Networking and Internetworking (Basics)

3.1 Introduction
3.2 Types of Network
3.3 Network Principles

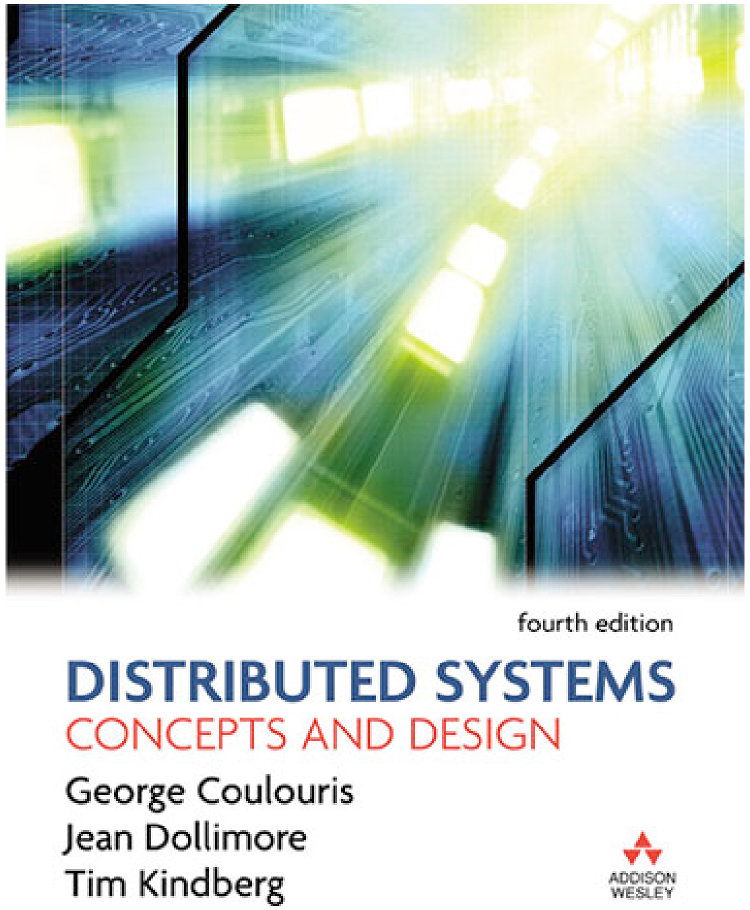

Introduction

- The networks used in distributed systems are built from a variety of
  - transmission media (wire, cable, fiber and wireless channels)
  - hardware devices (routers, switches, bridges, hubs, ...)
  - software components (protocol stacks, communication handlers and drivers)
- The resulting functionality and performance available to distributed system and application programs is affected by all these.

Some Terminology

- Communication subsystem: the collection of hardware and software components that provide the communication facilities for a distributed system.
- Host: computer or other device that use the network for communication purposes.
- Node: any computer or switching device attached to a network.
- Example: the Internet is a *single communication subsystem* providing communication between all of the hosts that are connected to it. The Internet is constructed from many *subnets*. *A subnet is a unit of routing*; it is a collection of nodes that can all be reached on the same physical network.

Network Performance Parameters

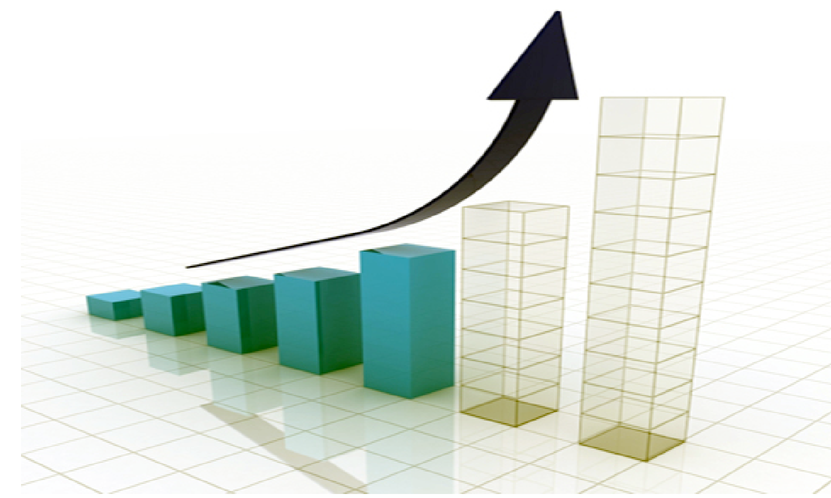

- The network performance parameters that are of primary interest are those affecting the *speed* with which individual messages can be transferred between interconnected computers. These parameters are:
  - Network Latency: *the delay that occurs after a send operation is executed before data starts to arrive at the destination computer*.
    - ✓ Measured as the time required to transfer an empty message
    - ✓ Determined primarily by software overheads, routing delays, ...
  - Data Transfer Rate: *the speed at which data can be transfered between two computers in the network once transmission has begun, usually in bits per second.*
    - ✓ Determined primarily by the network physical characteristics

Message Transmission Time

- Following the previous definitions, the time required for a network to transfer a message containing *length* bits between two computers is:

Message transmission time = latency + length/data transfer rate

- The above equation is valid for messages whose length does not exceed the maximum that is determined by the underlying network technology.

- Longer messages have to be segmented and the transmission time is the sum of the times for the segments.

- Many of the messages transferred between processes in distributed systems are small in size; *latency is therefore often of equal or greater significance than transfer rate in determining performance*.

Total System Bandwidth

- The *total system bandwidth* of a network is a measure of throughput: the total volume of traffic that can be transferred across the network in a given time.

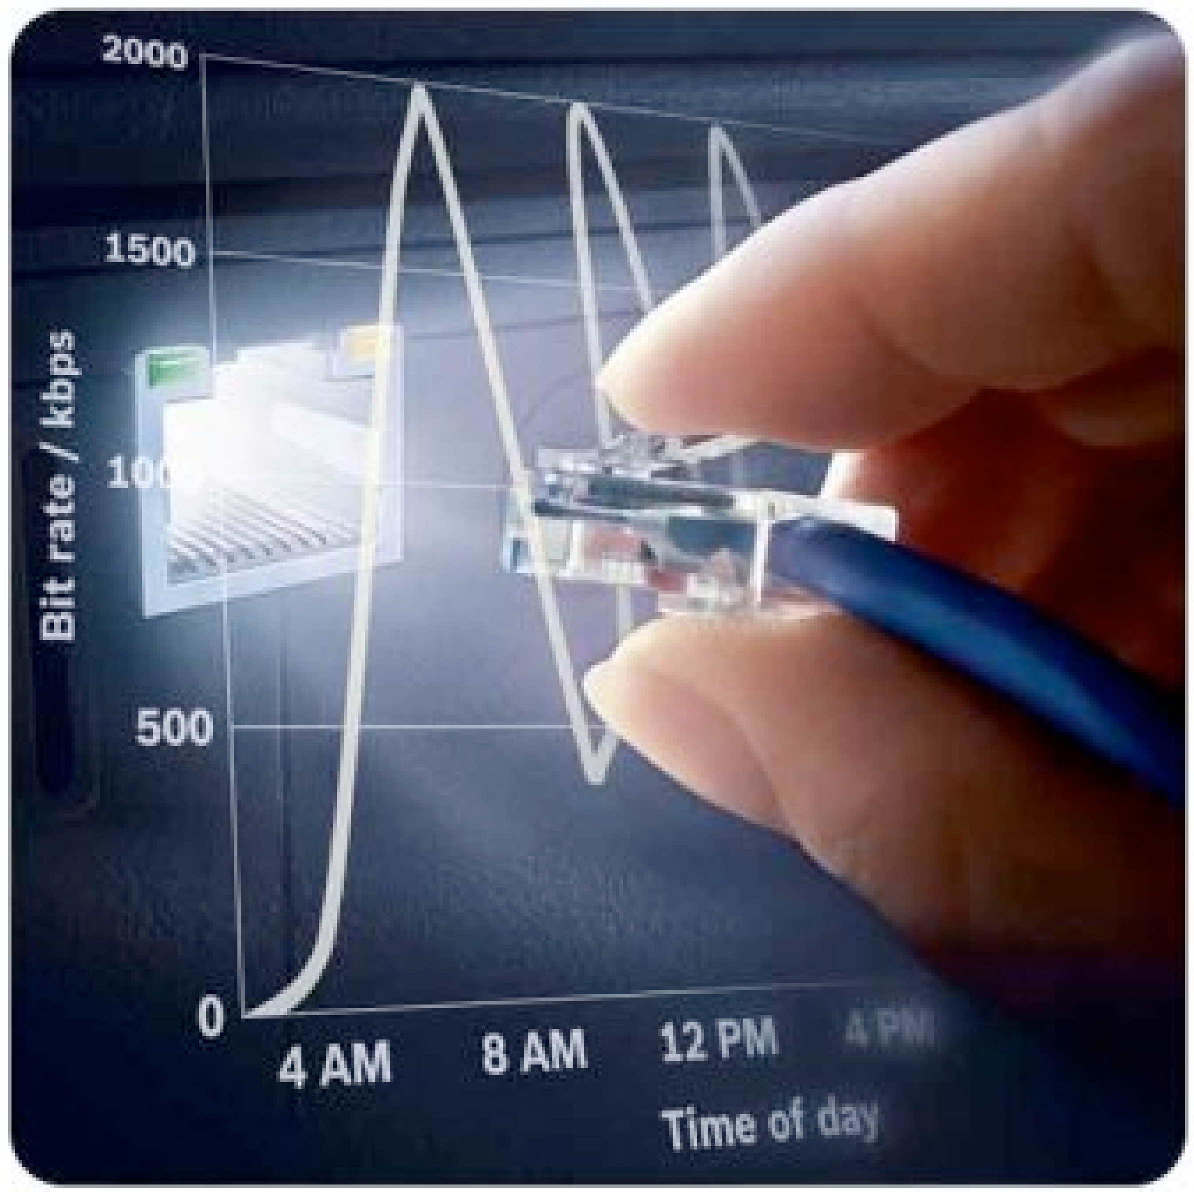

Networking and Internetworking (Basics)

3.1 Introduction
3.2 Types of Network
3.3 Network Principles

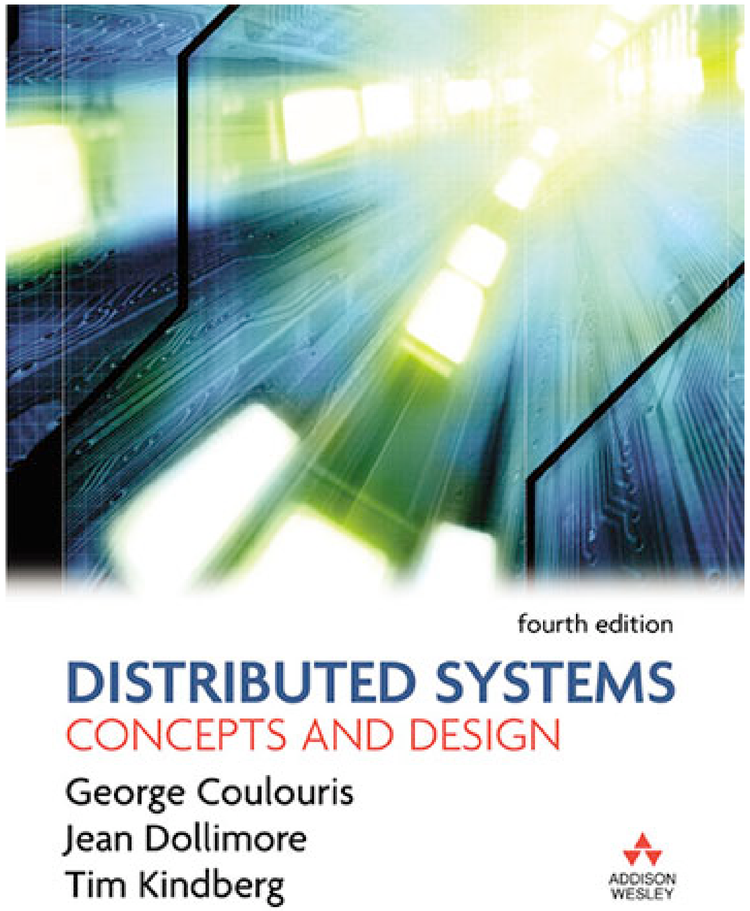

Types of Network

- Personal area networks, local area networks, wide area networks, metropolitan area networks.
- Internetworks: networks composed of many interconnected networks (of the above types), integrated to provide a single data communication medium.

	Example	*Range*	*Bandwidth (Mbps)*	*Latency (ms)*
Wired:				
LAN	Ethernet	1-2 kms	10-1000	1-10
WAN	IP routing	worldwide	0.010-600	100-500
MAN	ATM	250 kms	1-150	10
Internetwork	Internet	worldwide	0.5-600	100-500
Wireless:				
WPAN	Bluetooth (802.15.1)	10 - 30m	0.5-2	5-20
WLAN	WiFi (IEEE 802.11)	0.15-1.5 km	2-54	5-20
WMAN	WiMAX (802.16)	550 km	1.5-20	5-20
WWAN	GSM, 3G phone nets	worldwide	0.01-02	100-500

Personal Area Networks (PANs)

- PANs are a sub-category of local networks in which the *various digital devices carried out by a user are connected by a low-cost, low energy network.*

- Wireless Personal Area Networks (WPAN) are of increasing importance due to the number of personal devices such as mobile phones, PDA, digital cameras, music players and so on that are now carried by many people.

- Example: Bluetooth WPAN.

