Multicast and Quality of Service

Internet Technologies and Applications

Aims and Contents

- Aims
  - Introduce the multicast and the benefits it offers
  - Explain quality of service and basic techniques for delivering QoS

- Contents
  - Addressing and Delivery Mechanisms
  - IP Multicast
  - QoS
    - Service Differentiation (DiffServ)
    - Guaranteed QoS (IntServ)

Addressing and Delivery Mechanisms

Unicast

- One-to-one association between address and host
- The most common method of data delivery
- A host sends a packet to destination address X – the packet should be delivered to X (and only X will process the packet)
  - IP unicast: IP datagram sent to 125.70.16.3 is delivered only to the host with that address
  - MAC unicast: frame sent to 00:17:31:5A:E5:89 is processed only by the host with that address
- A host that receives a packet to which it is not the unicast destination will discard that packet

Broadcast

- One-to-many association between address and hosts
- Many = every host on the network
- Examples:
  - IP network broadcast address (all 1’s in host portion): a host with IP address 192.168.1.2/24 sends a datagram to the network broadcast address 192.168.1.255/24. All other hosts on the IP network will receive the datagram
  - MAC broadcast address (all 1’s): a host with MAC address 00:17:31:5A:E5:89 sends a frame to the broadcast address FF:FF:FF:FF:FF:FF. All other hosts on the network will receive the frame
- Applications:
  - Distributing control/management information, e.g. routing information, address requests
    - Used by many routing protocols, DHCP, ARP, …

Multicast

- One-to-many association between address and hosts
- Many = selected group of hosts
- A more general form of broadcast, where only a selected set of the hosts in the network receive the packet
- Requires management protocols for the hosts to “subscribe” or “join” the multicast destination group (much more complex than broadcast)
- Example:
  - IP multicast: a set of hosts on the Internet “subscribe” to the multicast group with address 225.70.8.20. When a host sends a datagram to 225.70.8.20, the datagram is delivered to all hosts subscribed to that multicast group
- Applications:
  - Multimedia and collaborative applications that involve many users
    - Audio/video/TV streaming, presentations, video/audio conferencing, shared document editing, ..

Anycast

- One-to-many association between address and hosts
- Many = any host in a group
- Similar to multicast, but a packet sent to an anycast address is delivered to only one (any one) of the subscribed hosts
- Applications:
  - Used in advanced DNS implementations: there are multiple, replicated copies of DNS servers in the world; a DNS query will be sent to an anycast address; anycast routing will deliver that query to one of the replicated DNS servers

