IP Multicasting

IP Multicast Principles and Applications, IGMP
DVRMP, MOSPF, PIM-SM, PIM-DM, MBone, RTP/RTPC

Agenda

- **IP Multicasting**
  - RFC 1112
  - IGMPv1, IGMPv2, IGMPv3
  - IGMP Snooping
- **IP Multicast Routing**
  - DVMRP
  - PIM-DM
  - PIM-SM
  - MOSPF
- **MBone**
- **Multicast Applications**
- **RTP/RTPC**

 IP_Multicasting, v4.4 2

IP-Multicast Service-Model Summary

- **According to RFC 1112 IP-Multicast architecture can be summarized:**
  - Senders send to IP multicast address
  - Receivers express an interest in an IP multicast address
    - register as group members for multicast address via IGMP
  - Routers conspire to deliver traffic from the senders to the receivers
    - building distribution-trees, duplicating packets,...
    - actually done by multicast routing protocols

 IP_Multicasting, v4.4 3

IP Multicasting

- **only difference between point-to-point IP packet and multicast IP packet**
  - group address in the destination field
- **IP addresses of class D are used as group address of multicast group**
  - class D address: 1110xxxx.xxxxxxxx.xxxxxxxx.xxxxxxxx
  - So 2^{28} = 268.435.456 possible groups
  - class D range: 224.0.0.0 up to 239.255.255.255
    - well known multicast addresses
    - transient addresses
      - assigned and reclaimed dynamically

 IP_Multicasting, v4.4 4

Multicast Service Models

- **Any Source Multicast (ASM)**
  - an end point indicates that it wishes to receive all multicast traffic sent to a certain multicast address, no matter where that traffic originated
  - other name Internet Standard Multicast (ISM)
- **Source Specific Multicast (SSM)**
  - an endpoint indicates that it wishes to receive multicast traffic sent to a certain multicast address, but only if it originates from a specific identified source address
  - other name Single-Source Multicast
  - coupled to IGMPv3

IP Multicast Addresses 1

- **some well known multicast addresses**
  - Local Network Control Block: 224.0.0.0 – 224.0.0.255
    - is used by network protocols on a local subnet segment with Time To Live (TTL) of 1

    - 224.0.0.1 … all systems on this subnet
    - 224.0.0.2 … all routers on this subnet
    - 224.0.0.4 … all DVMRP routers
    - 224.0.0.5 … all OSPF routers
    - 224.0.0.6 … all designated OSPF routers
    - 224.0.0.9 … all RIPv2 routers
    - 224.0.0.10 … all eIGRP routers
    - 224.0.0.11 … all mobile agents
    - 224.0.0.12 … all DHCP server/relay agents
    - 224.0.0.13 … all PIM routers
    - 224.0.0.18 … all VRRP routers
    - 224.0.0.22 … all IGMP queriers
    - 224.0.0.102 … all HSRP routers

Institute of Computer Technology - Vienna University of Technology

L71 - IP Multicasting

IP Multicast Addresses 2

- **some well known multicast addresses (cont.)**
  - Internetwork Control Block: 224.0.1.0 – 224.0.1.255
    - These addresses are similar to the Local Network Control Block except that they are used by network protocols when control messages need to be multicast beyond the local network segment

    - 224.0.1.1 … NTP Network Time Protocol
    - 224.0.1.8 … Sun NIS+ Information Service
    - 224.0.1.14 … IETF-2-AUDIO
    - 224.0.1.15 … IETF-2-VIDEO
    - 224.0.1.21 … DVMRP on MOSPF
    - 224.0.1.24 … Microsoft-ds
    - 224.0.1.32 … Mtrace
    - 224.0.1.33 … RVSP-encap-1
    - 224.0.1.34 … RVSP-encap-2
    - 224.0.1.39 … Cisco RP Announce
    - 224.0.1.40 … Cisco RP Discovery
    - 224.0.1.41 … Gatekeeper

IP Multicast Addresses 3

- ADHOC Block: 224.0.2.0 – 224.0.2.255
  - 224.0.18.0 – 224.0.18.255 … Dow Jones
  - 224.0.19.0 – 224.0.19.63 … Walt Disney

- SDP/SAP Block: 224.2.0.0 – 224.255.255.0
  - The multicast group range of 224.2.0.0 through 224.2.255.255 (224.2/16) is the SDP/SAP Multicast Block, which is reserved for applications that send and receive multimedia session announcements using the SAP described in RFC 2974

  - 224.2.0.0 – 224.2.127.253 … Multimedia Conference Calls
  - 224.2.127.254 … UDP Port 9875 SAPv1 / SDP
    - IPv4 Region (ttl=63) or World (ttl=127) in case of SDR tool
  - 224.2.128.0 – 224.2.255.255 … SAP Dynamic Assignments

IP Multicast Addresses 4

- Various/reserved:
  - 224.3.0.0 – 231.255.255.255
- Source Specific Multicast (SSM):
  - The multicast group range of 232.0.0.0 through 232.255.255.255 (232/8) is reserved for SSM. SSM is a new extension to PIM Sparse mode that eliminates the need for the Rendezvous Point and the Shared Tree and uses only the Shortest-Path Tree to the desired sources.
  - 232.0.0.0 – 232.255.255.255
- Reserved:
  - 234.0.0.0 – 238.255.255.255

IP Multicast Addresses 5

- Administratively Scoped Addresses:
  - These addresses are similar in nature to the reserved IP unicast ranges (such as 10.0.0.0/8) defined in RFC 1918 and will not be assigned by the IANA to any other group or protocol. This means that network administrators are free to use multicast addresses in this range inside of their domain without fear of conflicting with others elsewhere in the Internet. However, administratively scoped addresses should not be used when sourcing IP Multicast traffic to the Internet. This is because it has become a common practice to block multicast traffic in these ranges from entering or leaving an Autonomous Domain.
  - 239.0.0.0 – 239.255.255.255
  - "Private IP Multicast Addresses"
  - 239.192.0.0 – 239.251.255.255 … Organizational-Local Scope
  - 239.255.0.0 – 239.251.255.255 … Site-Local Scope
  - 239.255.255.255 … UDP Port 9875 SAPv1 / SDP
    - IPv4 Local Scope (ttl=15) in case of SDR tool
- RFC3171 (BCP) provides general guidance
  - on the use of the multicast address space and procedures for number allocation in certain blocks of the addresses

Institute of Computer Technology - Vienna University of Technology

L71 - IP Multicasting

IP Multicasting on Broadcast Networks

- **on broadcast networks an IP multicast packet should be sent**
  - either as layer 2 global broadcast
    - all stations will receive multicast packet on this network
    - 0xFF-FF-FF-FF-FF-FF (IEEE 802 LAN)
  - or as layer 2 multicast
    - only group of stations will receive multicast packet
    - mapping of layer 3 IP multicast address to layer 2 multicast address (hardware address) is necessary
    - multicast receivers of a group must be programmed to listen to this hardware address

Class D ⇨ LAN Multicasts Address

- **class D address - 28 bit**
- **IANA owns an Ethernet address block (vendor code)**
  - 0x00-00-5E (IEEE notation, Ethernet)
  - 0x01-00-5E … Ethernet multicast (I/G bit = 1)
  - 0x00-00-7A (Token Ring)
  - 0x10-00-7A … Token Ring Multicast (I/G bit = 1)
- **low-order 23 bits of class D are mapped**
  - to low-order 23 bits of the Ethernet multicast address 0x01-00-5E-00-00-00
  - done accordingly for other LANs
- **specified in RFC 1112**

Class D ⇨ LAN Multicasts Address

- **Example of**
  - Mapping Class D Address to Ethernet Multicast Address
    - Static mapping, no ARP
    - IANA has reserved range 0x 01005E000000- 0x 01005E7FFFFF

Class D Address: 226.139.9.7

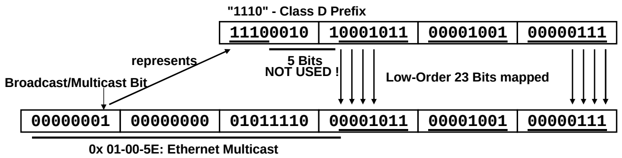

Ethernet Multicast Address: 0x 01-00-5E-0B-09-07

Class D ⇨ LAN Multicasts Address

- **5 bits cannot be mapped**
- **32:1 address ambiguity**
  - 32 different IP-Multicast-Groups have the same multicast MAC-Address
  - Filtering is needed by taking IP-Address into account
- **In IPv6 similar: (RFC 2464)**

IPv6-Multicast Address:0x FF02:0:1:FF68:12CB

Low-Order 32 bits are mapped

MAC-Address: 33-33-FF-68-12-CB

Institute of Computer Technology - Vienna University of Technology

L71 - IP Multicasting

IP Multicast Implementation

- **RFC 1112 "host extensions for IP multicasting"**
  - defines components of multicast architecture and tasks of components
    - model of host IP implementation
      - IP service interface
      - IP module
      - local network service interface
      - local network module
    - tasks of multicast sender
    - tasks of multicast receiver
    - tasks of multicast router
    - host-to-router protocol
      - Internet Group Management Protocol (IGMP)

RFC 1112 Highlights

- **IP multicasting is transmission of an IP datagram to a host group**
  - same best effort reliability as regular unicast IP datagram's
- **each group identified by a single class D address**
- **groups may be of any size, members could be located anywhere in the internet**
- **membership of a host group is dynamic**
  - hosts may join or leave groups at any time
- **a group could be permanent or transient**
  - permanent uses reserved "well known" IP addresses
  - transient may use all other IP addresses which are not reserved
- **multicast senders need not be member of groups**

RFC 1112 Highlights (cont.)

- **a host transmits an IP multicast datagram**
  - as a local network multicast, which reaches all immediately neighboring members of a destination group
  - default TTL=1 (higher value must be requested by application)
  - multicast datagram's which should reach remote members must be sent with a TTL higher than 1
  - TTL could be used to control the scope of multicast traffic
- **a host wanting to receive multicast traffic**
  - must join a group by preparing its local network interface for receiving of corresponding local network multicasts
  - must tell its membership to local multicast router with IGMP messages

RFC 1112 Highlights (cont.)