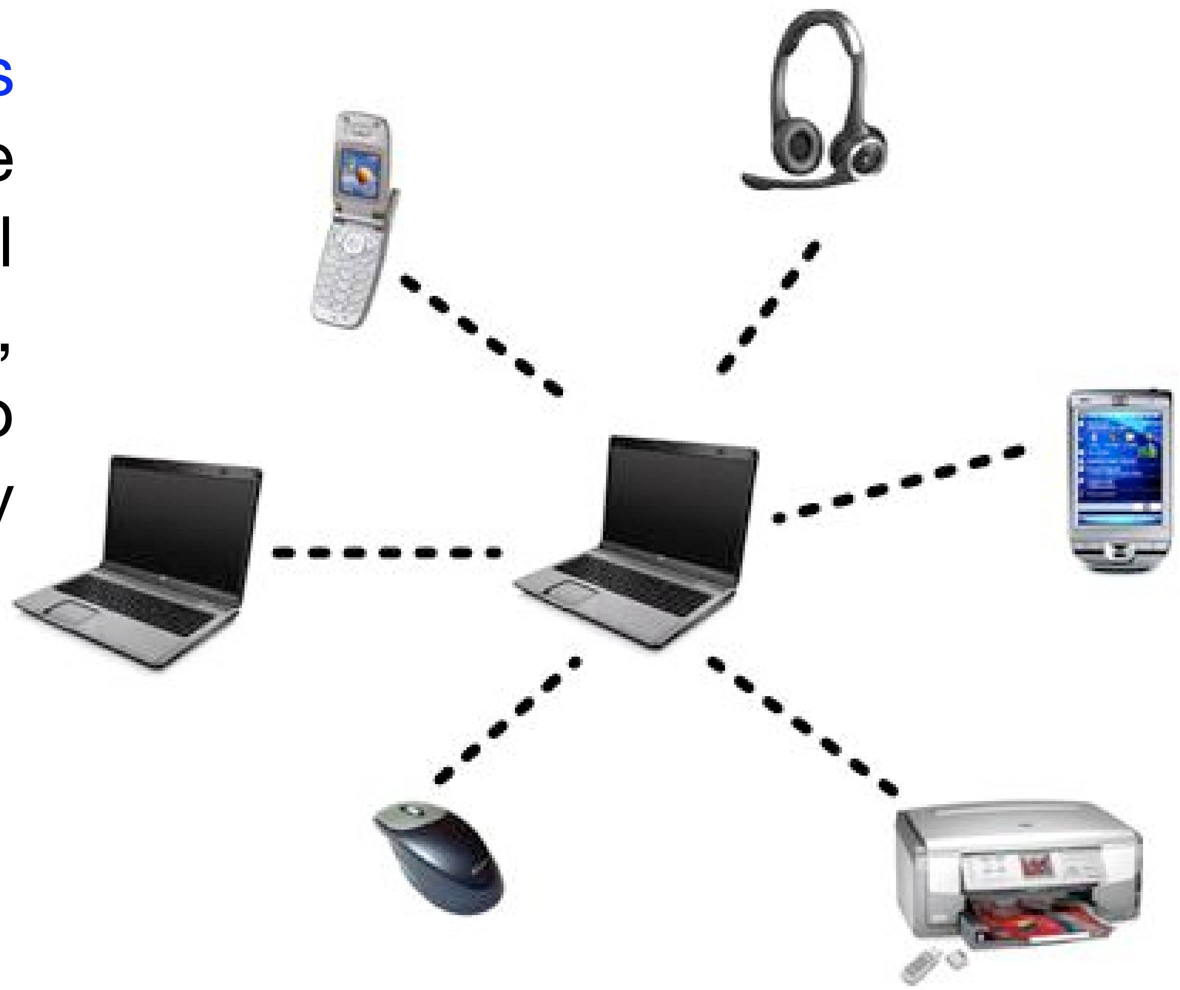

Local Area Networks (LANs)

- LANs *carry messages at relatively high speeds between computers connected by a single communication medium*, such as coaxial cable or optical fibre.
- A segment is a section of cable that serves a department or a floor of a building and may have many computers attached.
- No routing of messages is required within a segment, since the medium provides direct connections between all of the computers connected to it.

- Larger local area networks, such as whose that serve a campus or an office building, are composed of many segments interconnected by switches of hubs.

Wide Area Networks (WANs)

- WANs carry messages at lower speeds between nodes that are often in different organizations and may be separated by large distances. They may be located in different cities, countries and continents.

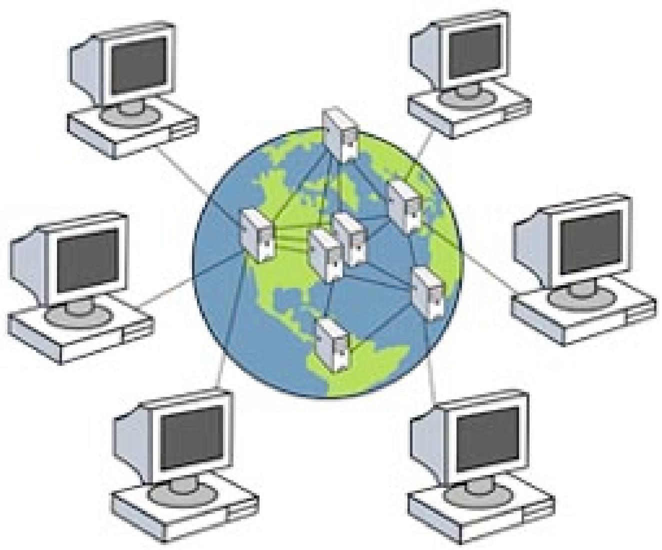

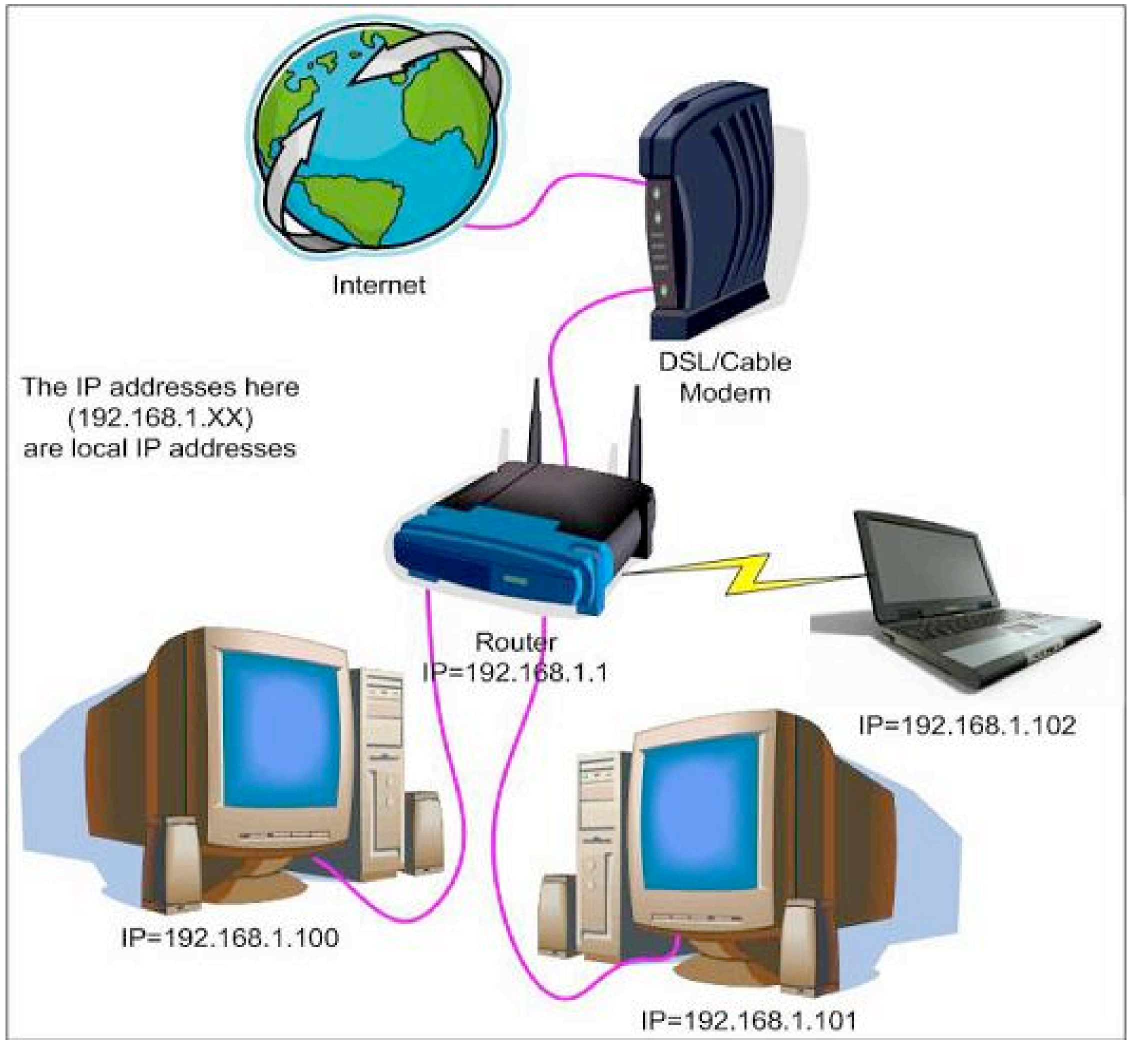

- The communication medium is a set of communication circuits linking a set of dedicated computers called routers.

- Routers manage the communication network and route messages or packets to their destinations.

Metropolitan Area Networks (MANs)

- This type of network is based on a high-bandwidth copper and fiber optic cabling installed in some towns and citied for the transmission of data over distances of up to 50 kilometers.

A MAN is optimized for a larger geographical area than a LAN, ranging from several blocks of buildings to entire cities. MANs can also depend on communications channels of moderate-to-high data rates. A MAN might be owned and operated by a single organization, but it usually will be used by many individuals and organizations. MANs might also be owned and operated as public utilities. They will often provide means for internetworking of local networks. Metropolitan area networks can span up to 50km, devices used are modem and wire/cable. IEEE 802-2001 standard

Networking and Internetworking (Basics)

Nicola Dragoni
Embedded Systems Engineering
DTU Informatics

- 3.1 Introduction
- 3.2 Types of Network
- 3.3 Network Principles

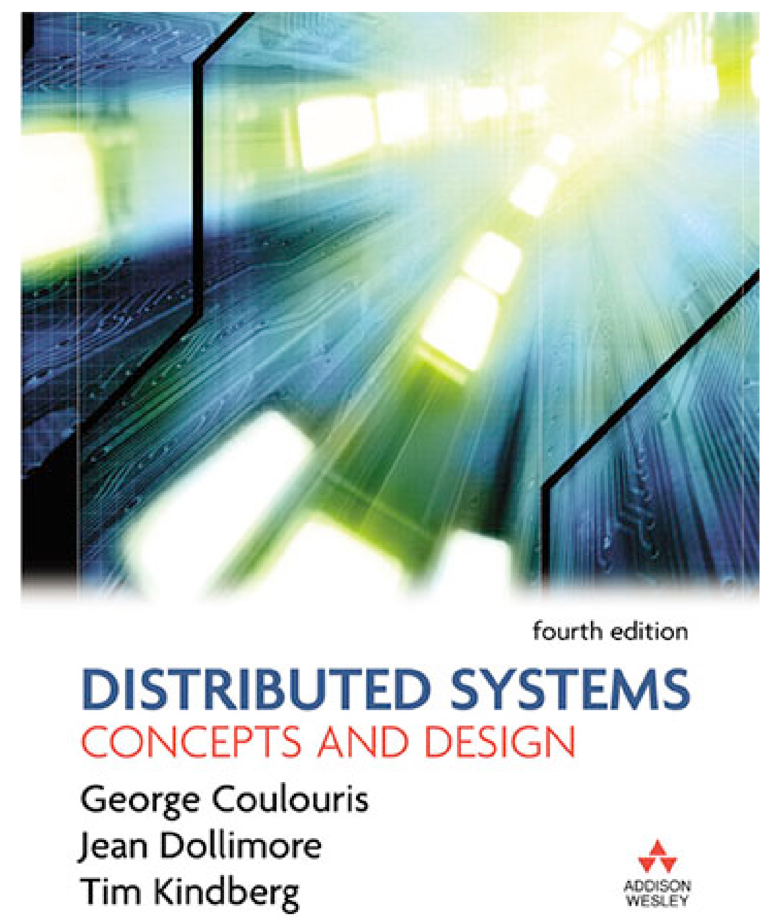

Packet Transmission

- Communication is asynchronous: messages arrive at their destination after a delay that varies depending upon the time that packets take to travel through the network.
- Messages: sequences of data items of arbitrary length.
- Before a message is transmitted it is subdivided into packets.
- The simplest form of packet is a sequence of binary data (an array of bits or bytes) of *restricted* length, together with addressing information sufficient to identify the source and destination computers.
- To transmit information between two arbitrary nodes, a *switching scheme* is required.

Switching Scheme: Broadcast

- Broadcasting is a transmission technique that involves no switching.
- Everything is transmitted to every node.
- It is up to potential receivers to notice transmissions addressed to them.
- Example: some LAN technologies, including Ethernet, are based on broadcasting.
- Example: wireless networking is necessarily based on broadcasting, but in the absence of fixed channels the broadcasts are arranged to reach nodes grouped in cells.

Switching Scheme: Circuit Switching

- If two processes want to communicate, a *permanent, dedicated and not-shared physical link (or **circuit**, or **path**)* is established between them.

- The packets are sent over this *dedicated* path.

- This link is allocated for the duration of the communication, and no other process can use that link during this period.

- Example: telephone system. Once a communication line has been opened between two parties (that is, party A calls party B), no one else can use this circuit until the communication is terminated explicitly (for example, when one of the parties hangs up).

Switching Scheme: Packet Switching

- Each packet is sent to its destination separately.
    - Each packet must include a source and destination address with its data.
- Each packet may take a different path through the network.
- The packets must be reassembled into messages as they arrive.

