Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

Internet Group Communication: IP Multicasting

- Introduction
- Why to Talk in Groups?
- Aspects of Group Communication
- IP-Multicasting
- Addressing
- The Internet Group Protocol

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

Motivation

Current Situation: Use of the Internet penetrates new areas:

- Multimedia Information Services
- Infotainment
- Gaming (MMORPGs)
- Synchronous Network Information Services
- Group Communication Tools

⇒ A Transport Infrastructur is needed for

Group Communication Services

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

Requirements

- Must fit into current IP Infrastructure
- Location Transparancy
- Efficiency in Data Distribution
- Utmost Independence of Specific Hardware
- Independent Open Standards
- Interoperability

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

Example: Video Streaming

- High Requirements of Bandwidth and Server Performance
- Continuous Data Streams
- Real-Time Synchronisation
- Global Distribution (Web-TV, IPTV, Video-Conferencing)

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

Issues:

- Internet Group Communication Fundaments

  Group Communication, Multicasting, Adressing, IGMP/MLD

- Layer 2 Multicasting

  Local Networks, Adress-Mapping, Framing, Discovery, ATM

- Multicast Routing

  Specialities, Algorithms, Protocols

- New Developments

  IPng, SSM, Multicast Mobility

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

Why to Talk in Groups?

Internet based communication steadily gains importance, quantitatively as well as qualitatively. New communication forms arise, old services spread rapidly:

- Multimedia Distribution
- ‚Broadcasting‘ - Offers
- Telecommunication Services

⇒ Scalable Communication Paths needed to Distribute Data in Parallel

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

Ineffective Group Communication

Unicast

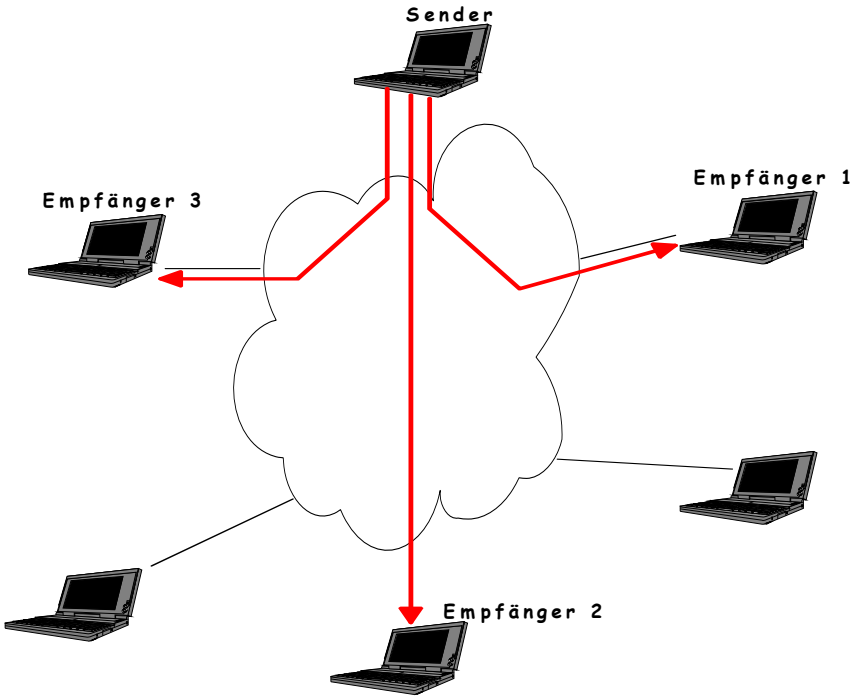

Broadcast

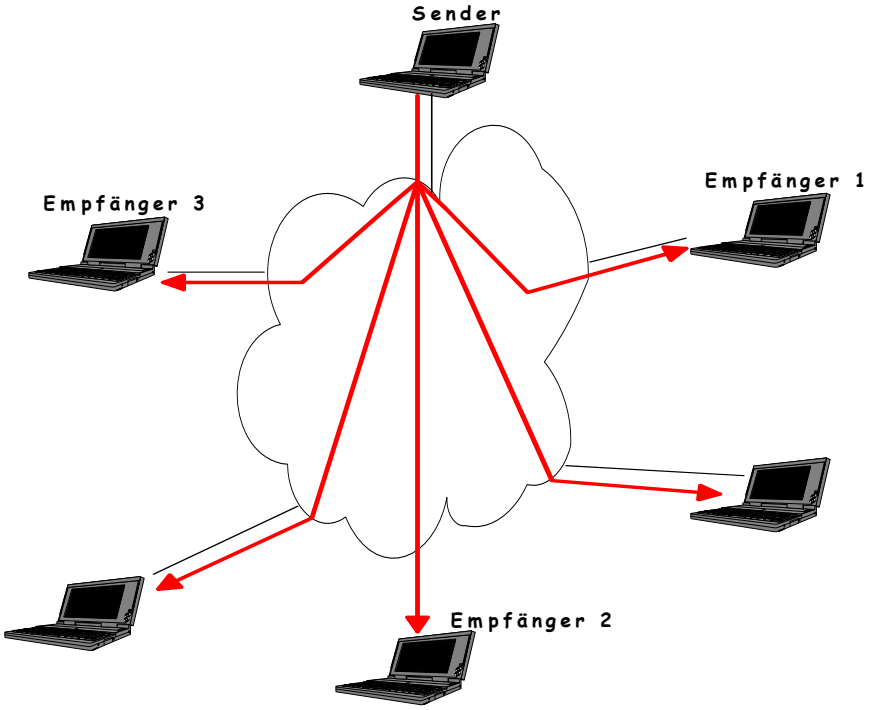

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

Effective Group Communication

Multicast

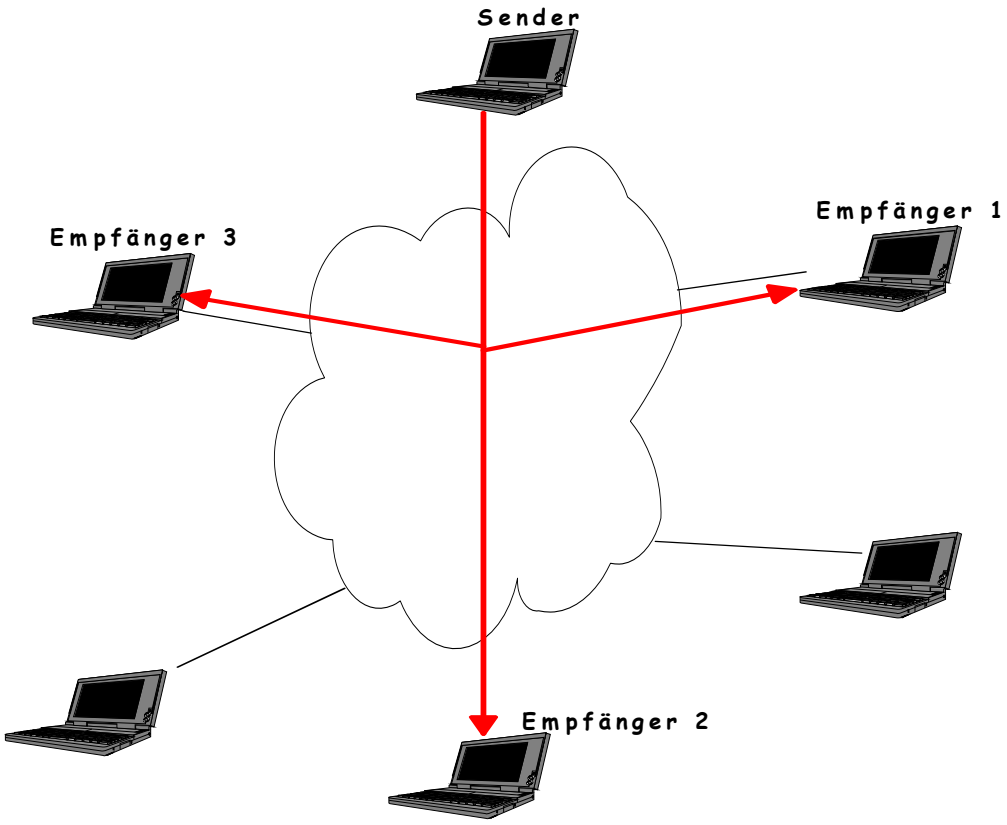

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

Group Communication Differs

Classical TCP/IP Communication Model:

- Client/Server Principle
- Individual Communication Channels
  - Initiated by client
  - Server answers individually
  - Server speaks on many point-to-point channels
- Exception for unspecific message distribution:

  Broadcasts

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

Examples

IRC – Client-to-Client Communication via Server

NTP – many Clients ask one Server

Routing (RIPv1) – Broadcast of Routing Tables

Multisource Webpage – Client asks many Server

Internet Server Farm – one Client asks one of many Servers

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

Group Communication Modes

Broadcast – one Sender to all Members of the Subnet

Concast – one Receiver of a Group of Senders

Multicast – one Sender Addresses a Group of Receivers

Multipeer – a Group of Senders to a Group of Receivers

Anycast – Communication Partner selected from a Group of potential Partners (Unicast)

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

Aspects of Group Communication

- Openness: Support of open and closed groups
- Dynamic: Change of group membership
- Reliability: Securing of data transport
- Flow Control: Adapting data streams to buffers
- Group Management: Mechanisms of addressing and membership control

Openness & Dynamic

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

Relevant Mechanisms

- Identification/Announcement in a group/ with the sender
- Authorisation in closed Groups
- Management of send/receive allowances
- Registration & deregistration, definition of group composition
- Definition of group lifetime

Reliability

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

A securing layer requests for some acknowledgements

ACK:

- Group members need to register with sender
- results in ACK-‘Implosion‘

NACK:

- Retransmission for one may disturb the entire group
- Last loss may be unseen

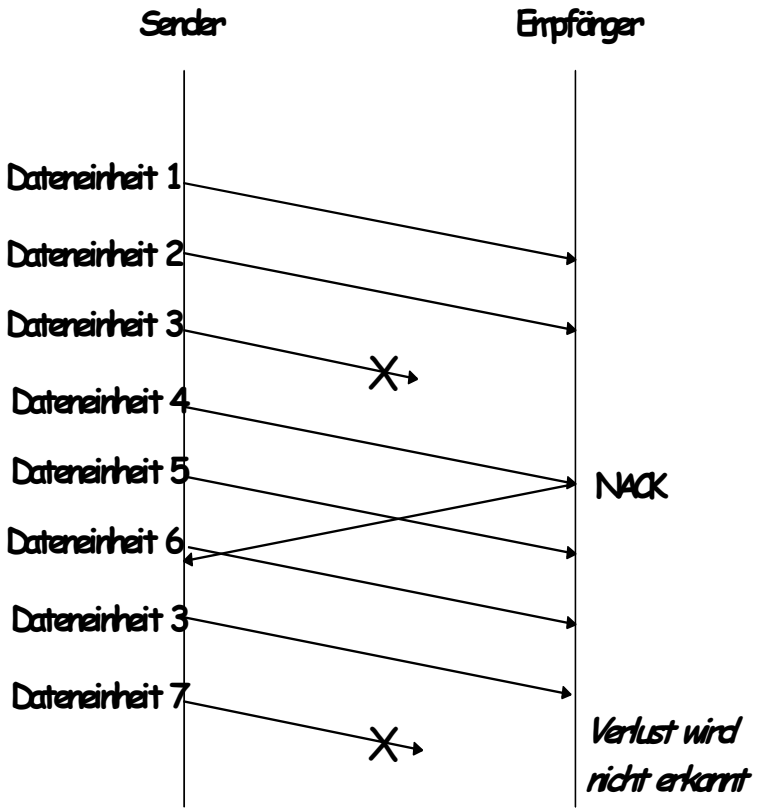

Ack Implosion

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

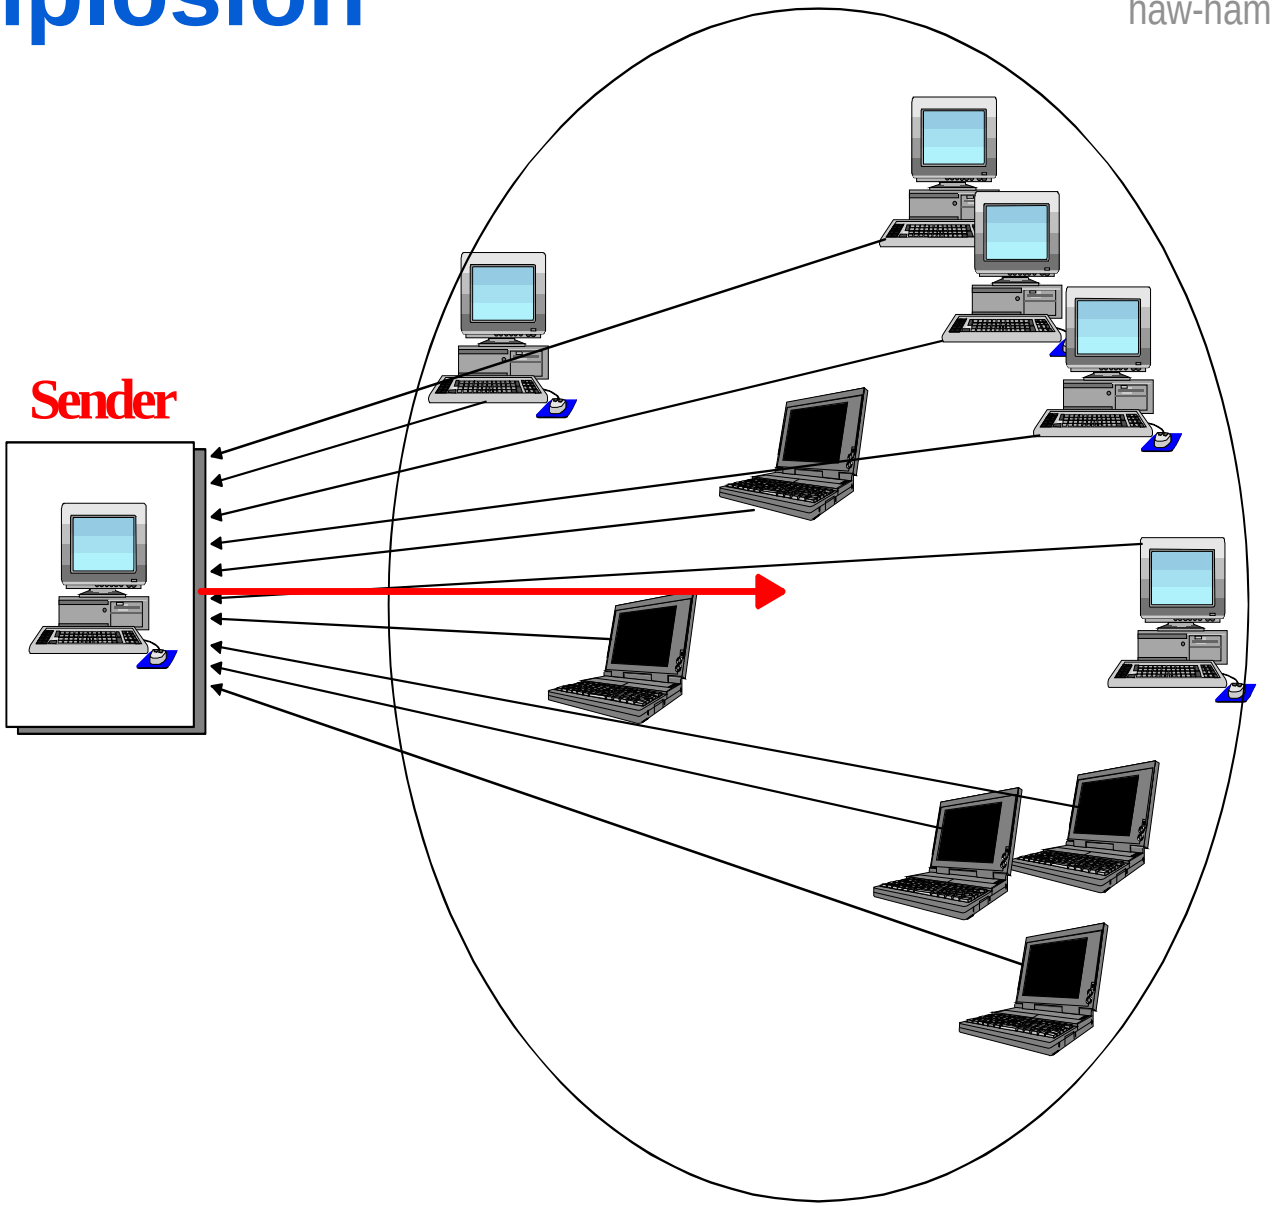

Flow Control

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

Window based:

- uses positive acknowledgements for sliding the window
  $\Rightarrow$ unsuitable

Rate based:

- Adjust source intensity (Burst Rate)
- May be announced by receiver with membership registration

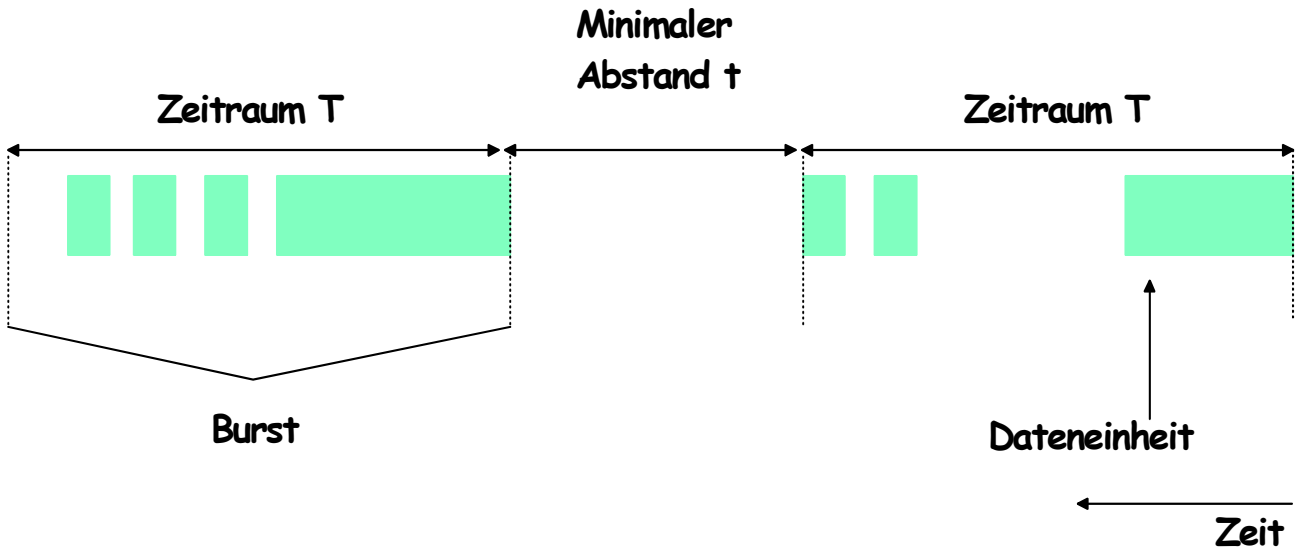

Group Management

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

Addressing

- Address scheme for a group
- Address allocation (centralised or decentralised)

Signaling

- Registration/Deregistration
- Member management (centralised or decentralised)

IP Multicasting

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

Method for Transfering IP Datagrams to Host-Groups

- Initially: RFC 1112 (S. Deering et.al., 1989)
- Addresses a host group by *one* group address
- Two kinds of multicast:
  - Any Source Multicast (ASM)
  - Source Specific Multicast (SSM)
- Client Protocol for registration (IGMP/MLD)
- Routing throughout the Internet (Multicast Routing)
- Address translation into Layer 2

Main Advantage of IP Multicasting

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

- Omits redundant network traffic
- Reduces network and server load

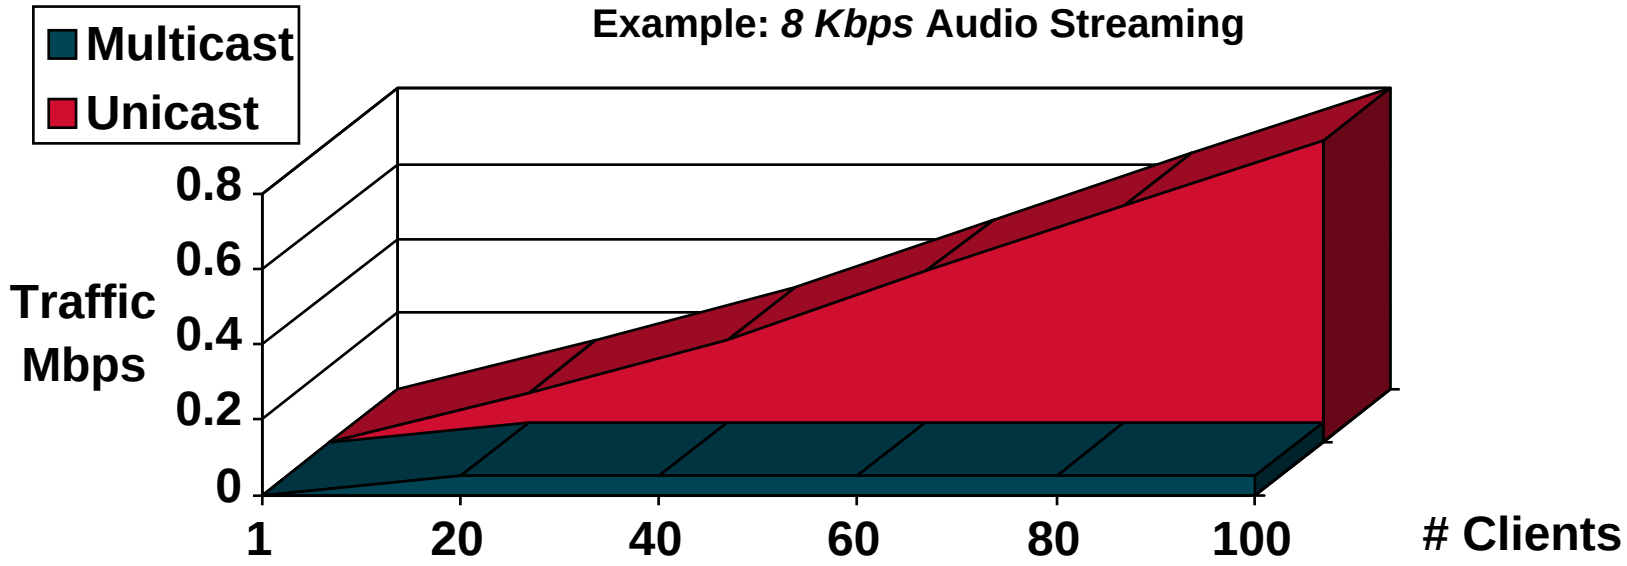

Aspects of IP Multicasting

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

- Offers to sender a data delivery service to a distributed unknown group of receivers (multipoint access)
- UDP-based
- Best Effort Transport
- Securing and flow control left to application
- No closed groups
- No restriction on senders
- Applications may react source address sensitive