Addressing and Delivery Mechanisms

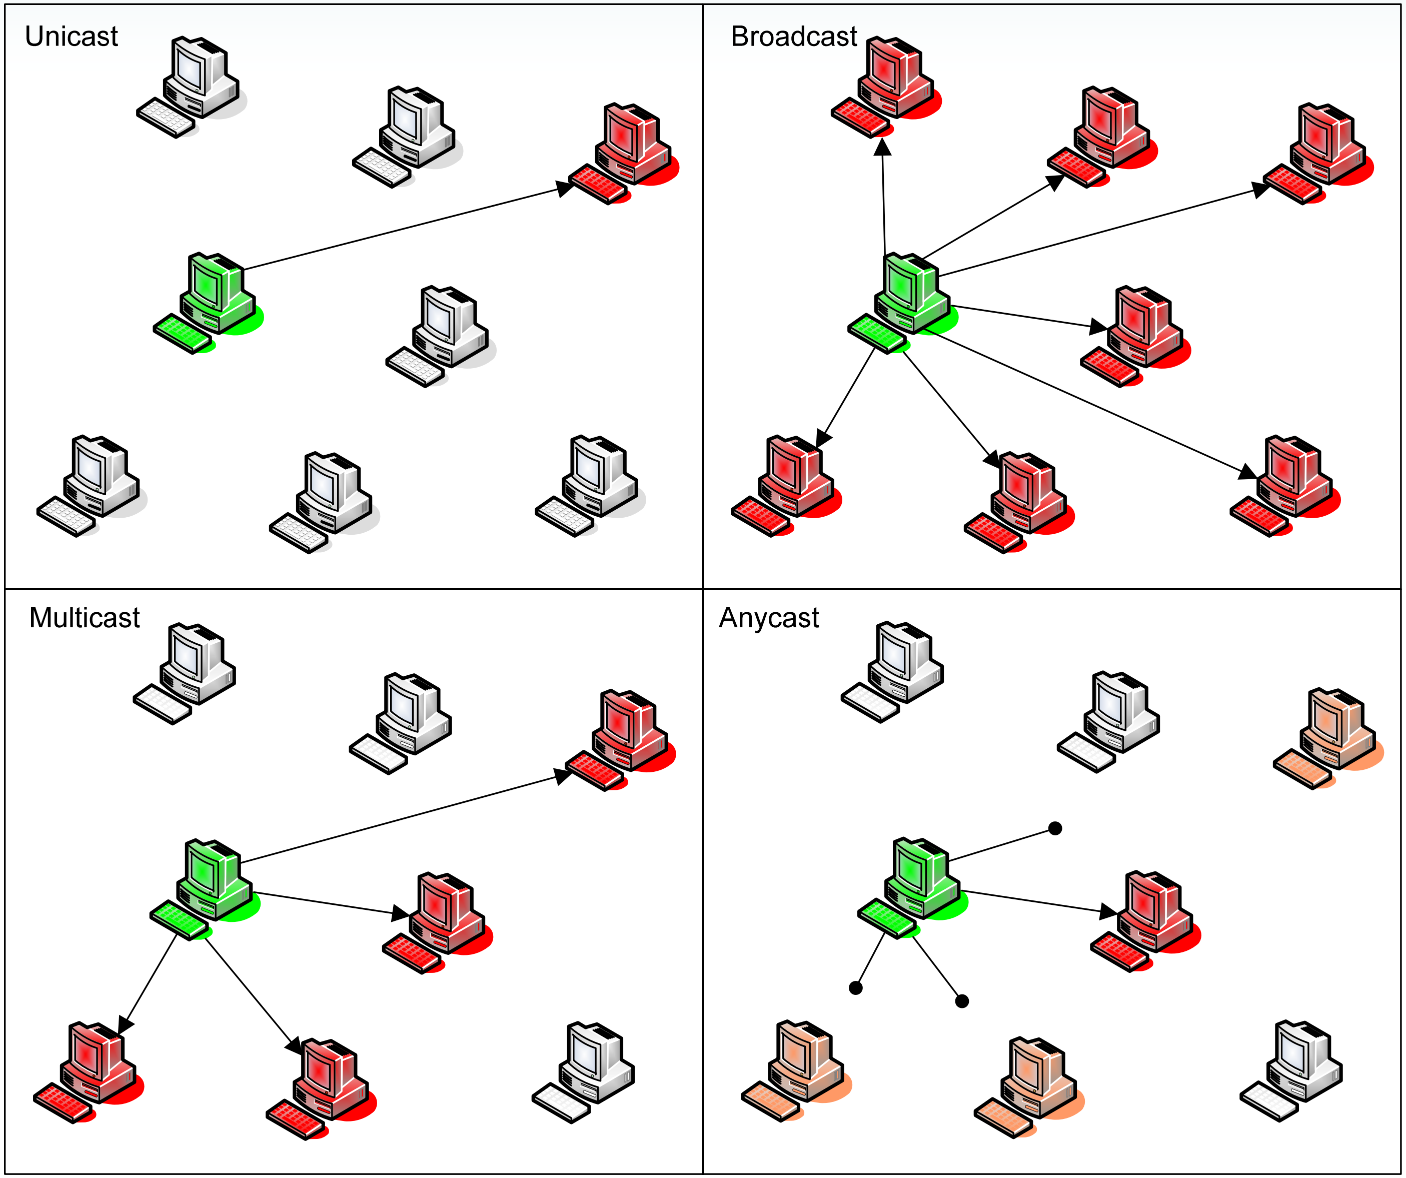

IP Multicast

Multicast Example

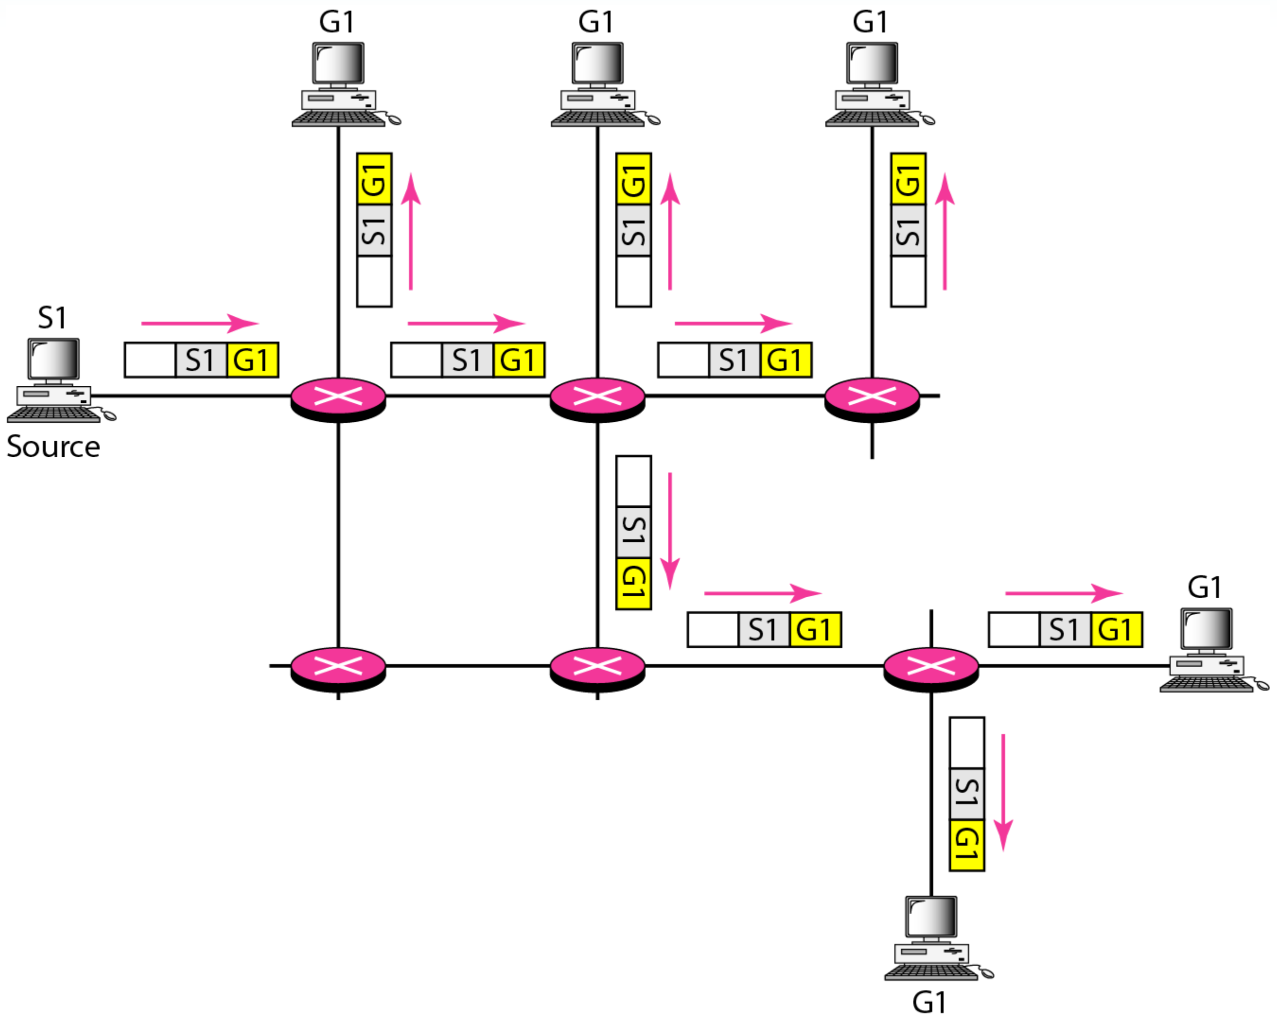

Multicasting vs Multiple Unicasting

