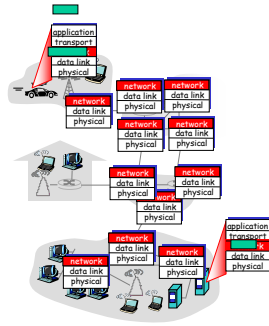


## Network layer

- ❖ transport segment from sending to receiving host
- ❖ on sending side encapsulates segments into datagrams
- ❖ on rcving side, delivers segments to transport layer
- ❖ network layer protocols in every host, router
- ❖ router examines header fields in all IP datagrams passing through it



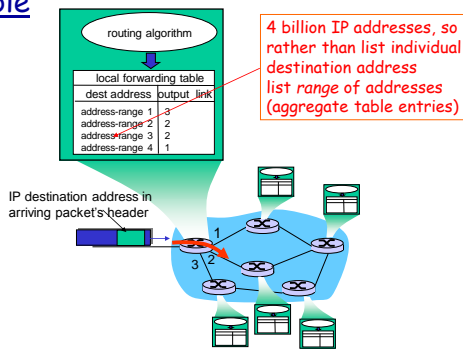
Network Layer 4-1

## Two Key Network-Layer Functions

- ❖ **forwarding**: move packets from router's input to appropriate router output
  - ❖ **routing**: determine route taken by packets from source to dest.
    - routing algorithms
- analogy:**
- ❖ **routing**: process of planning trip from source to dest
  - ❖ **forwarding**: process of getting through single interchange

Network Layer 4-2

## Datagram Forwarding table



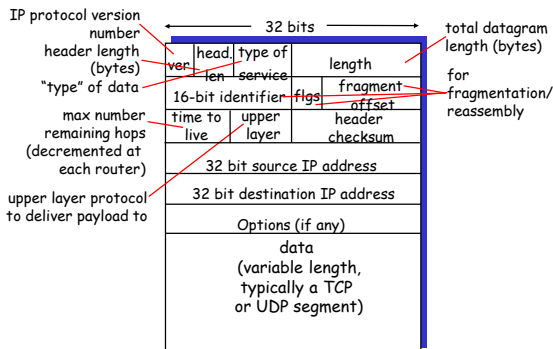
Network Layer 4-3

## Datagram Forwarding table

Destination Address Range	Link Interface
11001000 00010111 00010000 00000000 through 11001000 00010111 00010111 11111111	0
11001000 00010111 00011000 00000000 through 11001000 00010111 00011000 11111111	1
11001000 00010111 00011001 00000000 through 11001000 00010111 00011111 11111111	2
otherwise	3

Network Layer 4-4


## IP datagram format



Network Layer 4-5

## IP Addressing: introduction


- ❖ **IP address**: 32-bit identifier for host, router interface
- ❖ **interface**: connection between host/router and physical link
  - router's typically have multiple interfaces
  - host typically has one interface
  - IP addresses associated with each interface



Network Layer 4-6

## Subnets


- ❖ **IP address:**
  - subnet part (high order bits)
  - host part (low order bits)
- ❖ **What's a subnet ?**
  - device interfaces with same subnet part of IP address
  - can physically reach each other without intervening router



Network Layer 4-7

## Subnets

- ❖ **Recipe**
  - ❖ to determine the subnets, detach each interface from its host or router, creating islands of isolated networks
  - ❖ each isolated network is called a **subnet**.




Network Layer 4-8

## IP addressing: CIDR

### CIDR: Classless InterDomain Routing

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



Network Layer 4-9

## IP addresses: how to get one?

**Q:** How does a *host* get IP address?

- ❖ hard-coded by system admin in a file
  - Windows: control-panel->network->configuration->tcp/ip->properties
  - UNIX: /etc/rc.config
- ❖ **DHCP:** Dynamic Host Configuration Protocol: dynamically get address from as server
  - "plug-and-play"

Network Layer 4-10

## IP addressing: the last word...

**Q:** How does an ISP get block of addresses?

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

**Names and Numbers**

- allocates addresses
- manages DNS
- assigns domain names, resolves disputes

Network Layer 4-11


## IPv6

- ❖ **Initial motivation:** 32-bit address space soon to be completely allocated.
- ❖ Additional motivation:
  - header format helps speed processing/forwarding
  - header changes to facilitate QoS
- IPv6 datagram format:**
  - fixed-length 40 byte header
  - no fragmentation allowed

Network Layer 4-12

## IPv6 Header (Cont)

- Priority:** identify priority among datagrams in flow
- Flow Label:** identify datagrams in same "flow."  
(concept of "flow" not well defined).
- Next header:** identify upper layer protocol for data



