

Chapter 4: Network Layer

Chapter goals:

- r understand principles behind network layer services:
 - m routing (path selection)
 - m dealing with scale
 - m how a router works
 - m advanced topics: IPv6, multicast
- r instantiation and implementation in the Internet

Overview:

- r network layer services
- r routing principle: path selection
- r hierarchical routing
- r IP
- r Internet routing protocols
- r reliable transfer
 - m intra-domain
 - m inter-domain
- r what's inside a router?
- r IPv6
- r multicast routing


4: Network Layer 4a-1

Network layer functions

- r transport packet from sending to receiving hosts
- r network layer protocols in every host, router

three important functions:

- r *path determination*: route taken by packets from source to dest. *Routing algorithms*
- r *switching*: move packets from router's input to appropriate router output
- r *call setup*: some network architectures require router call setup along path before data flows



4: Network Layer 4a-2

Network service model

Q: What *service model* for “channel” transporting packets from sender to receiver?

- r guaranteed bandwidth?
- r preservation of inter-packet timing (no jitter)?
- r loss-free delivery?
- r in-order delivery?
- r congestion feedback to sender?

service

The most important abstraction provided by network layer:

virtual circuit
or
datagram?

4: Network Layer 4a-3

Virtual circuits

“source-to-dest path behaves much like telephone circuit”


- m performance-wise
- m network actions along source-to-dest path

- r call setup, teardown for each call *before* data can flow
- r each packet carries VC identifier (not destination host OD)
- r every router on source-dest path s maintain “state” for each passing connection
 - m transport-layer connection only involved two end systems
- r link, router resources (bandwidth, buffers) may be *allocated* to VC
 - m to get circuit-like perf.

4: Network Layer 4a-4

Virtual circuits: signaling protocols


- r used to setup, maintain teardown VC
- r used in ATM, frame-relay, X.25
- r not used in today's Internet



4: Network Layer 4a-5

Datagram networks: the Internet model

- r no call setup at network layer
- r routers: no state about end-to-end connections
 - m no network-level concept of "connection"
- r packets typically routed using destination host ID
 - m packets between same source-dest pair may take different paths



4: Network Layer 4a-6

Network layer service models:

Network Architecture	Service Model	Guarantees ?				Congestion feedback
		Bandwidth	Loss	Order	Timing	
Internet	best effort	none	no	no	no	no (inferred via loss)
ATM	CBR	constant rate	yes	yes	yes	no congestion
ATM	VBR	guaranteed rate	yes	yes	yes	no congestion
ATM	ABR	guaranteed minimum	no	yes	no	yes
ATM	UBR	none	no	yes	no	no

- r Internet model being extended: Intserv, Diffserv
- m Chapter 6

4: Network Layer 4a-7

Datagram or VC network: why?

Internet

- r data exchange among computers
 - m "elastic" service, no strict timing req.
- r "smart" end systems (computers)
 - m can adapt, perform control, error recovery
 - m simple inside network, complexity at "edge"
- r many link types
 - m different characteristics
 - m uniform service difficult

ATM

- r evolved from telephony
- r human conversation:
 - m strict timing, reliability requirements
 - m need for guaranteed service
- r "dumb" end systems
 - m telephones
 - m complexity inside network

4: Network Layer 4a-8


Routing

Routing protocol

Goal: determine "good" path (sequence of routers) thru network from source to dest.

Graph abstraction for routing algorithms:

- r graph nodes are routers
- r graph edges are physical links
 - m link cost: delay, \$ cost, or congestion level



r "good" path:

- m typically means minimum cost path
- m other def's possible

4: Network Layer 4a-9

Routing Algorithm classification

Global or decentralized information?

Global:

- r all routers have complete topology, link cost info
- r "link state" algorithms

Decentralized:

- r router knows physically-connected neighbors, link costs to neighbors
- r iterative process of computation, exchange of info with neighbors
- r "distance vector" algorithms

Static or dynamic?