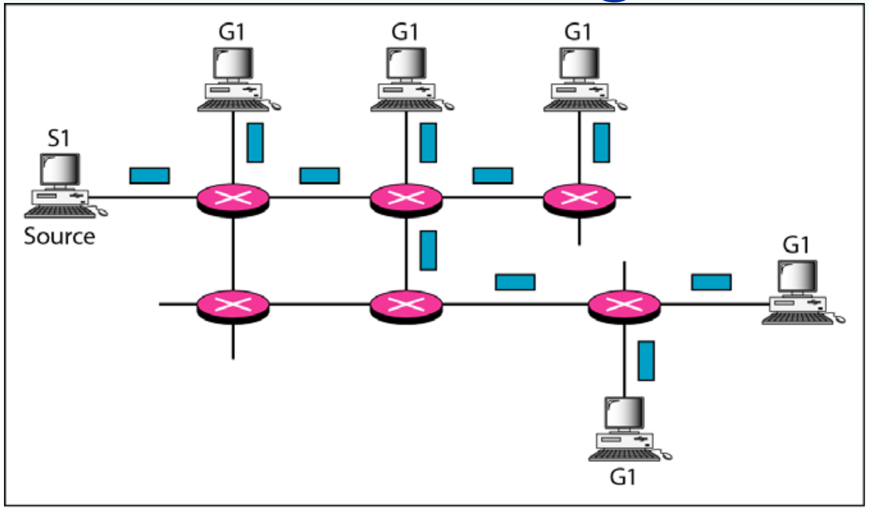

a. Multicasting

- Multicasting
  - A single packet is sent from source
  - Copies are made at multicast router where necessary

