


More on IP and the Network Layer



Dept. of Computer Science, University of Rochester



The Internet Protocol

- hierarchical addressing
- packet format
- hierarchical routing:
 - intra-AS (RIP, OSPF)
 - inter-AS (BGP)



IPv6

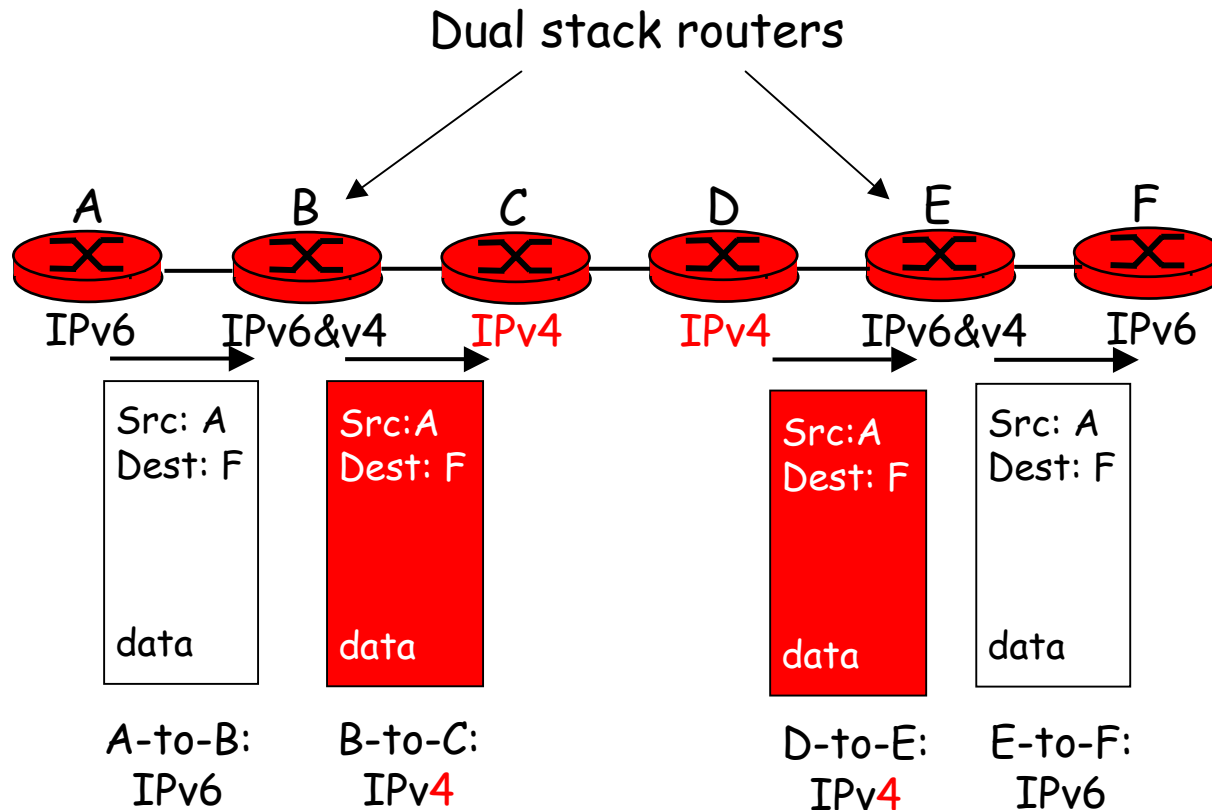
- **A key motivation:** address depletion \Rightarrow 32-bit address space soon to be used up.
 - 128-bit addresses
- **Additional motivation:**
 - simplification help speed processing/forwarding
 - fixed-length 40 byte header
 - no fragmentation allowed in intermediate routers
 - header changes to facilitate QoS
 - flow label
 - new "anycast" address: route to one of a set of several servers



