Multicast Routing

Issue: Some applications, such as a multiplayer game or live video of a sports event streamed to many viewing locations, send packets to multiple receivers. Unless the group is very small, sending a distinct packet to each receiver is expensive. On the other hand, broadcasting a packet is wasteful if the group consists of, say, 1000 machines on a million-node network, so that most receivers are not interested in the message (or worse yet, they are definitely interested but are not supposed to see it).

Multicast Routing

Thus, we need a way to send messages to well-defined groups that are numerically large in size but small compared to the network as a whole.

- Sending a message to such a group is called **multicasting**, and the routing algorithm used is called **multicast routing**.

Multicast Routing

- Multicast routing schemes build on the broadcast routing schemes we have already studied, sending packets along spanning trees to deliver the packets to the members of the group while making efficient use of bandwidth.
- However, the best spanning tree to use depends on whether the group is dense, with receivers scattered over most of the network, or sparse, with much of the network not belonging to the group.

11

Multicast spanning tree

- The solution explored by Deering and Cheriton (1990) is to prune the broadcast spanning tree by removing links that do not lead to members. The result is an efficient multicast spanning tree.

12

Multicast spanning tree

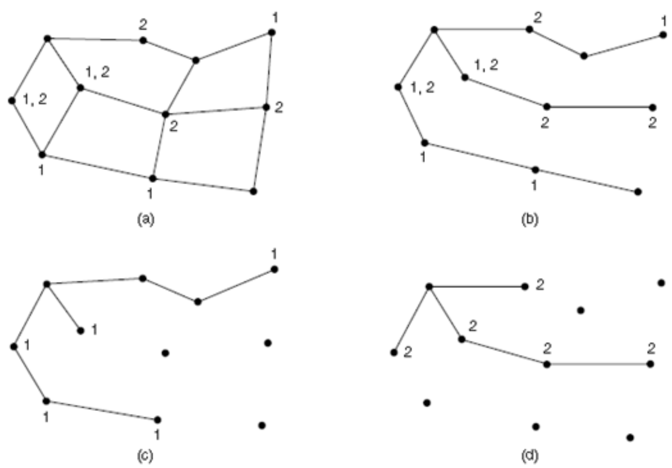

Figure 5-16. (a) A network. (b) A spanning tree for the leftmost router. (c) A multicast tree for group 1. (d) A multicast tree for group 2.

13

Multicast spanning tree

- Various ways of pruning the spanning tree are possible. The simplest one can be used if link state routing is used and each router is aware of the complete topology, including which hosts belong to which groups.
- Each router can then construct its own pruned spanning tree for each sender to the group in question by constructing a sink tree for the sender as usual and then removing all links that do not connect group members to the sink node.
- **MOSPF** (**Multicast OSPF**) is an example of a link state protocol that works in this way (Moy, 1994).

14

Multicast spanning tree

- With distance vector routing, a different pruning strategy can be followed. The basic algorithm is reverse path forwarding.
- Whenever a router with no hosts interested in a particular group and no connections to other routers receives a multicast message for that group, it responds with a PRUNE message.
- When a router with no group members among its own hosts has received such messages on all the lines to which it sends the multicast, it, too, can respond with a PRUNE message.
- **In this way, the spanning tree is recursively pruned.**

- **DVMRP** (**Distance Vector Multicast Routing Protocol**) is an example of a multicast routing protocol that works this way (Waitzman et al., 1988).

A disadvantage

Pruning results in efficient spanning trees that use only the links that are actually needed to reach members of the group.

- One potential disadvantage is that it is lots of work for routers, especially for large networks. Suppose that a network has *n* groups, each with an average of *m* nodes. At each router and for each group, *m* pruned spanning trees must be stored, for a total of *m.n* trees. When many large groups with many senders exist, considerable storage is needed to store all the trees.

Core-based trees

- An alternative design uses **core-based trees** to compute a single spanning tree for the group (Ballardie et al., 1993). All of the routers agree on a root (called the **core** or **rendezvous point**).
- To send to this group, a sender sends a packet to the core. When the packet **reaches the core, it is forwarded down the tree.**
- As a **performance optimization**, packets destined for the group do not need to reach the core before they are multicast. As soon as a packet reaches the tree, it can be forwarded up toward the root, as well as down all the other branches.

17

Core-based trees

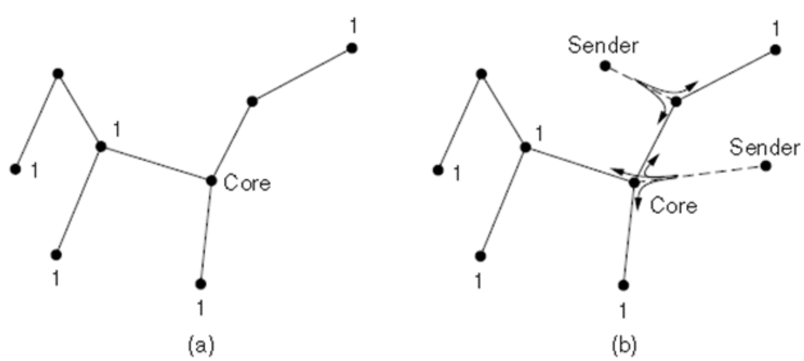

Figure 5-17. (a) Core-based tree for group 1. (b) Sending to group 1.

