

Multicast

Algorithmen und Programmierung V Netzprogrammierung

Volker Roth

Institut für Informatik
Freie Universität Berlin

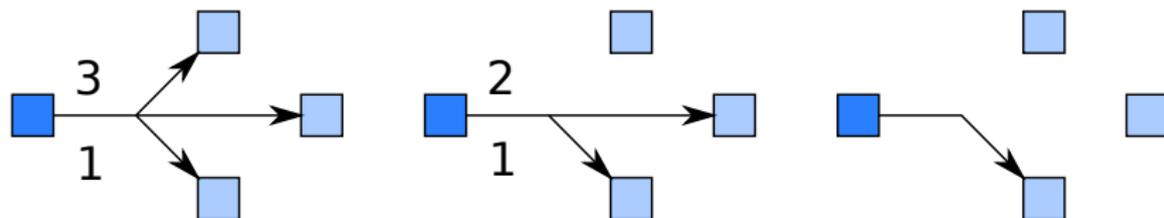
Wintersemester 2009/2010

Freie Universität



Berlin

Concepts



Sending information on a network:

- 1 : * Broadcast
- 1 : n Multicast
- 1 : 1 Unicast
- 1 : ? Anycast

Multicast

True broadcast and multicast:

- ▶ no redundant packets
- ▶ may require spanning tree

Multicast can live on different layers:

- ▶ data link layer (Ethernet)
- ▶ networking layer (IP)
- ▶ application layer (e.g., Overcast)

Connection-oriented vs. not, reliable vs. not.

Internet Protocol Multicast

Unreliable service based on UDP

Internet Group Management Protocol (IGMP)

- ▶ group membership management
- ▶ connect client to local multicast router

Protocol Independent Multicast (PIM)

- ▶ routing protocol (without topology discovery)
- ▶ sparse-mode most common
- ▶ other modes: dense, bidirectional, source-specific

Internet Protocol Multicast

Can be viewed as a multicast variant of UDP

- ▶ each destination address identifies a group of hosts
- ▶ available range is 224.0.0.1–239.255.255.255
- ▶ receivers join the group
- ▶ senders send to the group address
- ▶ all receivers receive, if things go well

Here is an explanation of the range choice:

- ▶ 0, class A addr
- ▶ 1, 0, class B addr (30)
- ▶ 1, 1, 0, class C addr (29)
- ▶ 1, 1, 1, 0, multicast addr (28)

Multicast Addresses

Some well-known multicast groups (off-limits)

224.0.0.1 all-hosts group (try pinging it, once)

224.0.0.2 all-routers group

224.0.0.5 OSPF routers

224.0.0.* local purposes (not routed)

239.*.*.* administrative scoping

Mapping to Ethernet Frames

Mapping of IP multicast addresses to Ethernet addresses

- ▶ Ethernet frame has 48 bits
- ▶ 01-00-5e-*-*-* (lowest 23 bits) reserved for multicast
- ▶ IP multicast has 28 bits
- ▶ 32:1 mapping, imperfect filter

Limiting the Distribution

Time-to-live (TTL)

- ▶ field in IP header with added significance in multicast
- ▶ router interfaces can be assigned TTL threshold
- ▶ forward & decrement TTL only if $TTL > \text{threshold}$

TTL thresholds and associated scopes

- 0 same host
- 1 same subnet
- <32 same site, org or department
- <64 same region
- <128 same continent

Interpretation is fuzzy...

Leaving and Joining

Per-host membership not per-process

- ▶ Leave group but remain bound on port receives packets if other processes on same host are still joined