Multicast Network

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

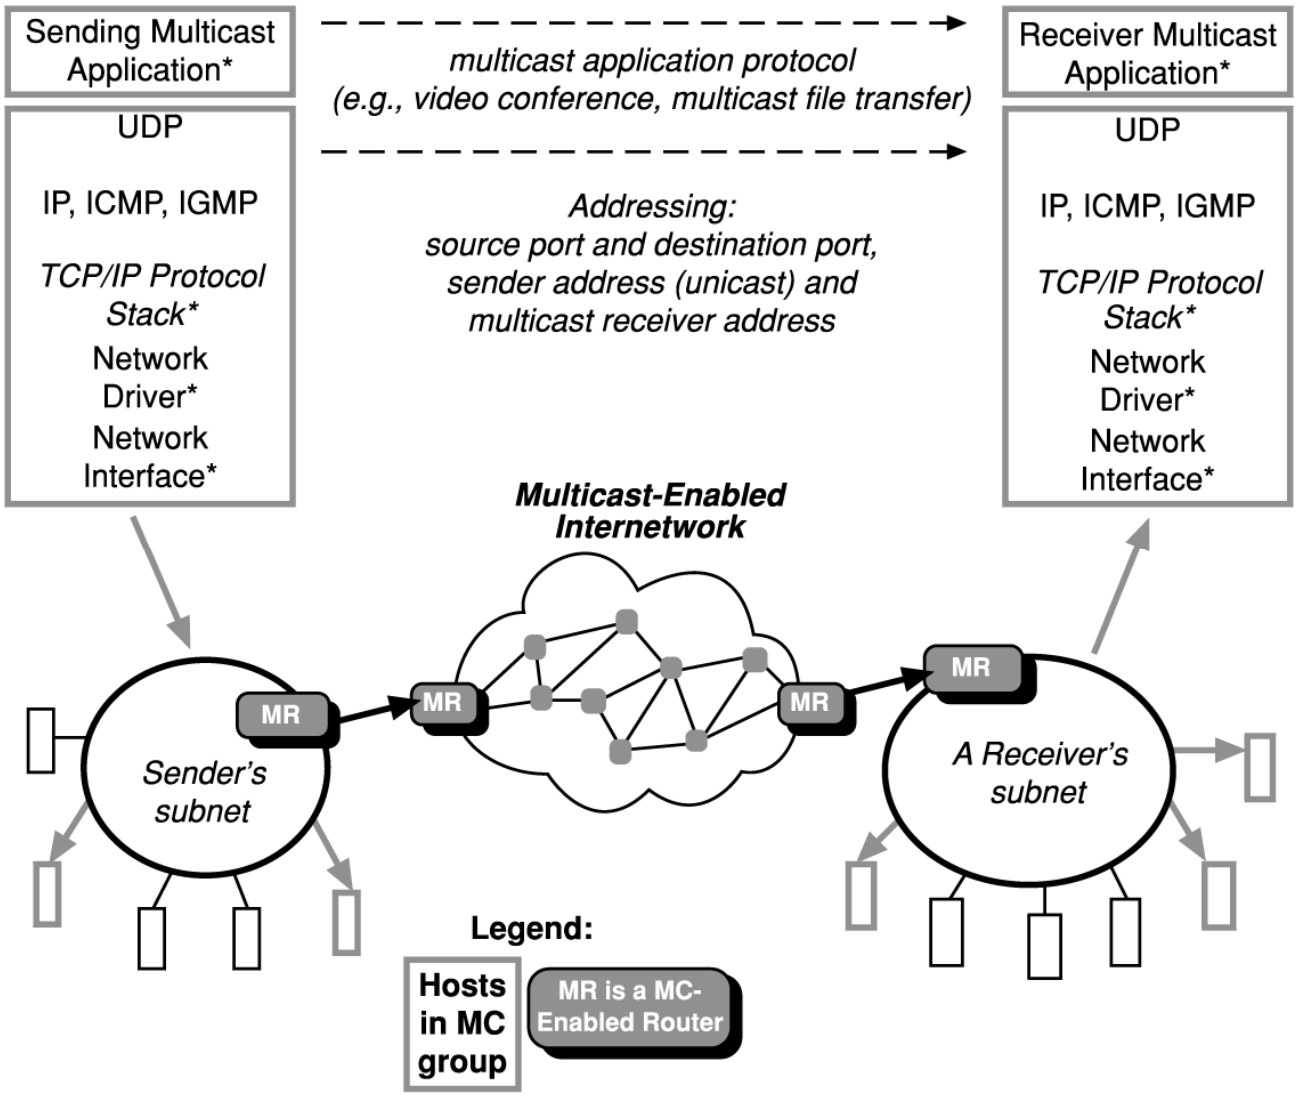

Applications of IP Multicasting

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

- Multimedia
  - Streaming video and audio (broadcasting)
  - Teleteaching
  - Conferencing
- Financial information services (stock price ticker,...)
- Netzwork information services
- Arbitrary data distribution services (Pusch Apps)

First Gobal Multicast Deployment: MBONE

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

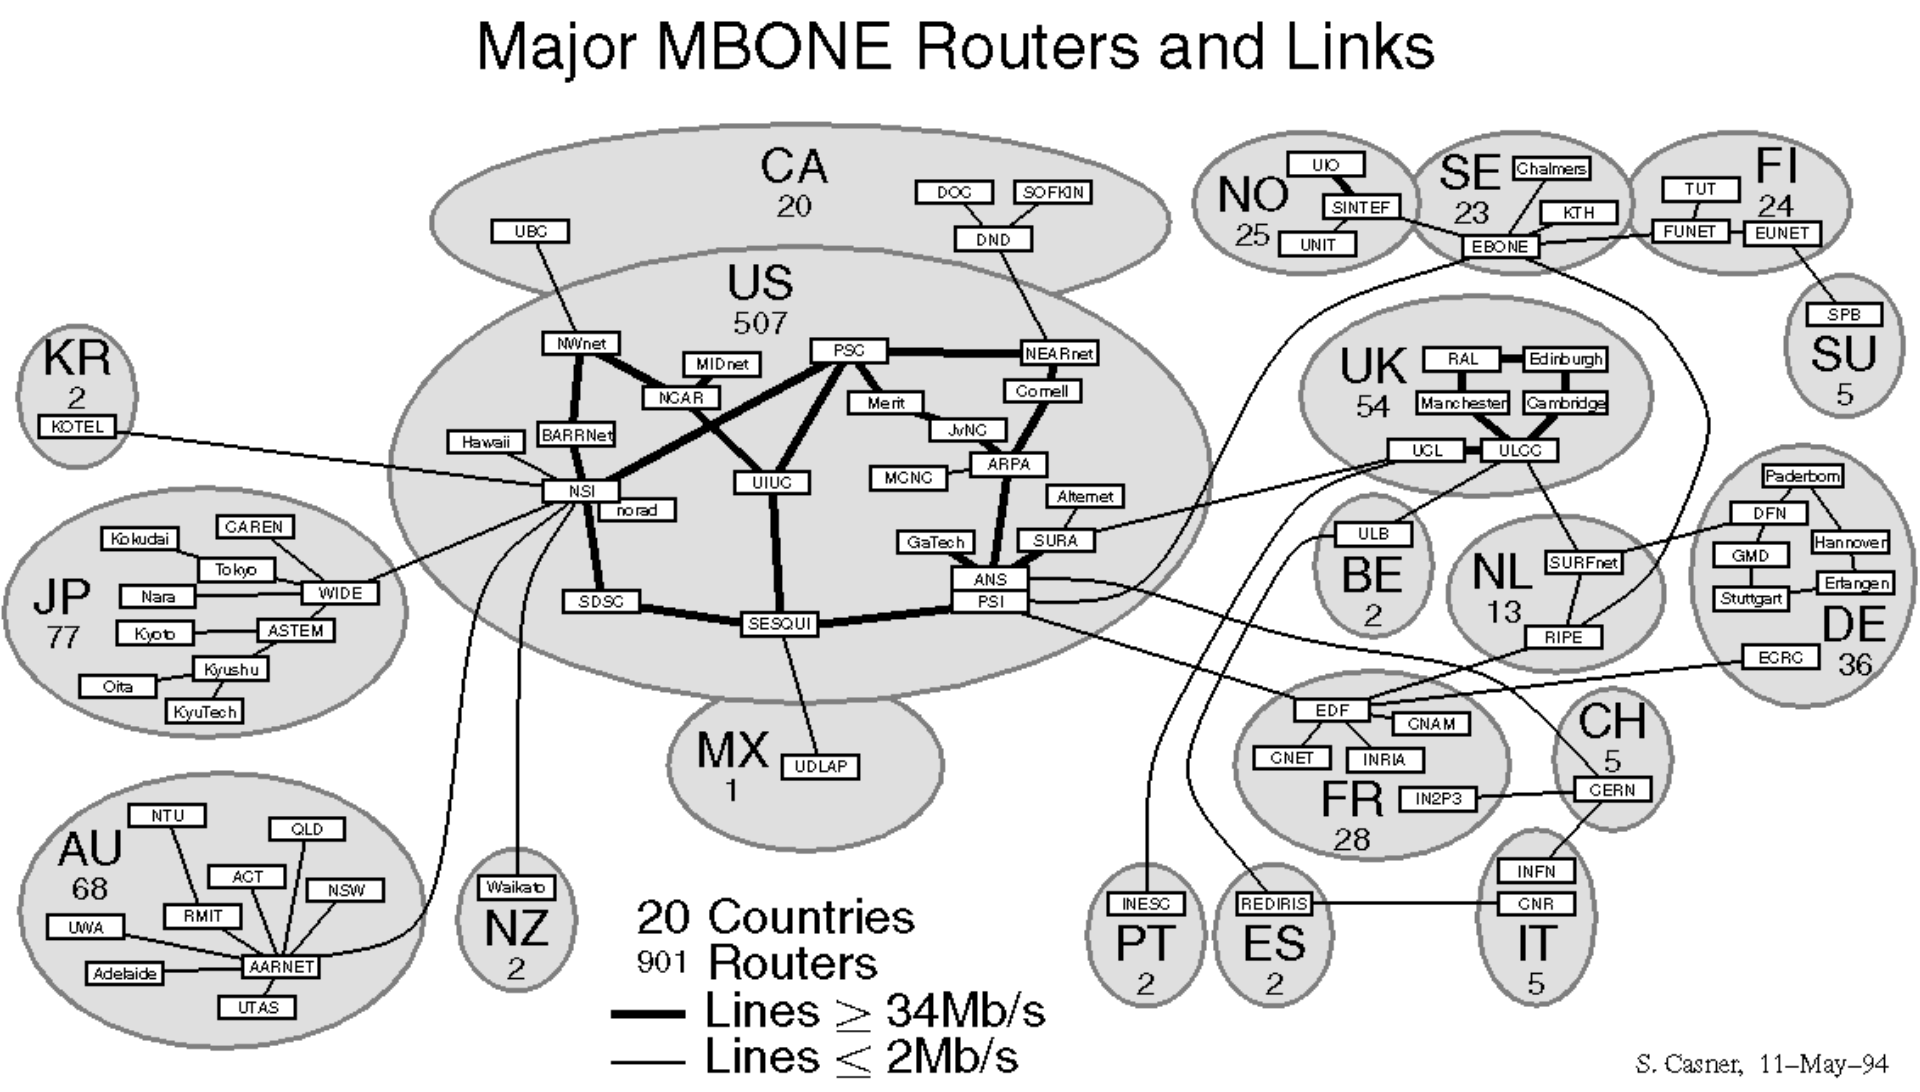

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

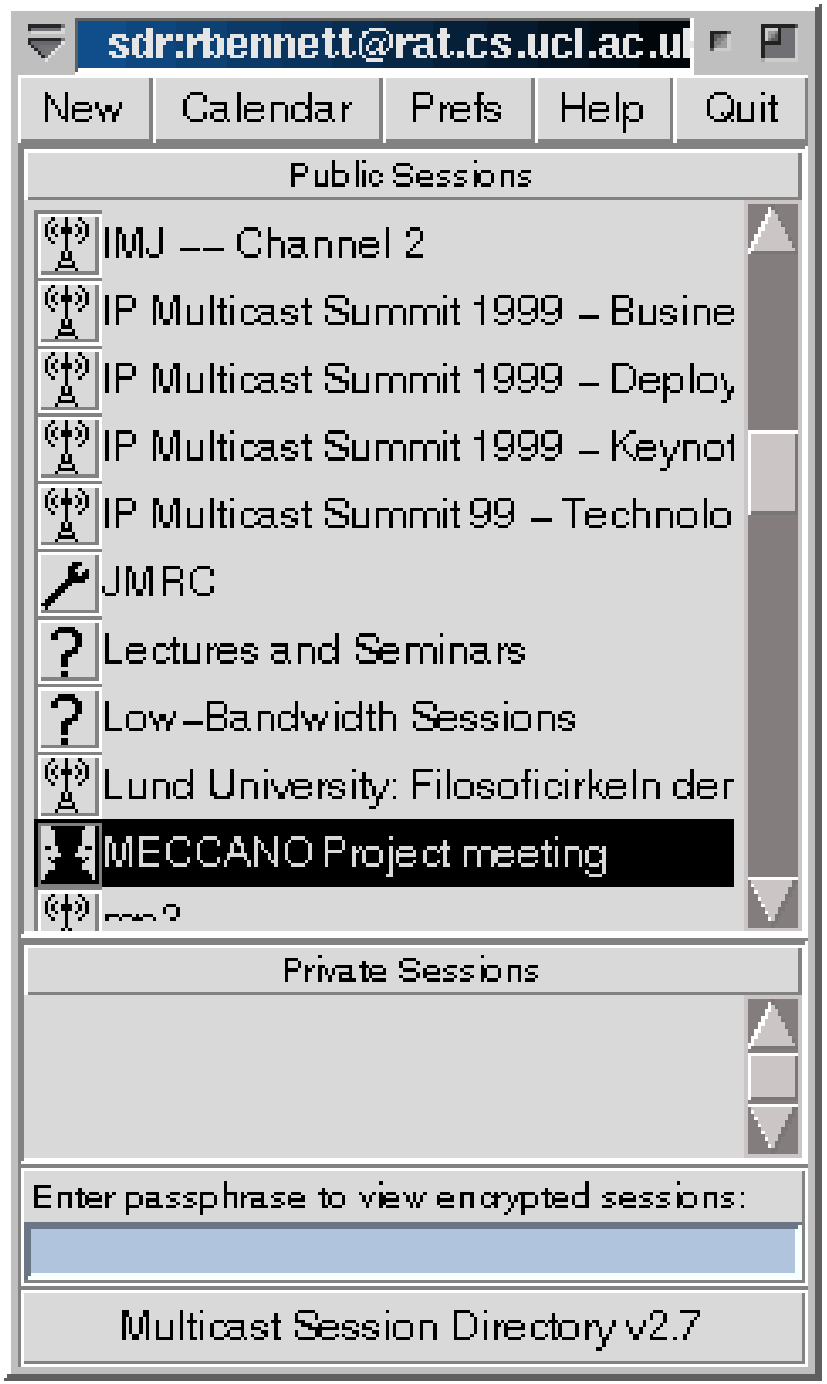

Example:

Mbone-Tools

SDR

Example: Mbone-Tools Rendez-Vous

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

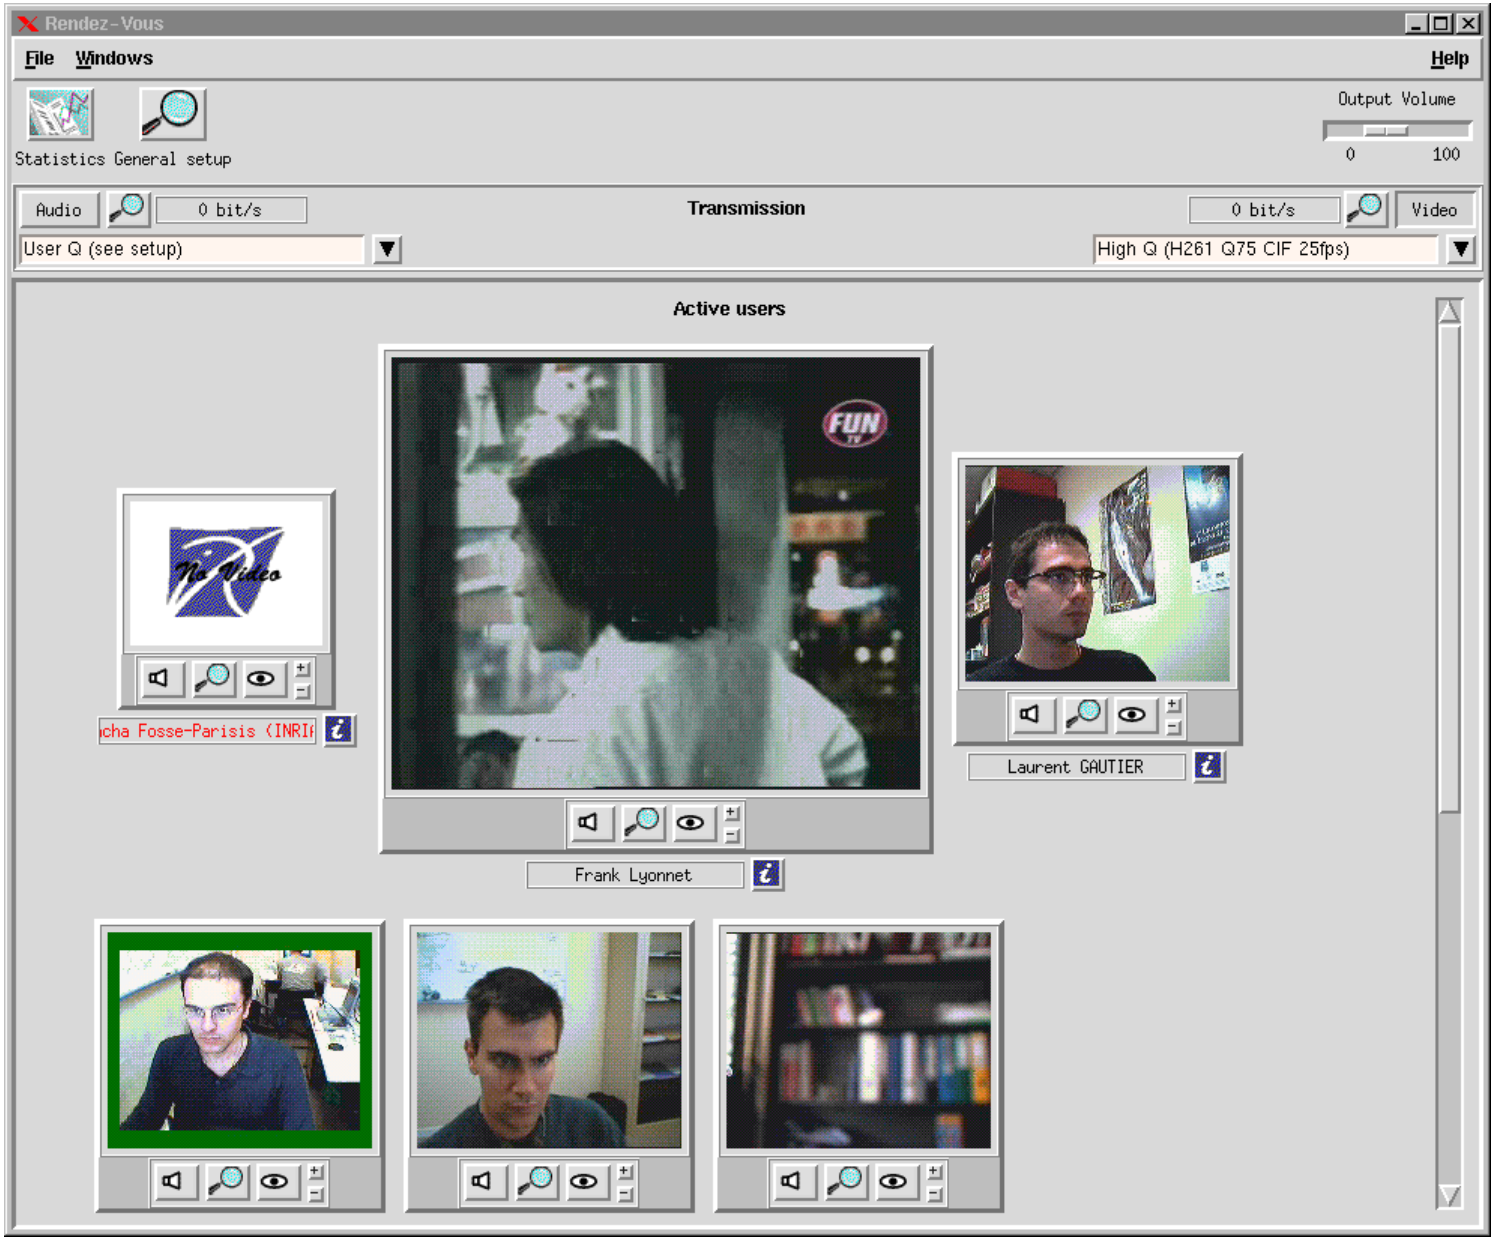

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

Multicast Addressing

- Denote delocalized addresses
- IPv4 Multicast Group addresses
  - 224.0.0.0–239.255.255.255
  - Class “D” Address Space
  - Special SSM block: 232.*.*.*
- IPv6: scoped multicast addresses
  - FF00::/8
  - Special SSM block: FF3x::/32
- Permanent Addresses assigned by IANA
  - RFC 1700: Assigned Adresses
  - *“http://www.iana.org/assignments/multicast-addresses ”* lists reserved addresses
- Dynamic Addresses
  - independent of local IP-address space (IPv4)
  - Unicast based Multicast addresses (IPv6)

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

Internet Address Classes

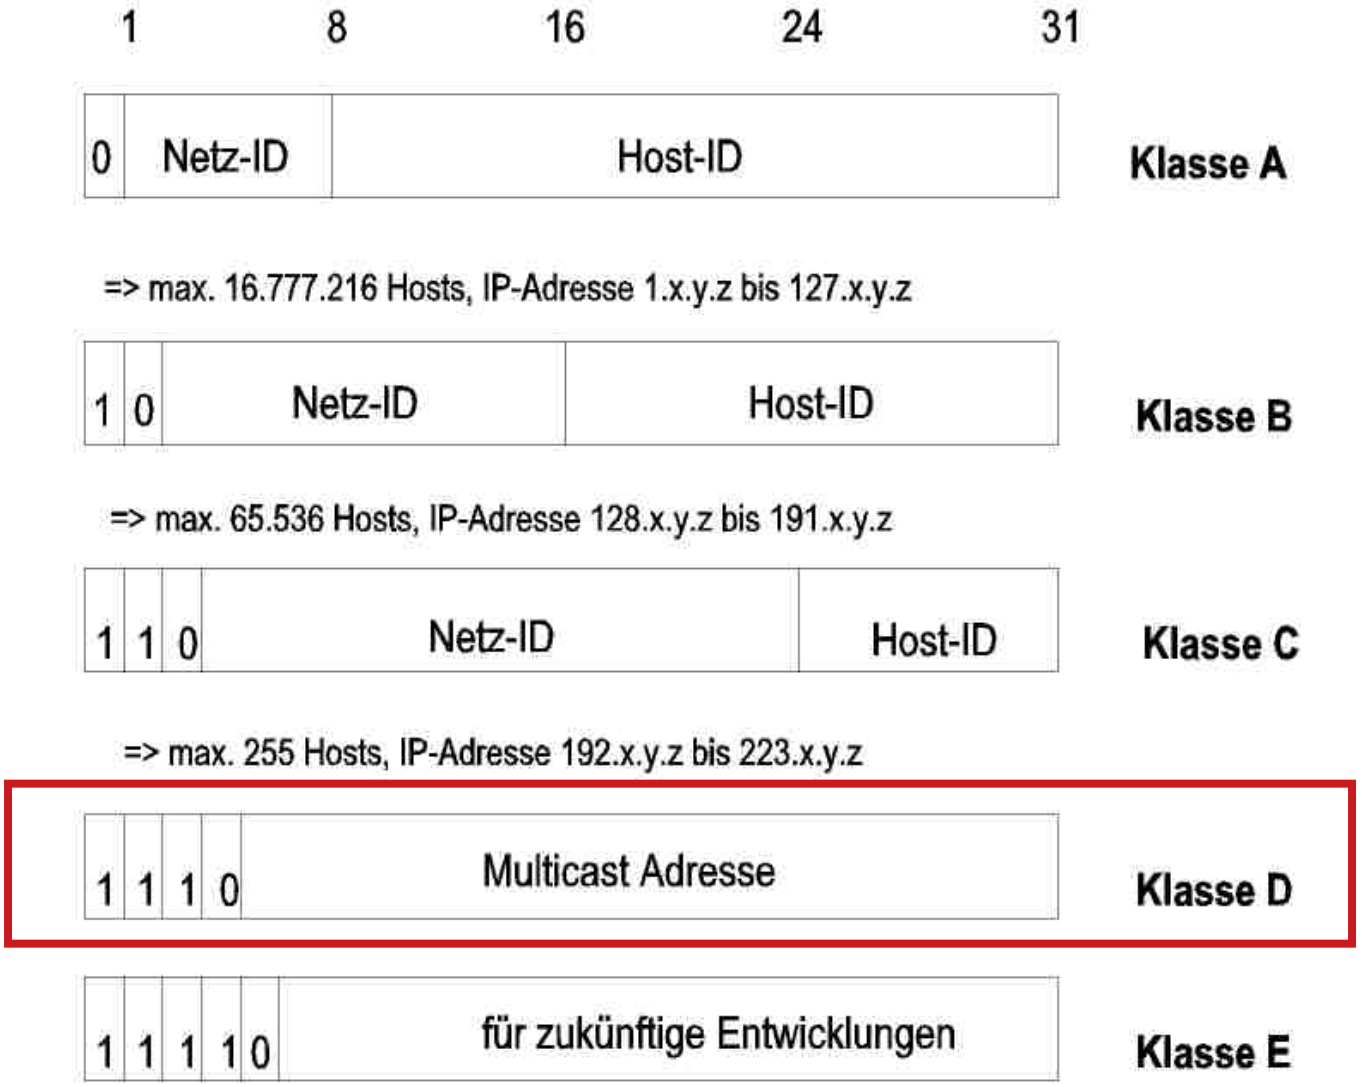

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

Private Multicast Addresses

- Officially not routed address range
  - 239.0.0.0–239.255.255.255
  - Private Address Space
    - Similar to RFC1918 Unicast Addresses
    - Unused for global Internet Traffic
    - Limits Multicast Traffic to own Institution
    - Same Addresses may be globally re-used
  - Example
    - Local range: 239.253.0.0/16
    - Organisation-wide range: 239.192.0.0/14