Transition From IPv4 To IPv6

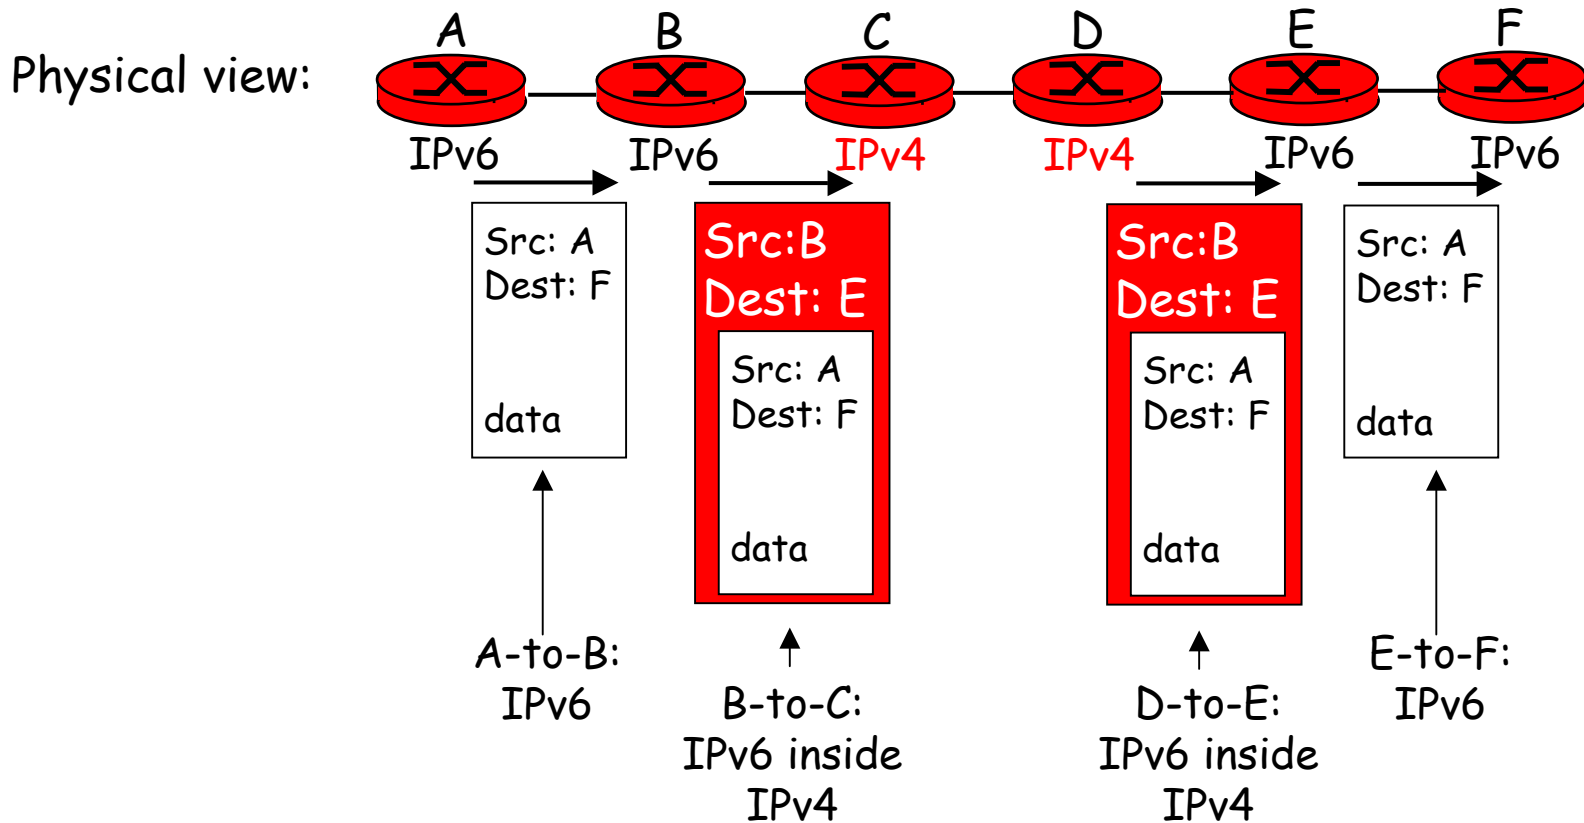

- It is difficult to change the network layer protocol
 - The large number of routers and administrative autonomy makes simultaneous change impossible.
 - Routers speaking different protocols do not naturally work together.
 - What should we do?
- Two proposed approaches:
 - Dual Stack: some routers with dual stack (v6, v4) can “translate” between formats
 - Tunneling: IPv6 carried as payload in IPv4 packet among IPv4 routers

Dual Stack Approach



Are there any problems?