Network Layer 4-13

Network Layer 4-14

## Other Changes from IPv4

- ❖ **Checksum:** removed entirely to reduce processing time at each hop
- ❖ **Options:** allowed, but outside of header, indicated by "Next Header" field
- ❖ **ICMPv6:** new version of ICMP
  - additional message types, e.g. "Packet Too Big"
  - multicast group management functions


## Transition From IPv4 To IPv6

- ❖ Not all routers can be upgraded simultaneously
  - no "flag days"
  - How will the network operate with mixed IPv4 and IPv6 routers?
- ❖ **Tunneling:** IPv6 carried as payload in IPv4 datagram among IPv4 routers

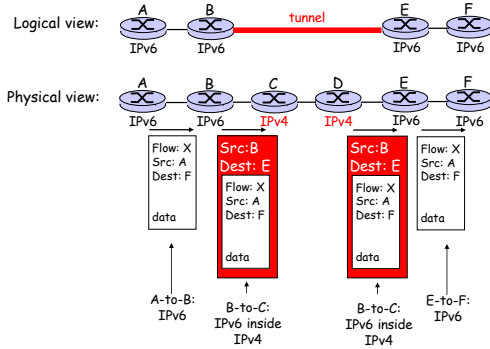
Network Layer 4-15

Network Layer 4-16

## Tunneling



## Tunneling




Network Layer 4-17

Network Layer 4-18

## Routingalgoritmer

Hur skapas innehållet i routingtabellerna??

## Graph abstraction



Graph:  $G = (N, E)$


$N$  = set of routers = { u, v, w, x, y, z }

$E$  = set of links = { (u,v), (u,x), (v,x), (v,w), (x,w), (x,y), (w,y), (w,z), (y,z) }

Remark: Graph abstraction is useful in other network contexts  
 Example: P2P, where  $N$  is set of peers and  $E$  is set of TCP connections

Network Layer 4-19

## Graph abstraction: costs



$c(x,x')$  = cost of link  $(x,x')$

- e.g.,  $c(w,z) = 5$

• cost could always be 1, or inversely related to bandwidth, or inversely related to congestion

Cost of path  $(x_1, x_2, x_3, \dots, x_p) = c(x_1, x_2) + c(x_2, x_3) + \dots + c(x_{p-1}, x_p)$

Question: What's the least-cost path between u and z ?

Routing algorithm: algorithm that finds least-cost path

Network Layer 4-20

## Routing Algorithm classification

### Global or decentralized information?

#### Global:

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

#### Decentralized:

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

### Static or dynamic?

#### Static:

- routes change slowly over time

#### Dynamic:

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

Network Layer 4-21

## A Link-State Routing Algorithm

### Dijkstra's algorithm

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

### Notation:

- $c(x,y)$ : link cost from node x to y;  $= \infty$  if not direct neighbors
- $D(v)$ : current value of cost of path from source to dest. v
- $p(v)$ : predecessor node along path from source to v
- $N'$ : set of nodes whose least cost path definitively known

Network Layer 4-22

## Dijkstra's Algorithm

- 1 **Initialization:**
- 2  $N' = \{u\}$
- 3 for all nodes v
- 4 if v adjacent to u
- 5 then  $D(v) = c(u,v)$
- 6 else  $D(v) = \infty$
- 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 shortest path cost to w plus cost from w to v \*/
- 14
- 15 **until all nodes in  $N'$**


Network Layer 4-23

## Dijkstra's algorithm: example

Step	$N'$	$D(v)$ p(v)	$D(w)$ p(w)	$D(x)$ p(x)	$D(y)$ p(y)	$D(z)$ p(z)
0	u	7,u	3,u	5,u	$\infty$	$\infty$
1	uw	6,w	6,u	11,w	$\infty$	
2	uwvx	6,w	11,w	14,x		
3	uwxv	6,w	10,v	14,x		
4	uwxvy	6,w	10,v	12,y		
5	uwxvzy					

### Notes:


- construct shortest path tree by tracing predecessor nodes
- ties can exist (can be broken arbitrarily)



Network Layer 4-24

## Dijkstra's algorithm: another example


Step	N'	D(v),p(v)	D(w),p(w)	D(x),p(x)	D(y),p(y)	D(z),p(z)
0	u	2,u	5,u	1,u	$\infty$	$\infty$
1	ux	2,u	4,x	2,x	$\infty$	$\infty$
2	uxy	2,u	3,y	4,y	4,y	$\infty$
3	uxyv	2,u	3,y	4,y	4,y	4,y
4	uxyvw	2,u	3,y	4,y	4,y	4,y
5	uxyvwz	2,u	3,y	4,y	4,y	4,y