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

Reserved Multicast Addresses

- Permanent IP Multicast Group Addresses
  - 224.0.0.0–224.0.0.255
  - Examples:
    - 224.0.0.1 All Systems of Subnet
    - 224.0.0.2 All Routers of Subnet
    - 224.0.0.4 All DVMRP Router
    - 224.0.0.5 All OSPF Router
    - 224.0.0.9 All RIP(v2) Router
    - 224.0.0.13 All PIMv2 Router
    - 224.0.1.1 NTP
    - 224.0.1.9 Multicast Transport Protocol (MTP)
- TTL – Standards in MBONE
    - TTL = 1: This Subnet
    - TTL = 15: This Site
    - TTL = 63: This Region
    - TTL = 127: This Internet

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

IPv6 Multicast Addresses

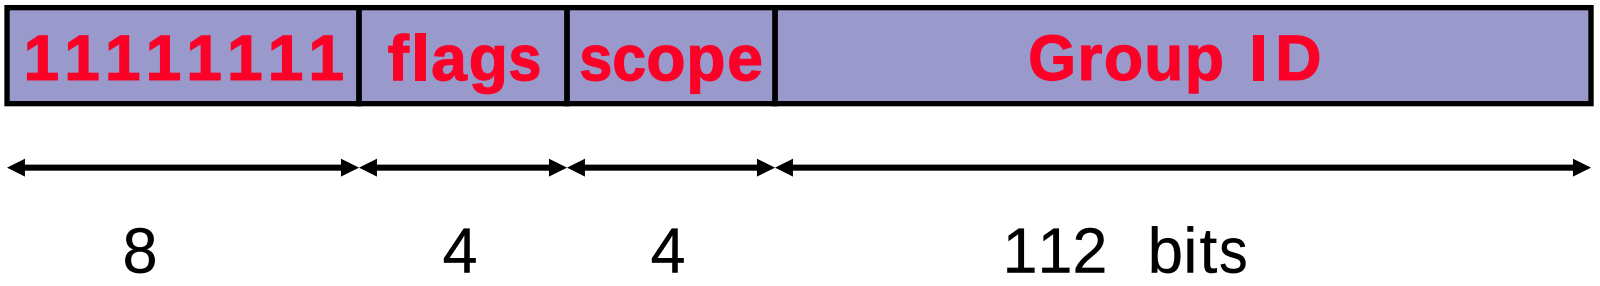

- Flag field: lower bit indicates permanent (=0) respectively transient (=1) group, rest is reserved (==0)
- Scope field:
  1 - node local
  2 - link-local
  5 - site-local
  8 - organisation local
  B - community-local
  E - global (other values reserved)

IPv6 Unicast Based Multicast Addresses (RFC 3306)

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

- **Solves the old IPv4 address assignment problem:**
  How can I get global IPv4 multicast addresses (GLOB, ..)
- **In IPv6, if you own an IPv6 unicast address prefix you implicitly own an RFC3306 IPv6 multicast address prefix:**

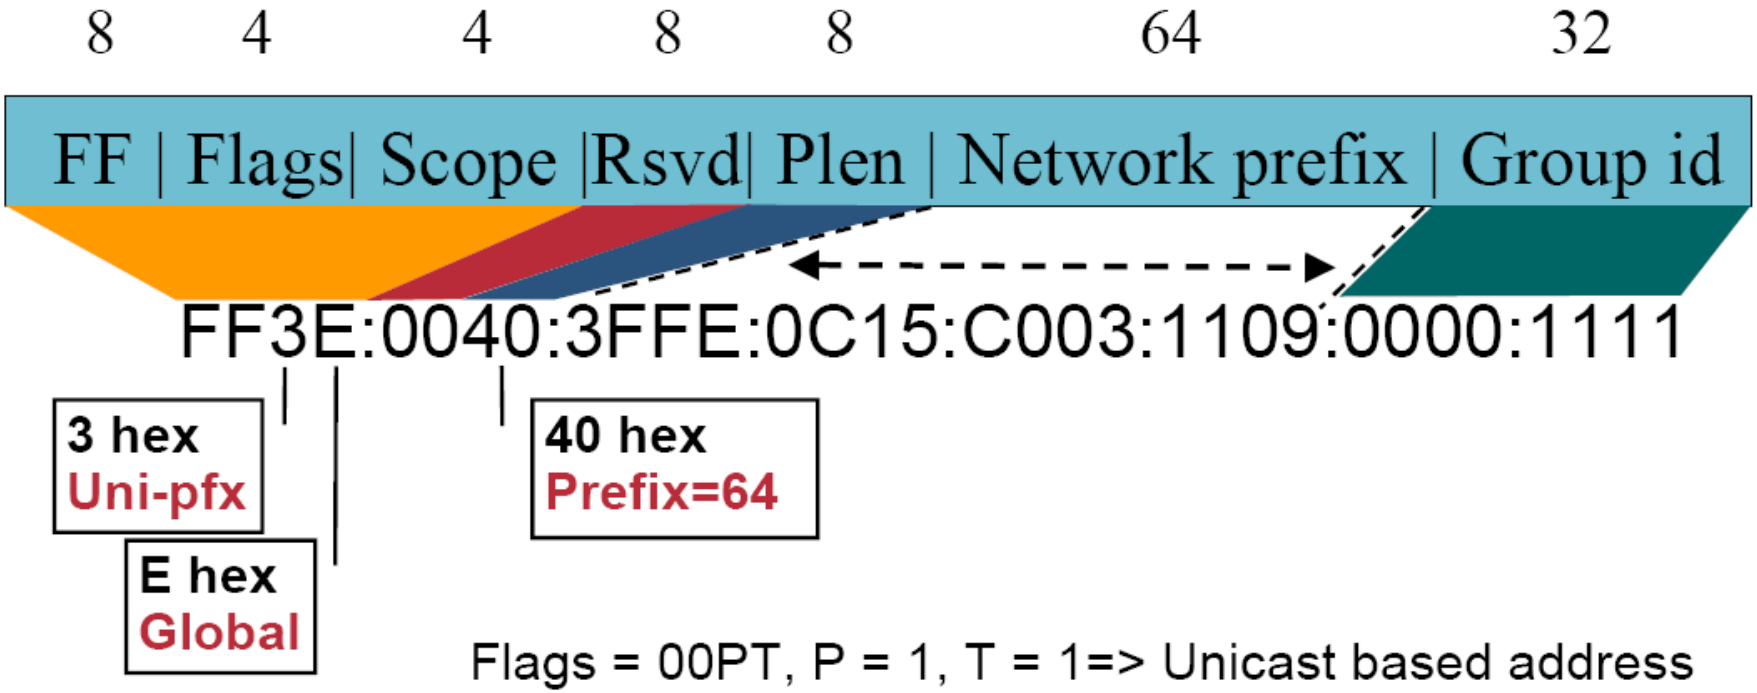

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

Dynamic Multicast Addressing

- Dynamic Assignment of Group addresses:
  - Until now: SDR Application
  - Sessions/Groups announced via well-known multicast groups
  - Address assignments and collisions are managed within initiation process
  - Brings up severe scaling issue
- Future Techniques and Planning:
  - Multicast Address Set-Claim (MASC)
    - Hierarchical, dynamical address allocation scheme
    - Difficult and far
  - MADCAP
    - Similar to DHCP
    - Needs own Protocol stack and application integration!

Internet Group Management

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

Internet Group Management Protocol (IGMP)

- Client Protocol to initiate, preserve and terminate group membership
- Local Router collect and monitor information
- IPv4: Internet Group Management Protocol (IGMP)
  - IGMP v1 RFC 1112
  - IGMP v2 RFC 2236 – implemented almost everywhere
  - IGMP v3 RFC 3376
- IPv6: Multicast Listener Discovery Protocol (MLD)
  - MLDv1 (RFC 2710) – analogue to IGMPv2
  - MLDv2 (RFC 3810) – starting from IGMPv3
- SSM Specialities: RFC 4604