Tunneling Approach





IPv6 Deployment

IPv6 deployment is slow:


- resistance to change \Leftrightarrow lack of critical mass
- immature technology: anycast

how about address depletion?

- DHCP (dynamic host configuration protocol)
- NAT (network address translation)

IPsec: Network Layer Security

- security services needed at the network layer:
 - source authentication
 - data integrity
 - data confidentiality
- source authentication & data integrity
 - signed message digests
- data confidentiality
 - encryption using a symmetric session key



- IPsec deployment?

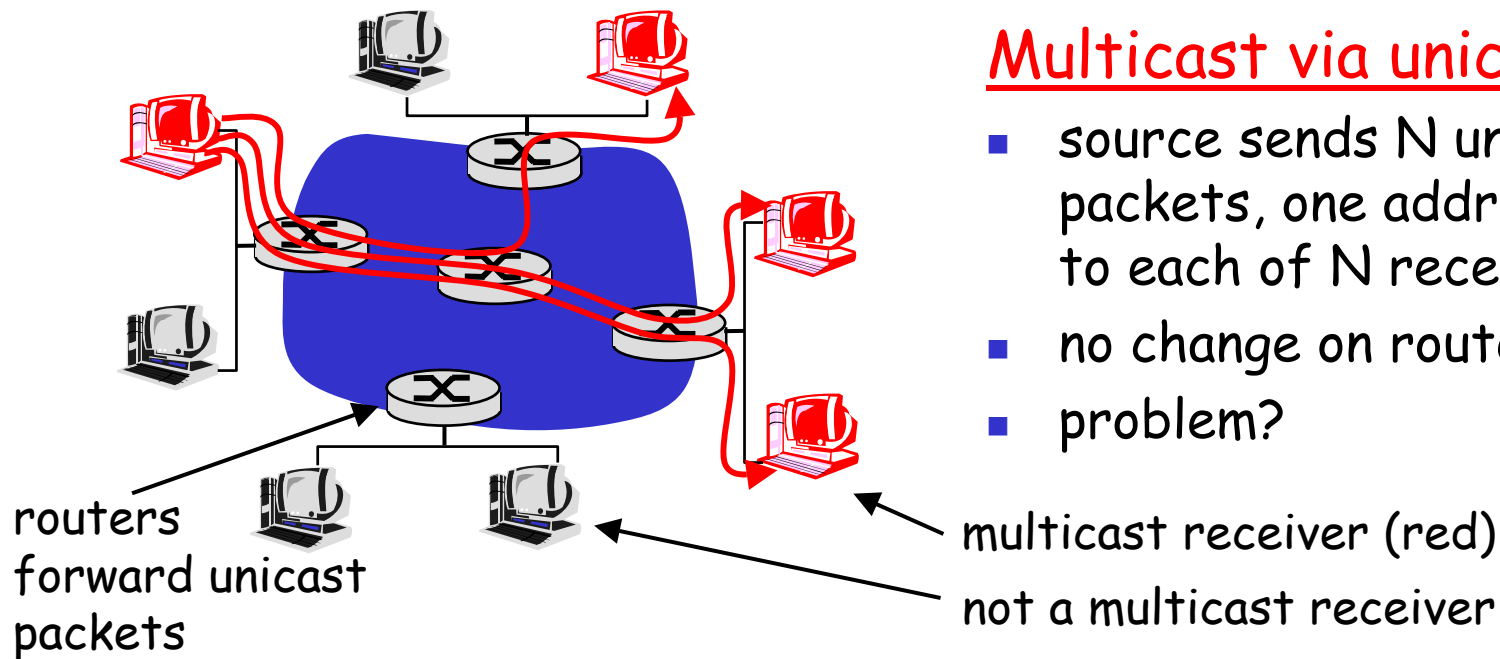


Outline

- IPv6
- IPsec
- Multicast routing
 - basic concept
 - Principle: basic approaches
 - Practice: protocols

Multicast

- Multicast: act of sending packet to multiple receivers with single "transmit" operation
- Applications??