- **forwarding of IP multicast datagram's is handled by multicast routers**
- **multicast routers will forward all received IP multicast datagram's with TTL value higher than 1**
  - towards all other networks that have members of the destination group
- **multicast routers attached to member networks that are reachable within TTL**
  - complete delivery by transmitting the IP multicast datagram as local network multicast
- **method of multicast routing**
  - defined in other RFCs (drafts)

Institute of Computer Technology - Vienna University of Technology

L71 - IP Multicasting

Agenda

- **IP Multicasting**
  - RFC 1112
  - IGMPv1, IGMPv2, IGMPv3
  - IGMP Snooping
- **IP Multicast Routing**
  - DVMRP
  - PIM-DM
  - PIM-SM
  - MOSPF
- **MBone**
- **Multicast Applications**
- **RTP/RTCP**

IGMPv1 (RFC 1112)

- **Internet Group Management Protocol (IGMP)**
  - protocol by which host reports its group membership to any immediately-neighboring multicast router
  - version (1) defined in RFC 1112
    - only host to router aspects
    - extension of IGMP for router-router communication in other documents (RFCs and drafts)
  - operates over broadcast LANs and point-to-point
  - IGMP for IP multicasts is an integral part of IP as ICMP for IP unicasts
  - IGMP messages are encapsulated in IP datagram's
    - IP protocol number = 2
    - multicast addresses are used as destination address
      - 224.0.0.1 ...all systems on this subnet

© 2009, D.I. Manfred Lindner

IGMP Format

0	4	8	16	31
Version	Type	reserved	Checksum	
Group Address				

- Version = 1
- Type = 1 Membership Query
  = 2 Membership Report
- Checksum … standard IP-style checksum
- Group Address … group being reported
  (all zero´s in Query messages)

IGMP Procedures

- **on each broadcast network**
  - one multicast router is elected the "querier"
  - election is outside the scope of RFC 1112
- **elected multicast router**
  - sends periodically IGMP Membership Query messages using multicast IP address 224.0.0.1 with TTL = 1 (Polling)
  - receives IGMP Membership Report messages in response to IGMP Membership Query
- **IGMP Report messages are sent by multicast hosts**
  - to refresh routers knowledge about group memberships present on local attached networks
  - with a destination address equal the group to be reported
- **multicast router must listen therefore to all multicast addresses**

IGMP Procedures (cont.)

- **a host wanting to receive multicast traffic for a certain group**
  - must prepare first its local network interface for receiving of corresponding local network multicasts
    - 224.0.0.1
    - certain class D for group (224.x.x.x - 239.255.255.255)
- **on receipt of IGMP Membership Query a host**
  - starts a report delay timer for each of its group memberships with a randomly chosen timeout
- **when report delay timer for group X expires**
  - IGMP Membership Report message is sent by the host with TTL= 1
  - destination address = group address to be reported

IGMP Procedures (cont.)

- **other members of group X**
  - will hear this message and hence will stop their timers
  - "implosion" of concurrent IGMP Reports is avoided
- **in normal case only one IGMP Report message per group will be sent in response to a query**
- **if host wants to join a new group**
  - one or two IGMP Membership Report are transmitted immediately instead of waiting for a IGMP Membership Query
- **IGMP timers**
  - Query Timer = typically 60 - 90 sec (min. 60 sec)
  - Report Delay Timer = 0 - 10 sec (max. 10 sec)

IGMPv1 Query-Response Process I

- **Allows router to determine which multicast groups are active**

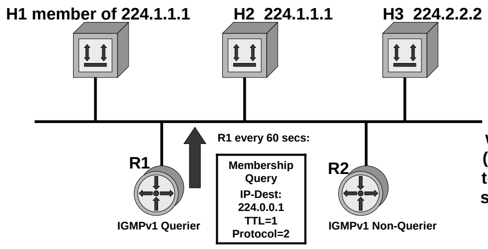

Hosts already listen to the multicast groups, they want traffic for (and also listen to 224.0.0.1 - all systems on this subnet)

IGMPv1 Query-Response Process II

- To avoid "Answer-burst" with identical information
  - Each Host sets timer randomly between 0 and 10 after reception
  - Host answers, when timer is 0 and no one else has answered before
  - This Report suppression mechanism helps reduce traffic

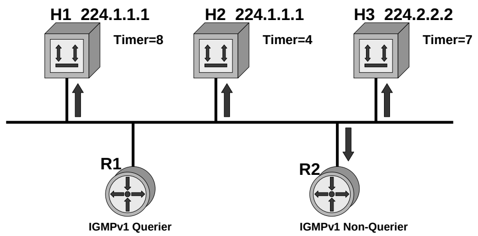

IGMPv1 Query-Response Process III

- 4 Seconds later:

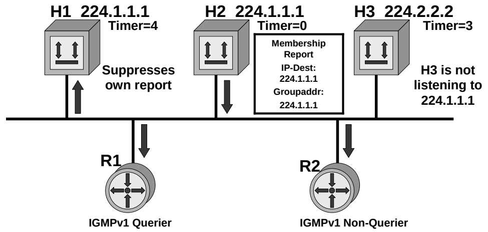

IGMPv1 Query-Response Process III

- Again 3 Seconds later:

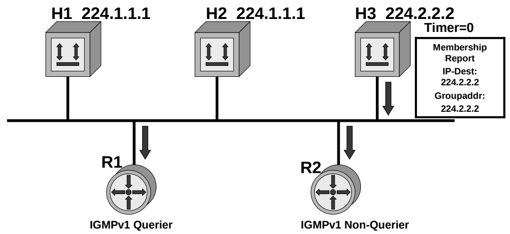

IGMPv1 Query-Response Process IV

– Final picture:

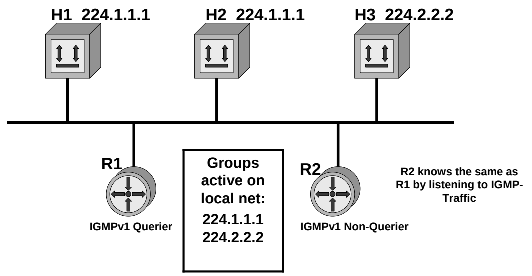

Note: In IGMPv1 R1 does not know number of listeners. Only groups, which are active

IGMPv1 Join Message

– If Host initially joins multicast group it sends an Unsolicited Report:

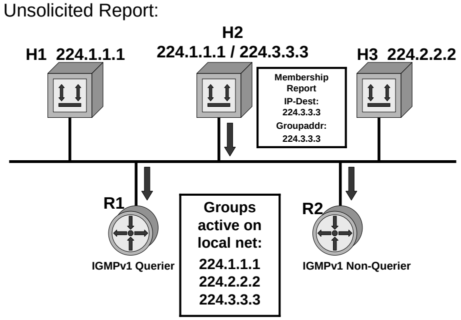

Institute of Computer Technology - Vienna University of Technology

L71 - IP Multicasting

IGMPv1 Leave Process

- **Router stops flooding Multicast traffic for a specific group into subnet, if no membership reports for this group are received any more**
- **Typical timeout interval is 3 times the Query Interval**
  - ~3 minutes
- **Long Leave Latency**
- **Can cause troubles**
  - Imagine high-bandwidth video application
  - User is "Channel-surfing"
  - Every time user leaves a group, traffic for this group is being delivered into subnet for another 3 minutes
- **Remedy: IGMPv2**

IGMPv2

- **RFC 2236**
- **backward compatible with version 1**
- **changes from version 1**
  - new messages and procedures to reduce latency for leaving a group
    - Leave Group Message
      - is sent by host to all routers address (224.0.0.2)
    - Group Specific Query
      - to detect if last member of a group has left
      - lower max. response timeout
  - standard querier election method
    - lowest IP address wins

IGMPv2 Format

0	4	8	16	31
Type		Max Response Time	Checksum	
Group Address				

- Type 0x11 = Membership Query
  - General Query (Group Address = all zero´s)
  - Group Specific Query (Group Address = specified Group)

  0x16 = v2-Membership Report

  0x17 = Leave Group

  0x12 = v1-Membership Report
- Max. Response Time
  - used only in Membership Query to specify maximum allowed timeout for report delay timer
  - in 1/10 second units

IGMPv3

- **RFC 3376**
- **adds "source filtering"**
  - ability to report interest in receiving packets only from a specific source addresses or from all but specific source addresses sent to a multicast address
  - reduces danger of Denial of Service Attack
    - ongoing multicast-session may be disturbed by someone sending junk-data to same multicast group
    - this source can be "turned off"
- **IGMP membership reports are sent**
  - to the well-known multicast address 224.0.0.22
    - eases the job of IGMP snooping for L2 switches because with IGMPv2 membership reports of listeners are sent to the multicast address of the wanted group address

Institute of Computer Technology - Vienna University of Technology

L71 - IP Multicasting

IGMPv3 - Source Filtering

- **Host A wants to receive from S=1.1.1.1 but not from S=2.2.2.2**

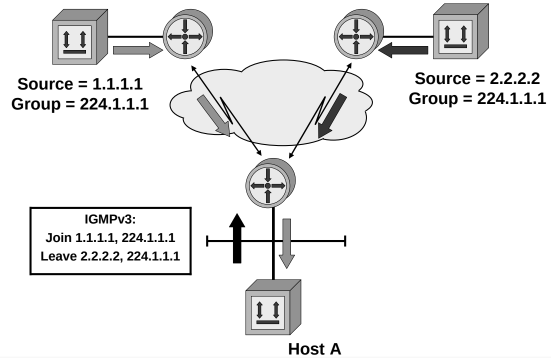

IGMP and IPv6

- **IGMP functions are integrated directly into IPv6 (ICMPv6)**
- **All IPv6 hosts are required to support multicasting**
- **In IPv4, multicasting and IGMP support is optional**

Agenda

- **IP Multicasting**
  - RFC 1112
  - IGMPv1, IGMPv2, IGMPv3
  - IGMP Snooping
- **IP Multicast Routing**
  - DVMRP
  - PIM-DM
  - PIM-SM
  - MOSPF
- **MBone**
- **Multicast Applications**
- **RTP/RTCP**

IGMP Snooping

- **Problem: Layer 2 Flooding of Multicast Frames**
  - Typical L2 Switches treat Multicast Traffic as Unknown or Broadcast und must "flood" the frame to every port
    - especially a problem when flat hierarchy with many members of broadcast domain
  - Bandwidth problem
  - Only members of groups should receive multicast traffic
  - How could switch know?
  - Solutions:
    - manual configuration of static entries in MAC-Table of Switch
    - CGMP - Cisco Group Multicast Protocol
      - Communication between Switch and Router - vendor specific
    - IGMP Snooping

IGMP Snooping

- Switch becomes "IGMP" aware
  - **must examine content of IGMP messages to determine which ports want what traffic**
    - **IGMP Membership Reports**
    - **IGMP Leave Messages**
- Impact:
  - **must process ALL layer 2 Multicast packets**
    - **IGMP packets have Class-D destination and Multicast MAC addresses like any other multicast traffic**