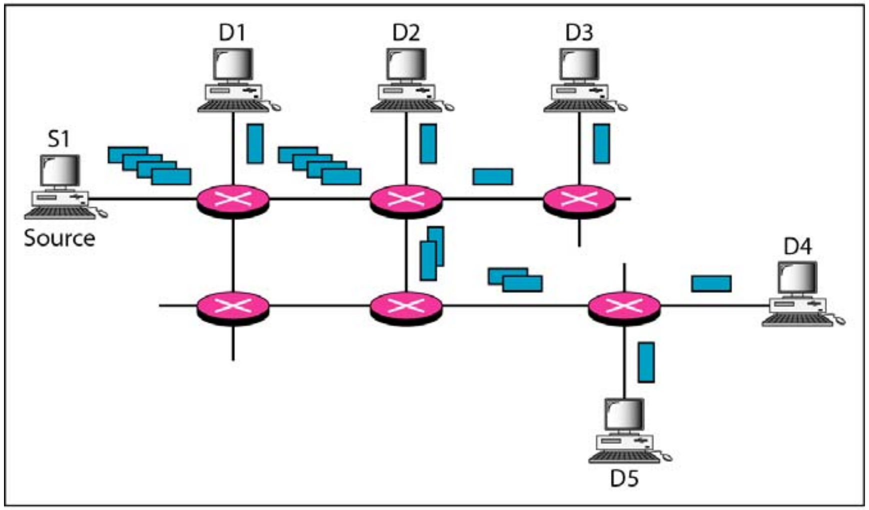

b. Multiple unicasting

- Multiple Unicasting
  - Use unicast to "emulate" multicast
  - Source uses unicast to send a copy of the same packet to each destination
  - Doesn't require routers/hosts to support multicast protocols
  - Very inefficient

IP Multicast Components

- IP Multicast is an optional feature of IP hosts and routers
  - Most hosts and routers *support* multicast today, but many ISPs have *not enabled* multicast

- When a group of computers want to communicate, they use a single IP multicast address
  - 32 bit IP addresses include a range for multicast addresses
- Routers must be notified as to which computers belong to which multicast group
  - Internet Group Management Protocol (IGMP)
- Routers must be able to delivery multicast datagrams in the correct direction
  - Forwarding principles are more complex than IP unicast forwarding
  - Multicast routing protocols such as DVMRP, PIM, MOSPF, PGM, …

IP Multicast Addresses

- IP multicast addresses range from:
  - 224.0.0.0 to 239.255.255.255 (there are some special cases)
- A single IP multicast address identifies a destination group of computers
  - Multicast address can only be a destination; never a source