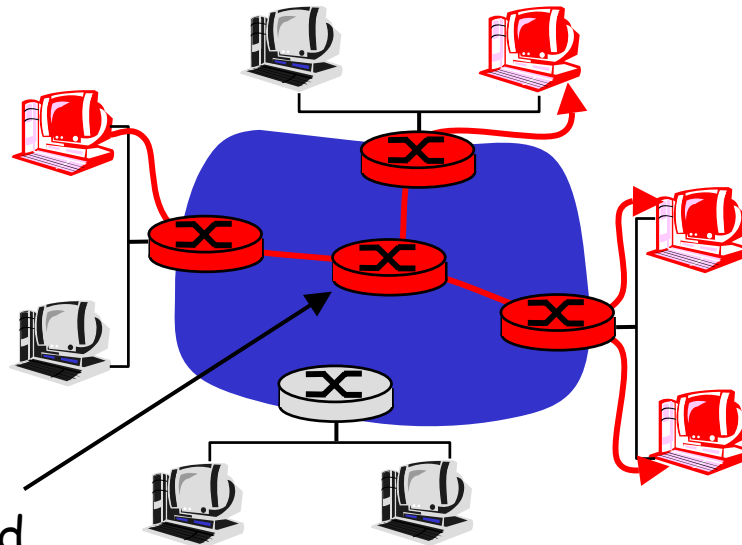


Multicast via unicast

- source sends N unicast packets, one addressed to each of N receivers
- no change on routers
- problem?


Ideal Network Multicast

Router actively participate in multicast, making copies of packets as needed and forwarding towards multicast receivers



Multicast routers
(red) duplicate and
forward multicast
packets

Internet Multicast Service Model

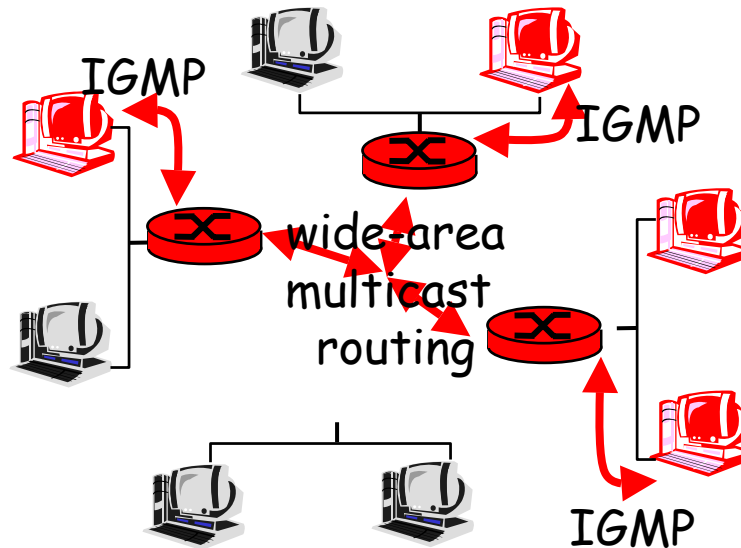


Multicast group concept: use of indirection

- hosts "join" multicast group
- hosts send IP packets with the multicast group address
- routers forward multicast packets to hosts that have "joined" that multicast group


Local and Wide-Area Multicast Management

- local: host informs local mcast router of desire to join group
 - IGMP (Internet Group Management Protocol)
- wide area: local router interacts with other routers to receive mcast packet flow




IGMP: Internet Group Management Protocol

- host: sends IGMP report when application joins mcast group
- local router: sends IGMP query at regular intervals
 - host belonging to a mcast group must reply to query
- host need not explicitly "unjoin" group when leaving