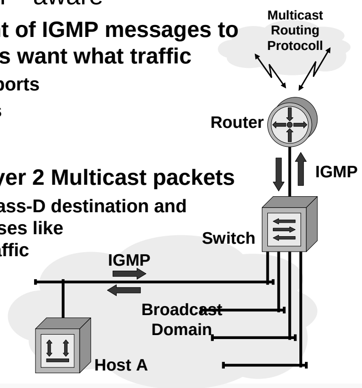

Agenda

- **IP Multicasting**
  - RFC 1112
  - IGMPv1, IGMPv2, IGMPv3
  - IGMP Snooping
- **IP Multicast Routing**
  - DVMRP
  - PIM-DM
  - PIM-SM
  - MOSPF
- **MBone**
- **Multicast Applications**
- **RTP/RTCP**

DVMRP

- **Distance Vector Multicast Routing Protocol**
  - first version
    - defined in RFC 1075 (experimental RFC)
  - actual version used today in the MBONE
    - defined in draft-ietf-idmr-dvrmp-v3-04.txt
- **DVMRP consists of two components**
  - conventional distance-vector routing protocol
    - to build a DVMRP routing table for all sources of multicast traffic
      - reverse path distances
    - using IGMP extensions for communication
  - truncated RPF with pruning and grafting
    - to forward multicast packets along truncated broadcast trees avoiding unnecessary branches of the tree

DVMRP

- **multicast routing protocol additional to unicast routing protocol is needed**
  - multicast routing process in parallel to unicast routing process
  - path from source S to multicast router is calculated and stored in multicast routing table using a very similar approach as RIP
    - On LANs neighbors discover each other by using DVMRP probe messages using the well-known DVMRP multicast address 224.0.0.4 as destination address
    - neighbors send each other route reports of the known networks and store this information in a separate DVMRP Route Table
    - on receipt of an advertised route with a hop-count greater than 32 (poison reverse), the router knows he is upstream for this network
    - like this, truncated broadcast trees to all hosts are created (truncated means that upstream router will send multicast only on an interface, if a poisened route was received before)

DVMRP- build RP routing table

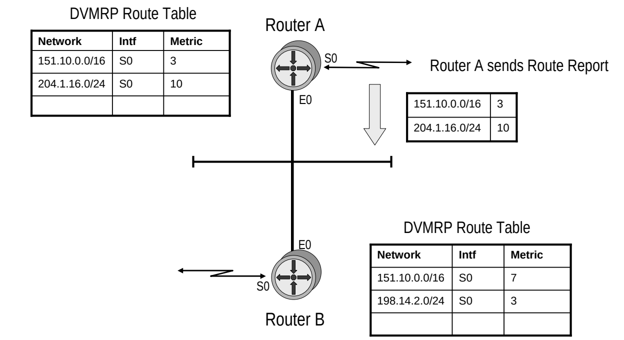

DVMRP- build RP routing table

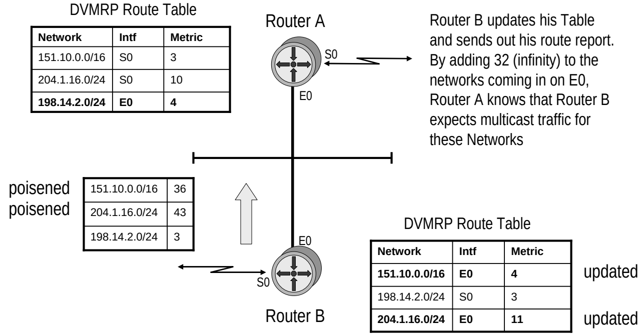

© 2009, D.I. Manfred Lindner

DVMRP

- **Inherent problem of first packet flooding and periodic flooding the whole network**
  - remember: flooding necessary to trigger pruning and to refresh states for any source/group relationship
- **also grafting previously pruned branches is possible**
- **DVMRP is not a protocol which scales well**
  - too much state information
  - max hop-count of 32
  - periodic flooding, route updates

Agenda

- **IP Multicasting**
  - RFC 1112
  - IGMPv1, IGMPv2, IGMPv3
  - IGMP Snooping
- **IP Multicast Routing**
  - DVMRP
  - PIM-DM
  - PIM-SM
  - MOSPF
- **MBone**
- **Multicast Applications**
- **RTP/RTCP**

PIM

- **Protocol Independent Multicast**
  - "IP routing protocol independent"
  - PIM-DM specified in draft-ietf-idmr-pim-dm-05.txt
  - PIM-SM
    - V1 defined in RFC 2117, V2 in RFC 2362
  - generic solution for multicast routing proposed by IDMR (Inter-Domain Multicast Routing) working group
  - emphasizes to do Internet-wide routing
    - e.g. tackle problems like explosion of states caused by large number of groups
  - two variants of PIM for different density of members
    - dense groups have large number of members in a large number of places (dense mode - DM)
    - spares groups have few members only in a small number of places (sparse mode - SM)

PIM-DM

- **PIM - Dense Mode**
  - RPF and prune method
    - flooding first packet, periodic flooding (3 minute cycle)
  - in principle very similar to DVMRP but no additional multicast routing protocol is used
    - relies on presence of any normal unicast routing protocol
    - unicast routing protocol will calculate path from router to destination networks (sources of multicast traffic)
    - unlike DVMRP's truncated broadcast trees PIM-DM only builds up broadcast trees.
      - trees are built on the fly when the first packet is flooded
      - only the first receipt of a given datagram will be forwarded
      - on receipt of the same datagram on another interface propagation will be stopped
    - path metric of unicast routing table will be used for RPF algorithm
    - cannot handle asymmetric metric
    - on LANs neighbors discover each other by using PIM hello messages using the well-known all-PIM-routers multicast address 224.0.0.13 as destination address (hello-time 30 seconds, hold-time 90 seconds)

Example 1: Multicast Traffic

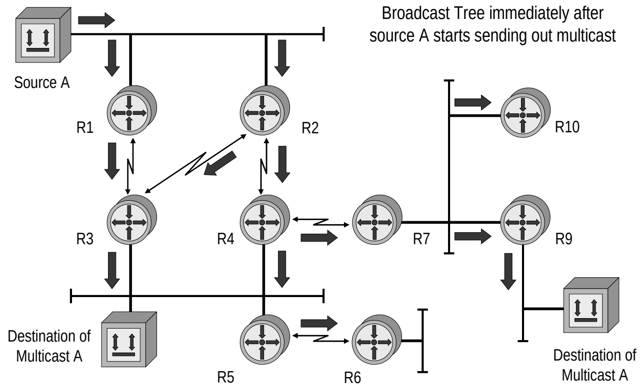

PIM-DM

- **PIM - Dense Mode (cont.)**
  - extremely simple protocol
    - memorizing (caching) of prune messages (source/group states)
    - in case of lacking memory resources pruning can be reduced to least recently used context
  - special procedures to solve pruning on broadcast networks and to handle equal-cost multi-paths
    - prune messages are sent to all-PIM-routers multicast address (224.0.0.13)
    - prunes are sent in the following cases
      - Traffic arriving on non-RPF point-to-point interfaces (example 1)
      - Leaf router and no directly connected receivers (example 2)
      - Router on a point-to-point link that has received a prune from its neighbor (example 3)
      - Router on a LAN segment that receives a prune from a neighbor on this segment and no one overrides the prune

Institute of Computer Technology - Vienna University of Technology

L71 - IP Multicasting

Example 1: pruning

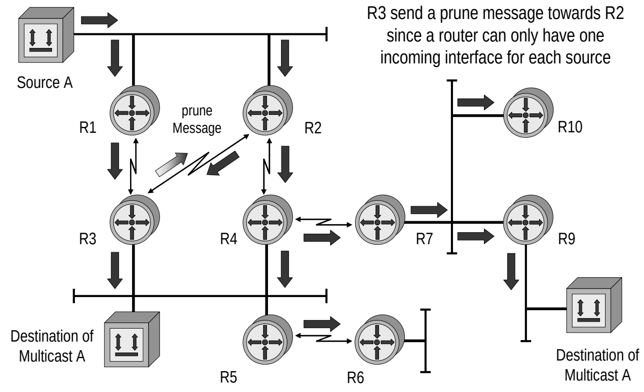

Example 2: pruning

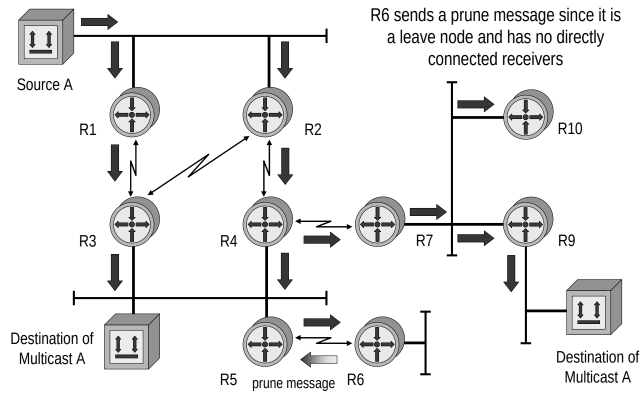

Example 3: pruning

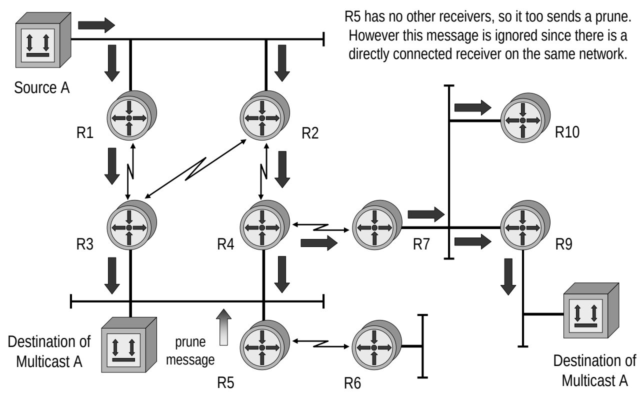

PIM-DM - Override

- **PIM - Dense Mode (cont.)**
  - In DVMRP a router remembers all its neighbors, so receiving a prune of one router on a shared LAN is no problem
    - Multicast traffic is sent out a LAN interface until all attached routers send a prune message
  - PIM-DM uses a special override algorithm instead
    - when an upstream router receives a prune message from a shared media segment, it starts a prune-delay timer (usually >3s)
    - since prune messages are sent to all-PIM routers every router on the segment receives the message
    - this gives other routers with active receivers the possibility to send a join message (244.0.0.13) to override the prune
    - this can lead to a delay accumulation