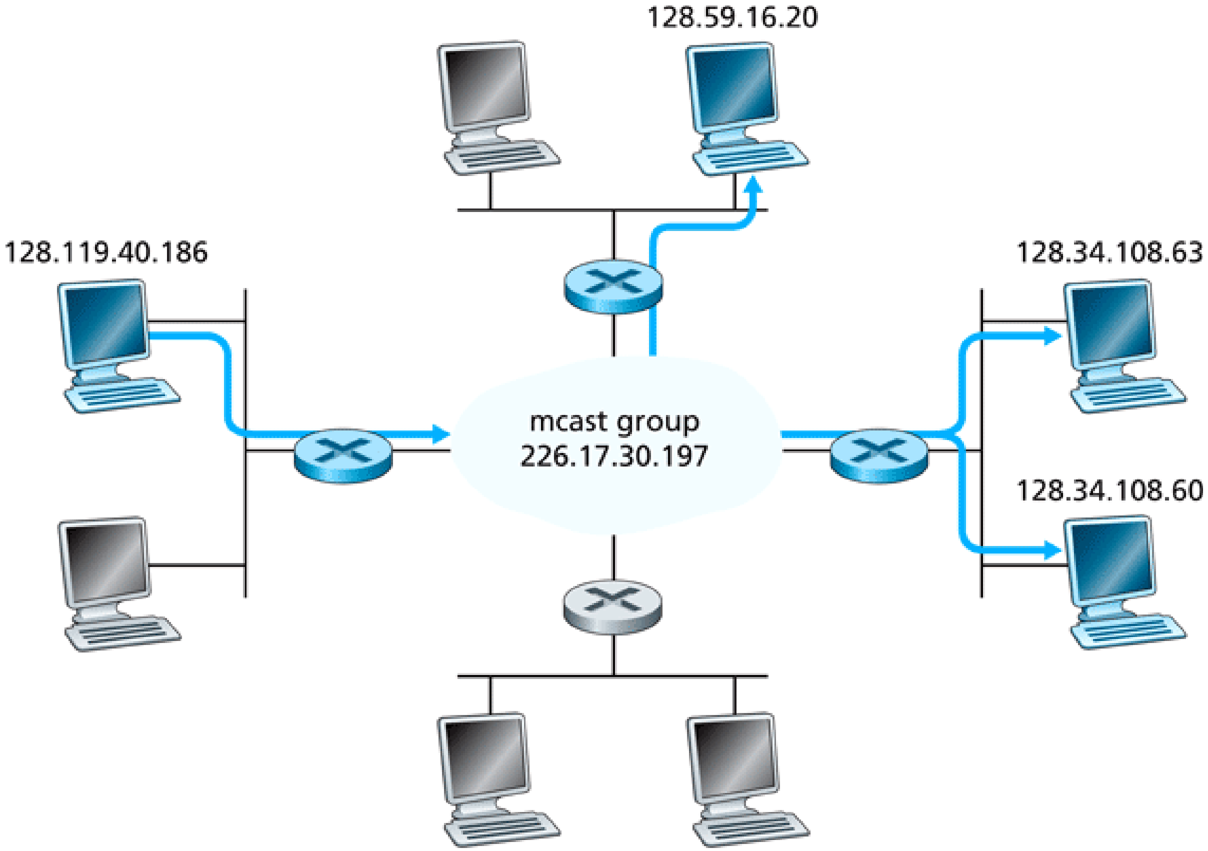

IGMP

- Any host can send a datagram to an IP multicast address
  - (Although in practice, security policies at routers may prevent this)
- The hosts registered with that address will receive the datagram
  - Those hosts are in the same multicast group
- IGMP allows hosts to join and leave a group
  - When a host wants to join a group it sends a IGMP Join message to the group multicast address
  - Local multicast routers receive the request and propagate the information to other multicast routers in the Internet
  - Local multicast routers periodically send messages to group members to confirm if the hosts are still part of the group
- With multicast routers managing group membership, a source does not need to know all members of the group

IP Multicast Forwarding and Routing

- When a host sends a datagram to a IP multicast address:
  - Sent to local multicast router
    - E.g. using Ethernet multicast or broadcast
  - The local multicast router must forward to the appropriate next router(s)
- Multicast routing protocol is used by multicast routers to determine the next router(s) to reach all destination
  - Several routing protocols in use: DVMRP, PIM, MOSPF, PGM, …
  - Trees are used to define the least-cost path from a source to all destinations in a group
  - A multicast router may need to send a copy of a datagram to multiple next routers
- IP Multicast forwarding and routing is much harder than unicast IP
  - E.g. least-cost routes may change rapidly (as hosts join/leave groups)
  - E.g. multicast delivery is often unreliable (cannot send ACKs back)

Quality of Service (QoS)

The Internet: Best Effort Service

- IP is connectionless protocol using packet switching
  - IP provides unreliable delivery
  - TCP adds reliability but still no guarantees on performance
- Some applications/users desire performance guarantees, e.g.
  - A user watching IPTV requires 6Mb/s sustained throughput
  - A VoIP call requires delay less than 100ms and jitter less than 5ms
  - A medical file transfer requires packet loss rate of less than 0.5%
  - A business requires 2Mb/s minimum and 3Mb/s average data rate
- How do we provide such guarantees in an IP network?

Quality of Service (QoS)

- QoS refers to statistical performance guarantees that a network system can make: Packet loss, delay, jitter, throughput are common measures
- There are numerous QoS mechanisms available
- QoS is only an issue with network utilisation is high
- Hard QoS (or guaranteed QoS):
  - Absolute performance guarantees
  - A user requests 1Mb/s from a network operator, then they will be guaranteed 1Mb/s
  - Reserve resources for users; control the number of users/traffic entering a network
- Soft QoS (or service differentiation):
  - Relative performance guarantees
  - A user requests 1Mb/s from a network operator; they will be delivered 1Mb/s when possible, but not guaranteed
  - Provide priority to users/traffic
- No QoS:
  - All users/traffic treated the same; no guarantees or prioritisation

Service Differentiation (Soft QoS)

- Give different priority to different types of traffic in the network
- May be applied at different levels of granularity:
  - Flow or Session level:
    - My VoIP call gets higher priority than your FTP download
  - Class (or Aggregate of Flows) level:
    - All VoIP calls get higher priority than all FTP downloads
    - All voice traffic gets higher priority than non-voice traffic
- In the Internet, DiffServ is the name of an architecture that provides service differentiation
  - Applied on the Class level
  - E.g. ISP treats all packets marked as voice with certain priority

Service Differentiation Example

- Simple network with two applications
  - Assume the LANs are 100Mb/s
  - The bottleneck link is router R1 to R2: 1.5Mb/s
- If sending rate of H1 and H2 exceeds 1.5Mb/s, then packet delay and packet loss will occur at R1

