

I. Introduction to Routing

Week 14

- **routing** is the process of moving data across an internetwork from a source to a destination
- why routing/routers, why not just use switches?
 - router is a layer 3 device, uses ? (IP) address
 - switch is a layer 2 device, uses ? (MAC) address
 - if we only had switches, each switch would need to know/store the location of every MAC address it was connected to
- routers identify all hosts on a network as a single entry, routers don't care where a node is, just the next hop to get the data there
- what do we need to route? routing tables
- who needs them? everybody on the network, e.g. each host and each router
- each **host** needs to know:
 1. how do I reach myself?
 - loopback address (127.0.0.1)
 - allows client & server to reside on same machine and not spew data across a NW
 - allows testing of client/server apps as well as self testing of NIC
 2. how do I reach my network?
 - how does a node know if an address is from/to their NW? by the subnet mask
 3. how do I reach everything else?
 - need to know the IP address of the default/gateway router

Note all 3 of these are manually configured!

- each **router** needs to know:
 1. how do I reach myself?
 - loopback address (127.0.0.1)
 2. how do I reach my network?
 - via the ports of the router
 3. how do I reach everything else?
 - via routing tables
 - a) static routing
 - manually configured via data files
 - must be configured for every router on NW
 - if NW changes, must be re-configured
 - b) dynamic routing
 - allows router to configure itself
 - via default route (0.0.0.0)
 - if not in table (no match found), behave like this route

II. General Routing Goals


- **optimality** – selecting the best route based upon routing metrics, i.e. # hops, delay, congestion, etc.
- **simple & efficient** – minimum overhead
- **robust & stable** – good performance in adverse conditions
- **rapid convergence** – agreement of all routers on optimal routes
- **flexible** – adaptable to network changes

III. Routing Algorithm Types and Techniques

- static vs. dynamic
- single-path vs. multi-path
- host intelligent vs. router intelligent
 - host intelligent (also called source routing) is where the host specifies the route
 - most routers drop
- intra-domain (IGP) vs. inter-domain (BGP, was EGP)
- distance vector vs. link state

IV. Distance Vector Routing

- RFC 1058
- DVP share information about how far away NW are in hops, ignoring distance
- lower hop counts are route of choice
- used by RIP, RIP-2, IGRP (Cisco)
- broadcast routing tables to all neighbors periodically (e.g. RIP – 30 sec)
- DVP “route by rumor”



R1 sends to R2: I have a route to NW 1
R2 sends to R3: I have a route to NW 1

- require very little in overhead and processor power
- routing loops avoided by setting an infinity value, e.g. RIP = 16

- employs split horizon: routing information cannot be sent back the way it was received
- see also poison reverse
- see Bellman-Ford algorithm (below)

V. Link State Routing

- routers send each other info about the links they have established with other routers
- information sent as Link State Advertisements (LSAs)
- as the LSAs are propagated throughout the NW, each router builds its own picture of the NW
- link state routers use this to build a forwarding table based upon cost, i.e. status and connection type (and thus speed)
- require more processing power than D-V
- used by OSPF (open shortest path first) which uses Dijkstra's shortest path first algorithm

VI. Other Routing Protocols

- Hybrid
 - features properties from both DVP and L-S routing
 - example: EIGRP – Enhanced Interior Gateway Routing Protocol (from Cisco)
- Path Vector
 - subset of DVP
 - uses path-vectors to make metric decisions
 - example: BGP – Border Gateway Protocol (from Cisco)

TABLE 7.1 Sample Routing Table Information

Destination	Gateway	Flags ^a	Ref ^b	Use ^c	Interface
localhost	localhost	UH	0	33106	lo0
215.103.16.227	187.96.25.13	UGHD	29	102	le0
215.103.16.141	187.96.25.35	UGHD	116	16128	le1
default	187.96.25.1	UG	0	2888104	
187.96.25.0	187.96.25.2	U	210	29024	le0

^a U = Route is up and operational; G = Packet must pass through at least one router;
^b H = Route is to a specific host and not a network; D = Route was created dynamically
^c Current number of routes that share the same link layer address
^d Number of packets sent using this route

Routing Algorithms

Two general algorithms are available for computing metric information: *distance-vector* and *link-state*. The goal of both types of algorithms is to route a packet from one point in the network to another point in the network through some set of intermediate routers without “looping,” a situation in which a packet is forwarded across the same link several times. The primary difference between distance-vector and link-state algorithms is the manner in which they collect and propagate routing information throughout the network. Let’s examine these two algorithms separately.

Distance-Vector Algorithms A *distance-vector routing algorithm* determines the distance (hence the name) between source and destination nodes by calculating the number of router hops a packet traverses en route from the source network to the destination network. An example of a distance-vector algorithm is the Bellman-Ford algorithm, which is described in Box 7.1. Two distance-vector-based routing protocols are RIP and RIP-2, which exchange routing tables with their neighbors every 30 seconds. RIP and RIP-2 also support a maximum of 15 hops. Thus, if the number of router-to-router hops between source and destination nodes is greater than 15, then the network to which the destination node is connected is considered “unreachable.” This limitation restricts the size of an internetwork to 15 consecutively connected networks.