Example 4: Override

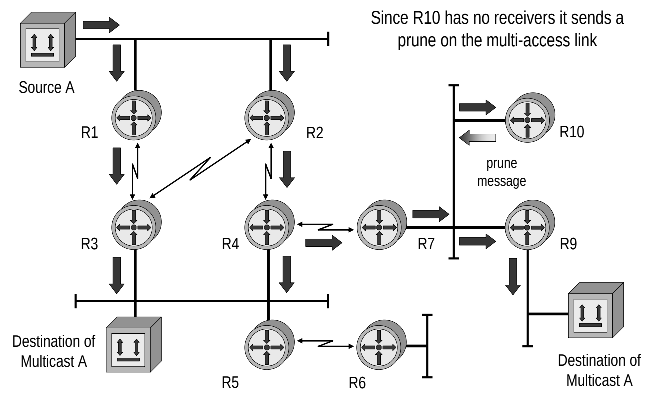

Example 5: Override

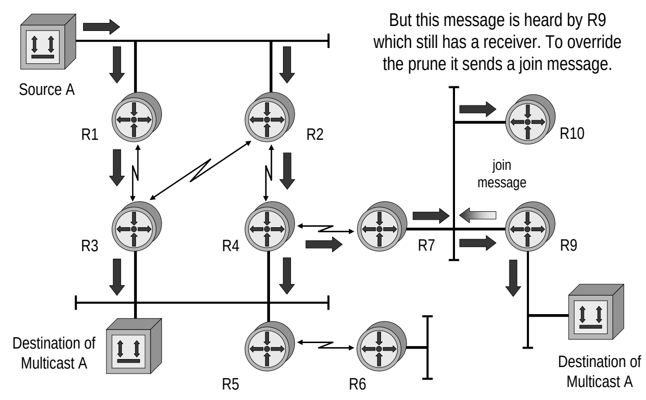

PIM-DM - Asserts

- **PIM - Dense Mode (cont.)**
  - unlike DVRMP, which builds up truncated broadcast trees, PIM-DM has to handle redundant paths
  - it does this with assert messages
    - if a router gets a multicast packet from an outgoing interface, it sends out an assert message (224.0.0.13) to resolve which router will continue
    - the message includes the routers administrative distance and the metric towards the source
    - if the values are the same the higher IP address wins

Example 6: Assert

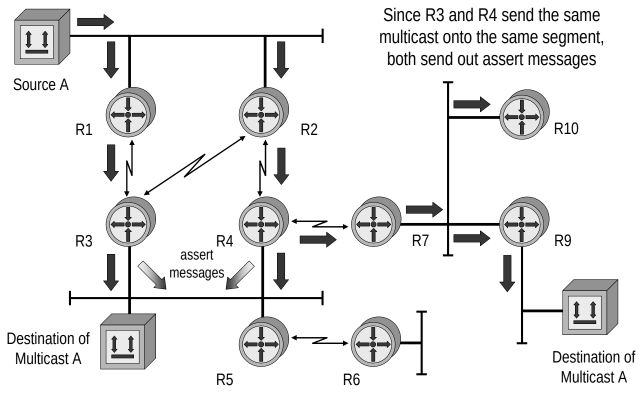

Example 7: Assert

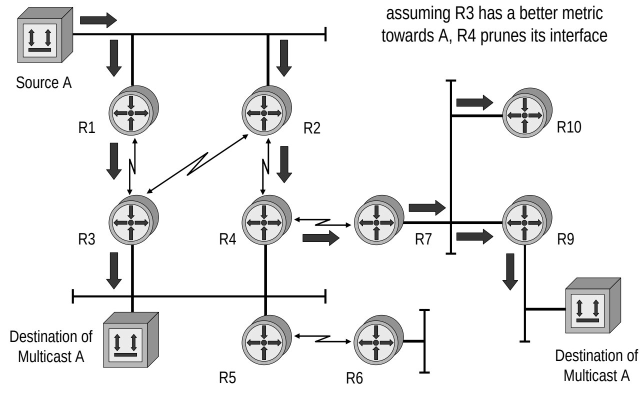

PIM-DM – State-Refresh

- **PIM - Dense Mode (cont.)**
  - Every pruned state has a timeout of about 3 minutes
  - This means that every 3 minutes multicast packets are flooded throughout the whole tree
  - In order to tackle this problem an interesting proposal was issued by the IETF called state-refresh
    - if a source is still sending traffic, a state-refresh message is sent down the whole broadcast tree resetting the timer of the pruned states
    - by including the same information as in assert messages, those states are refreshed as well

Example 8: State-Refresh

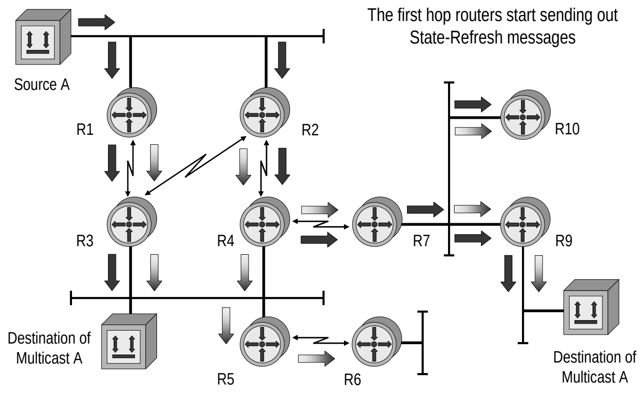

PIM-DM: Summery

- **PIM-DM is best used in high-speed networks, where periodic flooding is not a big issue**
  - If the unicast network is well structured PIM-DM will scale much better then DVRMP
    - especially if state-refresh becomes a fully supported feature
  - However the amount of routing states can become an issue in a network with many active source/groups

PIM-DM - Messages

- Hello
  - PIM neighbor discovery
  - default hello interval: 30 seconds
  - default hold time: 3,5 times hello interval
  - not used for designated router on LAN by PIM-DM
- Join/Prune
  - tree pruning after flooding first multicast datagram
  - "implicit joining" but explicit pruning
  - prune soft state times out after 3 minutes meaning that flooding will occur after timeout
- Assert
  - clarifying who is the forwarder on a LAN
- Graft / Graft-ACK
  - tree establishment of a pruned link
- Bootstrap (not used by PIM-DM)
- Candidate-RP-Advertisement (not used by PIM-DM)
- Register (not used by PIM-DM)
- Register-STOP (not used by PIM-DM)

Agenda

- **Introduction Multicasting**
- **Multicast Routing Principles**
- **IP Multicasting**
- **IP Multicast Routing**
  - DVMRP
  - PIM-DM
  - PIM-SM
  - MOSPF
- **MBone**
- **Multicast Applications**
- **RTP/RTCP**

PIM-SM

- **PIM - Spares Mode**
  - variant of shortest path tree (SPT) method
    - packets are not flooded to the whole network
    - multicasts are forwarded only to members that have explicitly joined the group along the (*, G) tree
  - differences to basic SPT method
    - core is called rendezvous point (RP)
    - tree is not bidirectional
    - RP can set up a source based (S, G) tree towards the source to avoid encapsulation / decapsulation of multicast packets
      - remember: encapsulation is done by first-hop router of a given source
    - routers can build source base trees towards the sending source
      - hence the name RP; the groups only meet temporarily at the RP before they build their own trees

PIM-SM

- **Joining works very similar as SPT**
  - A host sends a (*,G) join message to his first-hop router
    - if a router already has a state for this group, it simply adds the new interface to its (*,G) state entry
    - if this router has no other receivers it creates a new (*,G) entry and sends a join message upstream which is handled the same way
    - using the same messages a router can also join an (S,G) tree by exchanging the address of the RP with the address of the source
  - Unlike other sparse-mode protocols, PIM only creates soft states
    - therefore states have to be refreshed before they time out
    - routers have to send join messages every minute to the upstream router as long as the have state entries
      - with one join message a whole list of groups can be refreshed

PIM-SM

- **Source registration**
  - In PIM-SM it is possible for receivers to join a group, without a sending source and for a source to start sending multicast without having active receivers
  - to handle this special register messages are used
    - these messages notify the RP about an active source
    - and deliver the first multicast packets
  - When a source starts sending traffic its first-hop router encapsulates the packets into a register message which is unicast towards the RP
    - the RP decapsulates the packets to forward them to the group
    - and sends an (S,G) join message towards the source, hence normal multicast traffic will eventually arrive at the RP

PIM-SM

- **Source registration (cont.)**
  - Once the first packets arrive the RP will send a register stop message towards the source, causing the first-hop router to stop its register messages
    - The RP will send a register-stop message right away, if there are no active receivers in the group

PIM-SM: Registration

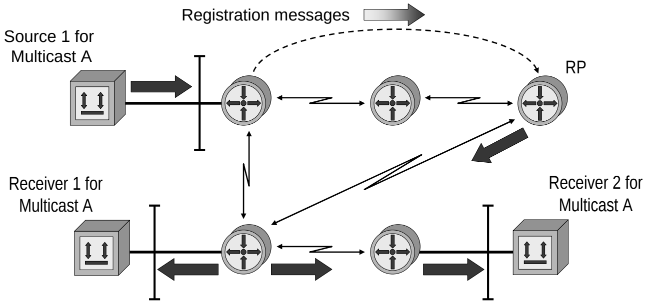

Source 1 starts sending multicast. The first-hop router encapsulates the packets and unicasts them towards RP, who then decapsulates the packets and forwards them along the (*,G) tree.