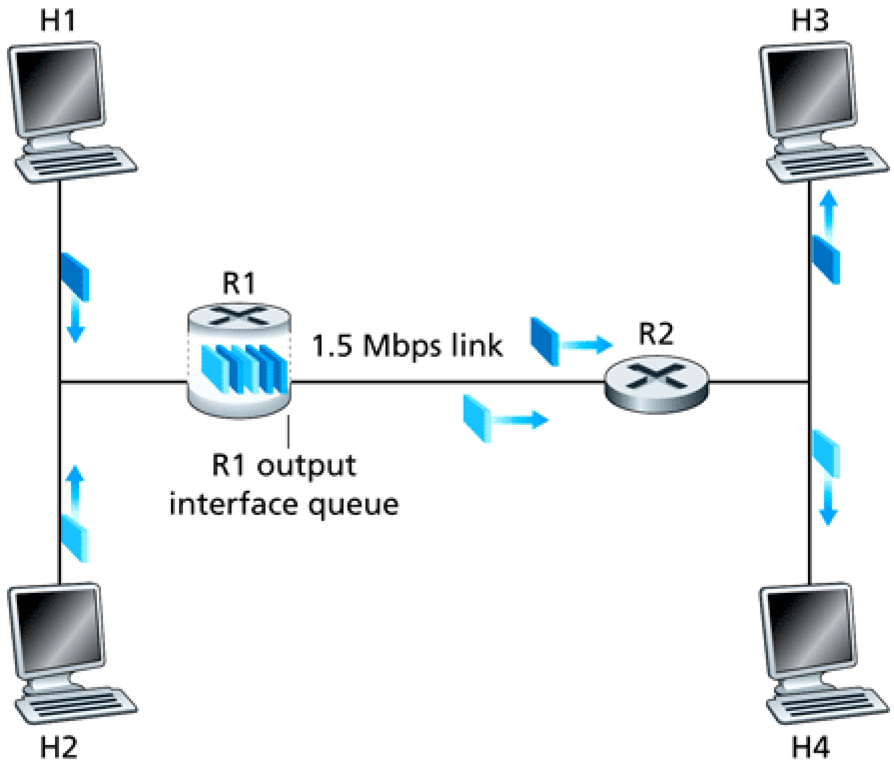

Service Differentiation: Marking

- Two applications:
  - 1Mb/s video application: requires low delay and packet loss
  - FTP file transfer: no strict time constraints
- How can we give priority to video application?
  - Queue at router R1 can treat video packets with priority
    - Whenever a video packet arrives, it is sent, even if there are FTP packets waiting in the queue
  - Need Packet Marking: IP packets must be marked such that routers can recognise the class of service they should be allocated
    - In IPv4, the ToS or DiffServ (DSCP) field in the header

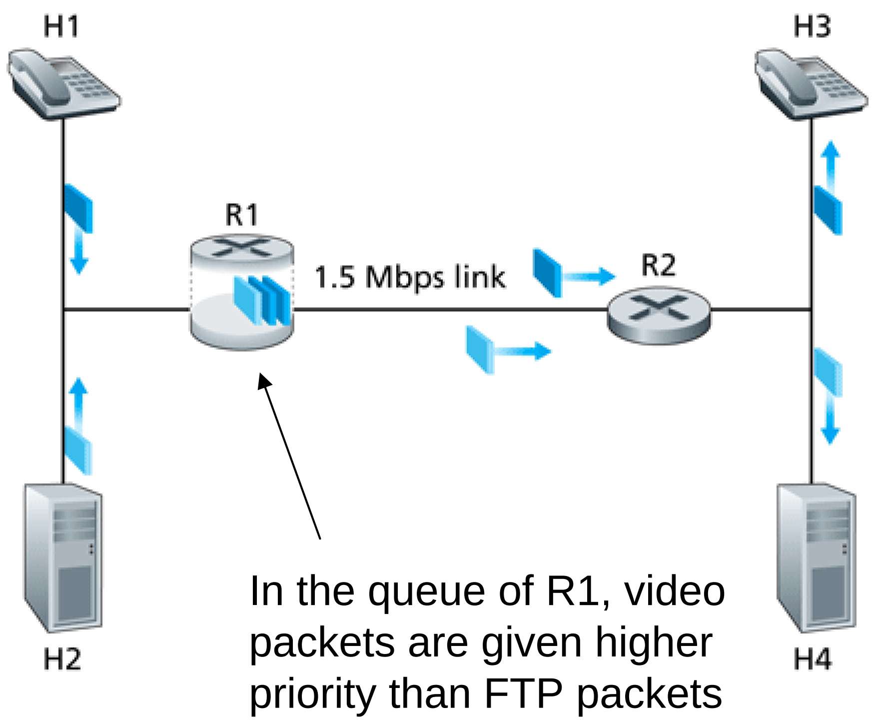

Service Differentiation: Classification

- What if FTP application belongs to a “premium” user?
  - User paid ISP more money than the video user
  - Now may want to give priority to FTP packets

- Packet Classification
  - Classify packets based on some criteria
    - Common criteria:
      - Class (e.g. Control, Video, Voice, Data, Background): DSCP in IP header
      - Source/Destination address (identify the user): IP addresses and port numbers
  - Packets are marked by a source host or some entry point into a network (e.g. ISP border router)
  - Routers determine the priority to give packets based on these markings
    - Achieved using queuing schemes: FIFO, Priority Queues, Round Robin, Weighted Fair Queuing (WFQ), …
  - The priority is a policy decision, e.g. made by ISP