Switching Scheme: Circuit vs Packet Switching

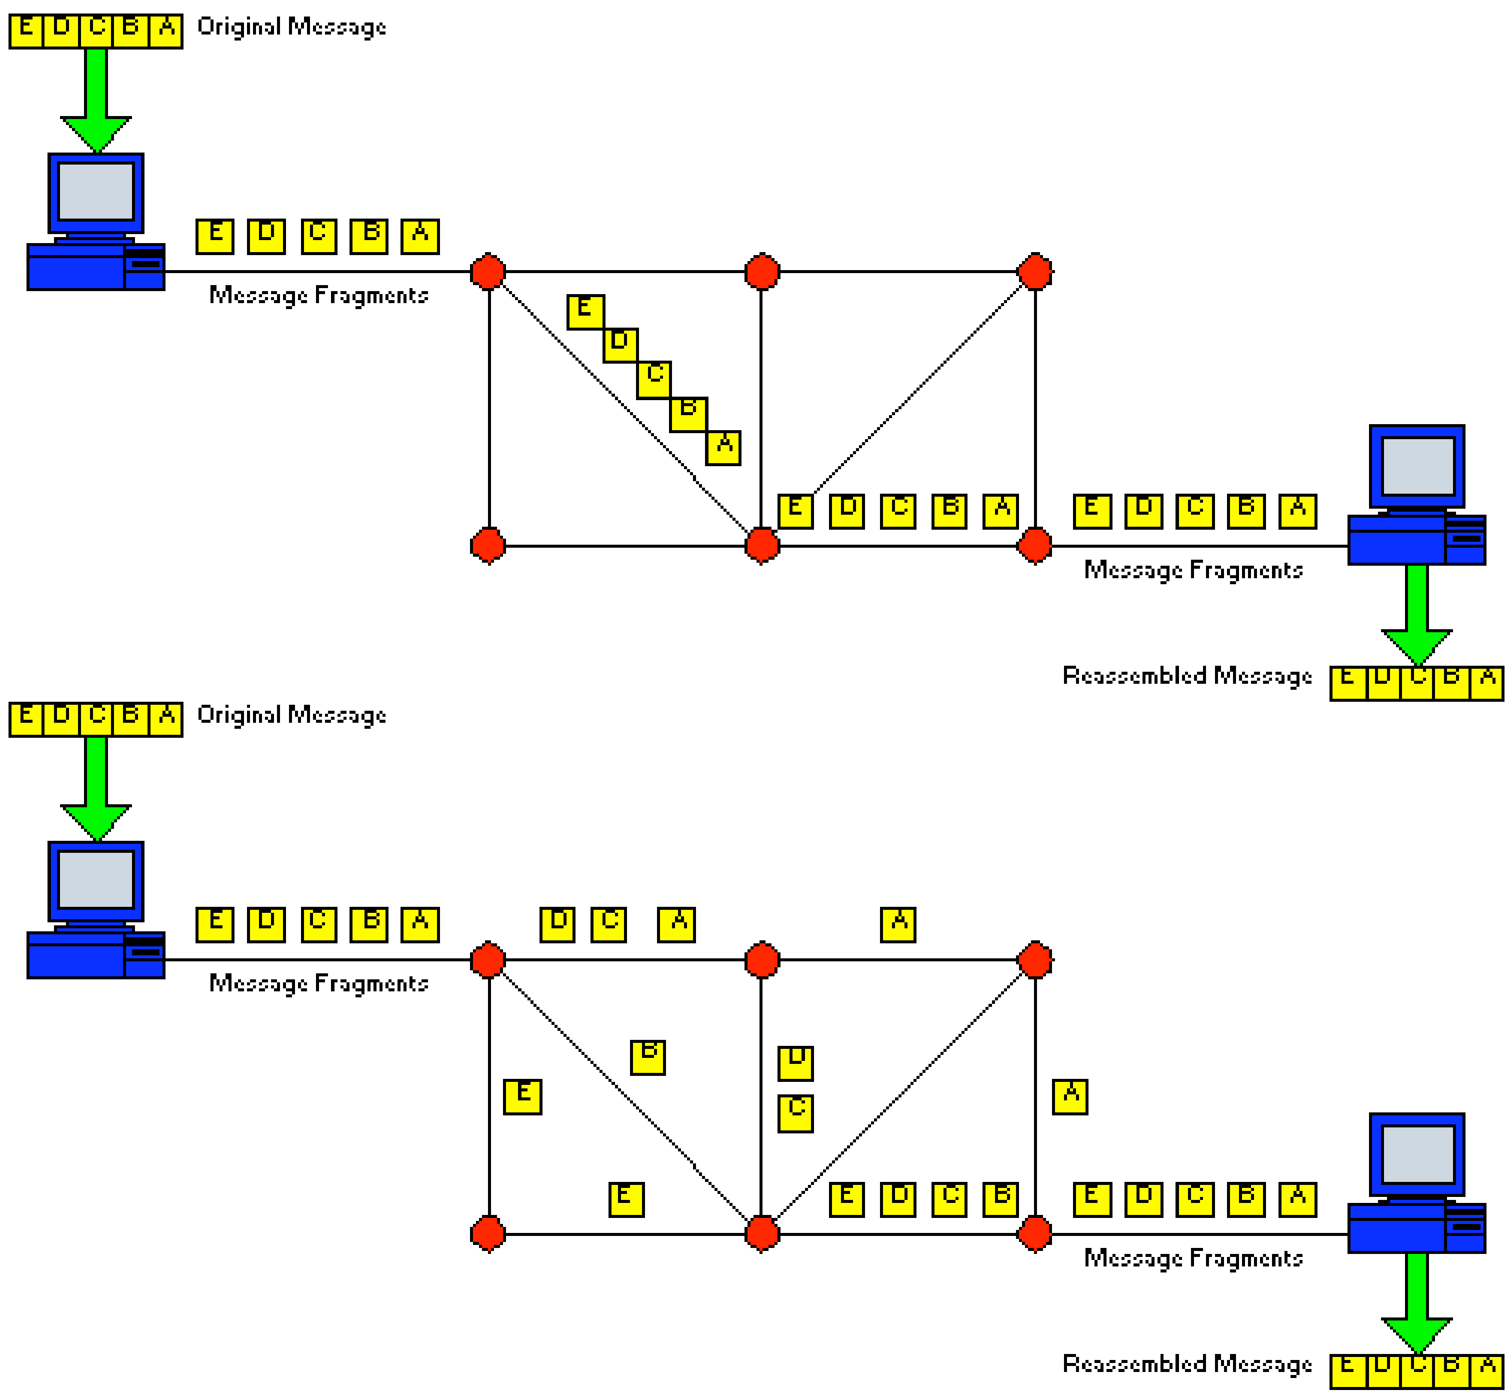

Protocols

- Protocol: *a well-known set of rules and formats to be used for communication between processes in order to perform a given task*.
- The definition of a protocol has two important parts to it:
  - a specification of the sequence of messages that must be exchanged
  - a specification of the format of the data in the messages
- Why protocols? The existence of well-known protocols enables the separate software components of distributed systems to be developed independently and implemented in different programming languages on different computers (having different data representations).

Protocol Layers

- Network software is arranged in a hierarchy of layers.
- Each layer presents an interface to the layers above it that *extends* the properties of the underlying communication system.
- A layer is represented by a module in every computer connected to the network.

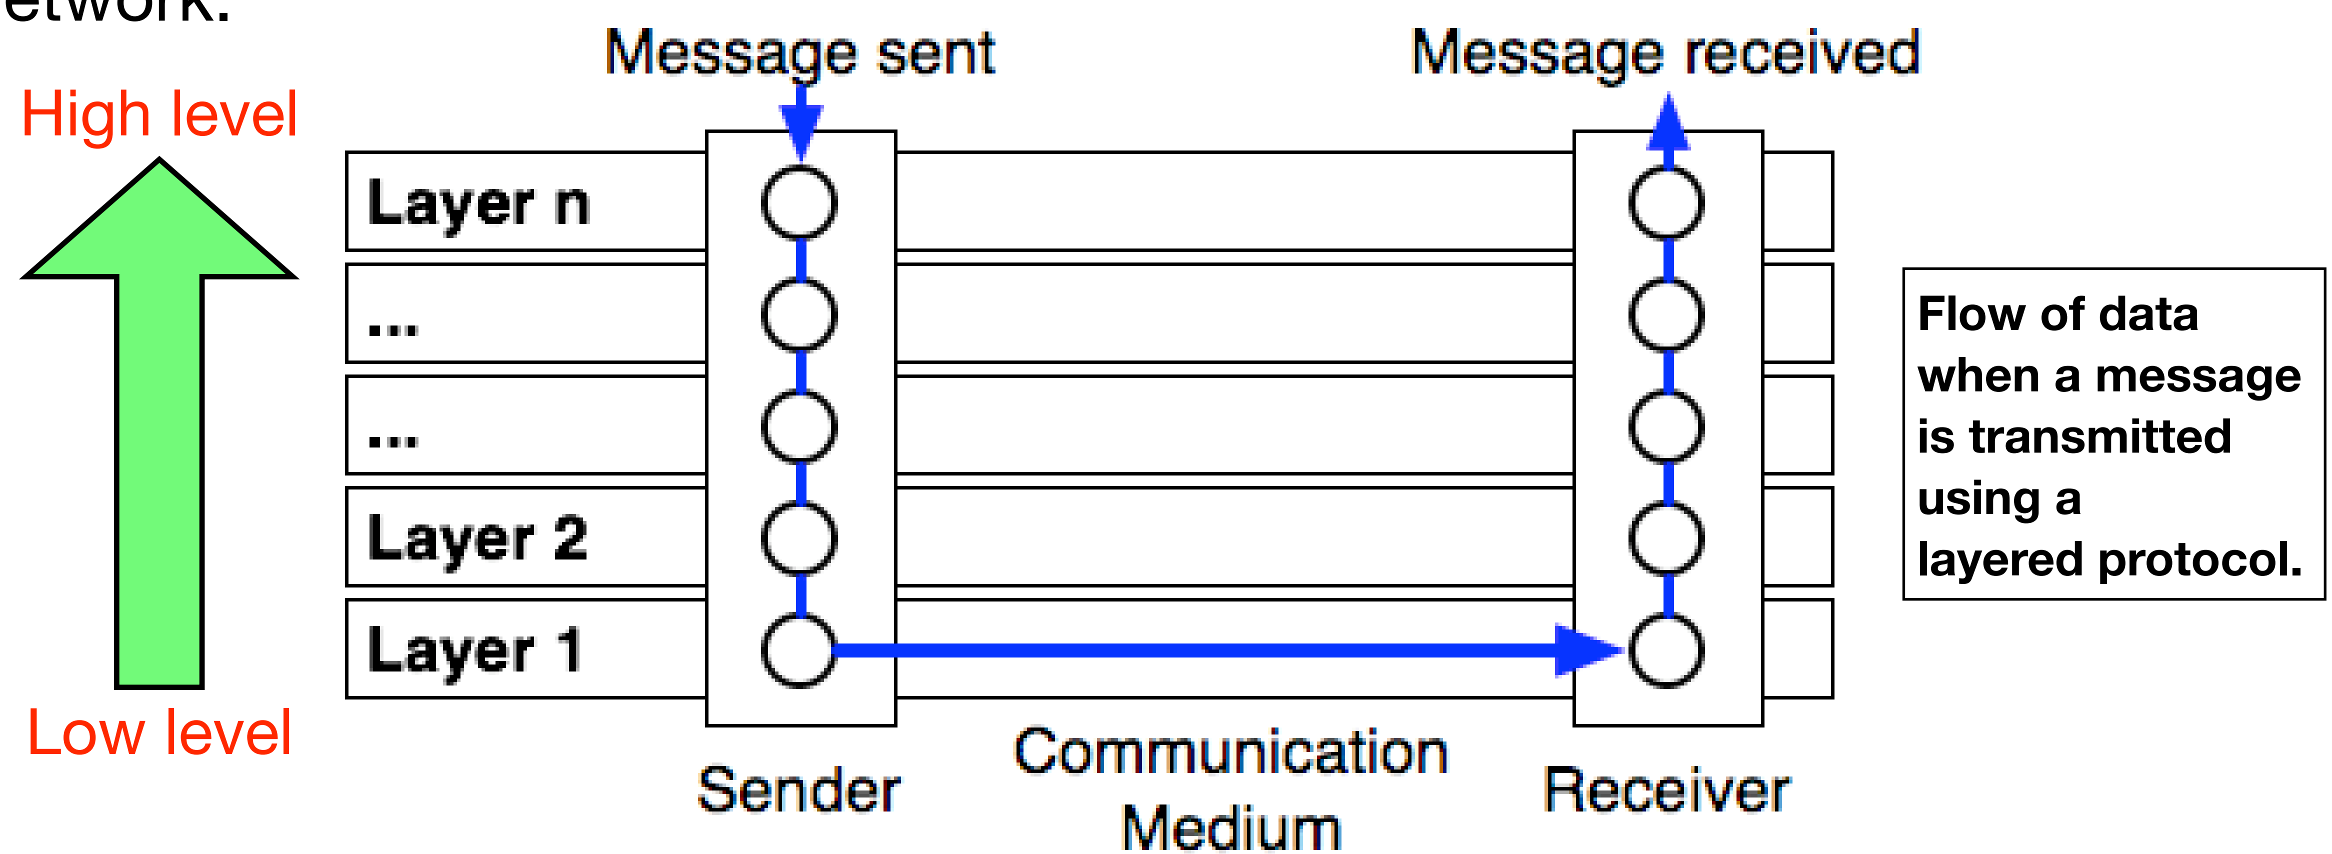

Flow of Data in a Layered Protocol

- Each module appears to communicate directly with with a module at the same level in another computer in the network...

- BUT in reality data is NOT transmitted directly between the protocol modules at *each level*.