PIM-SM: Registration

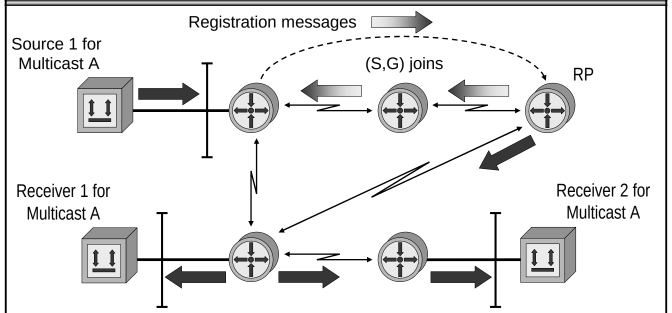

RP sends an (S,G) join towards the source, in order to build up an SPT from the source to the RP

PIM-SM: Registration

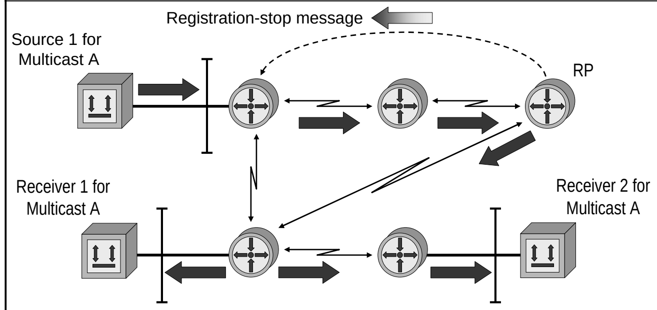

The first-hop router can now send multicast down the SPT towards RP, who in return sends a register-stop message to stop the registration messages

PIM-SM: SPT Switchover

- **First-hop multicast router of a multicast receiver can force building of source based tree (S, G)**
  - new join message on shortest path to source and prune to RP
  - this feature is triggered if a specified SPT-threshold in terms of bandwidth is exceeded
    - on CISCO routers by default 0
  - also any intermediate router can force building of a source-based distribution tree
    - e.g. if too much traffic from a given source coming along a suboptimal path is recognized
    - remember: multicast router are able to calculate shortest path to the RP and hence to any other network address

PIM-SM: SPT Switchover

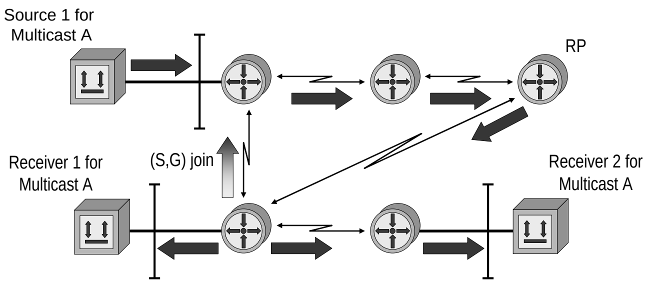

The first-hop router from receiver 1 sends out an (S,G) join towards the source

PIM-SM: SPT Switchover

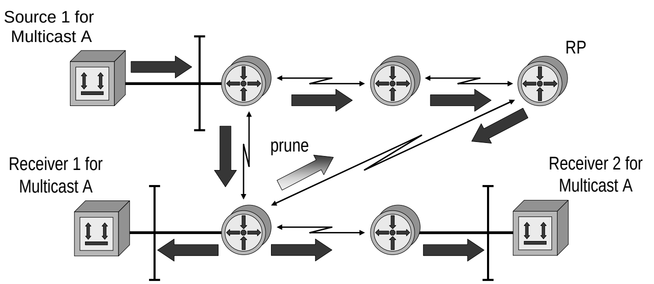

The join message will flow up to the source, building a SPT on the way.
Now there are two flows reaching the same router. A special prune message is triggered which is directed towards the RP

PIM-SM: SPT Switchover

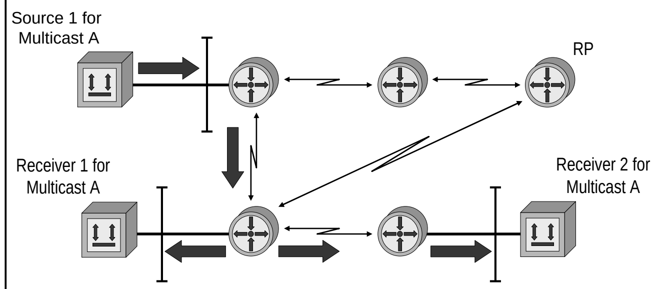

Since in our example RP has no other interfaces, it too will send a prune towards the source.
Resulting in the final picture shown above

PIM – SM: Summary

- **PIM-SM is better suited for multicast networks that have potential members at the end of WAN links.**
- **Being able to set the SPT-threshold helps engineers to configure there domain**
  - It helps to control the amount of state a multicast creates inside a network
    - e.g. high threshold results in a star-like topology
    - setting it very low reduces the importance of the RP to the very beginning

PIM-SM - Messages

- Hello
  - PIM neighbor discovery
  - default hello interval: 30 seconds
  - default hold time: 3,5 times hello interval
  - clarifying who is designated router on a LAN
- Join/Prune
  - explicit join to establish hard states
  - message interval is 60 seconds
  - hold-time is 180 seconds
- Assert
  - clarifying who is the forwarder on a LAN
- Graft / Graft-ACK (not used in PIM-SM)
- Bootstrap
- Candidate-RP-Advertisement
- Register
- Register-STOP

Agenda

- **IP Multicasting**
  - RFC 1112
  - IGMPv1, IGMPv2, IGMPv3
  - IGMP Snooping
- **IP Multicast Routing**
  - DVMRP
  - PIM-DM, PIM-SM
  - MOSPF
- **MBone**
- **Multicast Applications**
- **RTP/RTCP**

MOSPF

- **Multicast extensions to OSPF**
  - defined in RFC 1584
  - RFC 1585: MOSPF Analysis and Experience
  - provides multicast routing within in an AS of limited size
  - emphasizes on efficient route computation
- **MOSPF**
  - network map is complemented with new type of link state record ⇨ group membership LSA
    - created by multicast routers responsible for a subnet
    - summarized by area-border routers
    - special considerations for equal-cost paths
    - together with router and network LSAs a SPT for each source-**network**/group is calculated
      - Results in a (S/m, G) tree (source/subnet-mask, group)

MOSPF

- **MOSPF (cont.)**
  - RPF and prune computations
    - can be done locally in memory without flooding first packet or periodic flooding
  - no additional multicast routing protocol needed
    - will rely on normal unicast routing protocol OSPF
    - assumes symmetrical metric
  - problem of computation of shortest path for every source
    - could saturate even most powerful CPU´s of multicast routers
  - therefore computation on demand
    - when first packet of a given source and group arrives
      - but for a sufficient amount of active sources even this could lead to an over-saturation of CPU bandwidth, especially if receivers turn themselves on and off frequently

MOSPF – Interarea Routing

- **MOSPF Interarea Routing**
  - In order for MOSPF to work in between two areas the ABR's must be capable of multicast
    - MABR
  - An MABR will summarize multicast information and send these LSA's into area 0
    - note: area 0 will not forward information the other direction
  - this method is sufficient when sources are in area 0
  - to solve this dilemma MABR's will include a wildcard bit inside their router LSA's directed towards a non-backbone area
    - this specifies the MABR as a member for every multicast group
    - they will pass the information into area 0, acting as a source
  - This can lead to unwanted traffic, if there are no "real" sources

Agenda

- **IP Multicasting**
  - RFC 1112
  - IGMPv1, IGMPv2, IGMPv3
  - IGMP Snooping
- **IP Multicast Routing**
  - DVMRP
  - PIM-DM, PIM-SM
  - MOSPF
- **<u>MBone</u>**
- **Multicast Applications**
- **RTP/RTCP**

History of the MBONE

- **DARTNet**
  - U.S. government formed the DARPA Testbed Network
    - established in the early 1990s
    - playground for researchers to test and evaluate new tools and technologies without affecting the production Internet
    - initially composed of T1 lines connecting various sites
    - sites used Sun SPARCstations running mrouted as the DVMRP multicast routing daemon
    - hence DARTNet had native IP Multicast support between all sites
    - first uses provided audio multicasting of IETF meetings

The DARPA Test-bed Network (DARTNet)

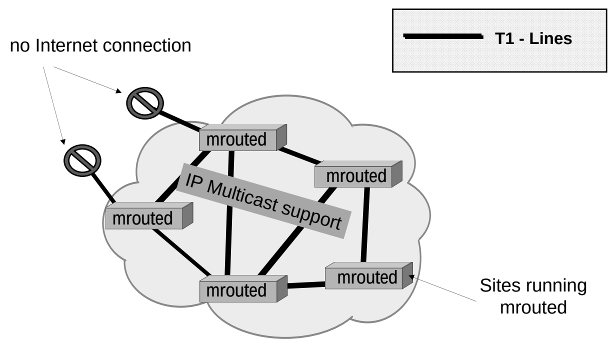

History of the MBONE (cont.)

- **DARTNet cont. / MBone**
  - at that time, 20 sites were connected to the backbone
  - two years later, simultaneous audio and video transmissions were distributed to more than 500 participants located in 15 countries
  - participants were not directly connected to the DARTNet backbone
  - to get the multicast traffic back into the DARTNet backbone DVMRP tunnels had to be built
  - these tunnel sites, coupled with DARTNet – the initial multicast core network – were soon called MBone (Multicast Backbone)
  - today, the MBone is still operational, but multicast connectivity is natively included in many Internet routers
  - efforts continue to integrate the MBONE directly into the Internet infrastructure

DVMRP - Tunneling

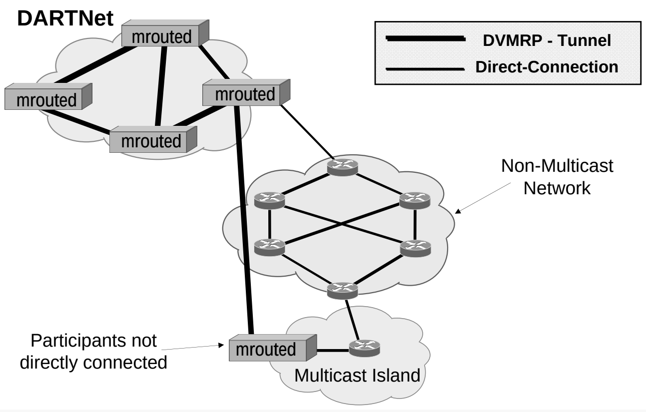

Institute of Computer Technology - Vienna University of Technology

L71 - IP Multicasting

DVMRP - Tunneling ⇨ MBone

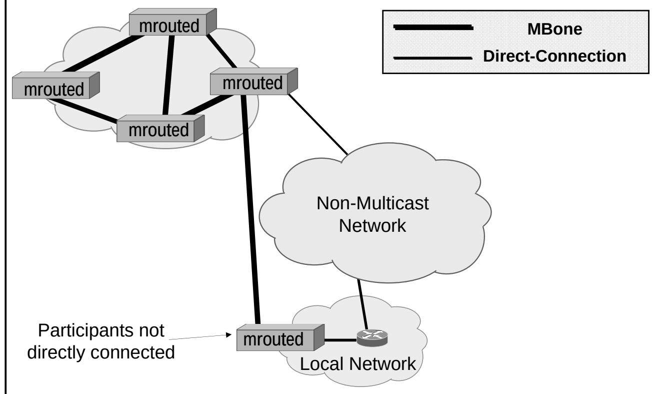

MBONE Routing

- **The Multicast Backbone today**
  - structure
    - interconnected set of multicast routers
    - virtual network overlaid on the Internet
    - overlay network using its own multicast routing protocols
  - multicast island
    - consists of multicast hosts and multicast routers
    - provides multicast service for testing multimedia applications
  - multicast islands are connected via tunnels
    - multicast routers encapsulate the packets in unicast packets
    - encapsulated packets are transmitted through the standard Internet routers
    - destination address contained in the unicast packets is the endpoint of the tunnel