Service Differentiation: Policing

- Assume the router knows it should give priority to the 1Mb/s video application
- What if the video application sends at a higher rate than agreed?
  - Error in the host, or malicious behaviour
- Traffic Metering and Policing is required
  - Routers measure the characteristics of arriving traffic
    - Average rate, peak rate, burst size
  - If application traffic exceeds some agreed values, routers may drop or delay packets
    - Achieved using Leaky Bucket

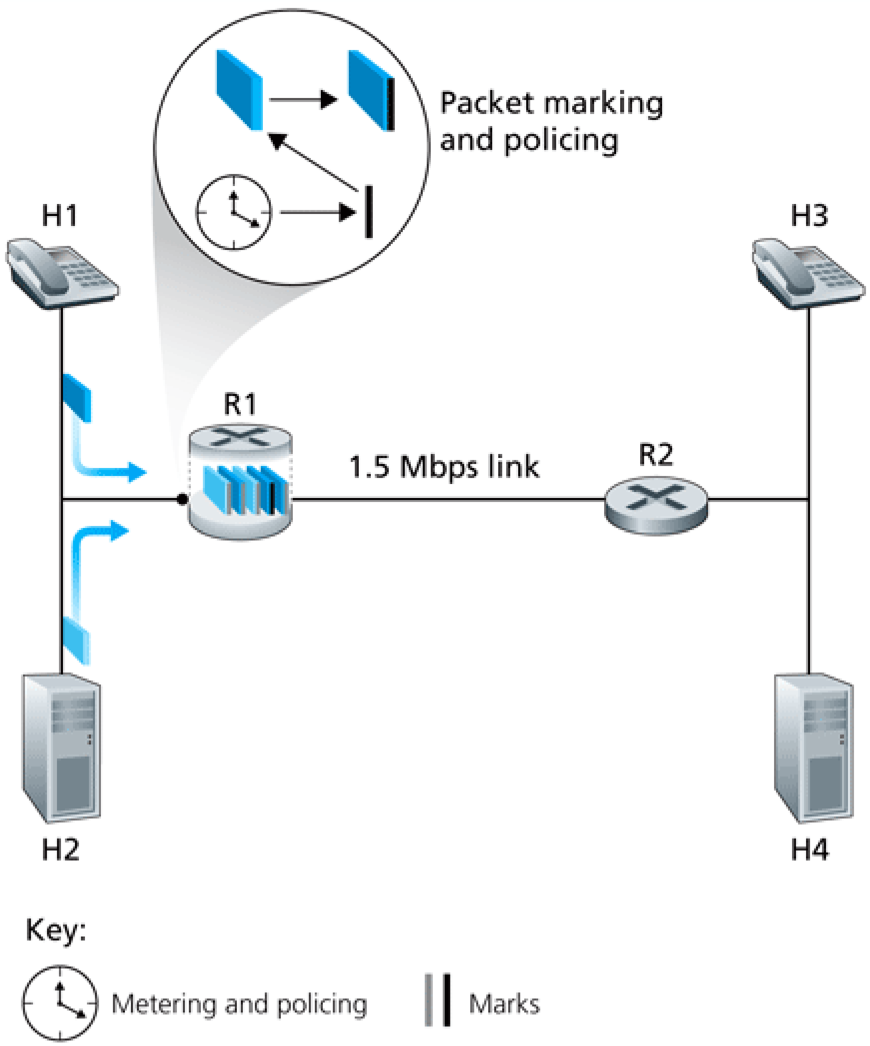

Limitation of Service Differentiation

- Service differentiation on its own canNOT deliver guaranteed performance
  - Two 1Mb/s video applications
  - Both given equal (high) priority, but the full 1Mb/s cannot be delivered to either
    - Packet delays and loss will occur at R1