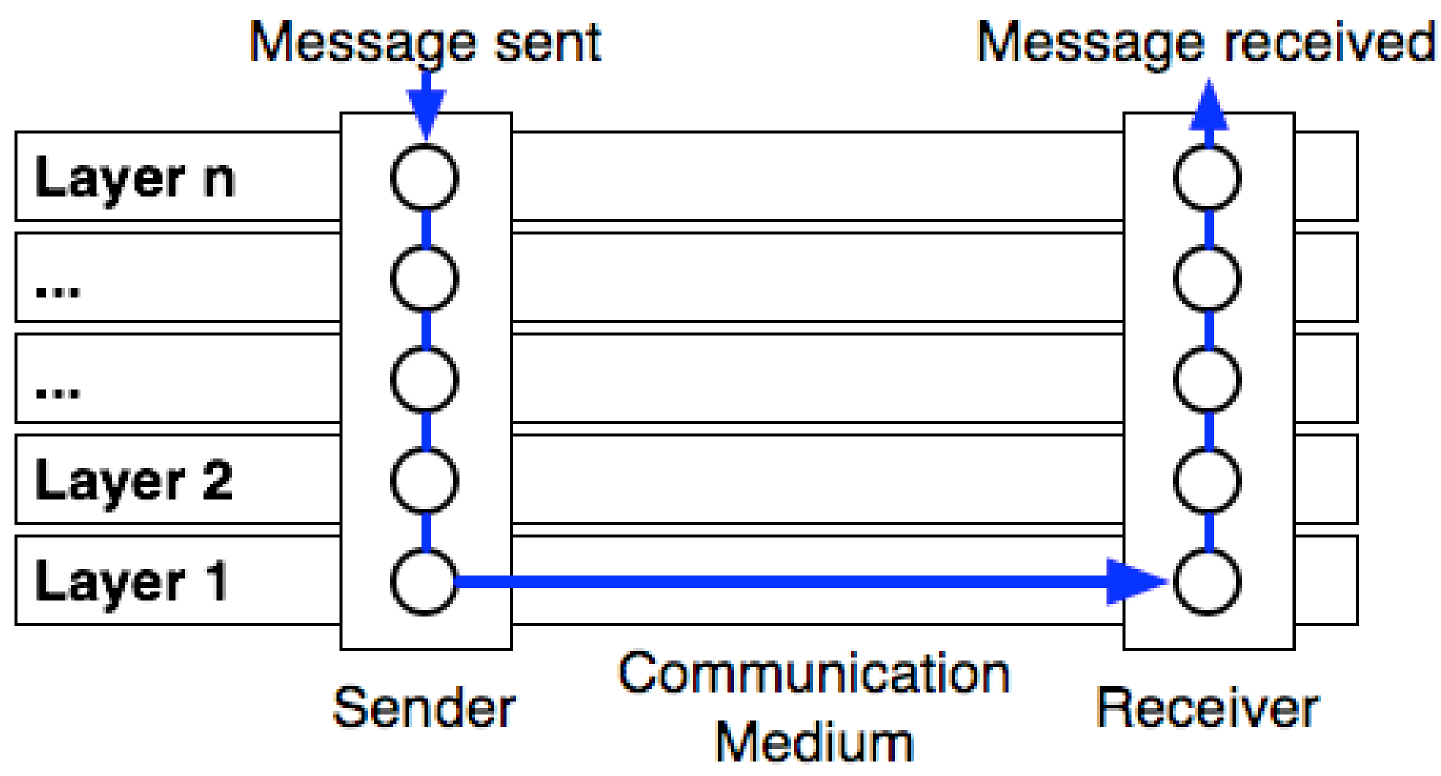

- Each layer of network software communicates by local procedure calls with the layers above and below is.

Flow of Data in a Layered Protocol: Sending Site

- Each layer (except the topmost):
  - accepts items of data in a specified format from the layer above it
  - applies transformations to encapsulate the data in the format specified for that layer before passing it to the layer below for further processing.

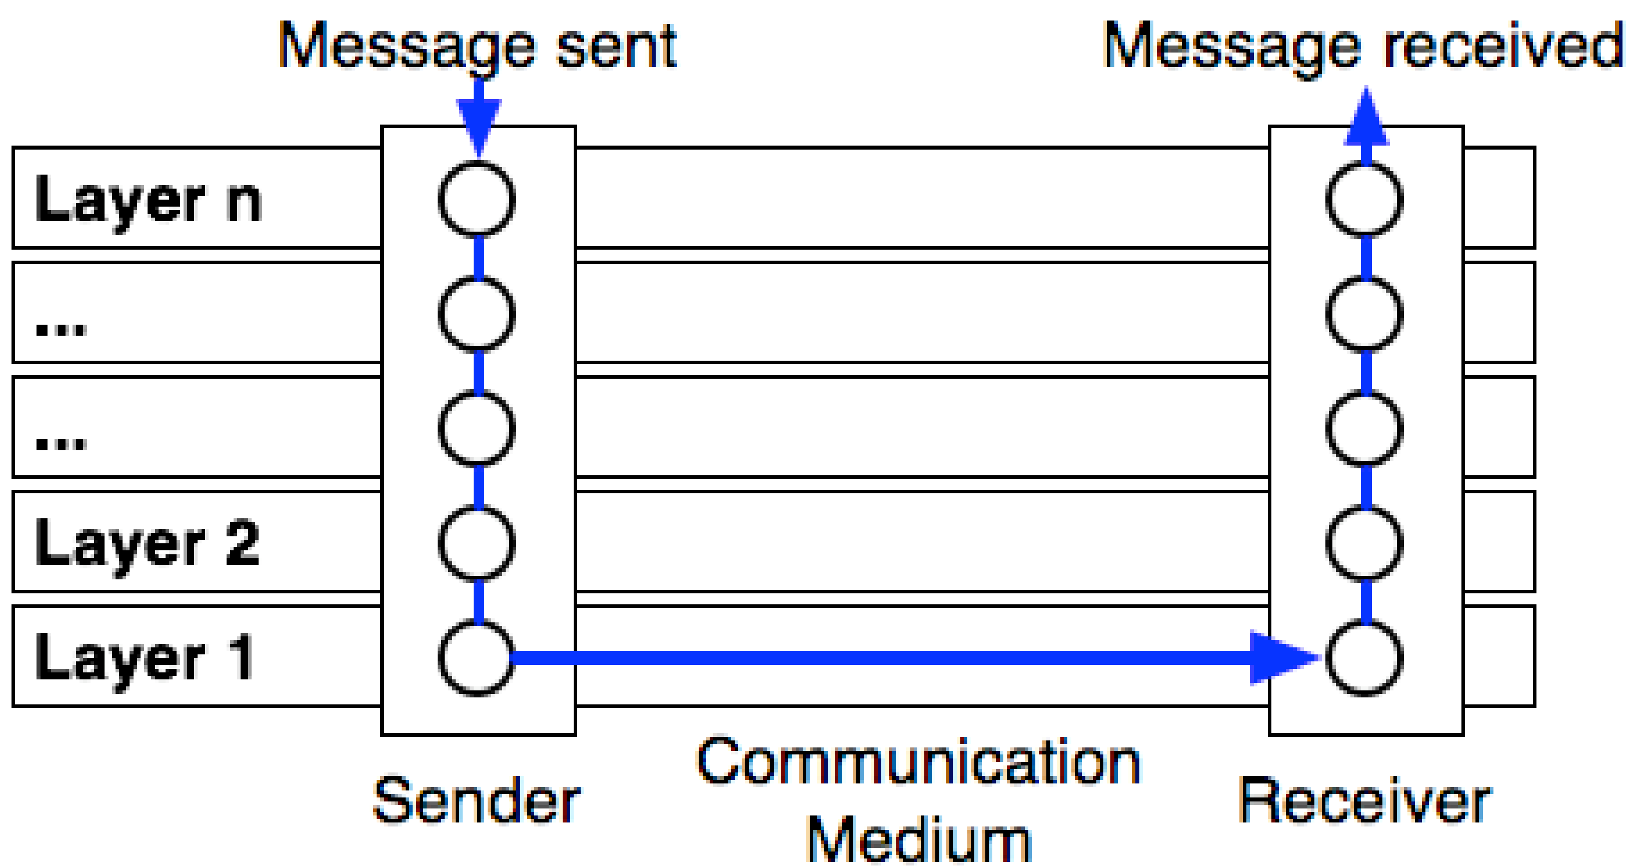

Flow of Data in a Layered Protocol: Receiving Site

- The converse transformations are applied to data items received from the layer below before they are passed to the layer above.

- The protocol type of the layer above is included in the header of each layer, to enable the protocol stack at the receiver to select the correct software components to unpack the packets.

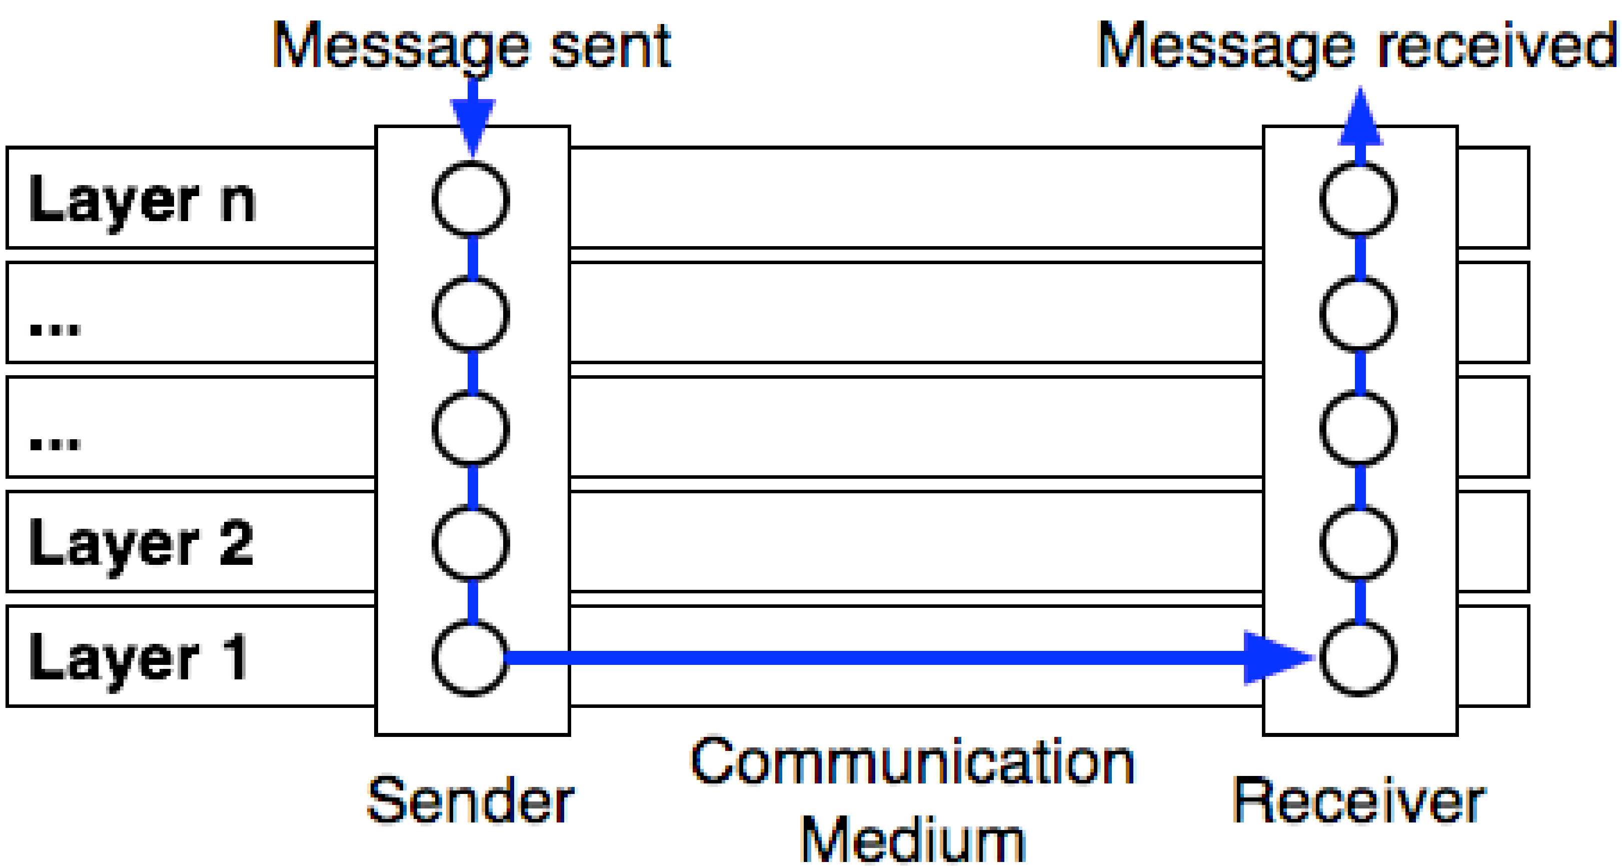

Protocol Layer as a Service

- Each layer provides a service to the layer above it and extends the service provided by the layer below it.

- At the bottom is a *physical layer*, implemented by a communication medium (copper or fiber optic cables, satellite communication channels, radio transmission, ...).

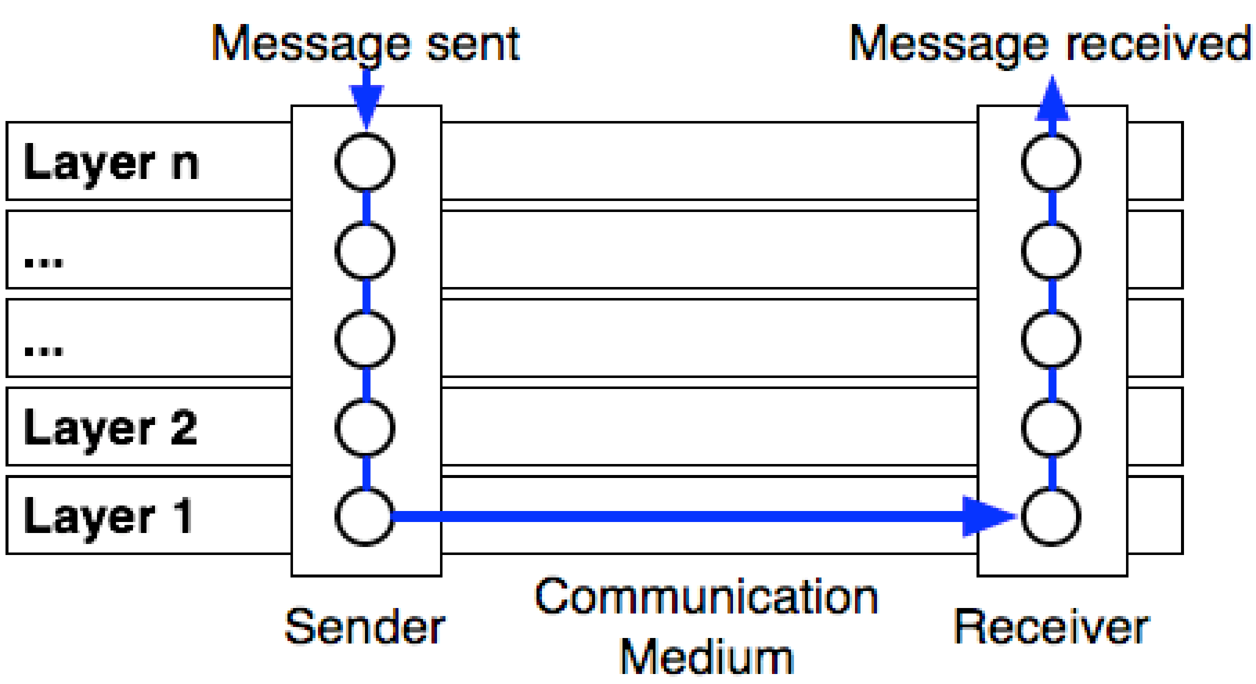

- At receiving nodes, data items are received and passed upwards through the hierarchy of software modules, transformed at each stage until they are in a form that can be passed to the intended recipient process.