IGMP

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

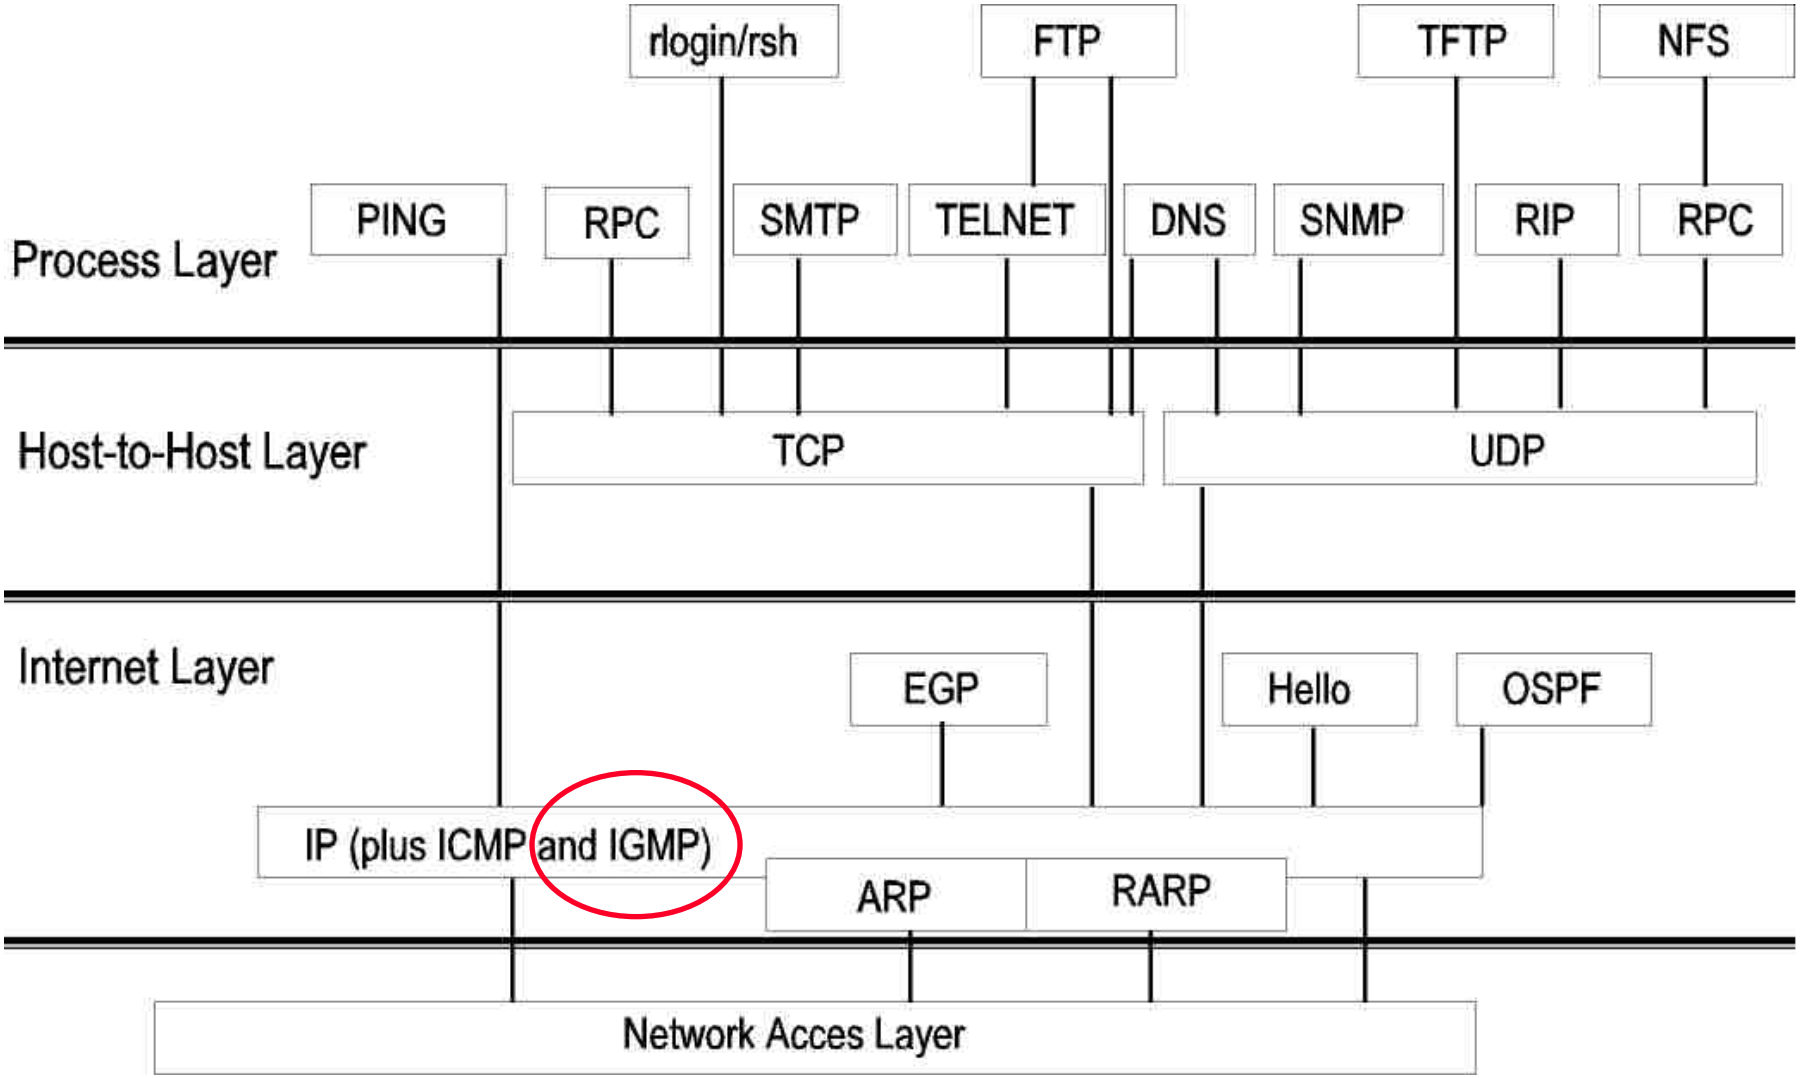

IGMP Protocol Architecture

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

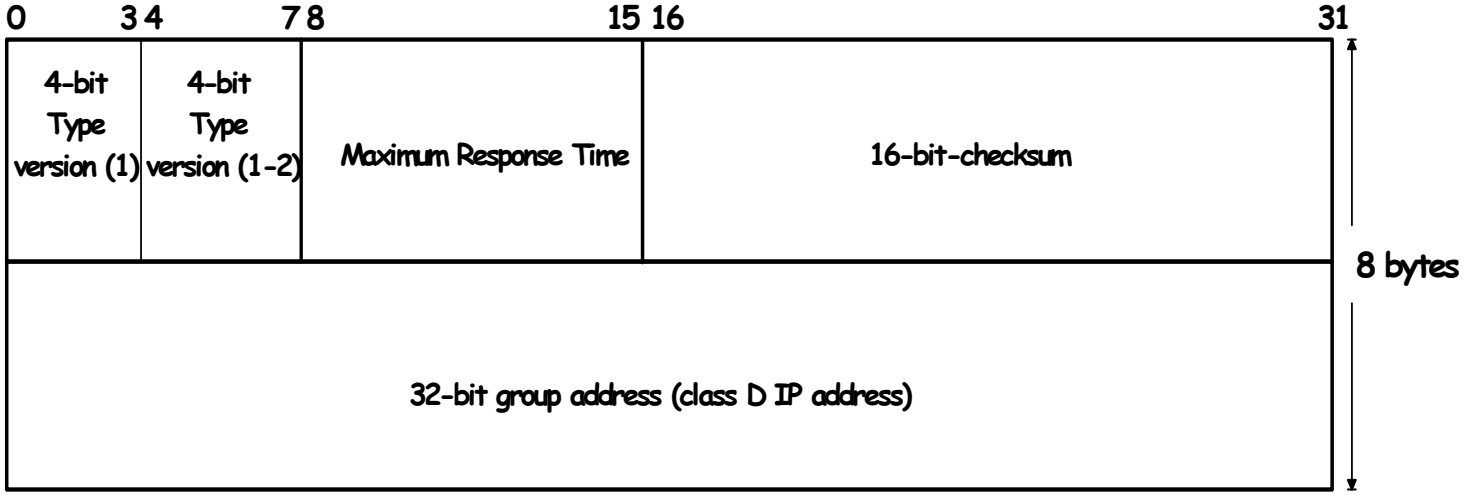

IGMP works like ICMP with Queries:

- General Membership
- Group specific Membership
- Version 2 Membership Report
- Leave Query
- Version 1 Membership Report

IGMP Communikation

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

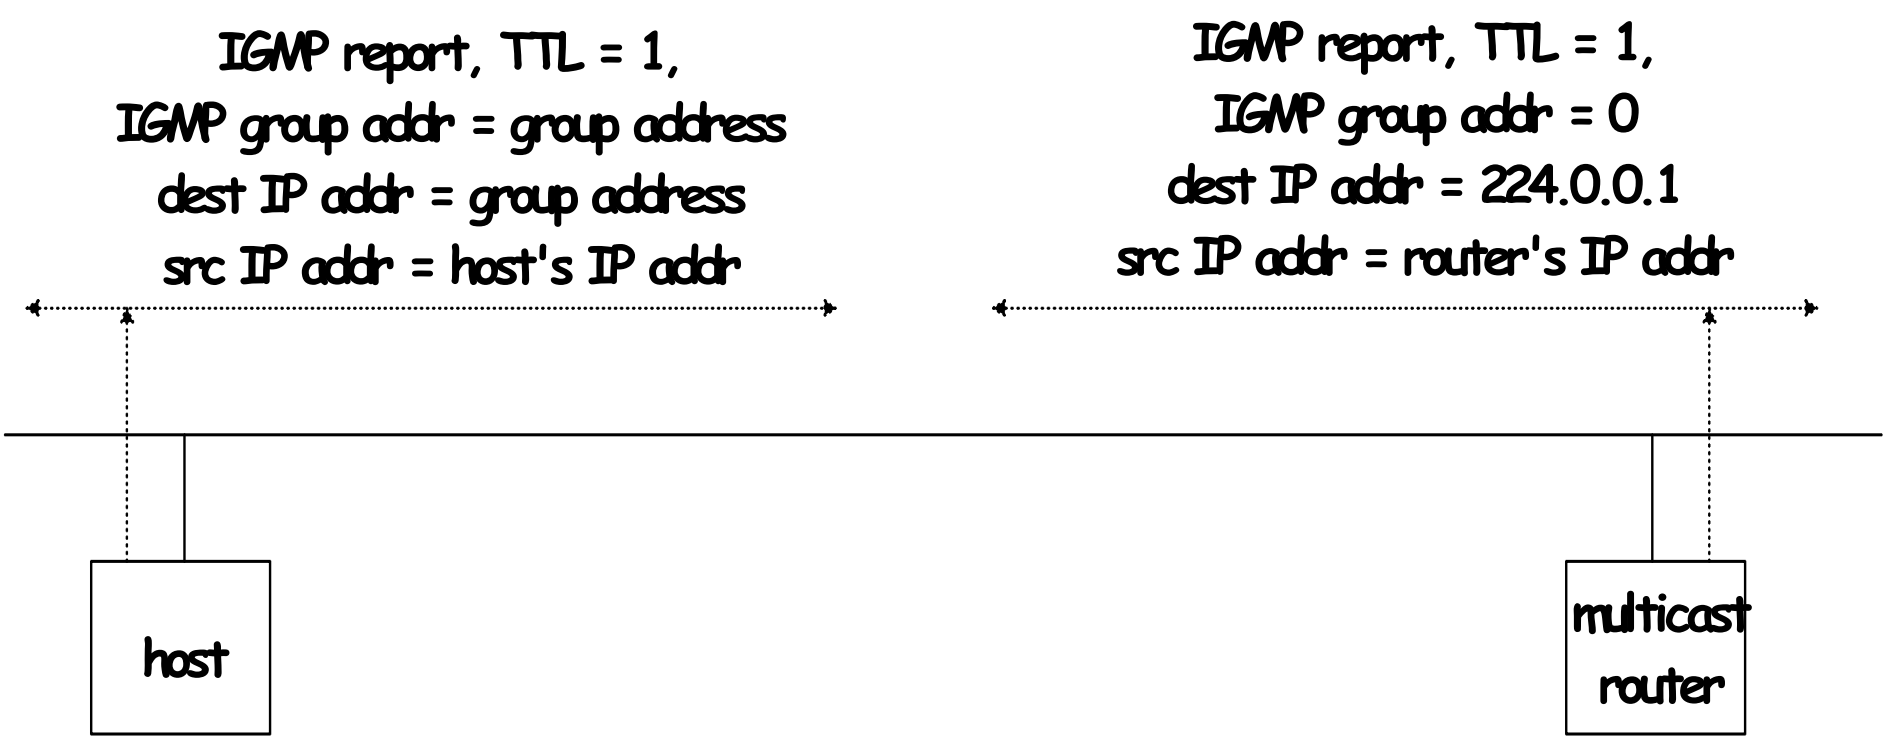

IGMP Host-Router Signalling

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

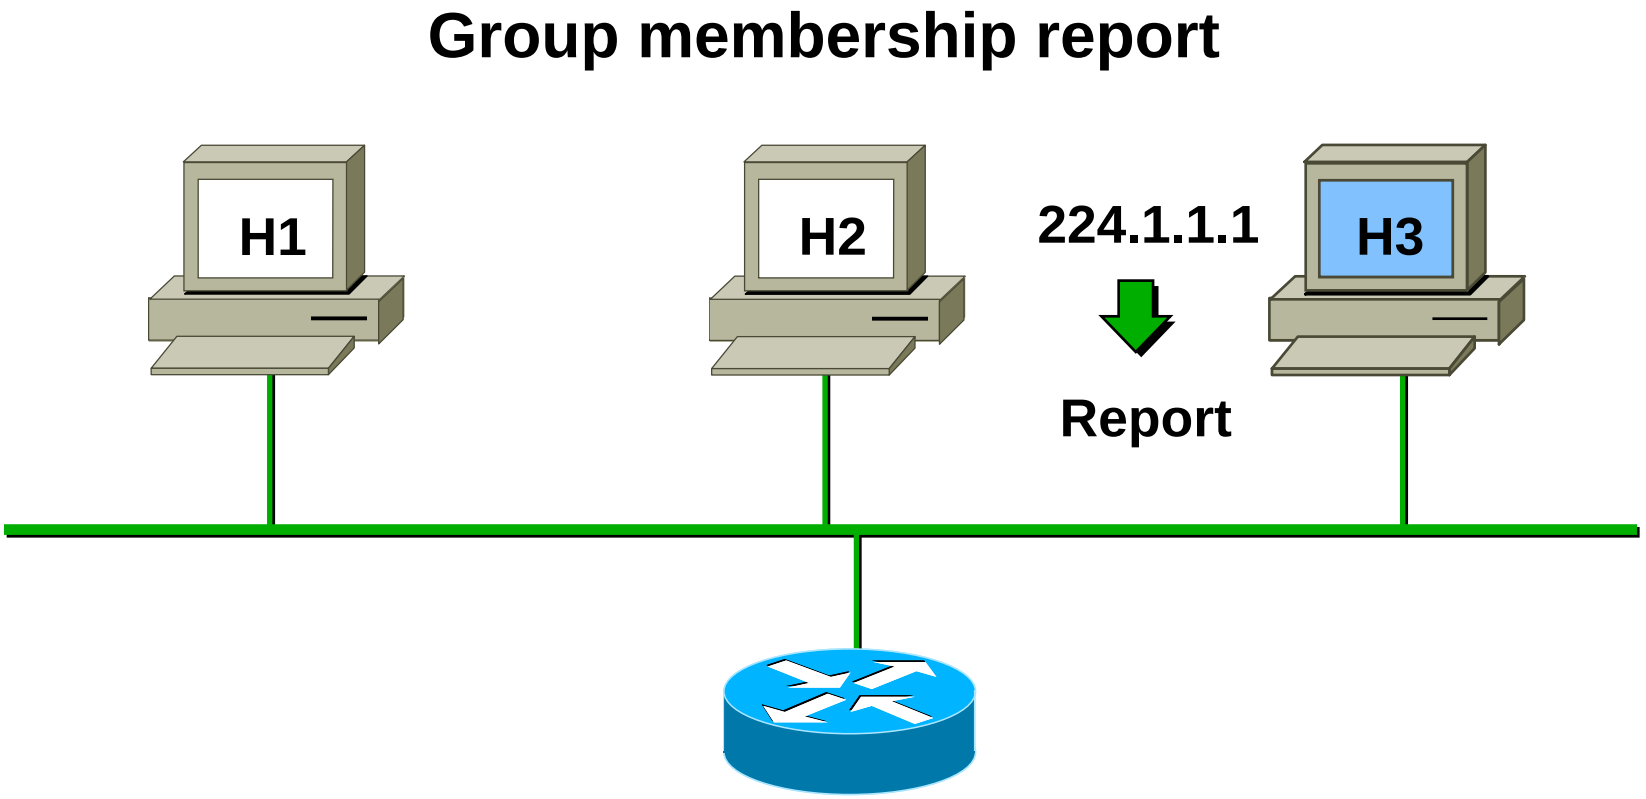

Members joining a group do not have to waited for a query to join.
They send in an unsolicited report indicating their interest

