

CSc 461/561

Multimedia Systems

Multicast




Jianping Pan
Spring 2006

2/10/06

CSc 461/561

1

What is multicast?

- Unicast: one sender, one receiver 
- Broadcast
 - one sender, all (possible) receivers
- Multicast
 - one sender, many willing receivers
 - many senders, many receivers
- Anycast
 - one sender, one of many possible receivers



2/10/06

CSc 461/561

2

Why multicast?

- Some applications involve many parties
 - e.g., webcasting, video conferencing
- Multicast can be emulated by
 - a bunch of unicast streams
 - too much replicated traffic!
- A better approach
 - replicate only when it is really necessary
 - IP multicast, application-layer multicast




2/10/06

CSc 461/561

3

IP multicast

- Receiver-oriented model
 - willing receivers with a common identity
 - IP multicast address (class D, 224-239.X.X.X)
 - receivers join a multicast group *explicitly*
 - senders does not care
 - senders: even no need to be a member of the group
 - send packets with destination to the group
 - best-effort network: bring packets to receivers
 - flow/error/congestion control: upper layer!



2/10/06

CSc 461/561

4

IP multicast addresses

- Class D IP address
 - 224.0.0.0 - 239.255.255.255
 - 224.0.0.0 - 224.0.0.255
 - local network control block
 - 224.0.0.1: all hosts; 224.0.0.2: all routers
 - 224.0.0.22: IGMP
 - 224.0.1.0 - 224.0.1.255
 - inter-network control block
 - <http://www.iana.org/assignments/multicast-addresses>


2/10/06

CSc 461/561

5

How multicast is done?

- Shared media (i.e., multicast-capable)
 - e.g., Ethernet
 - multicast address: 01-00-5e-0xxxxxxx-X-X
 - mapped from IP multicast address as a filter
- Non-shared media (by replication)
 - e.g., routers receive packets from one interface
 - replicate and send through all other interfaces with reachable *active* multicast receivers




2/10/06

CSc 461/561

6

IGMP




- Internet group management protocol
 - IGMP Join/Report sent by host
 - “I am here and forward me multicast packets”
 - IGMP Query sent by router with timer
 - “are you still there?”
 - membership refreshed by Report sent by host
 - IGMP Leave sent by host
 - explicit in IGMPv2: “I am leaving”
 - IGMPv3: source-specific joining and leaving

2/10/06

CSc 461/561

7

Flooding



- Flooding
 - receive from one interface
 - send through all other interfaces
 - broadcast storm problem
- Controlled flooding
 - no packets appear on the same link twice
 - if the packet received before, drop it
 - only forward packet from *reverse* shortest path


2/10/06

CSc 461/561

8

Pruning

- Spanning tree
 - no packet received by any node twice!
 - one path from source to any receiver
- If no member in downstream networks
 - by IGMP etc
 - remove such branches
 - periodical flooding or explicit join necessary



2/10/06

CSc 461/561

9

Multicast routing

- Essentially, building a tree
 - sender-specific tree; or
 - shared tree, i.e., shared by many senders
- Ideally, the tree should be
 - low cost: overall (for sender) or for receiver
 - balanced among members
 - stable: w.r.t. node joining and leaving
 - it is a HARD problem!

2/10/06

CSc 461/561

10

Sender-specific tree

- Shortest path from sender to receivers
 - Dijkstra's algorithm
 - build the tree from the sender
 - add a node (not in the tree) closest to the sender
 - check whether this node can also reduce the sender's distance to other nodes in the tree
 - repeat until all nodes are in the tree
 - node distance should be non-negative

2/10/06

CSc 461/561

11

Shared tree

- Minimal spanning tree
 - global knowledge of network topology
 - begin with the link of minimal cost
 - construct tree gradually by adding more links of increasing cost until a tree is formed
- Center-based tree: one node as the “center”
 - Join message sent to the center
 - possibly meet a node already in the group

2/10/06

CSc 461/561

12

Multicast routing protocols

- DVMRP
 - extension to DVRP (distance vector-based)
 - flooding and pruning
 - sender-specific tree with reverse shortest path
- MOSPF
 - extension to OSPF (link state-based)
 - global knowledge of network topology, cost etc
 - shortest path tree

2/10/06

CSc 461/561

13

PIM

- Independent of underlying routing protocol
- Dense mode: densely packed members
 - sender-specific tree
 - flood and prune
 - reverse shortest path
- Sparse mode: widely scattered members
 - shared tree
 - center-based tree at rendezvous point


2/10/06

CSc 461/561

14

Mbone: multicast backbone

- Multicast tunneling
 - traverse “not multicast-capable” networks
 - with tunnel between multicast-capable nodes
 - routers or hosts
- Mbone applications
 - vat: visual audio tool
 - wb: white board
 - sd: session directory



2/10/06

CSc 461/561

15

Application-layer multicast

- Inter-domain multicast is difficult
 - not only in technology, but also in policies
 - E.g., MBGP: multicast BGP (shared tree)
- Application-layer multicast
 - end-hosts act as multicast “router” to build trees
 - application-aware flow/error/congestion control
 - not as efficient as IP multicast
 - but provide more flexibility and easy-to-deploy

2/10/06

CSc 461/561

16

This lecture

- Multicast
 - IP multicast; IGMP
 - multicast routing
 - shortest-path tree
 - minimal spanning tree
- Bonus question
 - discuss the pros and cons of “forward first received only” and “forward received on reverse shortest path only” in controlled flooding

2/10/06

CSc 461/561

17

Next lecture

- RTP/RTCP
 - REF: <http://www.cs.columbia.edu/~hgs/rtp/>

2/10/06

CSc 461/561

18