Link-State Algorithms In a *link-state routing algorithm* every router of a network does not send every other router its routing table. Instead, routers send each other information about the links they have established to other routers. This information is sent via a *link-state advertisement* (LSA), which contains the names and various cost metrics of a router’s neighbors. LSAs are flooded throughout an entire router’s domain. (An example of how this is done is described later in our discussion of OSPF.) Routers also store the most recent LSA they receive, and destination routes are calculated using LSA information. Thus, rather than storing actual paths, which is the case with distance-vector algorithms, link-state algorithms store the information needed to generate such paths. An example of a link-state algorithm is Dijkstra’s shortest path algorithm, which iterates on length of path to determine a shortest route. Link-state-based routing protocols include OSPF, OSI’s IS-IS, and Netware’s Link Services Protocol (NLSP). Box 7.2 illustrates Dijkstra’s shortest path algorithm.

lu, 30 hops max, 40

follows from a source
 om www.fit.edu to
 destination node, not a
 d the round-trip time in
 en intermediate routers.


rs, the destination ad-
 work interface associ-
 packet, it looks at the
 searches its routing ta-
 wards the packet to the
 shows sample routing
 187.96.25.2. The com-

uting. This entry indi-
 5.0) will be forwarded
 this context, the local
 table. This entry indi-
 are forwarded to the
 ounter in turn will have
 mately reach its desti-
 h destination address
 dress is 187.96.25.35.
 entry of the table ref-
 he address of a default
 wn destination address
 he router forwards the
 eives a packet with the
 ket to the router whose
 ost’s routing table (Ta-
 ormation with neighbor
 ie frequency of routing

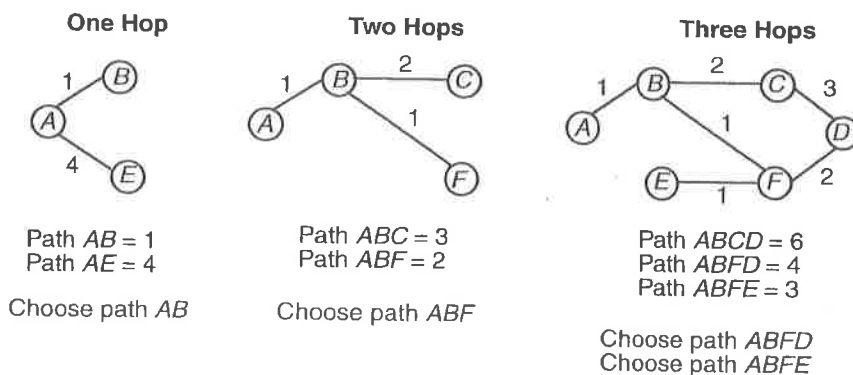
ted route that is entered
 configuration program.
 Although static routes
 dynamically to compen-

BOX 7.1 Bellman-Ford Algorithm

The Bellman-Ford routing algorithm is distance-vector-based and iterates on the number of hops a source node is from a destination node. To illustrate this algorithm, consider the following undirected graph, which depicts a sample network. The vertices *A*, *B*, *C*, *D*, *E*, and *F* may be thought of as routers, and the edges connecting the vertices are communication links. Edge labels represent an arbitrary cost. Our goal is to find the shortest path from *A* to *D* using the number of hops as the basis for our path selection.




We examine the costs of all paths leading from *A* to each node on a hop-by-hop basis.



In the last step (three hops), two paths are selected. The first path, *ABFD*, represents the least-cost path from *A* to *D* based on the hops metric. The second path, *ABFE*, is selected because it represents the least-cost path from *A* to *E*.

The final result of the Bellman-Ford algorithm yields a tree that represents the least cost incurred from the source node to every node of the network. Similar trees can be generated for every node of the network. Node *A*'s least-cost tree for our example is as follows:



- From Node *A*:
- the least-cost path to *B* is *AB* = 1
 - the least-cost path to *C* is *ABC* = 3
 - the least-cost path to *D* is *ABFD* = 4
 - the least-cost path to *E* is *ABFE* = 3
 - the least-cost path to *F* is *ABF* = 2

BOX 7.2 Dijkstra's Shortest Path

Dijkstra's SPF routing algorithm uses a "closest nodes" search.

Given a source node *S*, a node *N* directly connects to *S* if there is a path of one or more closest nodes between *S* and *N*.

Consider the following graph. The nodes *A*, *B*, and *F* may be thought of as source nodes. Edge labels represent distance.

To implement this algorithm, we start with the source node *A*. We examine the zero closest nodes to *A* and search for the success node.