IGMP Host-Router Signalling

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

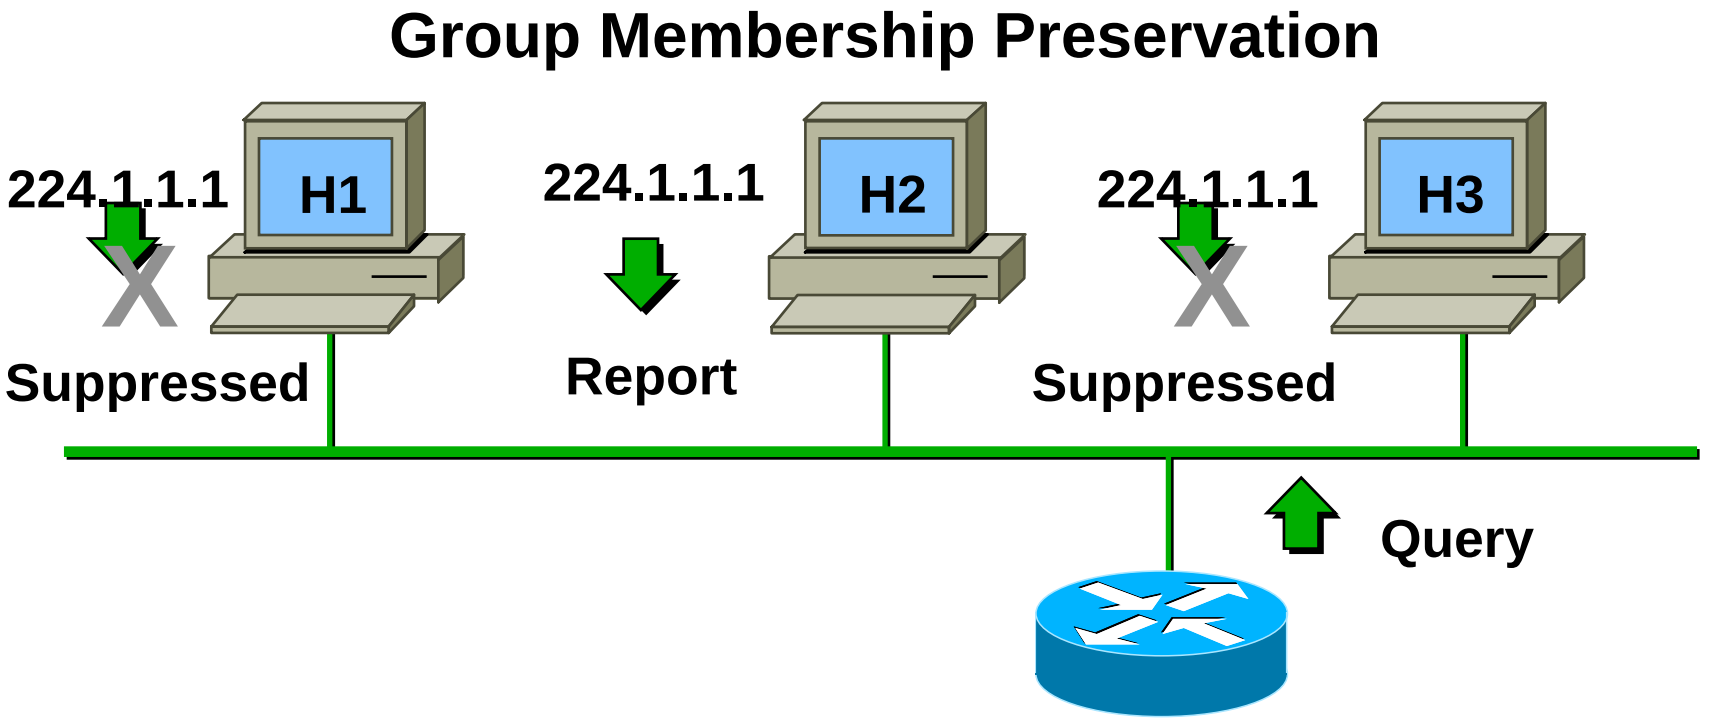

- Router sends periodic queries to 224.0.0.1
- One group member per subnet answers
- Others suppress answer

IGMP Host-Router Signalling

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

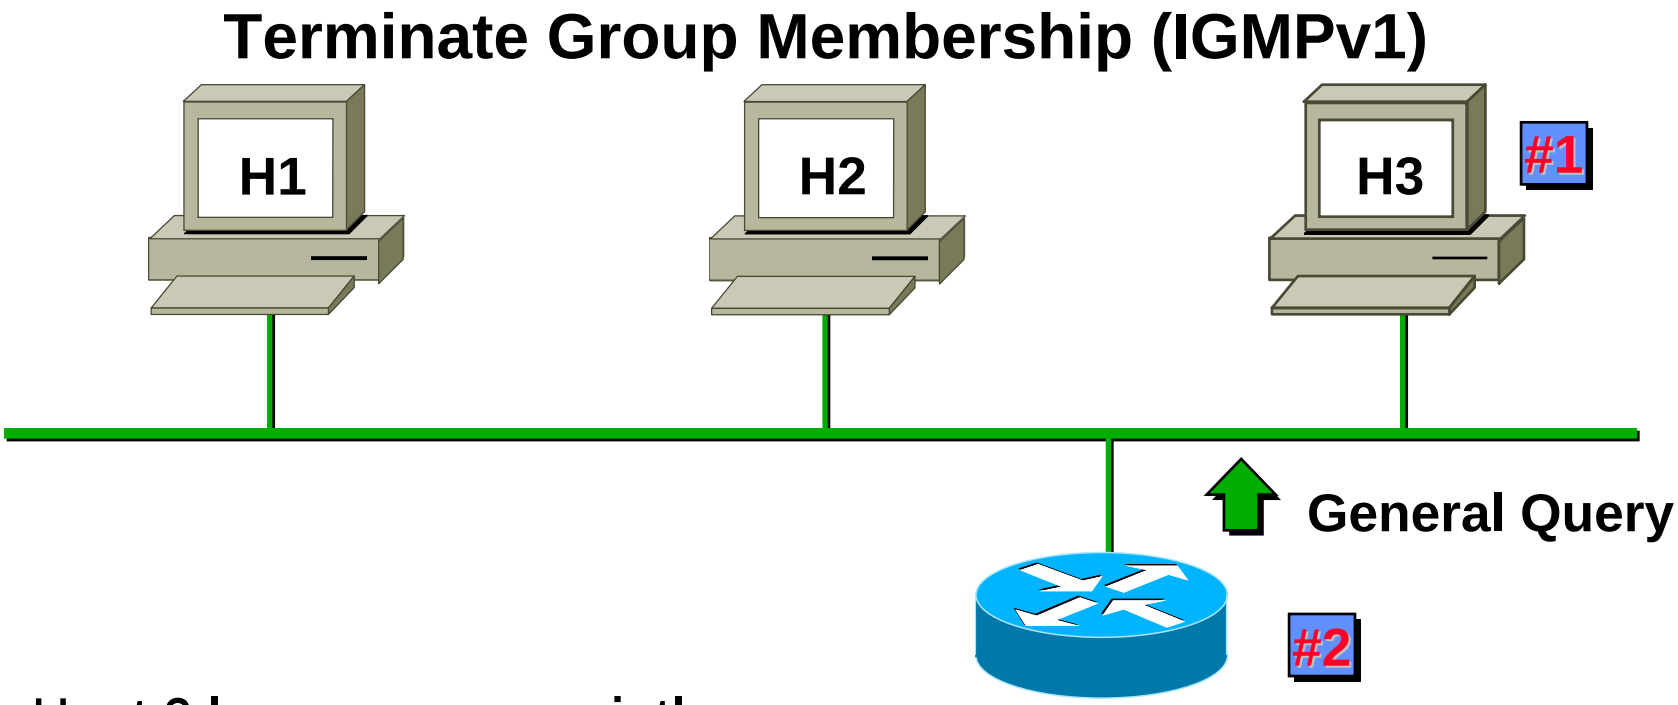

- Host 3 leaves group quietly
- Router queries remain unanswered
- Group terminate on timeout (up to 3 min)

IGMP Host-Router Signalling

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

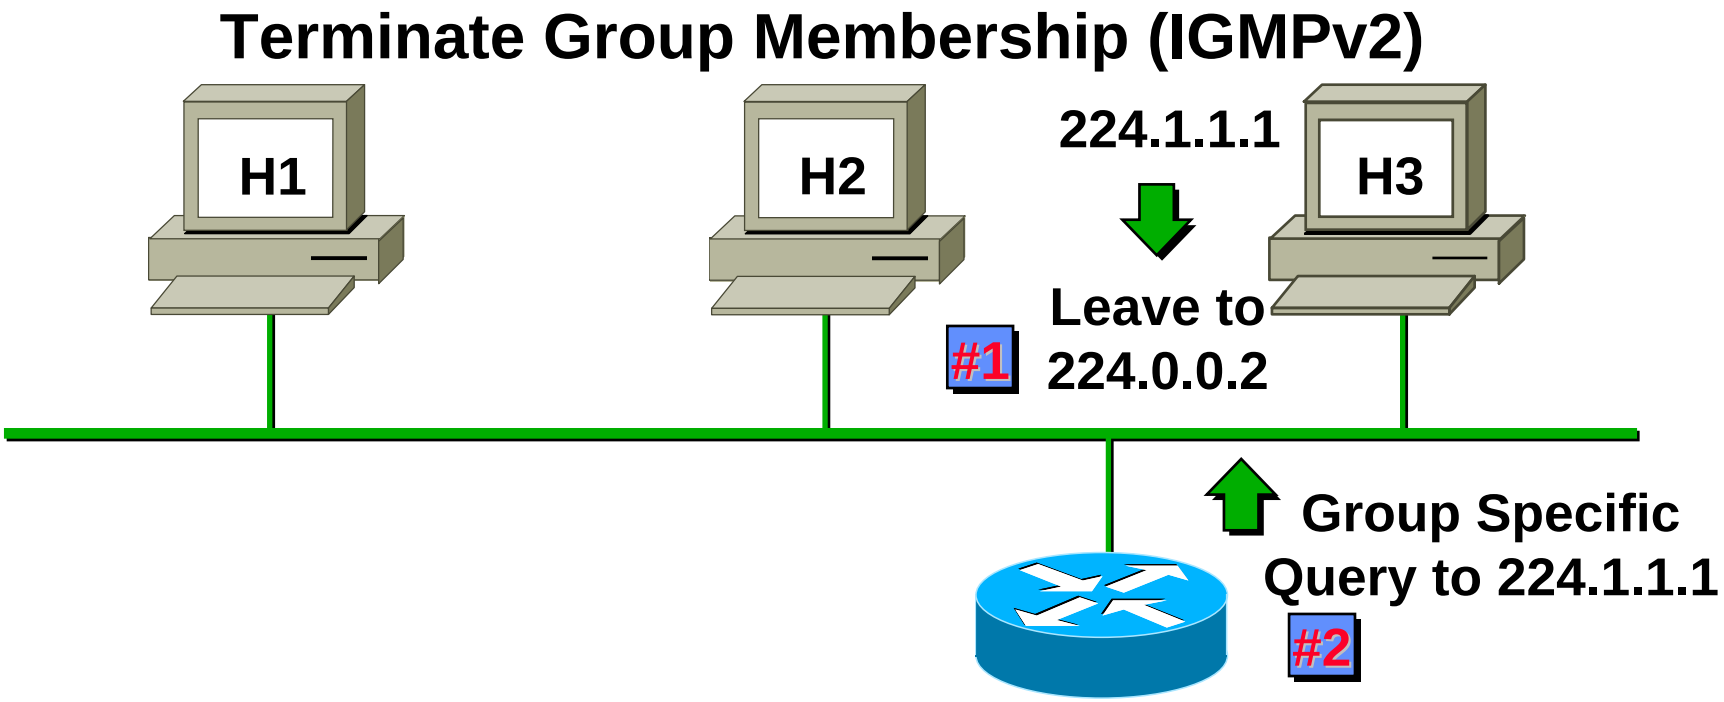

- Host sends Leave Message to 224.0.02
- Router sends group query to 224.1.1.1
- Timeout ~ 3 seconds for group 224.1.1.1