MBONE Tunnel

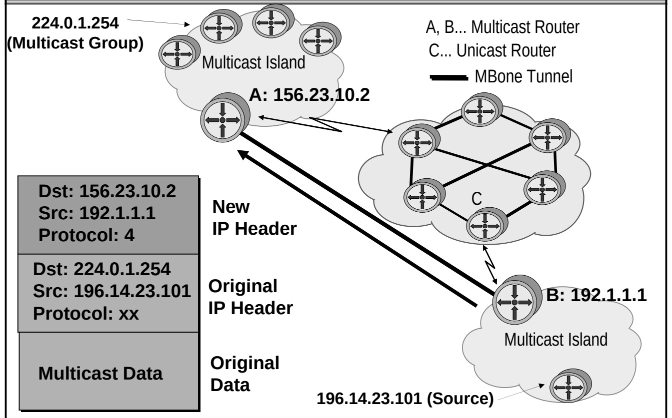

MBONE Routing (cont.)

- **MBone tunnels**
  - router at the remote end of the tunnel removes the encapsulation header and forwards the multicast packets to the receiving devices
  - tunnels have associated metric and threshold parameters
  - metric parameter is used as a cost in the multicast routing algorithm
  - routing algorithm uses this value to select the best path through the network (see next slide)
  - multicast packet sent from router A to router B should not use the tunnel directly connecting router A and router B
  - because cost of the alternate path using router D and router C is 5 (1 + 2 + 2)
  - threshold parameters also limit the distribution of multicast packets
  - specify a minimum TTL for a multicast packet forwarded into an established tunnel (TTL decremented by 1 at each router)

MBONE Tunnel Metric

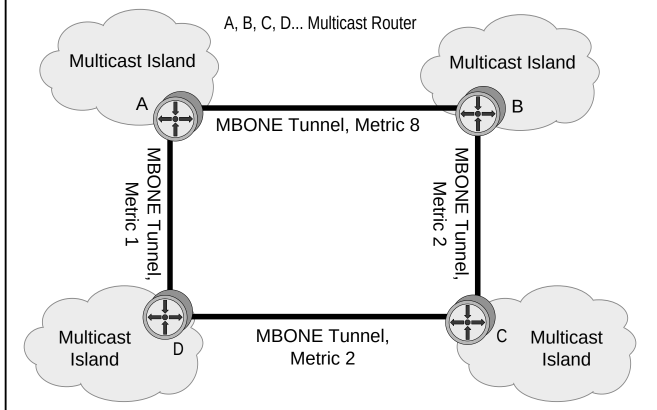

MBONE Future

- **Most Internet routers will provide direct support for IP multicast**
  - will eliminate the need for multicast tunnels
  - current MBone implementation is only a temporary solution
  - will become obsolete when multicasting is fully supported in every Internet router

Agenda

- **IP Multicasting**
  - RFC 1112
  - IGMPv1, IGMPv2, IGMPv3
  - IGMP Snooping
- **IP Multicast Routing**
  - DVMRP
  - PIM-DM
  - PIM-SM
  - MOSPF
- **MBone**
- **<u>Multicast Applications</u>**
- **RTP/RTCP**

IP Multicast Applications

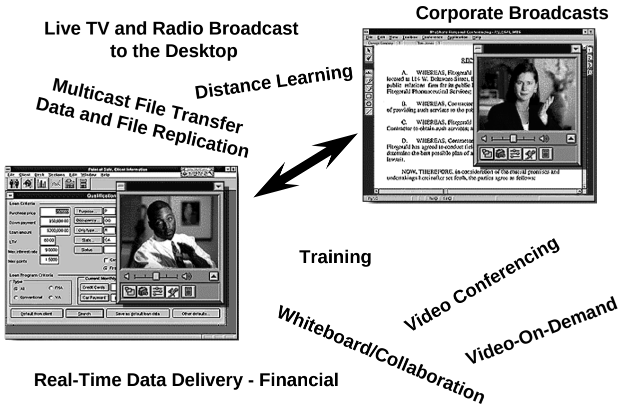

Session Information

- **SAP/SDP**
  - before joining a multimedia session, information about multicast address and port of this session is needed
  - furthermore time when session is going to be active and kind of application should be known
    - Session Announcement Protocol SAP is a protocol for advertising multicast conferences and sessions
    - SAP clients announce their conference sessions periodically by multicasting SAP packets containing session information
    - to a appropriate well-known multicast address and port
    - Session Description Protocol SDP is used to encode the actual session information
    - information may be encrypted optionally to avoid being read by unauthorized parties

Multicast Applications

- **Can be divided into 4 categories**
  - Multimedia Conferencing
  - Data Distribution
  - Gaming and Simulations
  - Real-time data multicasts

Multimedia Conferencing 1

- Multimedia Conferencing
  - video is only one of many IP multicast applications
  - some tools were developed for the use over the MBone
  - most widely used multimedia conferencing applications are the MBone freeware applications
  - SDR (Session Directory Tool)
  - SDR is listening for incoming SAP packets
  - if there are active multicast sessions on the network the application encodes session information using SDP
  - and shows the results on the screen
  - now the specific application starts using this session information
  - the multimedia conference can begin

Example of MBone Applications (SDR)

Multicast Sessions

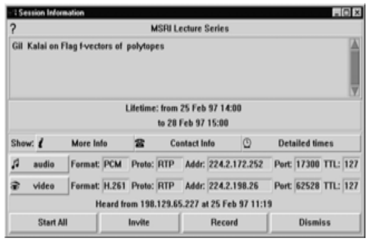

application can start using session information

MBone tools can be found at: http://www-mice.cs.ucl.ac.uk/multimedia/software

Institute of Computer Technology - Vienna University of Technology

L71 - IP Multicasting

Multimedia Conferencing 2

- Multimedia Conferencing (cont.)
  - VAT (MBone Multimedia Audio Tool)
  - VIC (MBone Multimedia Video Tool)
  - many-to-many audio-only or audio-video communication
  - WB (Shared Whiteboard Tool)
  - electronic whiteboard that members of the multicast session may share
  - in contrast to audio/video it uses a reliable multicast protocol
  - otherwise participants may not see the same things on their whiteboards
  - amount of time to retransmit lost audio/video packets would be too long

Example of MBone Applications (VAT, VIC)

VAT:

VIC:

MBone tools can be found at: http://www-mice.cs.ucl.ac.uk/multimedia/software

Example of MBone Applications (wb)

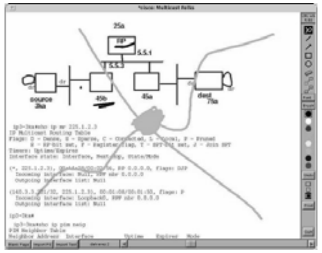

MBone tools can be found at: http://www-mice.cs.ucl.ac.uk/multimedia/software

Multimedia Conferencing 3

– Multimedia Conferencing (cont.)
- but there are also some commercial tools
- after the novelty of video wears off and high bandwith consumption become apparent
- audio-only conference coupled with a whiteboard application will be an extremely powerful form of multimedia conferencing which doesn`t consume much bandwith

Institute of Computer Technology - Vienna University of Technology

L71 - IP Multicasting

Data Distribution

– Data Distribution
- data replication is another multicast area becoming very popular
- MFTP (multicast form of ftp, “NAK only” protocol)
- one ore more files may be sent simultaneously with ftp to a group of nodes in the network by using IP multicasting
- for example, a central site can efficiently push updated data files to each district office

Gaming and Simulations

– Gaming and Simulations
- applications are already available
- integration of multicast services allow the applications to scale to a large number of users
- multicast groups can represent different sections of the game or simulation
- as users move from one section to the next, they exit and join different multicast groups

Real-time Data Multicasts

- Real-time data multicasts
  - applications distribute real-time data to a large number of users
  - for example, stock ticker information can be provided to sets of workstations on the trading floor
  - special applications and protocols are needed because of time-critical information

Multicast Applications

- **TCP?**
  - all of these applications must run on top of UDP or interface directly to IP providing their own transport layer
    - TCP is a unicast (point-point) only transport protocol
    - with TCP reliability and flow control mechanisms have not been optimized for real-time broadcasting of multimedia data
    - the potential to lose a small percentage of packets is preferred to the transmission delays introduced with TCP
    - hence multimedia streaming applications need a specialized transport layer
    - such as the Real-Time Transport Protocol RTP which operates over UDP in the application layer with the application
    - see next chapter

Operation over UDP or IP

- Multicast applications must run on top of UDP (e.g. RTP; left picture) or interface directly to IP providing their own customized transport layer (right picture)

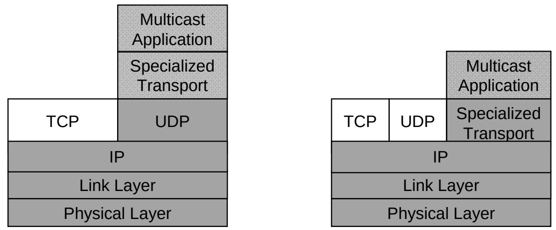

Real-time Applications based on RTP/RTPC

- **In addition to the mentioned MBone applications there are some others using RTP/RTCP**
  - Quick Time (Apple)
    - provides digital video and media streaming
  - Real Audio and Real Video (RealNetworks)
    - high quality audio and video streaming
  - NetMeeting (Microsoft)
    - provides IP telephony, white boarding, text chats and application and file sharing
  - CU-seeMe (CUseeMe Networks)
    - Internet video chat software supporting video, audio, text and whiteboard communications
  - IP/TV (Cisco Systems)
    - Live video, scheduled video, and video on demand

Agenda

- **IP Multicasting**
  - RFC 1112
  - IGMPv1, IGMPv2, IGMPv3
  - IGMP Snooping
- **IP Multicast Routing**
  - DVMRP
  - PIM-DM
  - PIM-SM
  - MOSPF
- **MBone**
- **Multicast Applications**
- **RTP/RTCP**