Static:

- r routes change slowly over time

Dynamic:

- r routes change more quickly
 - m periodic update
 - m in response to link cost changes

4: Network Layer 4a-10

A Link-State Routing Algorithm

Dijkstra's algorithm

- r net topology, link costs known to all nodes
 - m accomplished via "link state broadcast"
 - m all nodes have same info
- r computes least cost paths from one node ("source") to all other nodes
 - m gives **routing table** for that node
- r iterative: after k iterations, know least cost path to k dest.'s

Notation:

- r $c(i,j)$: link cost from node i to j. cost infinite if not direct neighbors
- r $D(v)$: current value of cost of path from source to dest. v
- r $p(v)$: predecessor node along path from source to v, that is next v
- r N : set of nodes whose least cost path definitively known

4: Network Layer 4a-11


Dijkstra's Algorithm

- 1 **Initialization:**
- 2 $N = \{A\}$
- 3 for all nodes v
- 4 if v adjacent to A
- 5 then $D(v) = c(A,v)$
- 6 else $D(v) = \text{infty}$
- 7
- 8 **Loop**
- 9 find w not in N such that $D(w)$ is a minimum
- 10 add w to N
- 11 update $D(v)$ for all v adjacent to w and not in N:
- 12 $D(v) = \min(D(v), D(w) + c(w,v))$
- 13 /* new cost to v is either old cost to v or known
- 14 shortest path cost to w plus cost from w to v */
- 15 **until all nodes in N**

4: Network Layer 4a-12

Dijkstra's algorithm: example

Step	start N	D(B),p(B)	D(C),p(C)	D(D),p(D)	D(E),p(E)	D(F),p(F)
→0	A	2,A	5,A	1,A	infinity	infinity
→1	AD	2,A	4,D		2,D	infinity
→2	ADE	2,A	3,E			4,E
→3	ADEB		3,E			4,E
→4	ADEBC					4,E
5	ADEBCF					



4: Network Layer 4a-13

Dijkstra's algorithm, discussion

Algorithm complexity: n nodes


r each iteration: need to check all nodes, w, not in N

r $n*(n+1)/2$ comparisons: $O(n^2)$

r more efficient implementations possible: $O(n \log n)$

Oscillations possible:

r e.g., link cost = amount of carried traffic



4: Network Layer 4a-14

Distance Vector Routing Algorithm

iterative:

- r continues until no nodes exchange info.
- r *self-terminating*: no "signal" to stop

asynchronous:

- r nodes need *not* exchange info/iterate in lock step!

distributed:

- r each node communicates *only* with directly-attached neighbors

Distance Table data structure


- r each node has its own
- r row for each possible destination
- r column for each directly-attached neighbor to node
- r example: in node X, for dest. Y via neighbor Z:

$$D^X(Y,Z) = \text{distance from } X \text{ to } Y, \text{ via } Z \text{ as next hop}$$

$$= c(X,Z) + \min_w \{D^Z(Y,w)\}$$

4: Network Layer 4a-15

Distance Table: example



$$D^E(C,D) = c(E,D) + \min_w \{D^D(C,w)\}$$

$$= 2+2 = 4$$

$$D^E(A,D) = c(E,D) + \min_w \{D^D(A,w)\}$$

$$= 2+3 = 5 \text{ loop!}$$

$$D^E(A,B) = c(E,B) + \min_w \{D^B(A,w)\}$$

$$= 8+6 = 14 \text{ loop!}$$

		cost to destination via		
$D^E()$	A	B	D	
A	1	14	5	
B	7	8	5	
C	6	9	4	
D	4	11	2	

no. tabs! (see)

4: Network Layer 4a-16

Distance table gives routing table

		cost to destination via				
$D^E()$		A	B	D	Outgoing link to use, cost	
no i tani tseal	A	1	14	5	A	A,1
	B	7	8	5	B	D,5
	C	6	9	4	C	D,4
	D	4	11	2	D	D,4

Distance table → Routing table

4: Network Layer 4a-17

Distance Vector Routing: overview

Iterative, asynchronous:


each local iteration caused by:

- r local link cost change
- r message from neighbor: its least cost path change from neighbor

Distributed:

- r each node notifies neighbors *only* when its least cost path to any destination changes
 - m neighbors then notify their neighbors if necessary

Each node:



4: Network Layer 4a-18

Distance Vector Algorithm:

