The UNIVERSITY *of* EDINBURGH
SCHOOL *of* COMPUTER SCIENCE

CS4/MSc

Distributed Systems

Jane Hillston

jeh@dcs.ed.ac.uk

Room 1604

(Lecture 2: Networking, 10th January 2002)

Communication Subsystem

The hardware and software within a distributed system which provides the communication facilities is known as the *communication subsystem*. It consists of:

Transmission media: providing the physical connectivity, e.g. wire, cable, fibre and wireless channels;

Hardware devices: providing the linkage, e.g. routers, bridges, hubs, repeaters, network interfaces and gateways;

Software components: managing the communication, e.g. protocol stacks, communication handlers and drivers.

Impact on Distributed Computing

The impact of the communication on a distributed computation will be one of delay — the delay introduced by the message passing. The delay experienced by each individual message can be broken down into two factors: *latency* and *transmission delay*.

Latency is the time which is necessary to set up the communication, i.e. it is the delay incurred from the time the message is sent until it starts to arrive at the destination.

The transmission delay, determined by the length of the message and the *data transfer rate*, is the time it takes to transfer the data of the message.

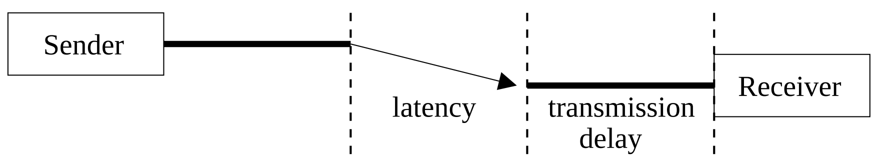

$$\text{message transmission time} = \text{latency} + \text{length}/\text{data transfer rate}$$

Types of Network

	Range	*Bandwidth(Mbps)*	*Latency(ms)*	*Reliability*
LAN	1-2 km	10–1000	1–10	high
Ethernets, token rings over twisted copper wire, coaxial cable or optical fibre; usually serves a single organisation; no switching, messages are broadcast.				
WAN	worldwide	0.010–600	100–500	high
dedicated computers act as routers and are the source of delay; transmission speeds approach the speed of light				
MAN	2-50 km	1–150	10	high
Ethernets or ATM networks over high-bandwidth copper and fibre optic cabling				
Wireless LAN	0.15-1.5 km	2–11	5–20	low
e.g. infra-red links for mobile devices; connection limited to the immediate vicinity				
Wireless WAN	worldwide	0.010–2	100–500	low
e.g. mobile phone technologies GSM and CDPD				

Switching schemes

Broadcast All messages are sent to all connected computers. Each computer is responsible extracting messages addressed to itself. This is the approach used in Ethernet and wireless networks and it requires no switches.

Circuit switching This is the approach taken in the telephone system where a physical link is established between the sender and the receiver (POTS).

Packet switching Otherwise known as *store-and-forward*. At each switching node (connection point) a computer manages the packets by reading each one into memory, examining its destination, and choosing an outgoing circuit appropriately.

Frame relay Reading in and storing the whole of each packet introduces a performance overhead which can become significant. In ATM networks a *frame* of fixed size is used in place of a packet and only its header needs to be examined. The remainder of the frame is simply relayed as a stream of bits.

Internetworking

Internetworks are networks built from subnetworks, which may themselves consist of subnetworks.... The Internet is the prime example.

When messages are passed across an internetwork, each constituent network may use a different transmission medium with differing switching requirements and reliability characteristics. However, it should appear as an intergrated network to the application programmer.

In most cases the problems associated with internetworking are solved by using IP addresses and the IP protocol.

Additionally specialised routers, known as *Internet Routers* are placed at points where subnetworks are connected. These machines will have distinct identities (IP addresses) within each subnetwork.

Bridges may also be used to link networks of different types, and the functionality of the bridge and the router may be combined in a single host. Similarly, routers often also serve as firewalls.