First Closest Node (*k* = 1)

The first closest node to *A* is *B*. The *AB* path has a small cost.

Second Closest Node (*k* = 2)

The second closest node to *A* includes the first closest node. The paths are *ABC* = 8, *ABE* = 7, or *AE* = 4. The closest node to *A* is *E*.

Third Closest Node (*k* = 3)

The third closest node to *A* includes the first two closest nodes. The paths are *ABCD* = 11, or *AEF* = 11. The closest node to *A* is *D*.

Fourth Closest Node (*k* = 4)

The fourth closest node to *A* includes the first three closest nodes. The paths are *ABCDE* = 11, or *ABFD* = 11. The closest node to *A* is *F*.

Fifth Closest Node (*k* = 5)

The fifth closest node to *A* includes the first four closest nodes. The paths are *ABCDEF* = 11, or *ABFE* = 11. The shortest path is *ABFE*.

Since *D* is the destination node, the algorithm stops.

I. Intro to Dynamic Routing

Routed vs. Routing Protocols

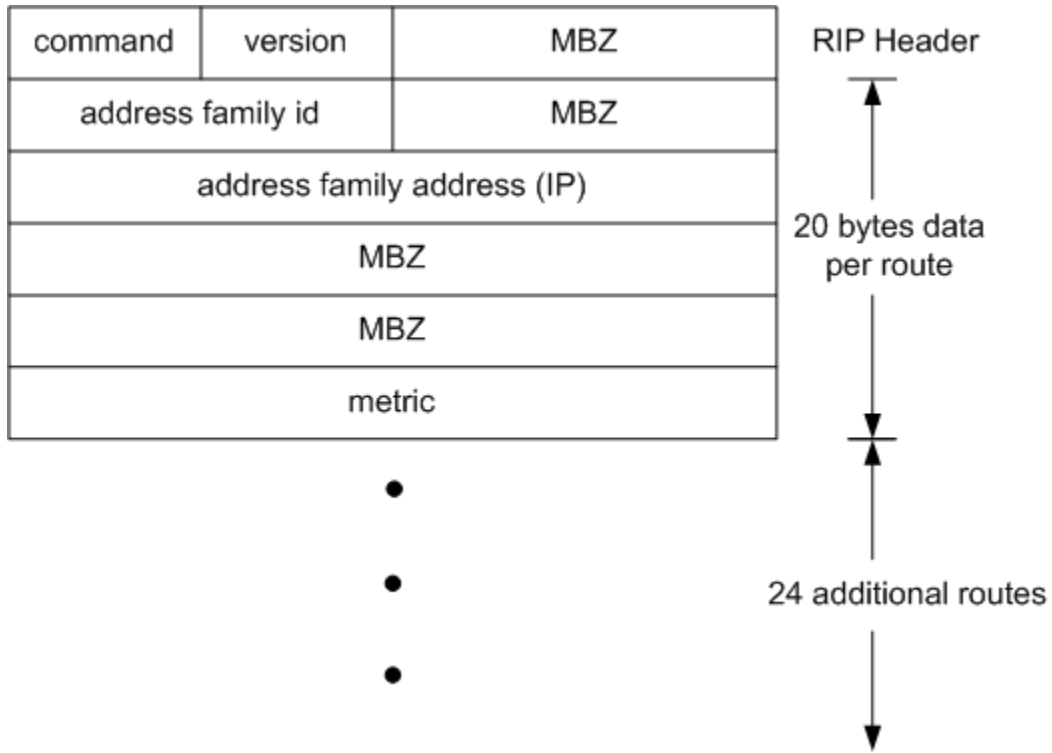
routed protocol - responsible for moving data across an internetwork, e.g. IP

routing protocol – responsible for exchanging routing information, e.g. RIP

RIP – Routing Information Protocol

- RFC 1058
- distance vector protocol (DVP), based upon hop count
- 2 versions, RIP & RIP-2
- originally from Xerox used in Xerox Network System – XNS (see http://www.tcpipguide.com/free/t_RIPOverviewHistoryStandardsandVersions.htm)
- IGP – thus does not scale well, mostly due to limit of 16 hops
- RIP uses UDP (specifically port 520)
- relatively slow convergence
- uses timers to trigger updates
 - broadcast (30 sec) – sends routing table to all neighboring routers
 - aging (180 sec) – if no message received from neighbor, drops from table
- other issues
 - cannot measure/calculate delay, throughput or reliability (like a L-S P)
 - single-path protocol
 - does not support variable length subnet masks (only natural)
 - not a secure protocol

II. RIP PDU



- max size: 4 (header) + 20 * 25 (routes) + 8 (UDP) + 20 (IP)
- **command:** 1 – request (for other routers' tables)
2 – reply (of this routers' table)
- **version:** 1 – RIPv1, 2 – RIPv2
- **AFI:** 2 for IP
- **IP addr:** addr of NW
- **metric:** # hops



IP dest addr = ? (broadcast)