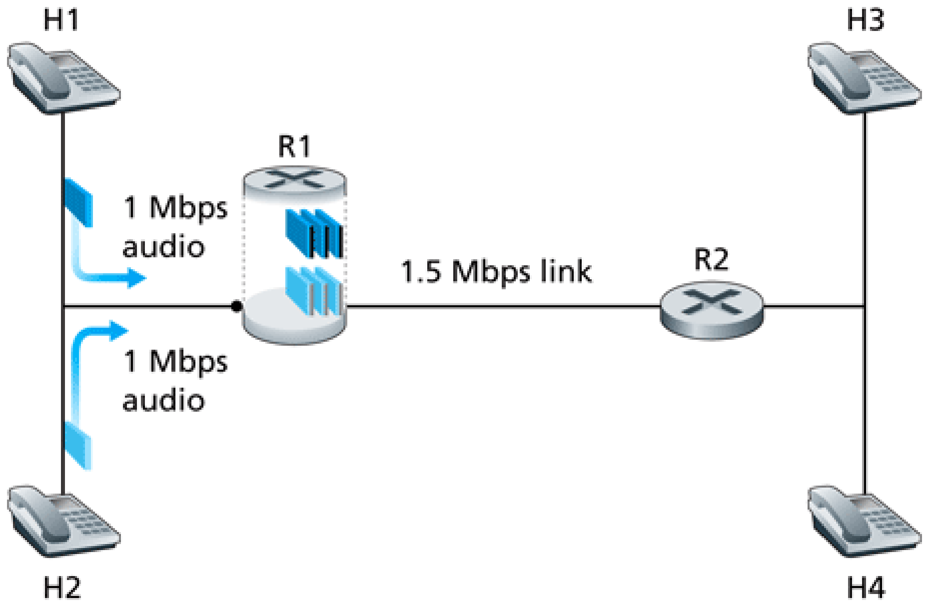

Guaranteed (Hard) QoS

- Resource Reservation:
  - Only way to guarantee that an application/user will have sufficient resources to reserve the resources (e.g. link bandwidth, queue space, processing time)
    - Cannot over-reserve
  - An application needs to be able to state the required resources
  - Resource Reservation Protocol (RSVP) is standardised for the Internet
- Admission Control
  - When application requests resources, the network (routers) must make a decision to accept or reject
    - Can only accept if sufficient resources to process the traffic

- In the Internet, IntServ is the name of the architecture that provides guaranteed QoS
  - Applied at the flow level
  - E.g. applications use RSVP to reserve resources along a path in the Internet

Resource Reservation Example

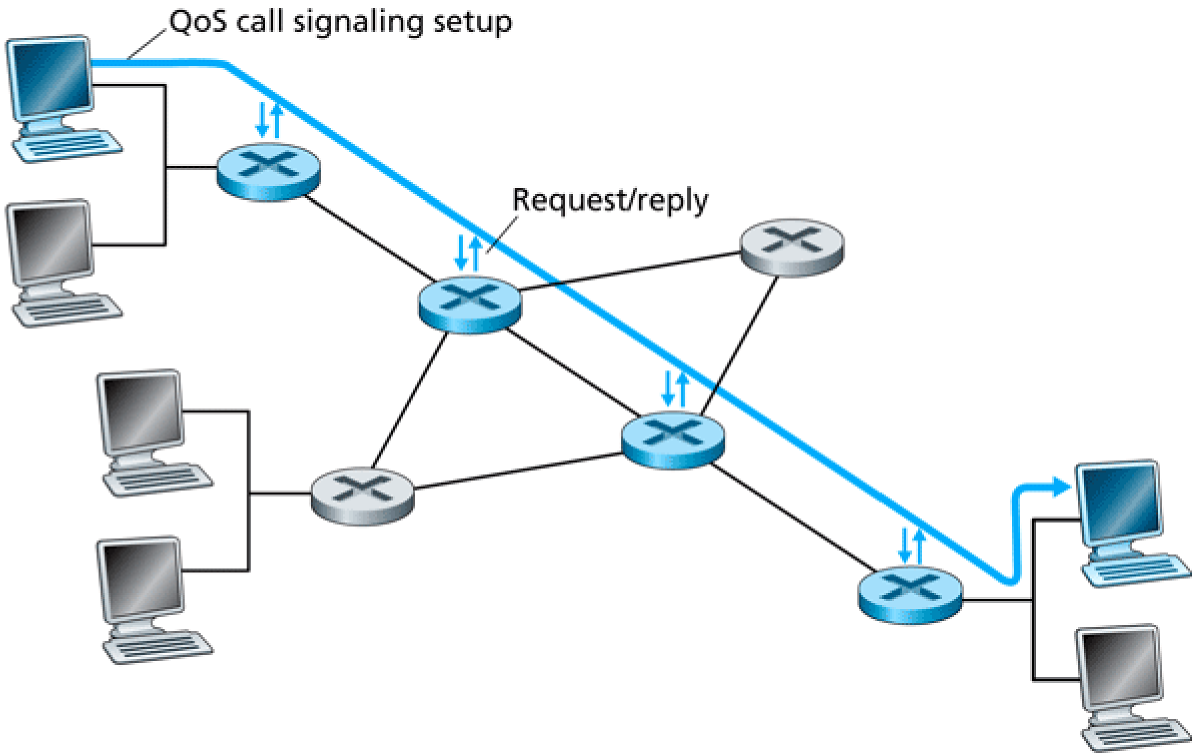

Soft vs Hard QoS

- Service Differentiation (including DiffServ)
  - Relative priorities given the class of traffic
  - Simpler to implement in large network
    - Do not need complex signalling, routers only prioritise amongst classes of traffic (not amongst each individual application flow)
  - No guaranteed QoS
- Guaranteed (Hard) QoS (including IntServ)
  - Absolute performance guaranteed to application flows
  - Same principles of traditional telephone networks
  - Complex (RSVP, routers) and requires applications to state their resource requirements
- Today, neither approach is widely deployed across entire Internet
  - Some ISPs use these approaches within their networks
  - There are combinations and variations (MPLS, traffic engineering)
  - End users do not have access to QoS mechanisms yet