Internetwork example - the DCS system

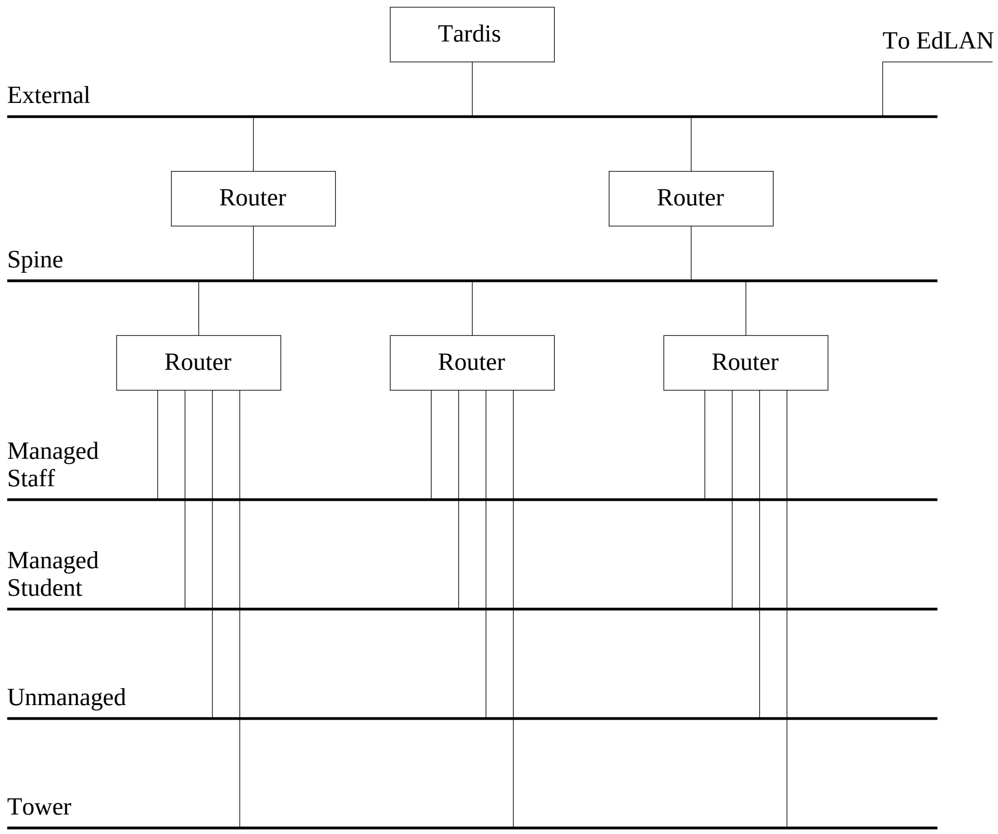

Internet Protocols

The Internet protocol stack works similarly to the ISO Open Systems Interconnections seven-layer reference model but is less rigidly defined, i.e. there is much more blurring between the layers. For example, the application layer protocol HTTP may be transported directly over TCP or the intermediate Secure Socket Layer (SSL) protocol may be introduced when needed.

IP datagrams form the basic form of communication over the Internet, but how these are actually sent over the network is not specified since this will depend on the technology of the underlying data links. For a single datagram this may change several times during transmission.

TCP and UDP provide two alternative transport layer protocols.

- TCP is designed to provide reliable, connection-oriented transmission and is quite sophisticated.
- UDP is much more basic, barely offering more than access to IP. The advantage of this is that it introduces very little overhead.

TCP/IP Layers

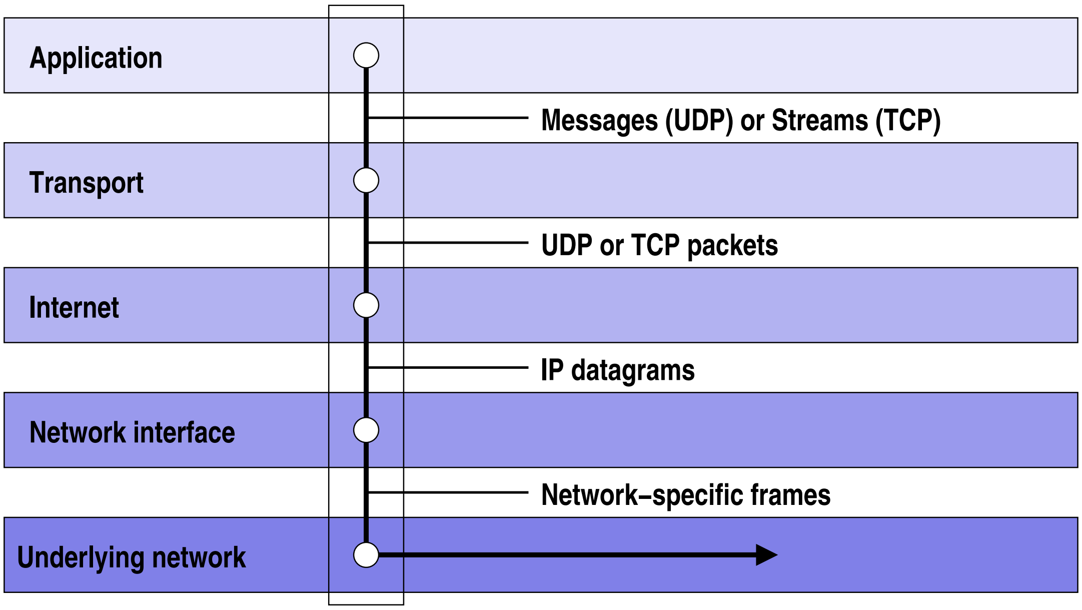

MobileIP

MobileIP allows IP communication to continue transparently with respect to the sever's current location.

- The mobile host is allocated a permanent IP address, corresponding to its "home" domain.
- When it is roaming, a *home agent* runs on a fixed machine in the home domain. Correspondingly a *foreign agent* running on a fixed machine at the temporary domain.
- The home agent keeps track of the current IP address of the mobile host and acts as a proxy during periods of disconnection.
- When the mobile machine is registered with the foreign agent, the foreign agent contacts the home agent, notifying it of the new temporary IP address.
- Requests for the server are captured by home agent and re-routed, embedded in MobileIP packets, to the foreign agent.