Protocol Stack

- A complete set of protocol layers is referred as a protocol stack (or suite).

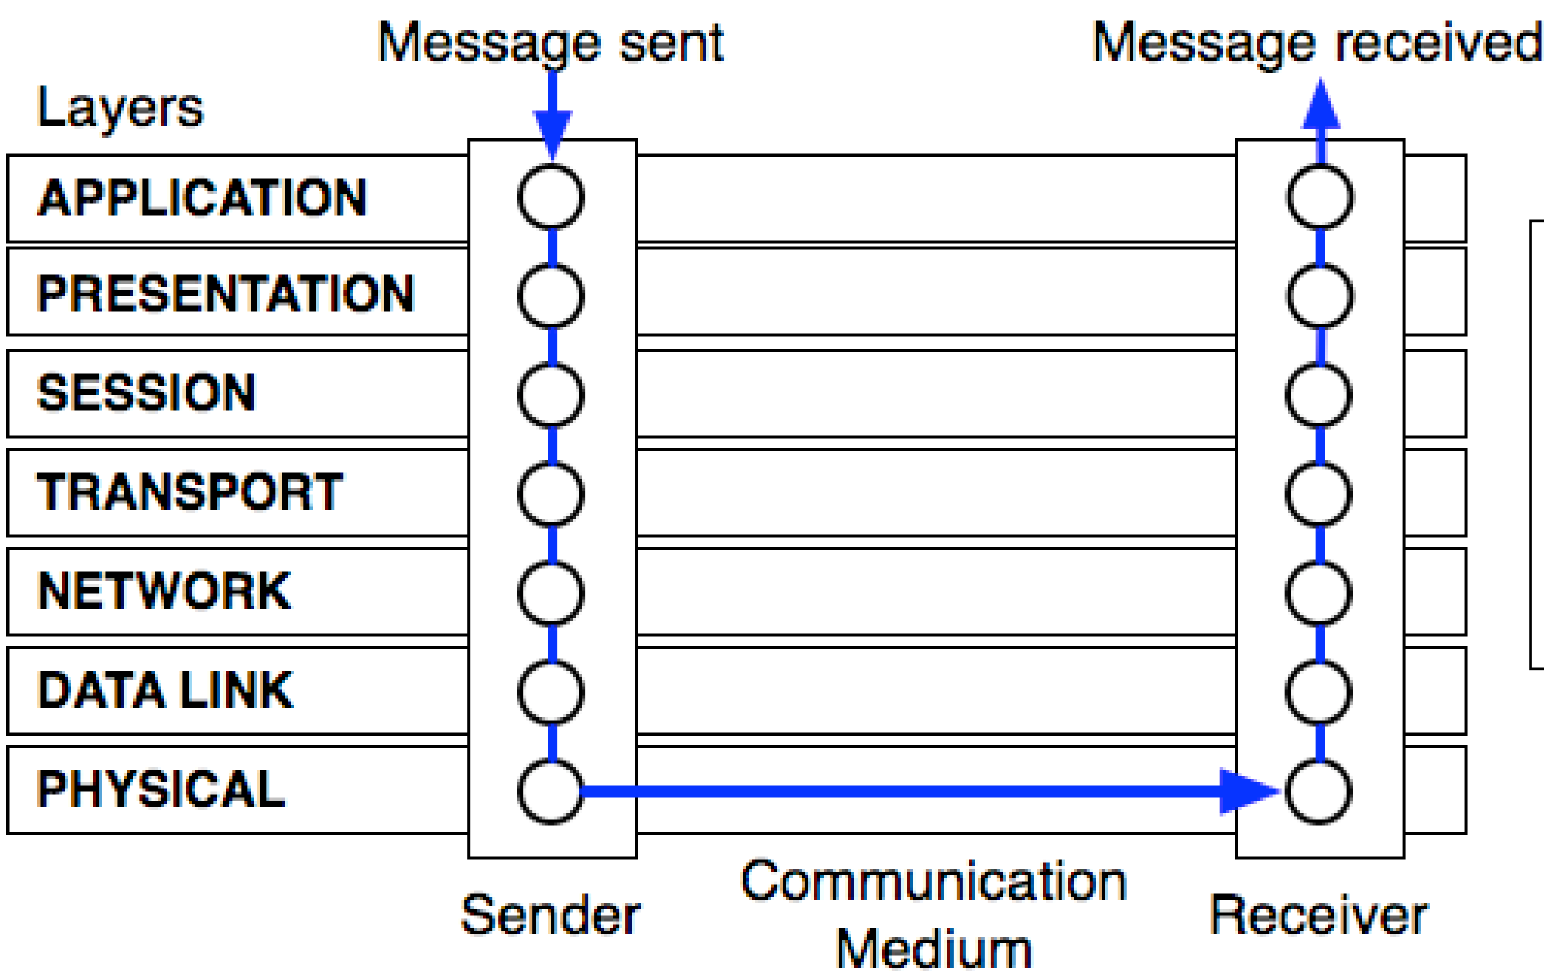

Protocol layers in the ISO (International Organization for Standardization) Open System Interconnection (OSI) Reference Model.

- OSI Reference Model: proposed by ISO to encourage the development of protocol standards that would meet the requirements of open systems.

OSI Protocol Summary

Layer	Description	Examples
Application	Protocols that are designed to meet the communication requirements of specific applications, often defining the interface to a service.	HTTP, FTP, SMTP
Presentation	Protocols at this level transmit data in a network representation that is independent of the representations used in individual computers, which may differ. Encryption is also performed in this layer, if required.	TLS security, CORBA Data Rep.
Session	At this level reliability and adaptation are performed, such as detection of failures and automatic recovery.	SIP
Transport	This is the lowest level at which messages (rather than packets) are handled. Messages are addressed to communication ports attached to processes, Protocols in this layer may be connection-oriented or connectionless.	TCP, UDP
Network	Transfers data packets between computers in a specific network. In a WAN or an internetwork this involves the generation of a route passing through routers. In a single LAN no routing is required.	IP, ATM virtual circuits
Data link	Responsible for transmission of packets between nodes that are directly connected by a physical link. In a WAN transmission is between pairs of routers or between routers and hosts. In a LAN it is between any pair of hosts.	Ethernet MAC, PPP
Physical	The circuits and hardware that drive the network. It transmits sequences of binary data by analogue signalling, using amplitude or frequency modulation of electrical signals (on cable circuits), light signals (on fibre optic circuits) or other electromagnetic signals (on radio and microwave circuits).	Ethernet base-band signalling, ISDN

Routing

- Routing is a function that is required in all networks except those LANs, such as the Ethernet, that provide direct connections between all pairs of attached hosts.

- In large networks, adaptive routing is employed: the best route for communication between two points in the network is *re-evaluated periodically*, taking into account the current traffic in the network and any faults such as broken connections or routers.

- The delivery of packets to their destinations in a network is the collective responsibility of the routers located at connection points.

- Unless the source and destination hosts are on the same LAN, the packet has to be transmitted in a series of hops, passing through router nodes.

Routing Algorithm

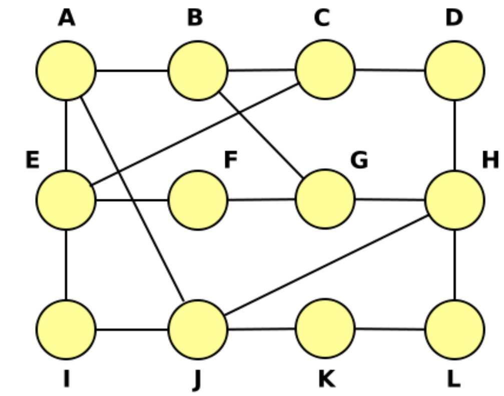

- The determination of routes for the transmission of packets to their destination is the responsibility of a routing algorithm.

- Implemented by a program in the network layer at each node.

- A routing algorithm has two parts:

  1. it must take decisions that determine the route taken by each packet as it travels through the network

  2. it must dynamically update its knowledge of the network based on traffic monitoring and the detection of configuration changes or failures.

     ➡ (Example: routing table construction and maintenance)