IGMP v3

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

Source = 1.1.1.1
Group = 224.1.1.1

Source = 2.2.2.2
Group = 224.1.1.1

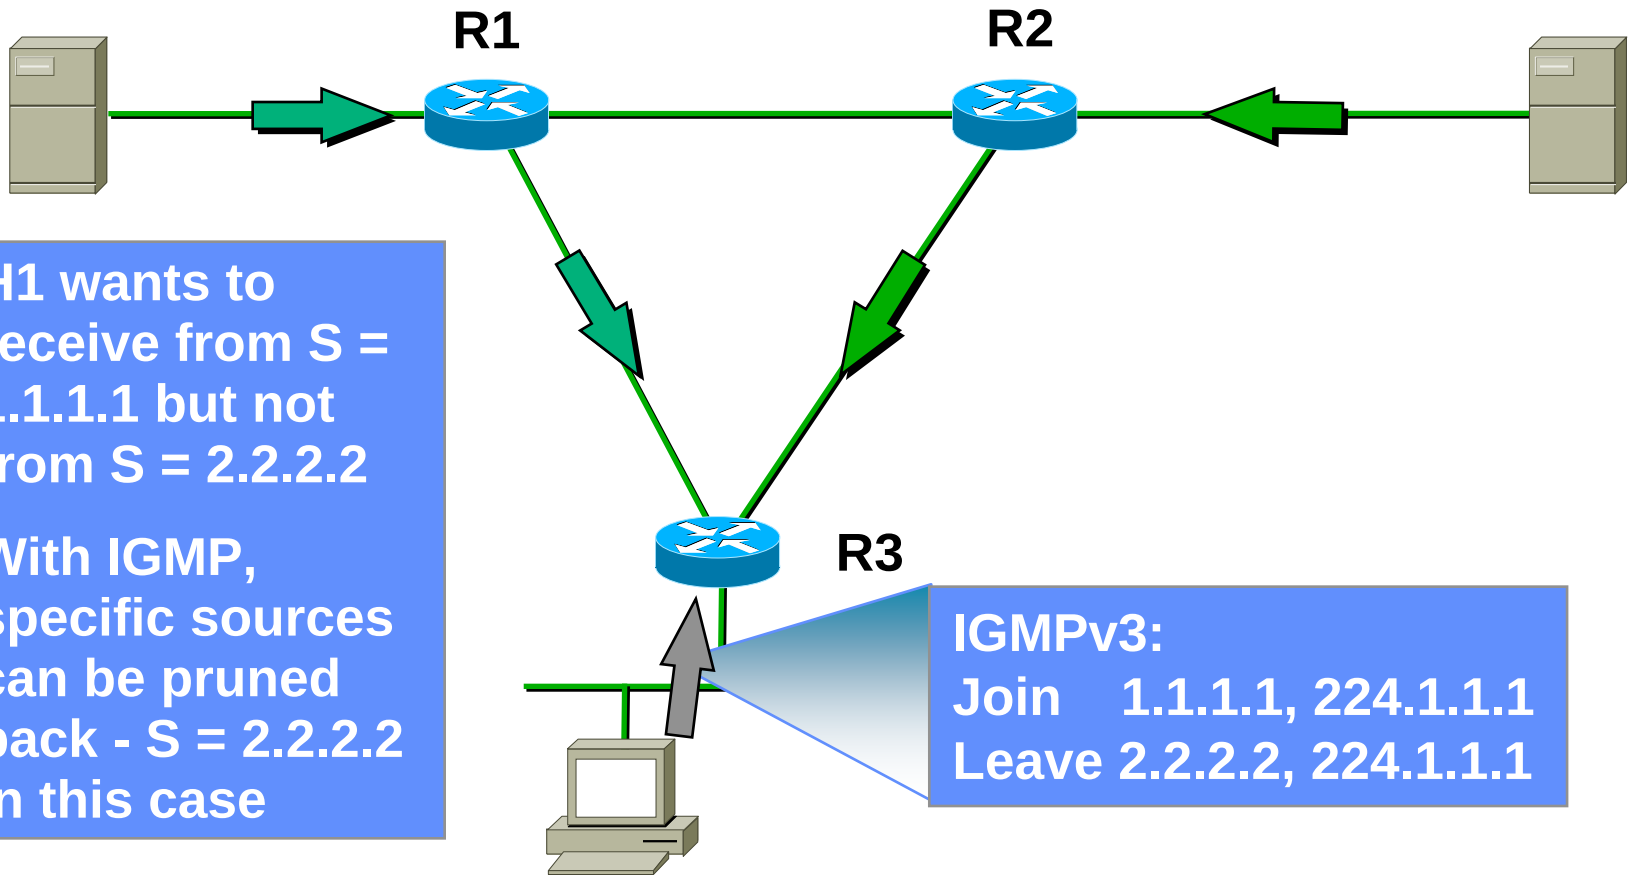

- H1 wants to receive from S = 1.1.1.1 but not from S = 2.2.2.2
- With IGMP, specific sources can be pruned back - S = 2.2.2.2 in this case

Limits of IGMP

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

IGMP Concept has no Group Directory

- Hosts not answering on Membership Queries remain unseen
- Closed groups impossible
- Undiscovered listener part of the concept

IGMP is relatively slow

- Time to reaction in the order of seconds
- Unsuitable for flow control or congestion avoidance
- Initiation or change of a non-local group tardy

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

API

Berkeley Sockets set/getsockopt():

- **IP_ADD_MEMBERSHIP** *to join a multicast group on a specific interface*
- **IP_DROP_MEMBERSHIP** *to leave a multicast group (no protocol action initiated with IGMP v1, but there is with IGMP v2)*
- **IP_MULTICAST_IF** *to set or get default interface for use with multicast sends*
- **IP_MULTICAST_LOOP** *to disable loopback of outgoing multicast datagrams*
- **IP_MULTICAST_TTL** *to set the IP time-to-live of outgoing multicast datagrams.*

API - Java

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

Package: java.net

Class MulticastSocket

with Methods

- *public void joinGroup(InetAddress mcastaddr)*
- *public void leaveGroup(InetAddress mcastaddr)*

IP Multicast in Java

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

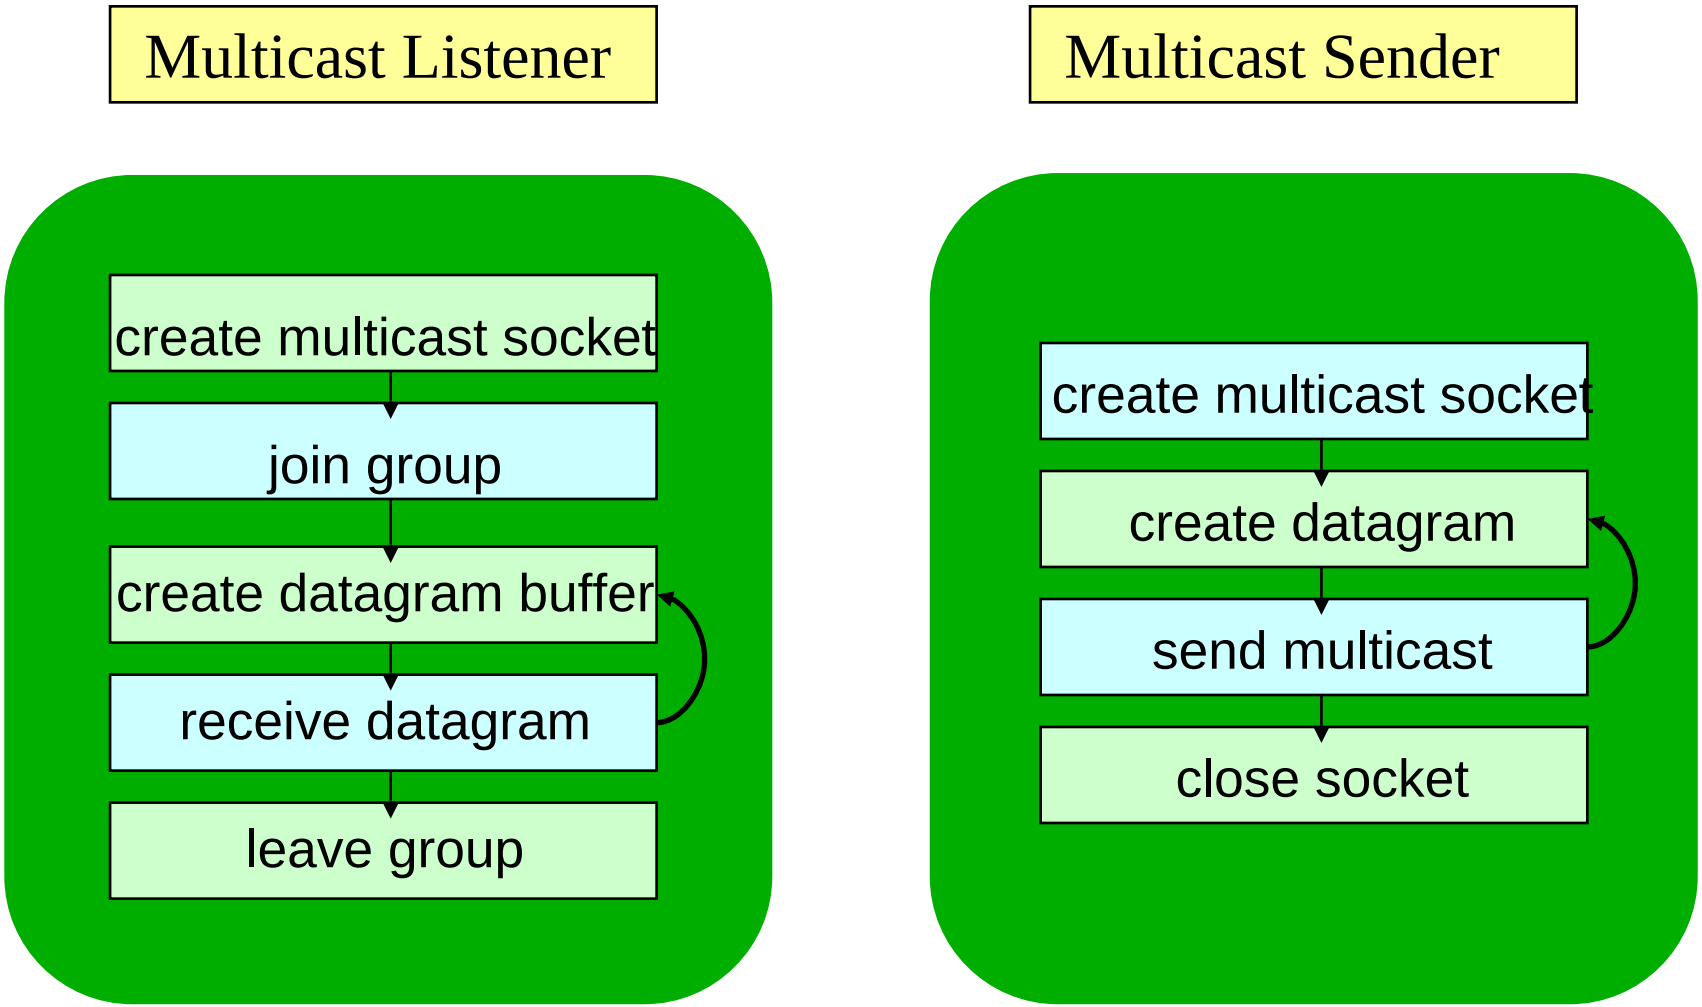

IP Multicast in Java

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

```
import java.net.*;
import java.io.*;
public class MulticastPeer{
   public static void main(String args[]){
   // args give message contents & destination multicast group
   // (e.g. "228.5.6.7")
      try {
         InetAddress group = InetAddress.getByName(args[1]);
         s = new MulticastSocket(6789);
         s.joinGroup(group);
         byte [] m = args[0].getBytes();
         DatagramPacket messageOut = new DatagramPacket(m, m.length, group, 6789);
         s.send(messageOut);
         // get messages from others in group
         byte[] buffer = new byte[1000];
         for(int i=0; i< 3; i++) {
            DatagramPacket messageIn = new DatagramPacket(buffer, buffer.length);
            s.receive(messageIn);
            System.out.println("Received:" + new String(messageIn.getData()));
            }
         s.leaveGroup(group);
         }catch (SocketException e){System.out.println("Socket: " +
e.getMessage());
         }catch (IOException e){System.out.println("IO: " + e.getMessage());}
      }  }   }
```

Reading

Thomas Schmidt
schmidt@informatik.
haw-hamburg.de

- R. Wittmann, M. Zitterbart: Multicast Communication, Morgan Kaufmann, 2001