18

Core-based trees

- Often it is reasonable when the core is in the middle of the senders. When there is only a single sender, as in a video that is streamed to a group, using the sender as the core is optimal
- Also of note is that shared trees can be a major savings in storage costs, messages sent, and computation. Each router has to keep only one tree per group, instead of *m* trees.

- For this reason, shared tree approaches like core-based trees are used for multicasting as part of popular protocols such as **PIM** (**Protocol Independent Multicast**) (Fenner et al., 2006).

19

Anycast Routing

- In anycast, a packet is delivered to the nearest member of a group (Partridge et al., 1993). Schemes that find these paths are called **anycast routing**. For example, anycast is used in the Internet as part of DNS.

- Luckily, we will not have to devise new routing schemes for anycast because regular distance vector and link state routing can produce anycast routes. This procedure works because the routing protocol does not realize that there are multiple instances of destination.

20

Anycast Routing

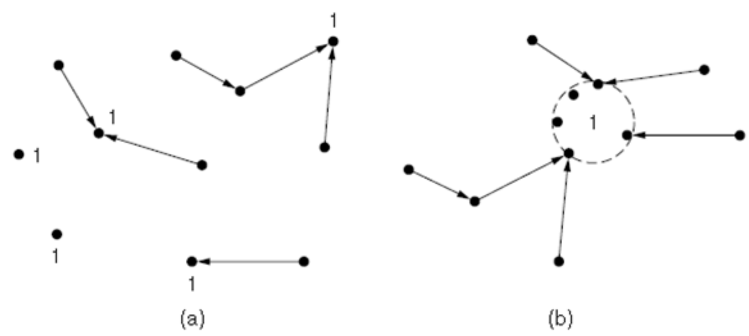

Figure 5-18. (a) Anycast routes to group 1. (b) Topology seen by the routing protocol.

21

Routing for Mobile Hosts

- Increasingly, people want to stay connected wherever in the world they may be, as easily as if they were at home. These mobile hosts introduce a new complication: to route a packet to a mobile host, the network first has to find it.

22

Recomputed routes

A model would be to recompute routes as the mobile host moves and the topology changes. We could then simply use the routing schemes described earlier in this section. However, with a growing number of mobile hosts, this model would soon lead to the entire network endlessly computing new routes.

Mobility in application layer

- Another alternative would be to provide mobility above the network layer, which is what typically happens with laptops today. When they are moved to new Internet locations, laptops acquire new network addresses. There is no association between the old and new addresses; the network does not know that they belonged to the same laptop.
- In this model, a laptop can be used to browse the Web, but other hosts cannot send packets to it (for example, for an incoming call), without, for example, signing into Skype *again* after moving. Moreover, connections cannot be maintained while the host is moving; new connections must be started up instead.

Mobility in network layer

- The model of the world that we will consider is one in which all hosts are assumed to have a permanent **home location** that never changes. Each hosts also has a permanent home address that can be used to determine its home location.
- The basic idea is for the mobile host to tell a host at the home location where it is now. This host, which acts on behalf of the mobile host, is called the **home agent**. Once it knows where the mobile host is currently located, it can forward packets so that they are delivered.

Mobility in network layer

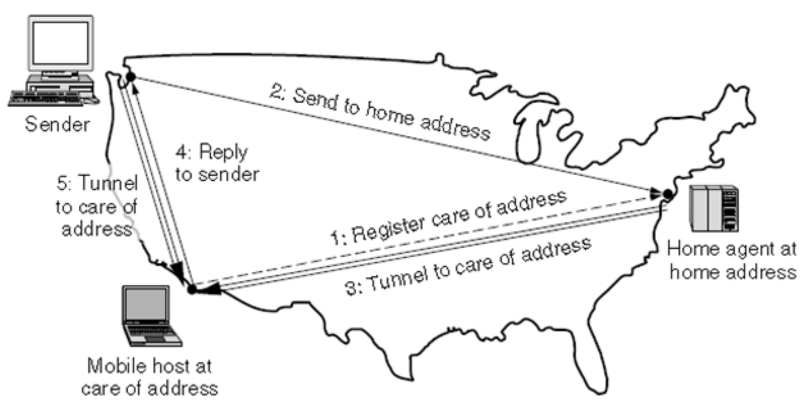

Figure 5-19. Packet routing for mobile hosts.

* Knowledge of the temporary location of the mobile host is limited to a small number of hosts (e.g., the mobile, home agent, and senders) so that the many routers in a large network do not need to recompute routes.