A Simple Routing Algorithm: Distance Vector

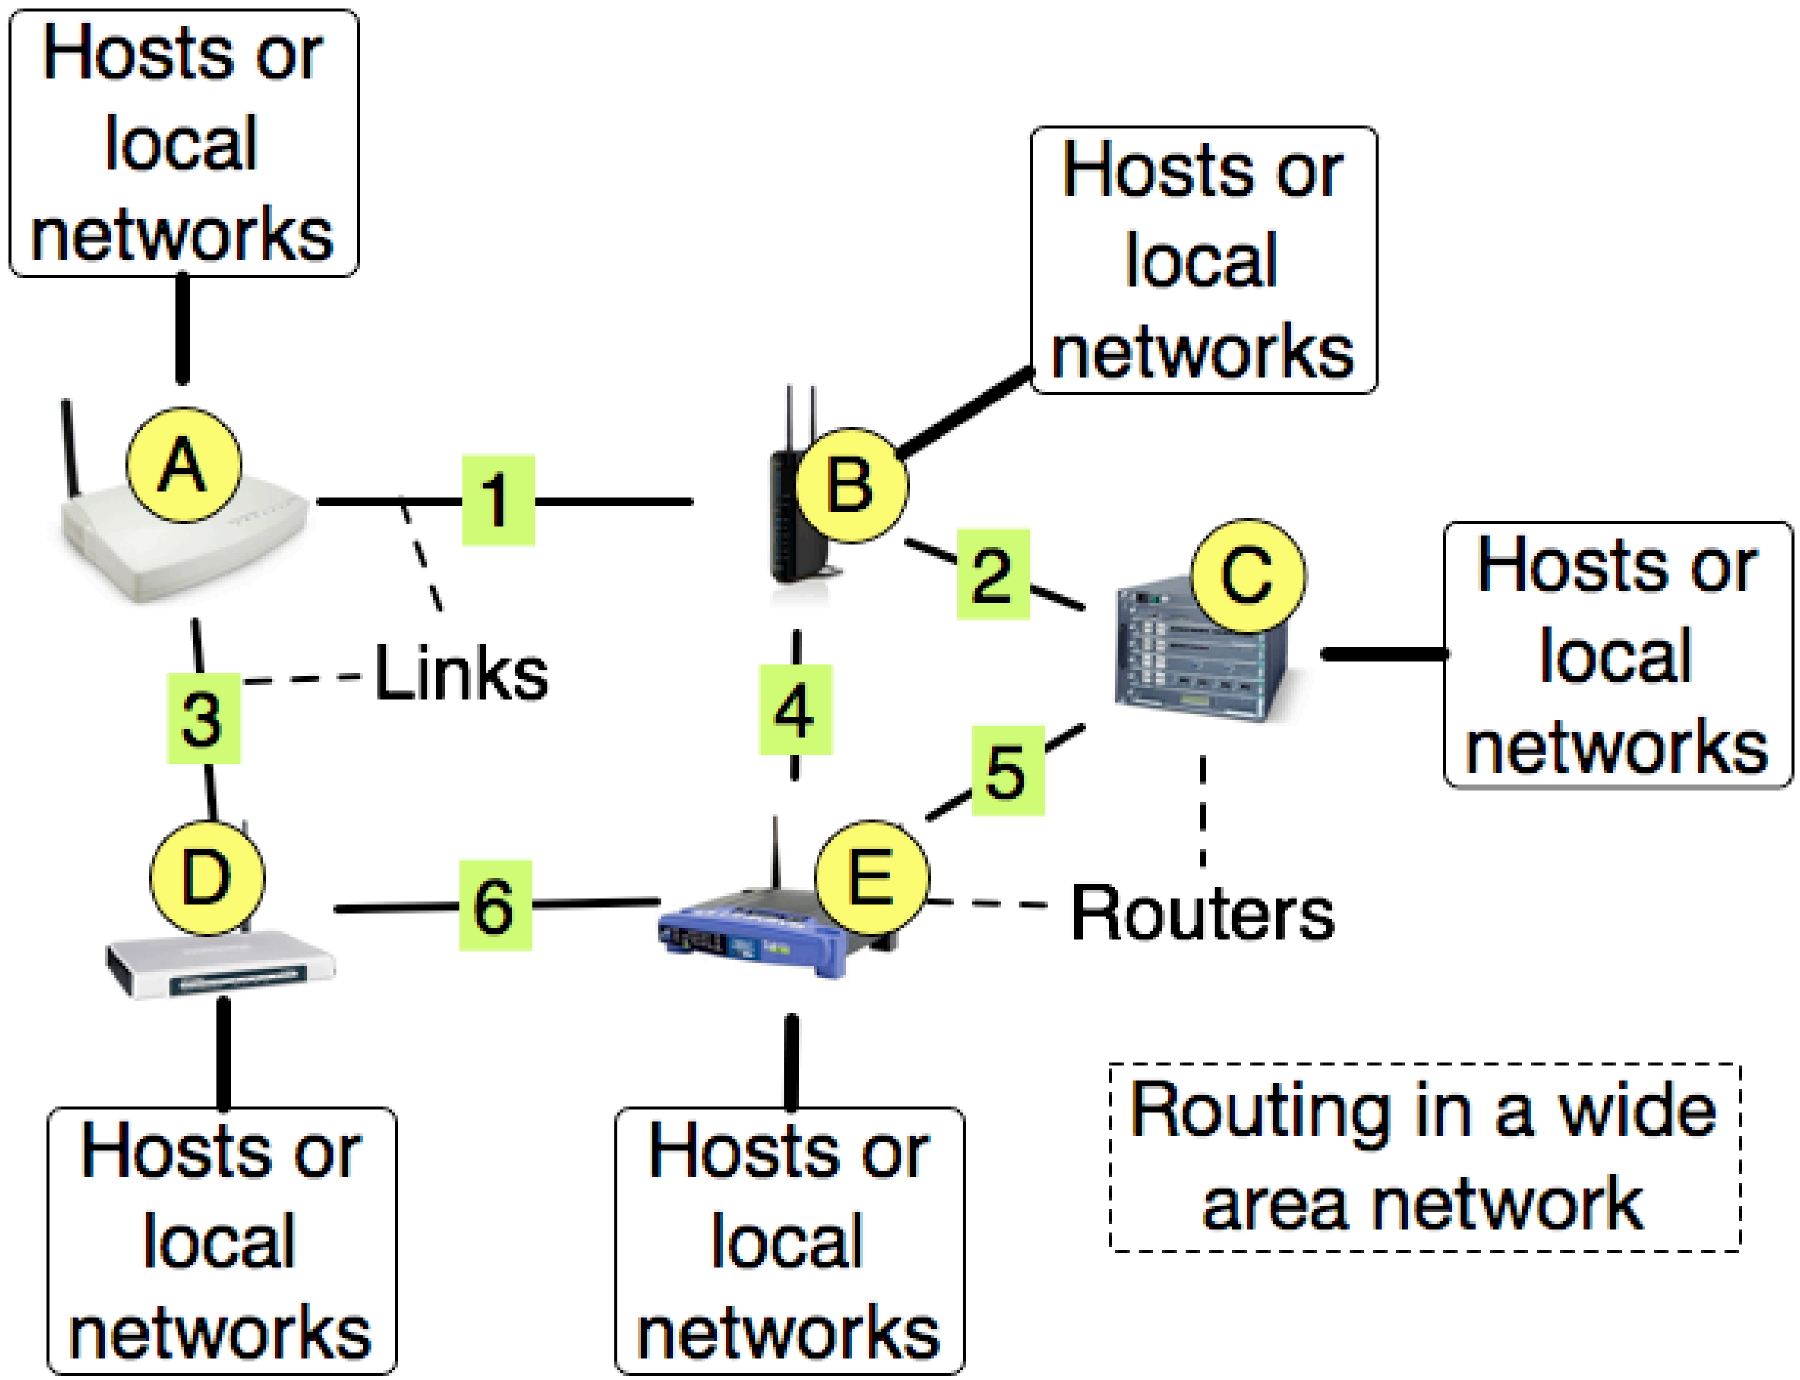

[Distance Vector Algorithm] Routing Tables

Routing Table

Routings from A

To	*Link*	*Cost*
A	local	0
B	1	1
C	1	2
D	3	1
E	1	2

Destination

Outgoing link for packets addressed to the destination

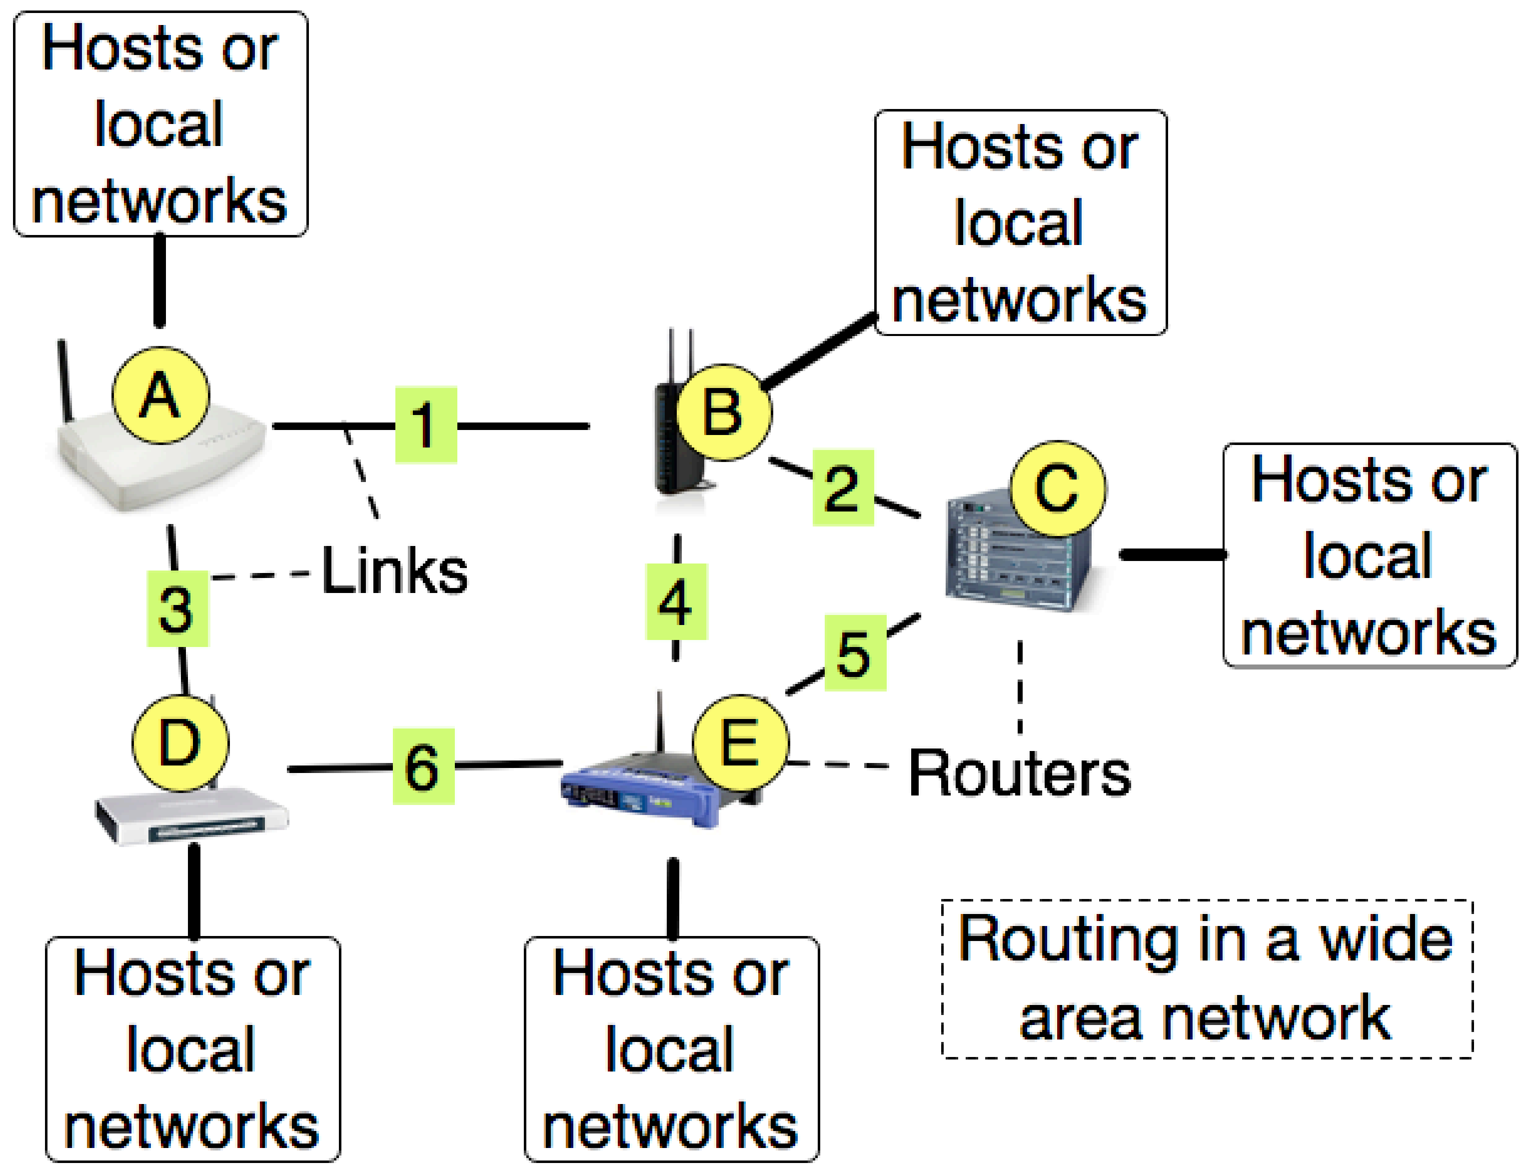

Calculation of the vector distance (number of hops to the given destination)