At all nodes, X:

```
1 Initialization:
2 for all adjacent nodes v:
3    $D^X(*,v) = \text{infty}$  /* the * operator means "for all rows" */
4    $D^X(v,v) = c(X,v)$ 
5 for all destinations, y
6   send  $\min_w D^X(y,w)$  to each neighbor /* w over all X's neighbors */
```


4: Network Layer 4a-19

Distance Vector Algorithm (cont.):

```
→ 8 loop
9   wait (until I see a link cost change to neighbor V
10    or until I receive update from neighbor V)
11
12   if (c(X,V) changes by d)
13     /* change cost to all dest's via neighbor v by d */
14     /* note: d could be positive or negative */
15     for all destinations y:  $D^X(y,V) = D^X(y,V) + d$ 
16
17   else if (update received from V wrt destination Y)
18     /* shortest path from V to some Y has changed */
19     /* V has sent a new value for its  $\min_w DV(Y,w)$  */
20     /* call this received new value is "newval" */
21     for the single destination y:  $D^X(Y,V) = c(X,V) + \text{newval}$ 
22
23   if we have a new  $\min_w D^X(Y,w)$  for any destination Y
24     send new value of  $\min_w D^X(Y,w)$  to all neighbors
25
26   forever
```

4: Network Layer 4a-20

Distance Vector Algorithm: example



		cost via	
		Y	Z
D ^X d e s t	Y	2	∞
	Z	∞	7

		cost via	
		X	Z
D ^Y d e s t	X	2	∞
	Z	∞	1

		cost via	
		X	Y
D ^Z d e s t	X	7	∞
	Y	∞	1


		cost via	
		Y	Z
D ^X d e s t	Y	2	8
	Z	3	7

		cost via	
		X	Z
D ^Y d e s t	X	2	8
	Z	9	1

		cost via	
		X	Y
D ^Z d e s t	X	7	3
	Y	9	1

4: Network Layer 4a-21

Distance Vector Algorithm: example



		cost via	
		Y	Z
D ^X d e s t	Y	2	∞
	Z	∞	7

		cost via	
		X	Z
D ^Y d e s t	X	2	∞
	Z	∞	1

		cost via	
		X	Y
D ^Z d e s t	X	7	∞
	Y	∞	1

		cost via	
		Y	Z
D ^X d e s t	Y	2	8
	Z	3	7

$$D^X(Y,Z) = c(X,Z) + \min_w \{D^Z(Y,w)\}$$

$$= 7 + 1 = 8$$

$$D^X(Z,Y) = c(X,Y) + \min_w \{D^Y(Z,w)\}$$

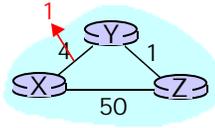
$$= 2 + 1 = 3$$

4: Network Layer 4a-22


Distance Vector: link cost changes

Link cost changes:

- r node detects local link cost change
- r updates distance table (line 15)
- r if cost change in least cost path, notify neighbors (lines 23,24)



“good news travels fast”





4: Network Layer 4a-23

Distance Vector: link cost changes

Link cost changes:

- r good news travels fast
- r bad news travels slow - “count to infinity” problem!





4: Network Layer 4a-24

Distance Vector: poisoned reverse

If Z routes through Y to get to X :

- r Z tells Y its (Z's) distance to X is infinite (so Y won't route to X via Z)
- r will this completely solve count to infinity problem?



4: Network Layer 4a-25

Comparison of LS and DV algorithms

Message complexity

- r **LS**: with n nodes, E links, $O(nE)$ msgs sent each
- r **DV**: exchange between neighbors only
 - m convergence time varies

Speed of Convergence

- r **LS**: $O(n^2)$ algorithm requires $O(nE)$ msgs
 - m may have oscillations
- r **DV**: convergence time varies
 - m may be routing loops
 - m count-to-infinity problem

Robustness: what happens if router malfunctions?

