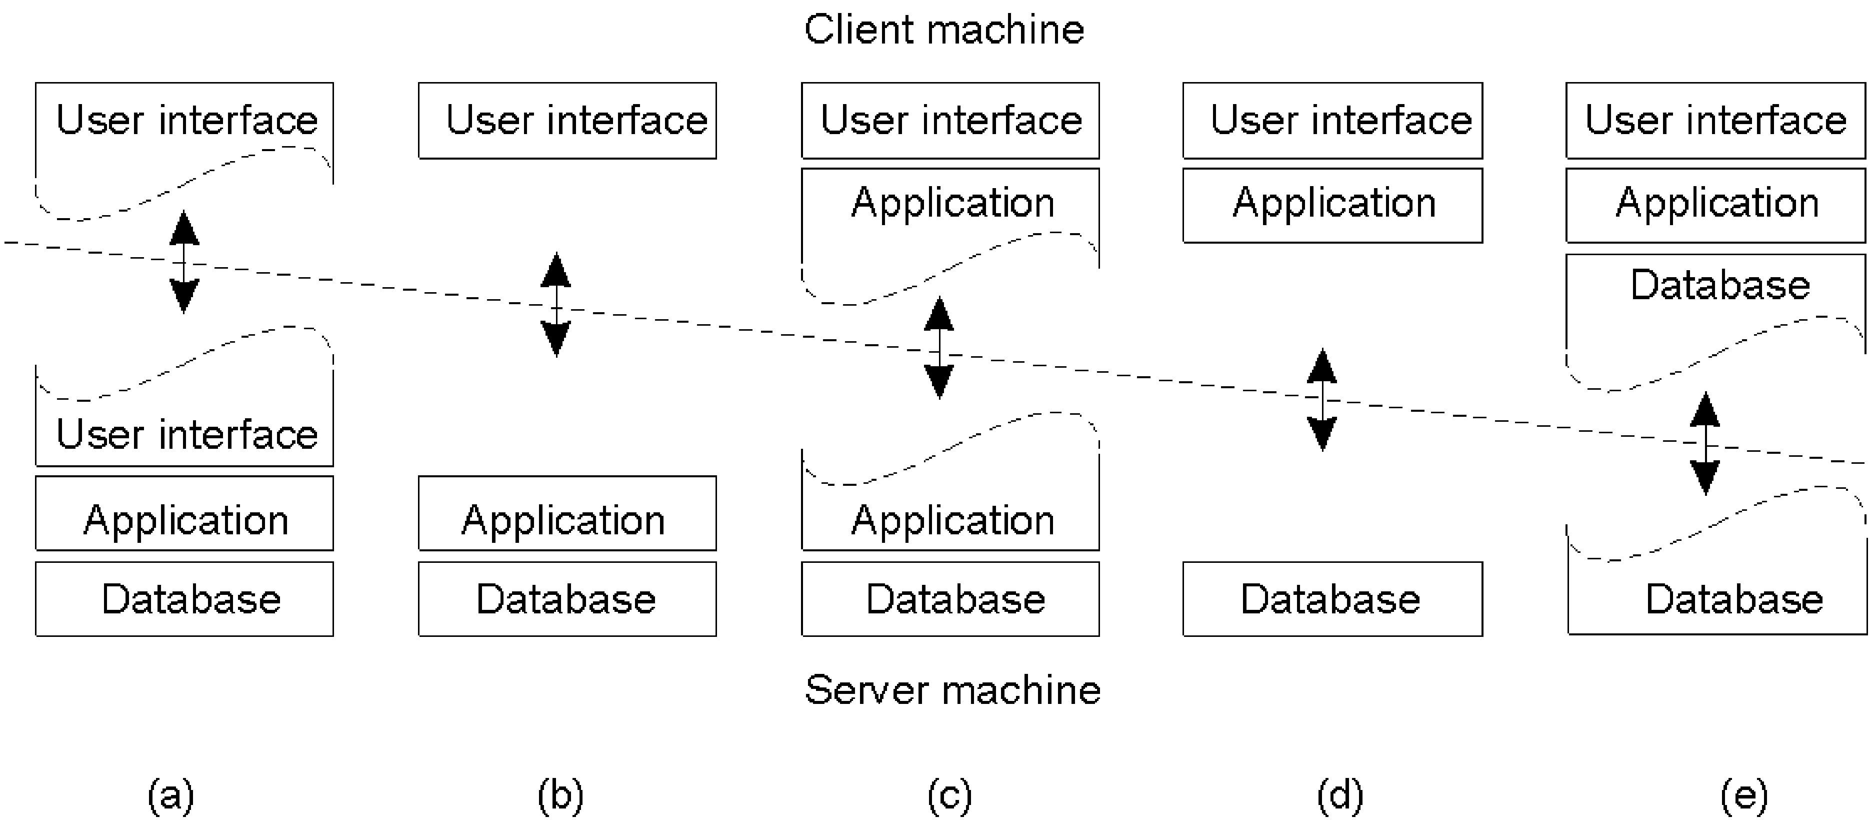

Multitiered Architectures (2)

- An example of a server acting as a client.