[Distance Vector Algorithm] Routing Actions Example

Part 1 of the routing algorithm

Routings from A

To	*Link*	*Cost*
A	local	0
B	1	1
C	1	2
D	3	1
E	1	2

Routings from B

To	*Link*	*Cost*
A	1	1
B	local	0
C	2	1
D	1	2
E	4	1

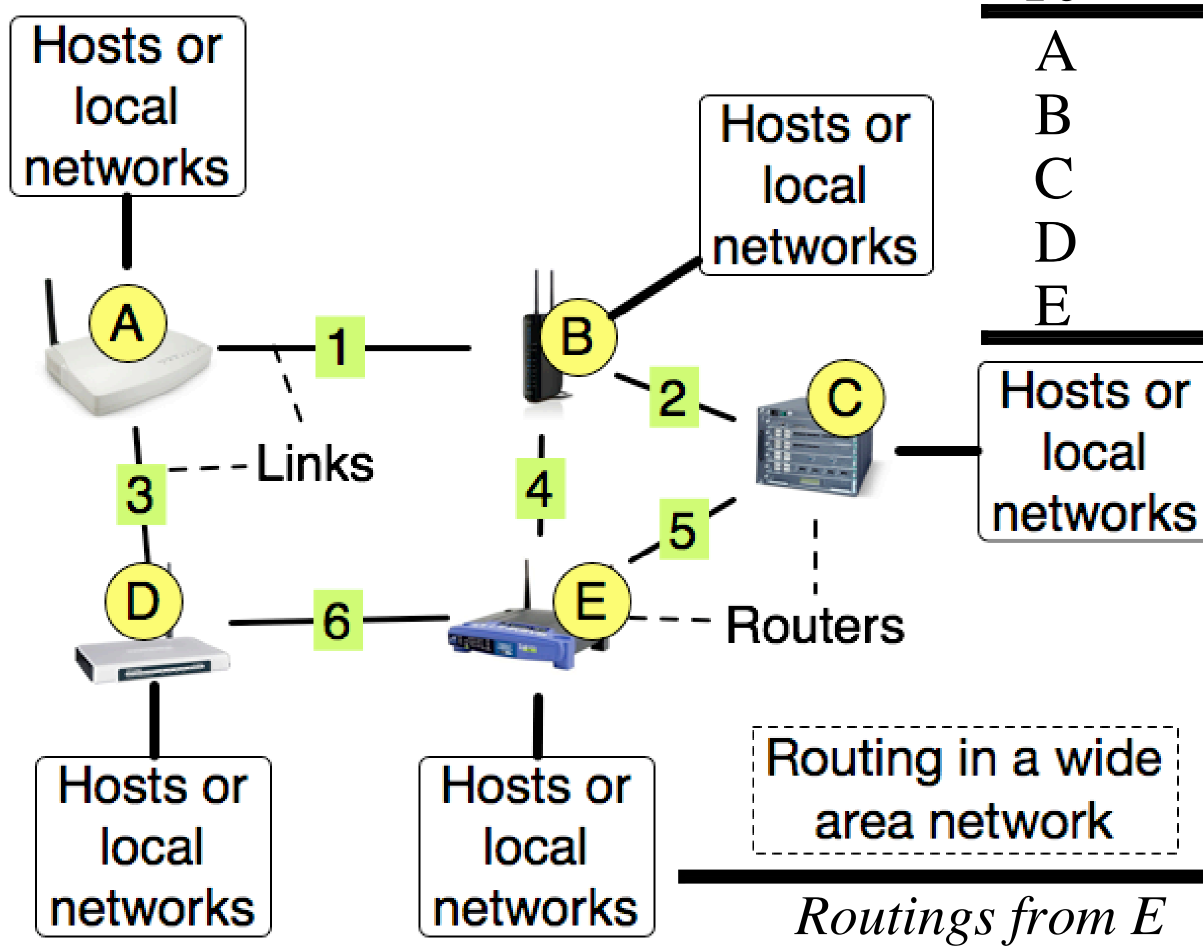

Routings from C

To	*Link*	*Cost*
A	2	2
B	2	1
C	local	0
D	5	2
E	5	1

Routings from D

To	*Link*	*Cost*
A	3	1
B	3	2
C	6	2
D	local	0
E	6	1

Routings from E

To	*Link*	*Cost*
A	4	2
B	4	1
C	5	1
D	6	1
E	local	0

[Distance Vector Algorithm] Routing Actions Example

Part 1 of the routing algorithm

Routings from A

To	*Link*	*Cost*
A	local	0
B	1	1
C	1	2
D	3	1
E	1	2

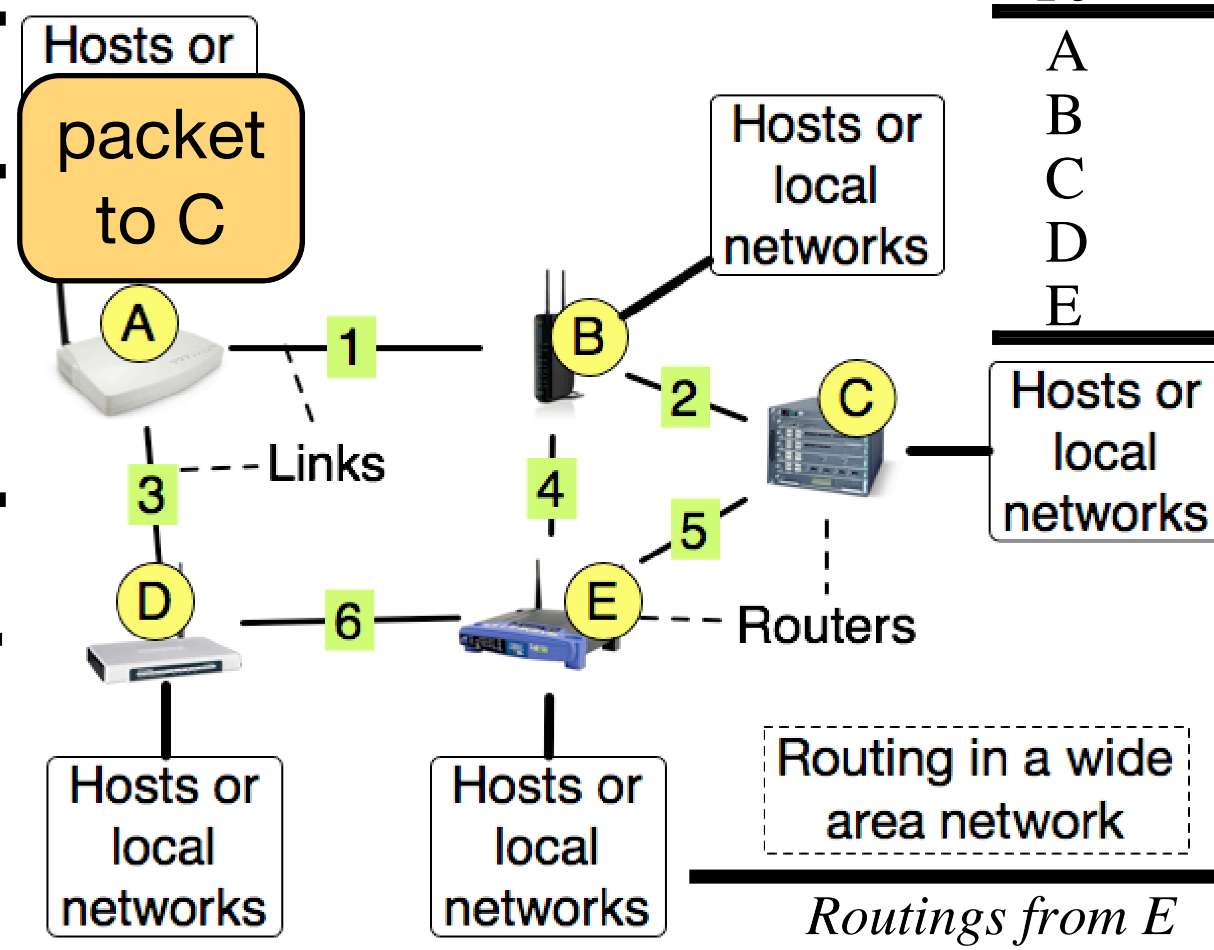

Routings from B

To	*Link*	*Cost*
A	1	1
B	local	0
C	2	1
D	1	2
E	4	1

Routings from C

To	*Link*	*Cost*
A	2	2
B	2	1
C	local	0
D	5	2
E	5	1

Routings from D

To	*Link*	*Cost*
A	3	1
B	3	2
C	6	2
D	local	0
E	6	1

Routings from E

To	*Link*	*Cost*
A	4	2
B	4	1
C	5	1
D	6	1
E	local	0

[Distance Vector Algorithm] Routing Actions Example

Part 1 of the routing algorithm

Routings from A

To	*Link*	*Cost*
A	local	0
B	1	1
C	1	2
D	3	1
E	1	2

Routings from D

To	*Link*	*Cost*
A	3	1
B	3	2
C	6	2
D	local	0
E	6	1

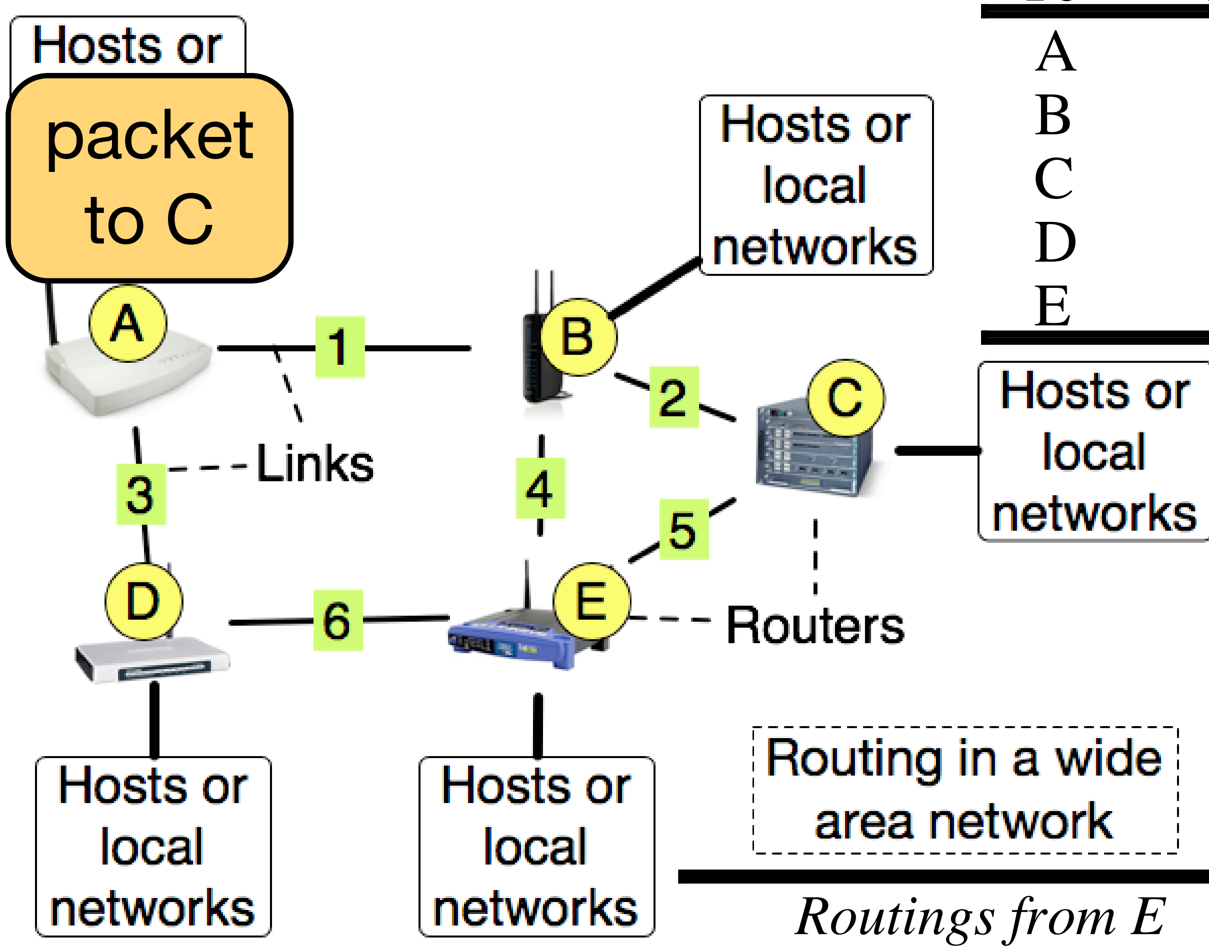

Routing in a wide area network

Routings from B

To	*Link*	*Cost*
A	1	1
B	local	0
C	2	1
D	1	2
E	4	1

Routings from C

To	*Link*	*Cost*
A	2	2
B	2	1
C	local	0
D	5	2
E	5	1

Routings from E

To	*Link*	*Cost*
A	4	2
B	4	1
C	5	1
D	6	1
E	local	0

[Distance Vector Algorithm] Routing Actions Example

Part 1 of the routing algorithm

Routings from A

To	*Link*	*Cost*
A	local	0
B	1	1
C	1	2
D	3	1
E	1	2

Routings from B

To	*Link*	*Cost*
A	1	1
B	local	0
C	2	1
D	1	2
E	4	1

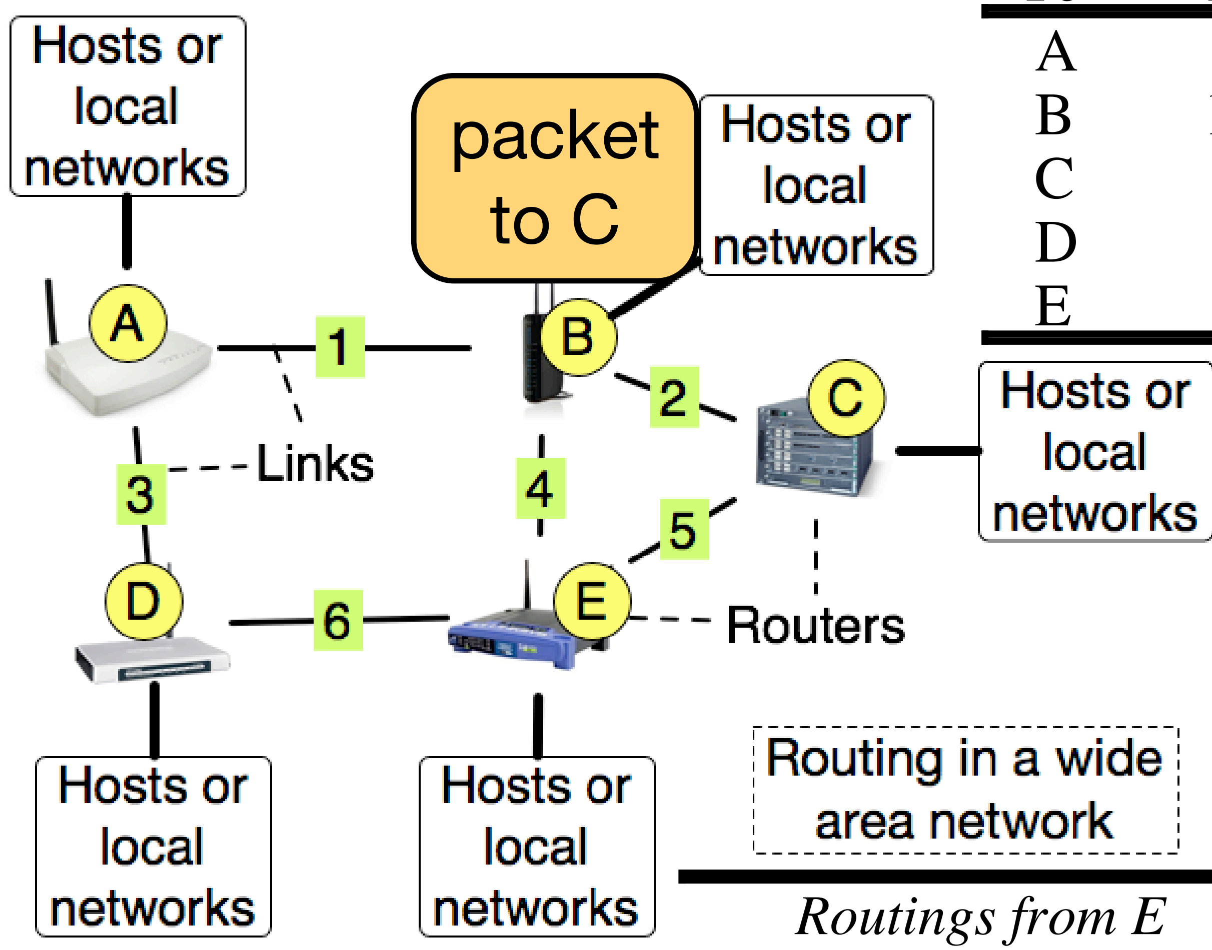

Routings from C

To	*Link*	*Cost*
A	2	2
B	2	1
C	local	0
D	5	2
E	5	1

Routings from D

To	*Link*	*Cost*
A	3	1
B	3	2
C	6	2
D	local	0
E	6	1

Routings from E

To	*Link*	*Cost*
A	4	2
B	4	1
C	5	1
D	6	1
E	local	0

[Distance Vector Algorithm] Routing Actions Example

Part 1 of the routing algorithm

Routings from A

To	*Link*	*Cost*
A	local	0
B	1	1
C	1	2
D	3	1
E	1	2

Routings from D

To	*Link*	*Cost*
A	3	1
B	3	2
C	6	2
D	local	0
E	6	1

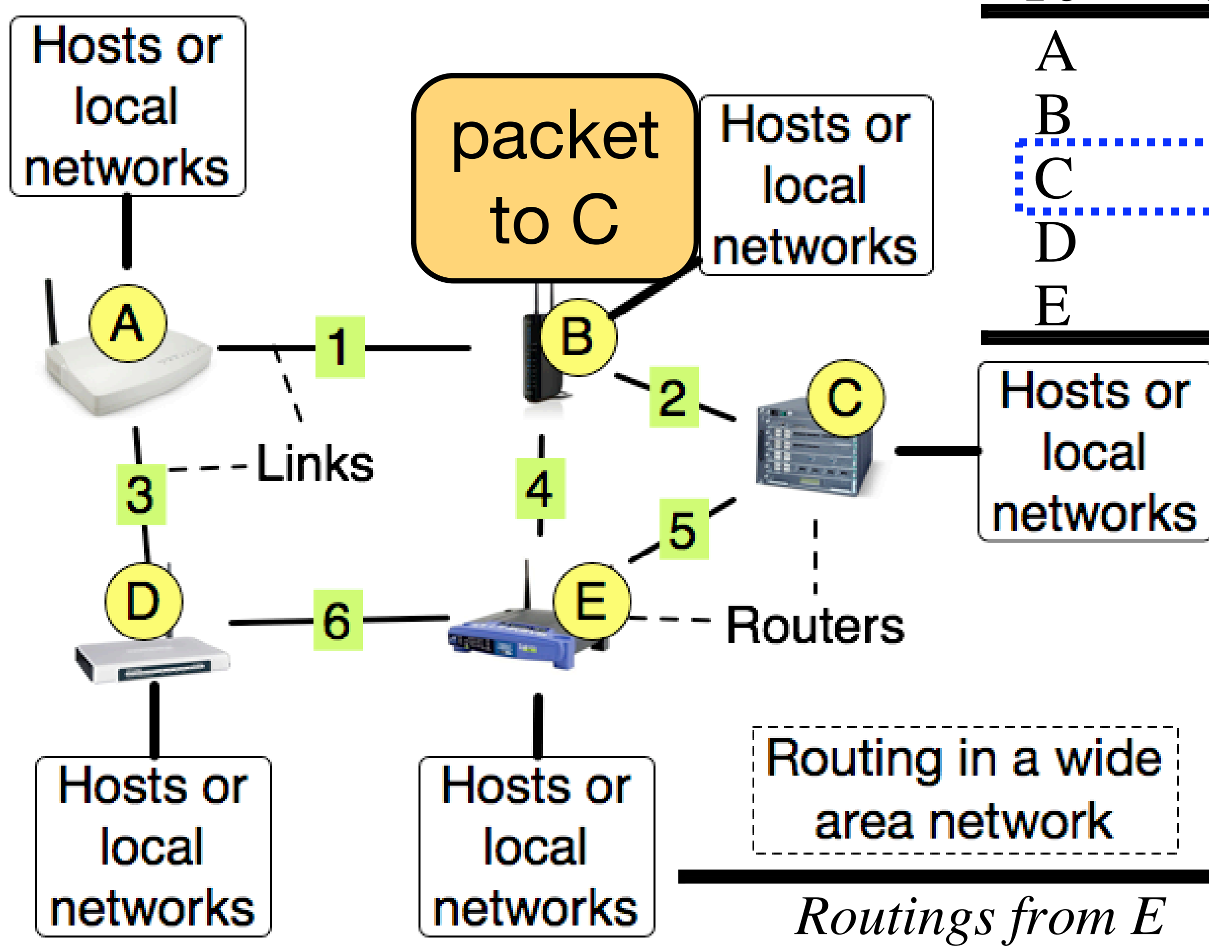

Routing in a wide area network

Routings from B

To	*Link*	*Cost*
A	1	1
B	local	0
C	2	1
D	1	2
E	4	1

Routings from C

To	*Link*	*Cost*
A	2	2
B	2	1
C	local	0
D	5	2
E	5	1

Routings from E

To	*Link*	*Cost*
A	4	2
B	4	1
C	5	1
D	6	1
E	local	0

[Distance Vector Algorithm] Routing Actions Example

Part 1 of the routing algorithm

Routings from A

To	*Link*	*Cost*
A	local	0
B	1	1
C	1	2
D	3	1
E	1	2

Routings from D

To	*Link*	*Cost*
A	3	1
B	3	2
C	6	2
D	local	0
E	6	1

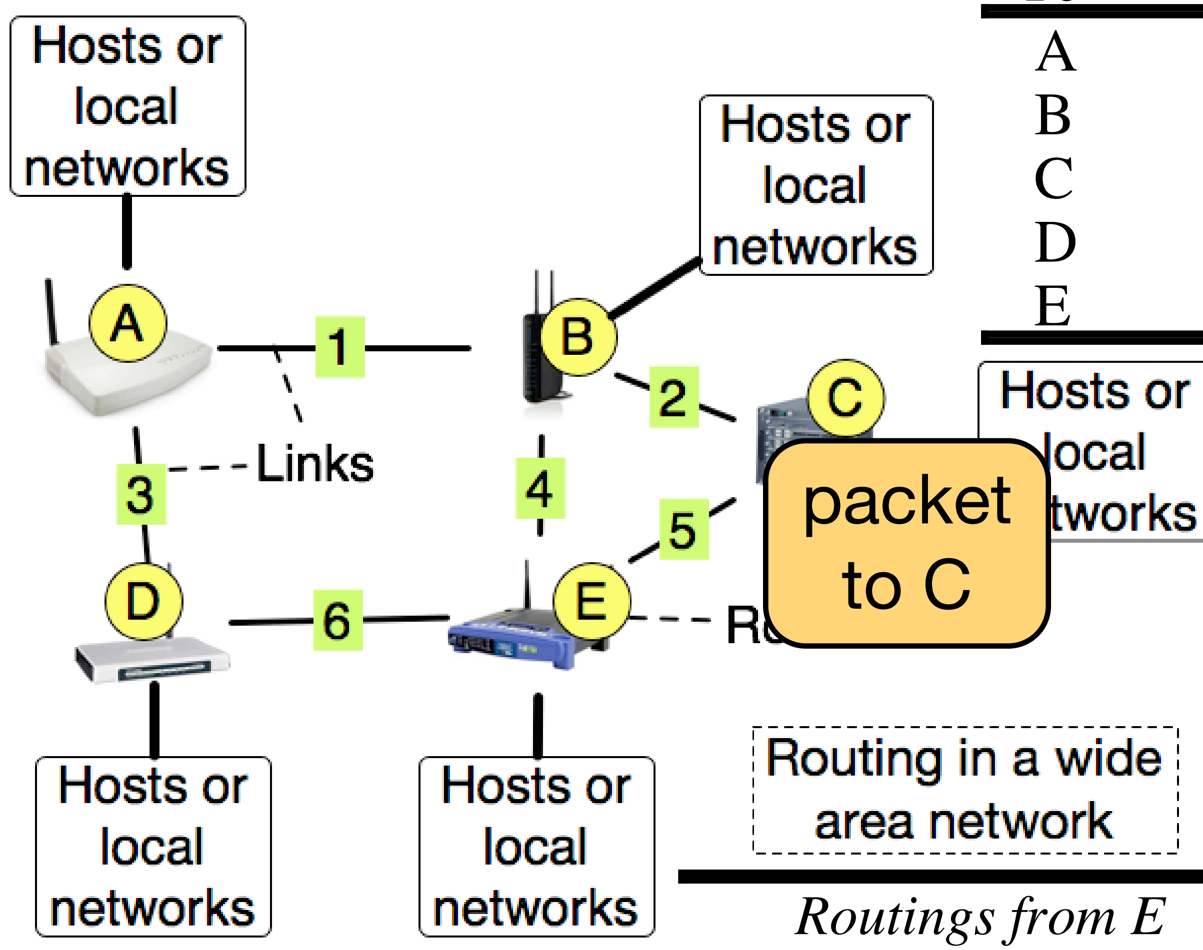

Routing in a wide area network

Routings from B

To	*Link*	*Cost*
A	1	1