RTP and RTCP

- **To use real-time services in an application two protocols must be implemented**
  - Real-Time Transport Protocol (RTP) provides the transport of real-time data packets
    - unicast and multicast transmissions
    - to accommodate new real-time applications the architecture was intentionally left incomplete
    - allows the protocol to easily adapt to new audio and video standards
  - the RTP Control Protocol (RTCP) monitors the quality of service provided to existing RTP sessions

RTP and RTCP

- **RTP**
  - implements the transport features needed to provide synchronization of multimedia data streams
    - RTP may be used to mark the packets associated with the individual video and audio streams
    - allows the streams to be synchronized at the receiving host
    - next slide shows the operation of RTP in a multimedia transmission
    - audio and video data are encapsulated in RTP packets
    - if the multimedia application does not utilize RTP services, the receiver may not be able to associate the corresponding audio and video packets
    - because of congestion or the varying performance conditions within the network

RTP

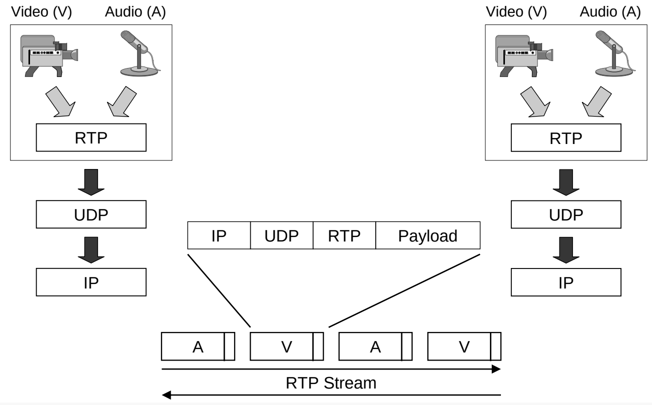

RTP and RTCP

- **RTP (cont.)**
  - because of performance issues of RTP many functions are not included
  - RTP protocol alone does not include any mechanism to provide guaranteed delivery or other quality of service functions
  - standard does not prevent out of sequence packet delivery nor does it assume that the underlying network is reliable and delivers packets in sequence
  - it also does not prevent the occurrence of network congestion
  - designers of applications must determine if these levels of service are acceptable

RTP Header Format

0 4 8 16 31

V	P	X	CSRC	M	Payload Type	Sequence Number
Timestamp						
Synchronization Source (SSRC) Identifier						
Contributing Source (CSRC) Identifiers						

- first 12 octets are required in every RTP packet
- V: Indicates the RTP version
- P: Contains the padding bit, used by encryption algorithms (bit is set)
- X: If this field is set a header extension follows the fixed header
- CSRC Count: This field contains the number of contributing source identifiers that follow the fixed header
- M: This field allows significant events to be marked in the packet stream (frame boundaries)

RTP Header Format (cont.)

0 4 8 16 31

V	P	X	CSRC	M	Payload Type	Sequence Number
Timestamp						
Synchronization Source (SSRC) Identifier						
Contributing Source (CSRC) Identifiers						

- SSRC identifier: All packets from the same source contain the same SSRC identifier. This enables the receiver to group packets for playback.
- CSRC identifiers: Contains a list of the sources for the payload in the current packet. This field is used when a mixer combines different streams of packets. (see later in this chapter)

RTP Header Format (cont.)

- **RTP protocol services**
  - RTP provides end to end transport services for applications transmitting real-time data
  - included in the RTP header
  - Payload type identification
    - an RTP packet can contain portions of either audio or video data streams
    - to differentiate between these streams, the sending application includes a payload type identifier within the RTP header
    - identifier indicates the specific encoding scheme used to create the payload
    - receiving application uses this identifier to determine the appropriate decoding algorithm

RTP Header Format (cont.)

- **RTP protocol services (cont.)**
  - Sequence numbering
    - sequence numbers are used by the receiving RTP host to restore the original packet order
    - the receiver is able to detect packet loss using the information in this field
  - Timestamping
    - time stamps are used in RTP to synchronize packets from different sources
    - timestamp represents the sampling (creation) time of the first octet in the RTP data packet
    - it is possible that several RTP packets may have the same time stamp
    - for example this can occur when a single video frame is transmitted in multiple RTP packets

RTP Timestamping

100…sequence number
JPEG…payload type

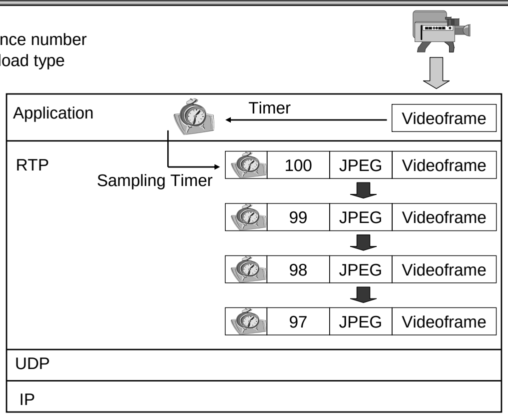

Institute of Computer Technology - Vienna University of Technology

L71 - IP Multicasting

RTP and RTCP

- **RTCP**
  - to manage real-time delivery many applications require feedback about the current performance of the network
    - primary function of RTCP is to provide feedback about the quality of RTP data distribution
    - RTCP is based on periodic transmission of control packets to all participants in a session
    - RTCP uses a UDP connection for communication
    - separate from any UDP connection used by the RTP protocol
  - RTCP architecture defines five types of control information used to report current performance

RTP and RTCP

- **Types of RTCP control information (cont.)**
  - Sender report:
    - sent out by the source of an RTP data stream (in intervalls)
    - provides the transmission and reception statistics observed by the sender
    - is sent as a multicast packet processed by all RTP session participants
  - Receiver report:
    - provides reception statistics for participants that are not active senders
    - is issued if the intervall times out and no data flows
  - Source description report:
    - used by an RTP sender to provide local capability information

RTP Translators and Mixers

- **RTP protocol supports the use of translators and mixers to modify the RTP packet stream**
  - these devices are used when some participants in a multimedia session need to receive data in different formats
- **RTP translators**
  - used to change the type of data in an RTP packet
  - in this example, three videoconferencing workstations are exchanging MPEG traffic over a high-speed LAN
  - each workstation is generating MPEG data (rate 1.5 Mbps)
  - another workstation connected via a lower-speed serial connection wishes to participate in the videoconference

RTP Translators and Mixers

- **RTP translators (cont.)**
  - bandwidth of this connection is not sufficient to support the video streams
  - one possible solution for this problem is changing all workstations to a video format, producing less traffic (e.g., H.261 with 256 Kbps)
  - but reducing data rate means reducing quality of video
  - an alternate solution uses RTP translation devices
  - each individual MPEG video stream is converted to an H.261 video stream with 256 Kbps which can be forwarded through the serial line
  - the three LAN attached workstations continue to use the higher quality MPEG format

Institute of Computer Technology - Vienna University of Technology

L71 - IP Multicasting

Videoconference without Translating

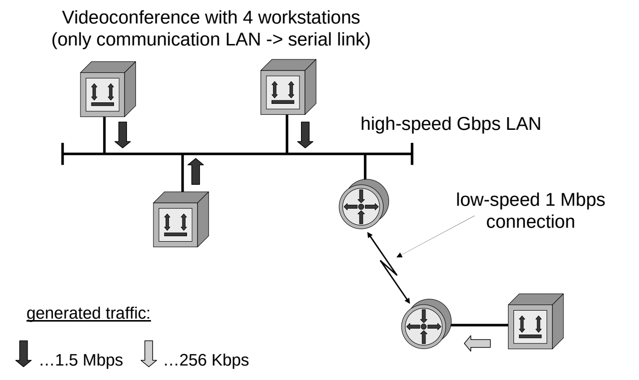

Videoconference without Translating

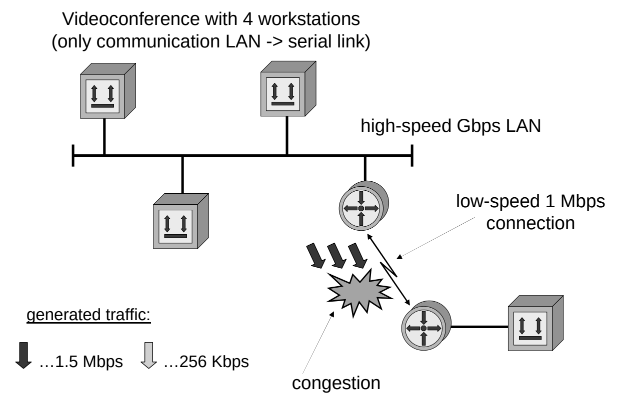

Videoconference with RTP Translating

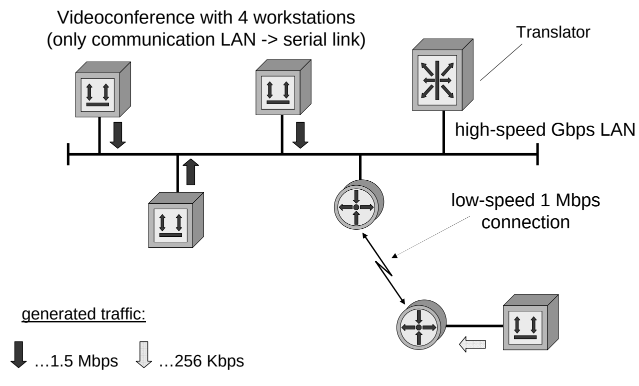

Videoconference with RTP Translating

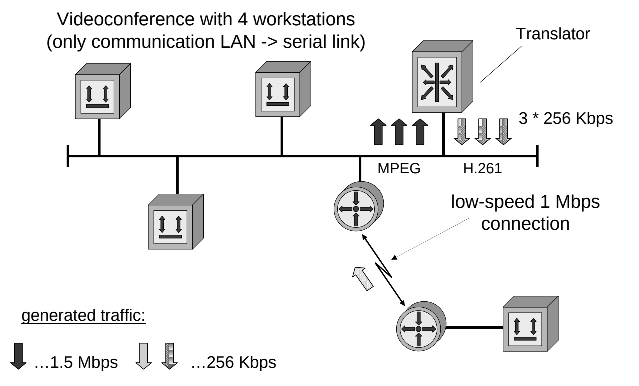

Videoconference with RTP Translating

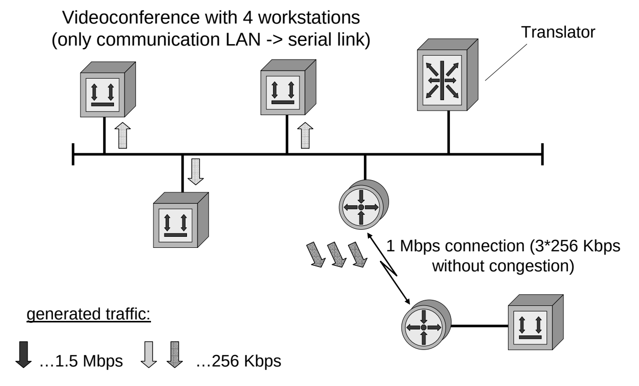

RTP Translators and Mixers

- **RTP translators (cont.)**
  - RTP translators are also used in case of firewalls which don`t pass multicast packets
  - two translators on each side of the firewall
  - one for secure tunneling the multicast packets
  - the second forwards information as multicast packets
- **RTP mixers**
  - RTP mixers are used to combine multiple data streams into a single RTP stream
  - these devices are used to support audio transmission applications where there are only one or two simultaneous speakers

