


Routing and the Network Layer

(ref: [Interconnections](#) by Perlman)



Netprog 2001 - Routing and the Network Layer

1

Network Service Types

- Two basic models of the services the *network* should provide:
 - Connectionless (*datagram*)
 - Connection-oriented (*virtual circuit*)

Netprog 2001 - Routing and the Network Layer

2

Connectionless Network

- Each packet is independently routed.
- Each packet includes the destination address.
- No guarantee that packets are kept in order.
- No guarantee that packets are not lost or duplicated.

Netprog 2001 - Routing and the Network Layer

3

Connection-Oriented Network

- A single path is first established for each new connection.
- The *network* guarantees that packets are delivered in order.
- No loss or duplication.
- If anything goes wrong the connection is broken.
- It is possible to limit the number of connections.

Connection-Oriented (cont.)

- The network can guarantee bandwidth at connect time.
- The network can refuse new connections.

Middle Ground

- It is possible to design service models that are somewhere in the middle:
 - connection-oriented, but without any bandwidth guarantee.
 - Routers take care of establishing a virtual circuit - hosts view the network as connectionless.

Advantages of Connectionless

- Connection-oriented requires duplication of service at the transport layer (to handle broken connections).
- Host software is much simpler at the network layer.
- Many applications do not require sequential delivery of packets (example: packet voice).

Netprog 2001 - Routing and the Network Layer

7

Advantages of Connectionless (cont.)

- Network traffic often comes in *bursts*, so reserving resources is wasteful.
- It is better to provide degraded service to everyone than to limit network access.
- Server (or router) could become overloaded managing too many connections.

Netprog 2001 - Routing and the Network Layer

8

Advantages of Connection-Oriented

- Most applications requires sequential packet delivery - the network should handle the complexity.
- Faster Routers. Once a connection is established each router can reference the connection via a small number.
- It is better to provide uniform service to a few than to degrade while handling everyone.

Netprog 2001 - Routing and the Network Layer

9

Advantages of Connection-Oriented (cont.)

- Life is easier for the Transport Layer
 - possible to calculate round-trip delay
 - possible to maximize packet size (it never needs to change).

Netprog 2001 - Routing and the Network Layer

10

IP Routing

- IP is a connectionless network layer.
- Each host has a routing table:
 - routes to specific hosts
 - routes to specific networks
 - default route

Netprog 2001 - Routing and the Network Layer

11

Host route determination

- Search for a matching host address
- Search for a matching network address
- Search for a default entry.

Netprog 2001 - Routing and the Network Layer

12

Routing table creation

- Static routes - Unix "route" command.
- ICMP Router Discovery
 - broadcast protocol that discovers routers on the local network.
- ICMP redirects.
- Run a routing daemon.

Netprog 2001 - Routing and the Network Layer

13

Dynamic Routing

- Routers need to tell each other about routes.
- Host routing tables can change over time by listening to routers.
- There are many dynamic IP routing protocols in use.

Netprog 2001 - Routing and the Network Layer

14

RIP Routing Information Protocol

- UDP based messages
- Each router sends out a broadcast (possibly a series of broadcasts) that contains the entire routing table of the router.
- Typically routers do this every 30 seconds or when something changes.

Netprog 2001 - Routing and the Network Layer

15

RIP routes

- Each RIP routing table entry includes:
 - IP address
 - metric (hop count 1-15).
 - timeout (seconds).
- directly connected networks have a metric of 1.
- If a route times-out the metric is set to 16 (no connection) and deleted after 1 minute.

Netprog 2001 - Routing and the Network Layer

16


Convergence

- When something changes (for example when a link or router goes down), it takes a while before the change is propagated to all affected routers.
- RIP suffers from slow convergence - there is not enough information in RIP routing tables to avoid this problem.

Netprog 2001 - Routing and the Network Layer

17

What happens when C dies?



A's Routing Table:

B is 1 hop away (directly connected)
C is 2 hops away (via B).

B's Routing Table:

A is 1 hop away (directly connected)
C is 1 hop away (directly connected)

Netprog 2001 - Routing and the Network Layer

18


OSPF vs. RIP

- *Open Shortest Path First* is an alternative IP routing protocol.
- RIP is a *distance-vector* protocol.
- OSPF is *link-state* protocol.
 - a router checks the condition of each of its connections (links) and reports this information to neighbors.

OSPF advantages

- Each router has a picture of the network topology.
- Faster convergence.
- support for independent routes for each IP *type-of-service*
- load balancing (distribute traffic among equal cost routes)

X.25 - Connection-Oriented Network Layer




DTE: Data Terminal Equipment (host).
DCE: Data Circuit-Terminating Equipment (router)

X.25 Network Service

- Allows a DTE to establish multiple simultaneous connections (over a single link to a DCE).
- Can also be used to connect 2 DTEs directly.
- Can support permanent connections.

X.25 Switched Virtual Circuit

- DTE A tells DCE A that it wants a connection to DTE B.
- DCE B tells DTE B that a new connection has been requested.



Virtual Circuit Numbers

- Each *Virtual Circuit* is assigned a number at setup time.
- A *virtual circuit number* identifies a connection between a DTE and its DCE only.
- The other end of the connection can use a different *virtual circuit number*.

Virtual Circuit Numbers

- Each data packet includes a virtual circuit number rather than a destination address.
- 12 bit identifier.
- Virtual circuit numbers are smaller than addresses and much faster for a router to process (just a table lookup).

Virtual Circuit Numbers

- 0 is reserved for control packets.
- non-overlapping ranges of numbers:
 - permanent virtual circuits
 - incoming connections
 - outgoing connections

Connection request


- DTE A selects an unused outgoing vc number.
- DTE A creates a *call request* packet and sends to DCE A.
- DCE A contacts DCE B (via an unspecified mechanism) and requests the connection.
- DCE A notifies DTE A when the connection is established.

Data Transfer

- Each connection (*call*) is full duplex.
- Each packet sent by a DTE includes:
 - virtual circuit number
 - sequence number (3 or 7 bits).
 - fragmentation information (M bit).
 - ACK sequence number.

Data Transfer

- Fragmentation can happen at DTE or DCE.
- Sequence numbers do not necessarily coincide at DTEs.



Flow Control

- Each packet includes an acknowledgement number (ACK).
- The *ack#* indicates that the send has received all packets with sequence number \leq *ack#*.
- Each sender has a window size *w* that determines how many unacknowledged packets can be outstanding.

DCE <-> DCE

- The interface between DCE is not part of the X.25 standard.
- Thought Exercise:
 - How could you design a network to route based on circuit numbers ?