B	local	0
C	2	1
D	1	2
E	4	1

Routings from C

To	*Link*	*Cost*
A	2	2
B	2	1
C	local	0
D	5	2
E	5	1

Routings from E

To	*Link*	*Cost*
A	4	2
B	4	1
C	5	1
D	6	1
E	local	0

[Distance Vector Algorithm] Routing Actions Example

Part 1 of the routing algorithm

Routings from A

To	*Link*	*Cost*
A	local	0
B	1	1
C	1	2
D	3	1
E	1	2

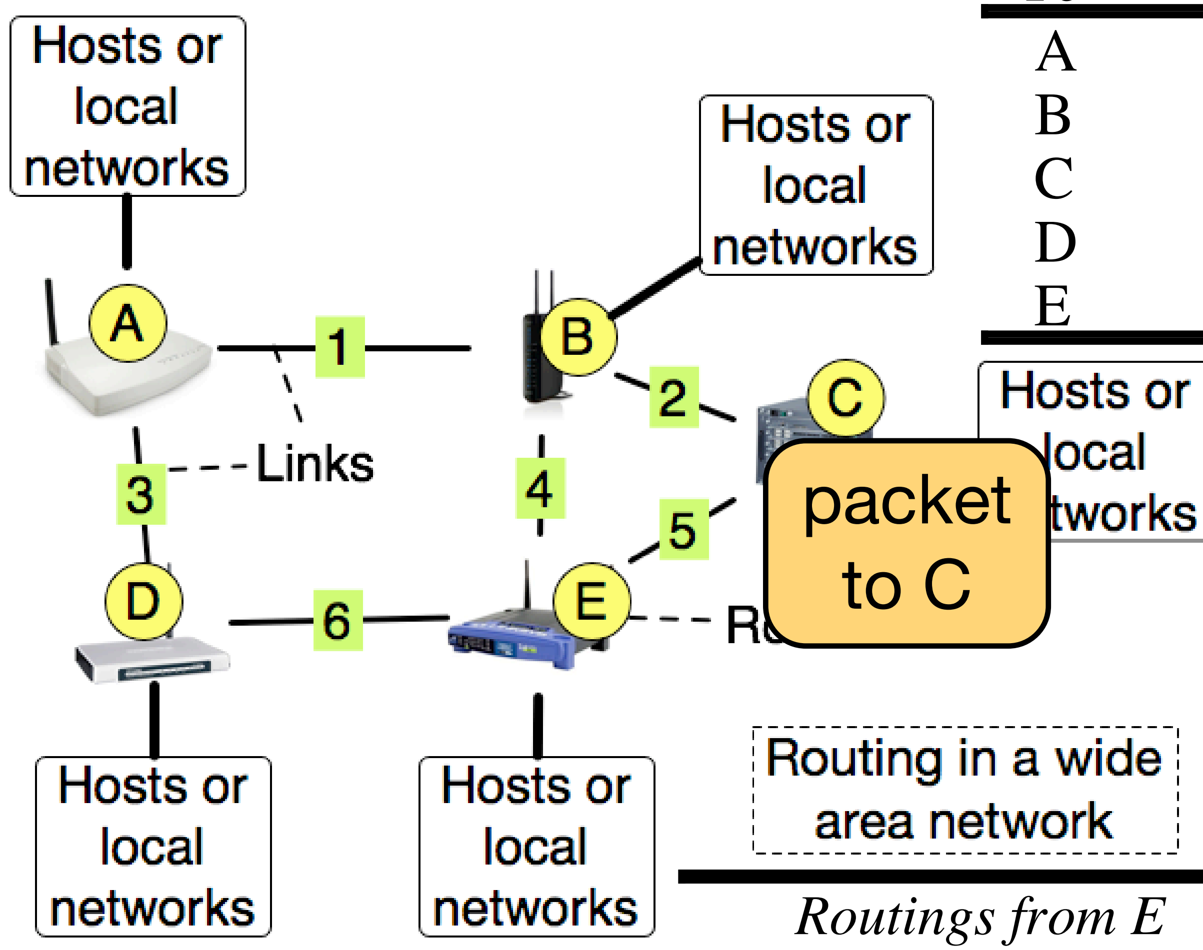

Routings from B

To	*Link*	*Cost*
A	1	1
B	local	0
C	2	1
D	1	2
E	4	1

Routings from C

To	*Link*	*Cost*
A	2	2
B	2	1
C	local	0
D	5	2
E	5	1

Routings from D

To	*Link*	*Cost*
A	3	1
B	3	2
C	6	2
D	local	0
E	6	1

Routings from E

To	*Link*	*Cost*
A	4	2
B	4	1
C	5	1
D	6	1
E	local	0

The packet should be delivered to a local host

Constructing and Maintaining Routing Tables

Part 2 of the routing algorithm

- Because *each routing table specifies only a single hop for each route*, the construction or repair of the routing information proceeds in a distributed fashion.

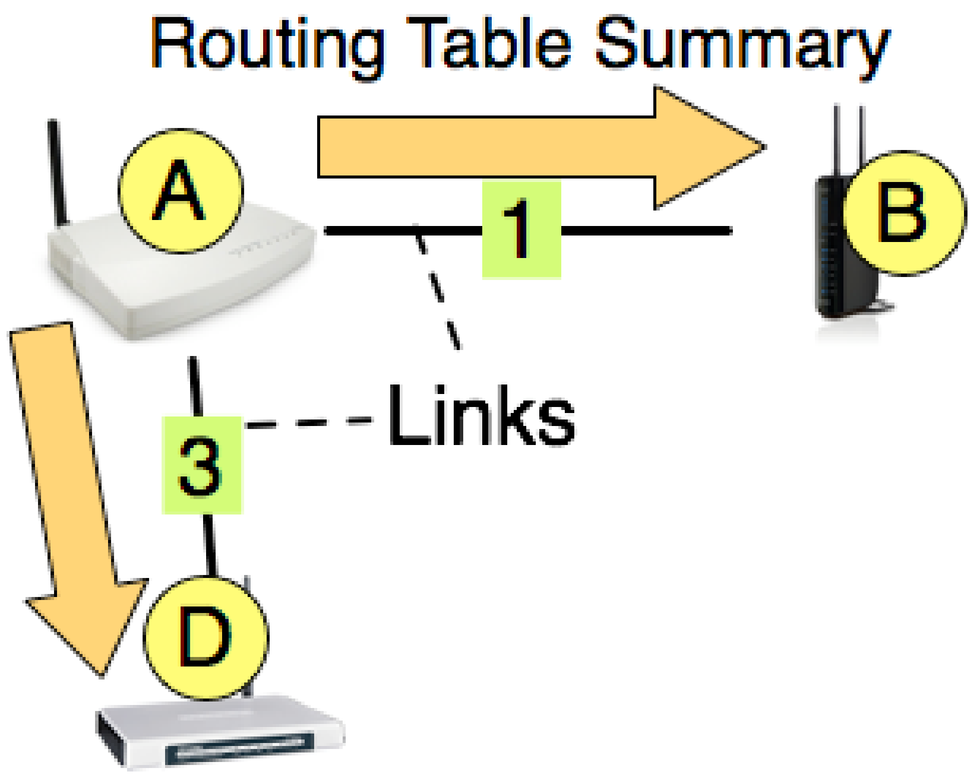

- A router exchanges information about the network with its neighbouring nodes by sending a *summary of its routing table.*

- This is done using a router information protocol (RIP).

Routing Information Protocol (RIP) - Intuition

Part 2 of the routing algorithm

- RIP actions performed at a router:

  1. *Periodically, and whenever the local routing table changes*, send the table (in a summary form) to all accessible neighbours.
     That is, send a RIP packet containing a copy of the table on each non-faulty outgoing link.

  2. *When a table is received from a neighbouring router*:
     IF the received table shows a route to a new destination, or a better (lower cost) route to an existing destination

     THEN update the local table with the new route

     IF the table was received on link n and it gives a different cost than the local table for a route that begin with link n

     THEN replace the cost in the local table with the new cost

Routing Information Protocol (RIP) - PseudoCode

Part 2 of the routing algorithm

Send: Each *t* seconds or when *Tl* changes, send *Tl* on each non-faulty outgoing link.

Receive: Whenever a routing table *Tr* is received on link *n*:

```
for all rows Rr in Tr {
    if (Rr.link ≠ n) {
        Rr.cost = Rr.cost + 1;
        Rr.link = n;
        if (Rr.destination is not in Tl) add Rr to Tl;
         // add new destination to Tl
        else for all rows Rl in Tl {
            if (Rr.destination = Rl.destination and
                    (Rr.cost < Rl.cost or Rl.link = n)) Rl = Rr;
            // Rr.cost < Rl.cost : remote node has better route
            // Rl.link = n : remote node is more authoritative
        }
    }
}
```

t adopted throughout the Internet: 30 secs.

Tl: local table
Tr: table received from another router