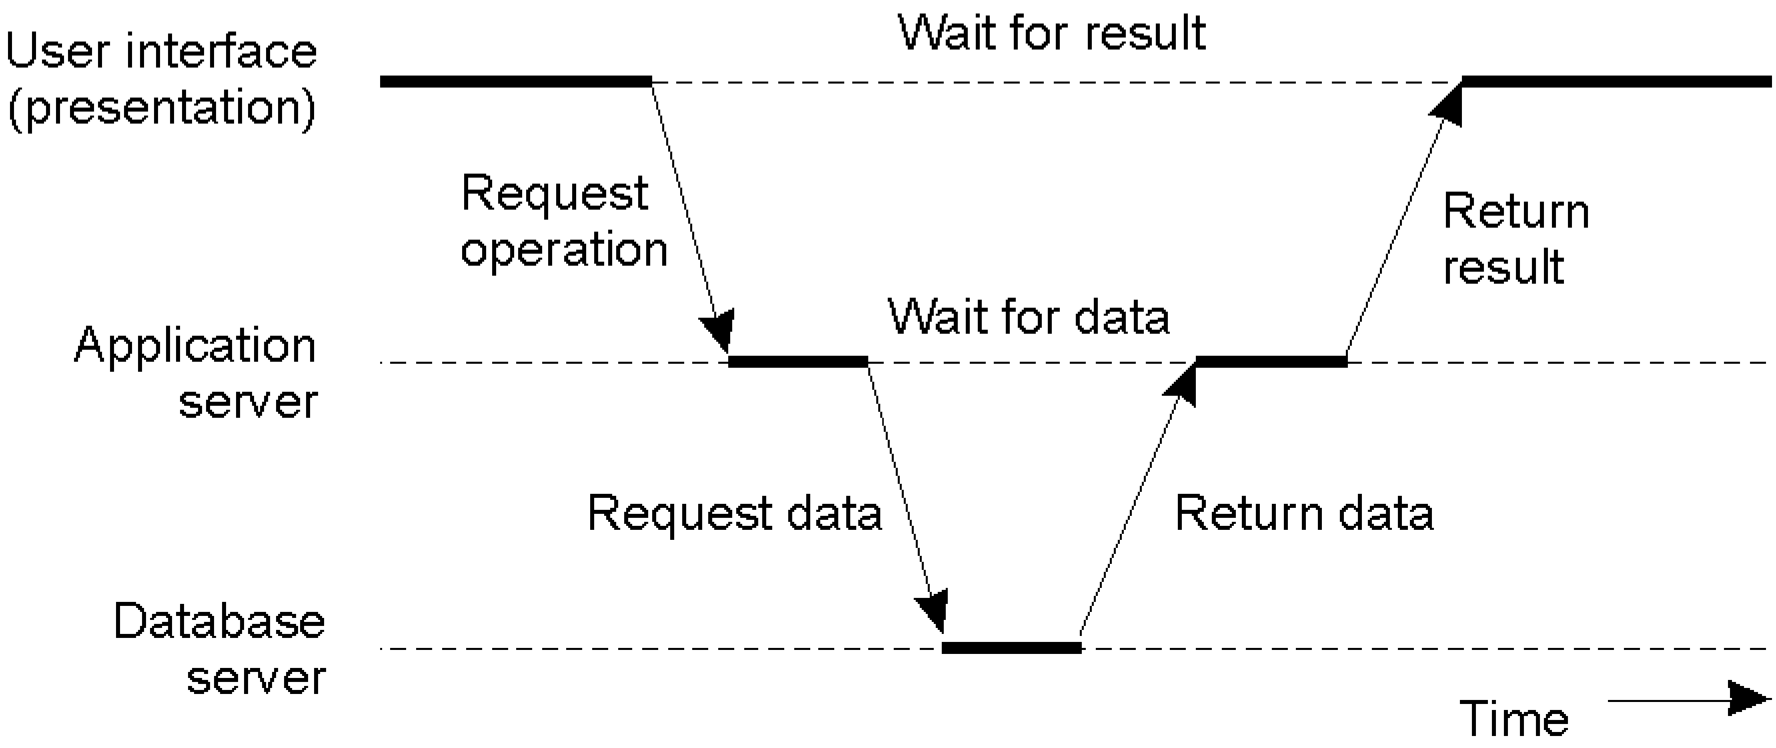

Modern Architectures

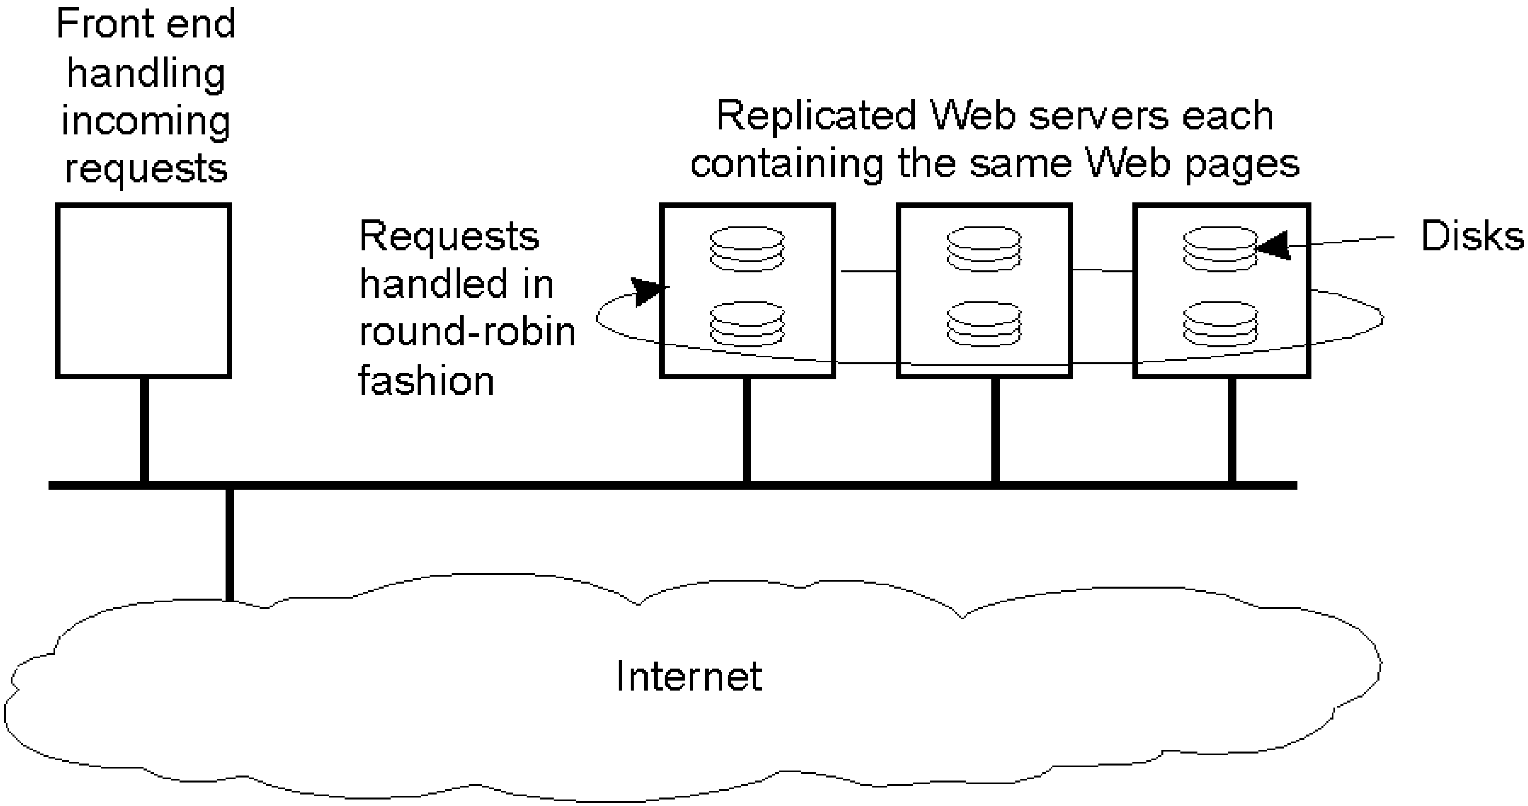

Vertical vs. Horizontal distribution

Fundamental models (CDK2.4): Synchronous distributed systems

Synchronous distributed systems: Hadzilacos and Toueg [1994] define a synchronous distributed system to be one in which the following bounds are defined:

- The time to execute each 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 a known bound.

another. The term *clock drift rate* refers to the rate at which a computer clock deviates from a perfect reference clock. Even if the clocks on all the computers in a distributed system are set to the same time initially, their clocks will eventually vary quite significantly unless corrections are applied.

Fundamental models (CDK2.4): Asynchronous distributed systems

Asynchronous distributed systems: Many distributed systems, such as the Internet, are very useful without being able to qualify as synchronous systems. Therefore we need an alternative model. An asynchronous distributed system is one in which there are no bounds on:

- Process execution speeds – for example, one process step may take only a picosecond and another a century; all that can be said is that each step may take an arbitrarily long time.
- Message transmission delays – for example, one message from process A to process B may be delivered in negligible time and another may take several years. In other words, a message may be received after an arbitrarily long time.
- Clock drift rates – again, the drift rate of a clock is arbitrary.