Network Layer 4-25

## Dijkstra's algorithm: example (2)

Resulting shortest-path tree from u:



Resulting forwarding table in u:

destination	link
v	(u,v)
x	(u,x)
y	(u,x)
w	(u,x)
z	(u,x)

Network Layer 4-26


## Dijkstra's algorithm, discussion

Algorithm complexity:  $n$  nodes

- ❖ each iteration: need to check all nodes,  $w$ , not in  $N$
- ❖  $n(n+1)/2$  comparisons:  $O(n^2)$
- ❖ more efficient implementations possible:  $O(n \log n)$

Oscillations possible:

- ❖ e.g., link cost = amount of carried traffic



Network Layer 4-27

## Distance Vector Algorithm

Bellman-Ford Equation (dynamic programming)

Define

$d_x(y)$  := cost of least-cost path from  $x$  to  $y$


Then

$$d_x(y) = \min_v \{ c(x,v) + d_v(y) \}$$

where min is taken over all neighbors  $v$  of  $x$

Network Layer 4-28

## Bellman-Ford example



Clearly,  $d_v(z) = 5$ ,  $d_x(z) = 3$ ,  $d_w(z) = 3$

B-F equation says:

$$\begin{aligned}
 d_u(z) &= \min \{ c(u,v) + d_v(z), \\
 &\quad c(u,x) + d_x(z), \\
 &\quad c(u,w) + d_w(z) \} \\
 &= \min \{ 2 + 5, \\
 &\quad 1 + 3, \\
 &\quad 5 + 3 \} = 4
 \end{aligned}$$

Node that achieves minimum is next hop in shortest path  $\rightarrow$  forwarding table

Network Layer 4-29

## Distance Vector Algorithm

- ❖  $D_x(y)$  = estimate of least cost from  $x$  to  $y$ 
  - $x$  maintains distance vector  $D_x = [D_x(y): y \in N]$
- ❖ node  $x$ :
  - knows cost to each neighbor  $v$ :  $c(x,v)$
  - maintains its neighbors' distance vectors. For each neighbor  $v$ ,  $x$  maintains  $D_v = [D_v(y): y \in N]$


Network Layer 4-30

## Distance vector algorithm (4)

### Basic idea:

- ❖ from time-to-time, each node sends its own distance vector estimate to neighbors
- ❖ when x receives new DV estimate from neighbor, it updates its own DV using B-F equation:
 
$$D_x(y) \leftarrow \min_v \{c(x,v) + D_v(y)\} \text{ for each node } y \in N$$
- ❖ under minor, natural conditions, the estimate  $D_x(y)$  converge to the actual least cost  $d_x(y)$

Network Layer 4-31



Network Layer 4-33

## Distance Vector Algorithm (5)

**Iterative, asynchronous:**  
each local iteration caused by:

- ❖ local link cost change
- ❖ DV update message from neighbor

**Distributed:**

- ❖ each node notifies neighbors *only* when its DV changes
  - neighbors then notify their neighbors if necessary


**Each node:**

wait for (change in local link cost or msg from neighbor)

recompute estimates

if DV to any dest has changed, notify neighbors

Network Layer 4-32




Network Layer 4-34

## Distance Vector: link cost changes

### Link cost changes:

- ❖ node detects local link cost change
- ❖ updates routing info, recalculates distance vector
- ❖ if DV changes, notify neighbors



"good news travels fast"

$t_0$ : y detects link-cost change, updates its DV, informs its neighbors.

$t_1$ : z receives update from y, updates its table, computes new least cost to x, sends its neighbors its DV.


$t_2$ : y receives z's update, updates its distance table. y's least costs do not change, so y does not send a message to z.

Network Layer 4-35

## Distance Vector: link cost changes

### Link cost changes:

- ❖ good news travels fast
- ❖ bad news travels slow - "count to infinity" problem!
- ❖ 44 iterations before algorithm stabilizes: see text



### Poisoned reverse:

- ❖ If Z routes through Y to get to X:
  - Z tells Y its (Z's) distance to X is infinite (so Y won't route to X via Z)
- ❖ will this completely solve count to infinity problem?

Network Layer 4-36

## Comparison of LS and DV algorithms

### Message complexity

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

### Speed of Convergence

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

### Robustness: what happens if router malfunctions?

#### LS:

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

#### DV:

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