LS:

- m node can advertise incorrect *link* cost
- m each node computes only its *own* table

DV:

- m DV node can advertise incorrect *path* cost
- m each node's table used by others
 - error propagate thru network

4: Network Layer 4a-26

Hierarchical Routing

Our routing study thus far - idealization

- r all routers identical
- r network "flat"
- ... *not* true in practice

scale: with 50 million destinations:

- r can't store all dest's in routing tables!
- r routing table exchange would swamp links!

administrative autonomy

- r internet = network of networks
- r each network admin may want to control routing in its own network

4: Network Layer 4a-27

Hierarchical Routing


- r aggregate routers into regions, "**autonomous systems**" (AS)
- r routers in same AS run same routing protocol
 - m "**intra-AS**" routing protocol
 - m routers in different AS can run different intra-AS routing protocol

gateway routers


- r special routers in AS
- r run intra-AS routing protocol with all other routers in AS
- r *also* responsible for routing to destinations outside AS
 - m run **inter-AS routing** protocol with other gateway routers

4: Network Layer 4a-28

Intra-AS and Inter-AS routing




Intra-AS and Inter-AS routing



r We'll examine specific inter-AS and intra-AS Internet routing protocols shortly

The Internet Network layer


Host, router network layer functions:



4: Network Layer 4a-31

IP Addressing: introduction

- r IP address: 32-bit identifier for host, router *interface*
- r *interface*: connection between host, router and physical link
 - m router's typically have multiple interfaces
 - m host may have multiple interfaces
 - m IP addresses associated with interface, not host, router



$$223.1.1.1 = \underbrace{11011111}_{223} \underbrace{00000001}_{1} \underbrace{00000001}_{1} \underbrace{00000001}_{1}$$

4: Network Layer 4a-32






Figure 4-3

Internet address classes




4: Network Layer 4a-35

IP Addresses

given notion of "network", let's re-examine IP addresses:

"class-full" addressing:

class




← 32 bits →

4: Network Layer 4a-36

Figure 4-4

Classes using decimal notation



4: Network Layer 4a-37

Table 4-1


Numbers in classes

<i>Class</i>	<i>Number of Networks</i>	<i>Number of Hosts</i>
A	$2^7 - 2 = 126$	$2^{24} - 2 = 16,777,214$
B	$2^{14} = 16,384$	$2^{16} - 2 = 65,535$
C	$2^{21} = 2,097,152$	$2^8 - 2 = 254$
D	Not Applicable	Not Applicable
E	Not Applicable	Not Applicable

4: Network Layer 4a-38

IP Addressing

- r IP address:
 - m network part (high order bits)
 - m host part (low order bits)
- r What's a network ? (from IP address perspective)
 - m device interfaces with same network part of IP address
 - m can physically reach each other without intervening router




network consisting of 3 IP networks (for IP addresses starting with 223, first 24 bits are network address)

IP Addressing

- How to find the networks?
- r Detach each interface from router, host
 - r create "islands of isolated networks"

Interconnected system consisting of six networks



Types of addresses

- r Introduced at layer 2 - ex. Ethernet MAC address
- r Addresses at layer 3 are IP
- r Packets to specific computer are *unicast*
- r Packets to groups of computers are *broadcast* or *multicast*
- r 4.3 deals with broadcasts

4: Network Layer 4a-41

Table 4-2

Special addresses


<i>Special Address</i>	<i>Netid</i>	<i>Hostid</i>	<i>Source or Destination</i>
Network address	Specific	All 0s	None
Direct broadcast address	Specific	All 1s	Destination
Limited broadcast address	All 1s	All 1s	Destination
This host on this network	All 0s	All 0s	Source
Specific host on this network	All 0s	Specific	Destination
Loopback address	127	Any	Destination

Why?

4: Network Layer 4a-42

Figure 4-6


Examples of network addresses



4: Network Layer 4a-43

Figure 4-7

Example of direct broadcast address



4: Network Layer 4a-44

Figure 4-8

Example of limited broadcast address




Figure 4-9

Example of "this" host on "this" network address





Figure 4-10


Example of specific host on "this" network



4: Network Layer 4a-47

Figure 4-11


Example of loopback address



4: Network Layer 4a-48

Figure 4-12

Sample internet



4: Network Layer 4a-49

IP addressing: CIDR

r classful addressing:

- m inefficient use of address space, address space exhaustion
- m e.g., class B net allocated enough addresses for 65K hosts, even if only 2K hosts in that network

r CIDR: Classless InterDomain Routing

- m network portion of address of arbitrary length
- m address format: **a.b.c.d/x**, where x is # bits in network portion of address



4: Network Layer 4a-50

IP addresses: how to get one?