Wide-area Multicast Routing

- **Goal:** find path/paths (tree) connecting routers having local mcast group members
- **source-based tree:** one tree per source
 - shortest path trees, reverse path forwarding
- **group-shared tree:** group uses one tree
 - minimal spanning (Steiner) , center-based trees




Source-based trees



Shared tree

Shortest Path Tree

- mcast forwarding tree: tree of shortest path routes from source to all receivers
 - Dijkstra's algorithm





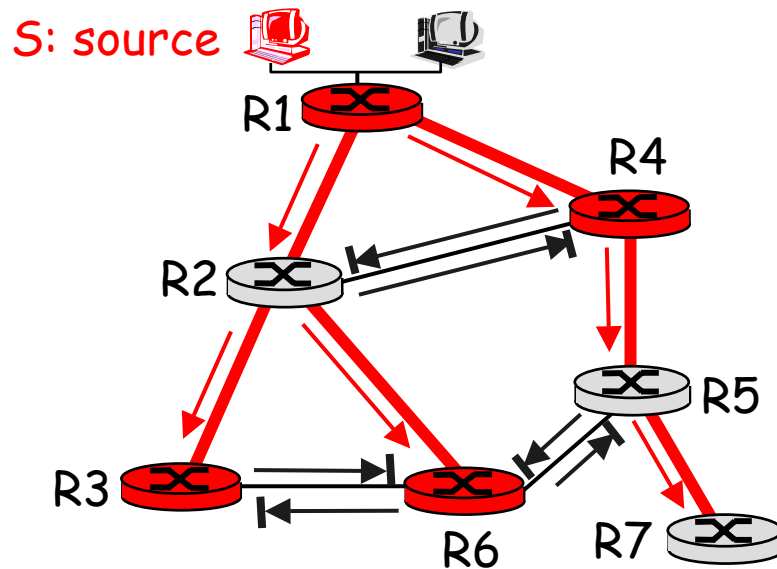
Reverse Path Forwarding

What if we don't have global network information?





- rely on router's knowledge of unicast routing table
- each router has simple forwarding behavior:

if (mcast packet received on incoming link on
shortest path back to source)
then flood packet onto all outgoing links
else ignore packet

Reverse Path Forwarding: Example




LEGEND

-  router with attached group member
-  router with no attached group member
-  packet will be forwarded
-  packet will not be forwarded





- result is a source-specific *reverse* SPT
- problem: reaching routers that are not in the group

Reverse Path Forwarding: Pruning

- forwarding tree contains subtrees with no mcast group members
 - "prune" msgs sent upstream by router with no downstream group members



LEGEND

-  router with attached group member
-  router with no attached group member
-  prune message
-  links with multicast forwarding



Shared-Tree

- Single delivery tree shared by all
- Assume complete global network information is available, derive minimum cost tree connecting all routers with attached group members
 - problem is NP-complete
 - heuristic approximation exists




A distributed solution: Center-based Tree




- One router identified as "*center*" of tree
- To join:
 - edge router sends unicast *join-msg* addressed to center router
 - *join-msg* "processed" by intermediate routers and forwarded towards center
 - *join-msg* either hits existing tree branch for this center, or arrives at center
 - path taken by *join-msg* becomes new branch of tree for this router

Center-based Trees: An Example

Suppose R6 chosen as center:



LEGEND

-  router with attached group member
-  router with no attached group member
-  path order in which join messages generated



Internet Multicasting Routing

- **DVMRP**: distance vector multicast routing protocol, based on RPF (reverse path forwarding) and pruning
 - first and the most commonly supported
- **PIM**: protocol independent multicast
 - in a **sparse** mode where DVMRP doesn't do well, it uses Center-based Trees
- IP multicast deployment?



Disclaimer

- Parts of the lecture slides contain original work of James Kurose, Larry Peterson, and Keith Ross. The slides are intended for the sole purpose of instruction of computer networks at the University of Rochester. All copyrighted materials belong to their original owner(s).