RTP Translators and Mixers

- **RTP mixers (cont.)**
  - RTP mixing is not usable in video application environments
  - in this example, two audio conferencing workstations produce PCM audio streams at a rate of 64 Kbps
  - another workstation connected via a lower speed serial connection wishes to participate in the audio conference
  - the bandwidth of this connection is not sufficient to support the combined 192 Kbps
  - an RTP mixer merges the two sender streams into a single 64 Kbps stream
  - this allows the new station to join the conference

RTP Mixing with 2 speakers

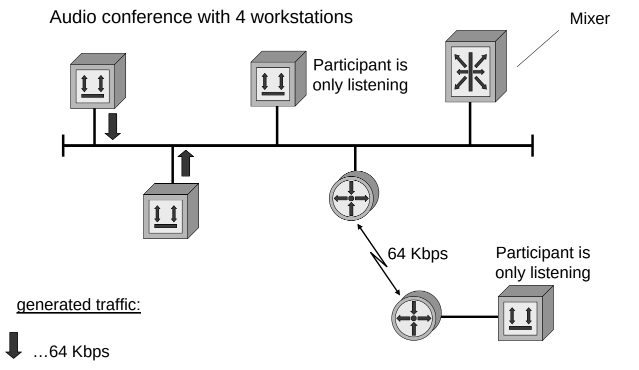

Institute of Computer Technology - Vienna University of Technology

L71 - IP Multicasting

RTP Mixing with 2 speakers

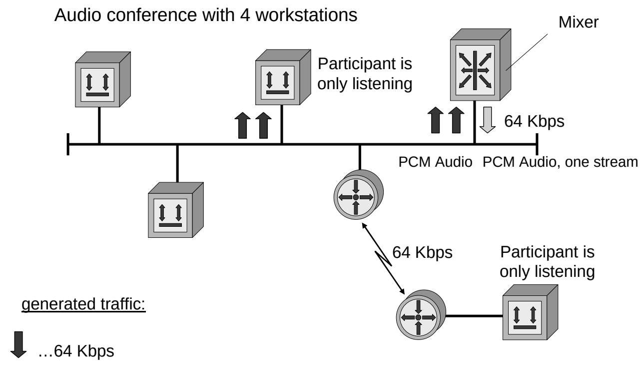

RTP Mixing with 2 speakers

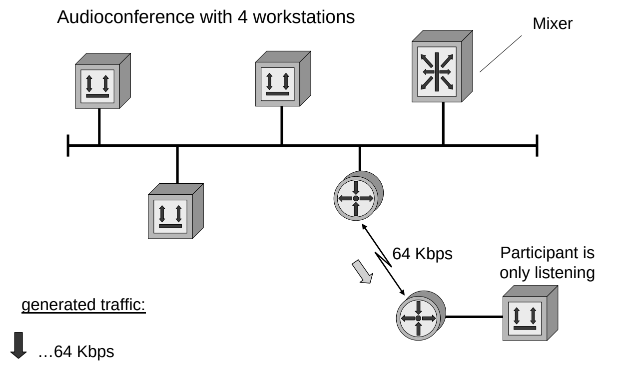

RTP Translators and Mixers

- **RTP mixers (cont.)**
  - payload type of the incoming and outgoing packets remain the same
  - it is possible to combine RTP mixing and RTP translating in the same environment
  - this would be required if the workstation is connected via a lower-speed link
  - payload format of the PCM stream must be changed to a lower bandwidth specification

RFCs of Interest 1

- RFC 1075 Distance Vector Multicast Routing Protocol, November 1988
- RFC 1112 Host Extensions for IP Multicasting, August 1989
- RFC 2236 Internet Group Management Protocol, Version 2, November 1997 (now obsolete)
- RFC 2327 SDP: Session Description Protocol, April 1998
- RFC 2362 Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification, June 1998 (Obsoletes 2117)
- RFC 2365 Administratively Scoped IP Multicast, July 1998
- RFC 2432 Terminology for IP Multicast Benchmarking, October 1998
- RFC 2588 IP Multicast and Firewalls, May 1999
- RFC 2614 An API for Service Location, June 1999
- RFC 2627 Key Management for Multicast: Issues and Architectures, June 1999
- RFC 2674 Definitions of Managed Objects for Bridges with Traffic Classes, Multicast Filtering and Virtual LAN Extensions, August 1999
- RFC 2715 Interoperability Rules for Multicast Routing Protocols, October 1999
- RFC 2729 Taxonomy of Communication Requirements for Large-scale Multicast Applications, December 1999

Institute of Computer Technology - Vienna University of Technology

L71 - IP Multicasting

RFCs of Interest 2

- RFC 2858 Multiprotocol Extensions for BGP-4, June 2000
- RFC 2887 The Reliable Multicast Design Space for Bulk Data Transfer, August 2000
- RFC 2932 IPv4 Multicast Routing MIB, October 200.
- RFC 2934 Protocol Independent Multicast MIB for IPv4, October 2000
- RFC 2974 Session Announcement Protocol, October 2000
- RFC 3031 Multiprotocol Label Switching Architecture, January 2001
- RFC 3048 Reliable Multicast Transport Building Blocks for One-to-Many Bulk-Data Transfer, January 2001
- RFC 3138 Extended Assignments in 233/8, June 2001
- RFC 3170 IP Multicast Applications: Challenges and Solutions, September 2001
- RFC 3171 IANA Guidelines for IPv4 Multicast Address Assignments, August 2001
- RFC 3180 GLOP Addressing in 233/8, September 2001.
- RFC 3208 PGM Reliable Transport Protocol Specification, December 2001.
- RFC 3228 IANA Considerations for IPv4 Internet Group Management Protocol (IGMP)
- RFC 3259 A Message Bus for Local Coordination, April 2002.

RFCs of Interest 3

- RFC 3261 SIP: Session Initiation Protocol, June 2002
- RFC 3353 Overview of IP Multicast in a Multi-Protocol Label Switching (MPLS) Environment, August 2002
- RFC 3376 Internet Group Management Protocol, Version 3, October 2002
- RFC 3446 Anycast Rendezvous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP), January 2003
- RFC 3488 Router-port Group Management Protocol (RGMP), February 2003
- RFC 3559 Multicast Address Allocation MIB, June 2003
- RFC 3569 An Overview of Source-Specific Multicast (SSM), July 2003
- RFC 3618 Multicast Source Discovery Protocol (MSDP), October 2003
- RFC 3678 Socket Interface Extensions for Multicast Source Filters, January 2004
- RFC 3740 The Multicast Group Security Architecture, March 2004
- RFC 3754 IP Multicast in Differentiated Services (DS) Networks, April 2004
- RFC 3913 Border Gateway Multicast Protocol (BGMP): Protocol Specification September 2004
- RFC 3918 Methodology for IP Multicast Benchmarking, October 2004
- RFC 3926 FLUTE - File Delivery over Unidirectional Transport, October 2004

RFCs of Interest 4

- RFC 3956 Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address, November 2004 (Updates RFC3306)
- RFC 3973 Protocol Independent Multicast – Dense Mode (PIM-DM): Protocol Specification (Revised), January 2005
- RFC 4045 Extensions to Support Efficient Carrying of Multicast Traffic in Layer-2 Tunneling Protocol (L2TP), April-20-2005
- RFC 4046 Multicast Security (MSEC) Group Key Management Architecture, April-29-2005
- RFC 4082 Timed Efficient Stream Loss-Tolerant Authentication (TESLA): Multicast Source Authentication Transform Introduction, June 2005
- RFC 4271 A Border Gateway Protocol 4 (BGP-4), January 2006 (Obsoletes 1771)
- RFC 4286 Multicast Router Discovery, December 2005
- RFC 4363 Definitions of Managed Objects for Bridges with Traffic Classes, Multicast Filtering, and Virtual LAN Extensions, January 2006 (Obsoletes RFC2674.

Other RFCs 1

- RFC 966 Host Groups: A Multicast Extension to the Internet Protocol, December 1985 (now obsolete)
- RFC 988 Host Extensions for IP Multicasting, July 1986 (now obsolete)
- RFC 1054 Host Extensions for IP Multicasting, May 1988 (now obsolete)
- RFC 1458 Requirements for Multicast Protocols, May 1993
- RFC 1584 Multicast Extensions to OSPF, March 1994
- RFC 1585 MOSPF: Analysis and Experience, March 1994
- RFC 1819 Internet Stream Protocol Version 2 (ST2) Protocol Specification-Version ST2+, August 1995
- RFC 2117 Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification, June 1997 (Obsoleted by 2362)
- RFC 2201 Core Based Trees (CBT) Multicast Routing Architecture, September 1997
- RFC 2429 RTP Payload Format for the 1998 Version of ITU-T Rec. H.263 Video (H.263+), October 1998
- RFC 2431 RTP Payload Format for BT.656 Video Encoding, October 1998
- RFC 2730 Multicast Address Dynamic Client Allocation Protocol (MADCAP), December 1999
- RFC 2770 GLOP Addressing in 233/8, February 2000 (obsoleted by RFC 3180)
- RFC 2771 An Abstract API for Multicast Address Allocation, February 2000

Other RFCs 2

- RFC 2776 Multicast-Scope Zone Announcement Protocol (MZAP), February 2000
- RFC 2907 MADCAP Multicast Scope Nesting State Option, September 2000
- RFC 2908 The Internet Multicast Address Allocation Architecture, September 2000
- RFC 2909 The Multicast Address-Set Claim (MASC) Protocol, September 2000