Hosts (host portion):

- r hard-coded by system admin in a file
- r **DHCP: Dynamic Host Configuration Protocol:** dynamically get address: "plug-and-play"
 - m host broadcasts "DHCP discover" msg
 - m DHCP server responds with "DHCP offer" msg
 - m host requests IP address: "DHCP request" msg
 - m DHCP server sends address: "DHCP ack" msg

4: Network Layer 4a-51

IP addresses: how to get one?

Network (network portion):


- r get allocated portion of ISP's address space:

ISP's block	<u>11001000</u>	<u>00010111</u>	<u>00010000</u>	00000000	200.23.16.0/20
Organization 0	<u>11001000</u>	<u>00010111</u>	<u>00010000</u>	00000000	200.23.16.0/23
Organization 1	<u>11001000</u>	<u>00010111</u>	<u>00010010</u>	00000000	200.23.18.0/23
Organization 2	<u>11001000</u>	<u>00010111</u>	<u>00010100</u>	00000000	200.23.20.0/23
...
Organization 7	<u>11001000</u>	<u>00010111</u>	<u>00011110</u>	00000000	200.23.30.0/23

4: Network Layer 4a-52

Hierarchical addressing: route aggregation


Hierarchical addressing allows efficient advertisement of routing information:



4: Network Layer 4a-53

Hierarchical addressing: more specific routes

ISPs-R-Us has a more specific route to Organization 1



4: Network Layer 4a-54

IP addressing: the last word...

Q: How does an ISP get block of addresses?

A: **ICANN**: Internet Corporation for Assigned

Names and Numbers

m allocates addresses

m manages DNS

m assigns domain names, resolves disputes

4: Network Layer 4a-55

Getting a datagram from source to dest.


IP datagram:

misc fields	source IP addr	dest IP addr	data
-------------	----------------	--------------	------

- r datagram remains unchanged, as it travels source to destination
- r addr fields of interest here

routing table in A

Dest. Net.	next router	Nhops
223.1.1		1
223.1.2	223.1.1.4	2
223.1.3	223.1.1.4	2



4: Network Layer 4a-56


Getting a datagram from source to dest.

misc fields	223.1.1.1	223.1.1.3	data
-------------	-----------	-----------	------

Starting at A, given IP datagram addressed to B:

- r look up net. address of B
- r find B is on same net. as A
- r link layer will send datagram directly to B inside link-layer frame
 - m B and A are directly connected

Dest. Net.	next router	Nhops
223.1.1		1
223.1.2	223.1.1.4	2
223.1.3	223.1.1.4	2



4: Network Layer 4a-57


Getting a datagram from source to dest.

misc fields	223.1.1.1	223.1.2.3	data
-------------	-----------	-----------	------

Starting at A, dest. E:

- r look up network address of E
- r E on *different* network
 - m A, E not directly attached
- r routing table: next hop router to E is 223.1.1.4
- r link layer sends datagram to router 223.1.1.4 inside link-layer frame
- r datagram arrives at 223.1.1.4
- r continued....

Dest. Net.	next router	Nhops
223.1.1		1
223.1.2	223.1.1.4	2
223.1.3	223.1.1.4	2



4: Network Layer 4a-58


Getting a datagram from source to dest.

misc fields	223.1.1.1	223.1.2.3	data
-------------	-----------	-----------	------

Arriving at 223.1.4,
destined for 223.1.2.2

- r look up network address of E
- r E on *same* network as router's interface 223.1.2.9
 - m router, E directly attached
- r link layer sends datagram to 223.1.2.2 inside link-layer frame via interface 223.1.2.9
- r datagram arrives at 223.1.2.2!!! (hooray!)

Dest. network	next router	Nhops	interface
223.1.1	-	1	223.1.1.4
223.1.2	-	1	223.1.2.9
223.1.3	-	1	223.1.3.27



4: Network Layer 4a-59