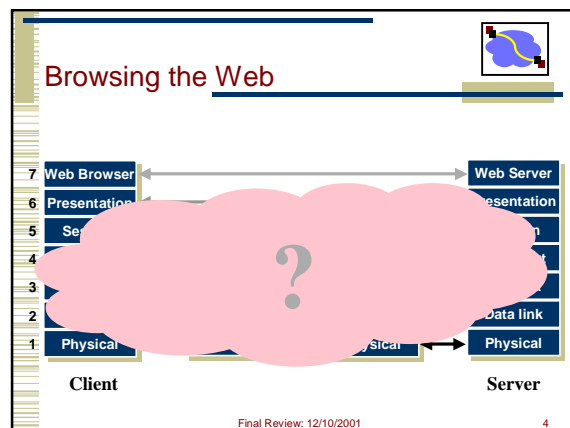
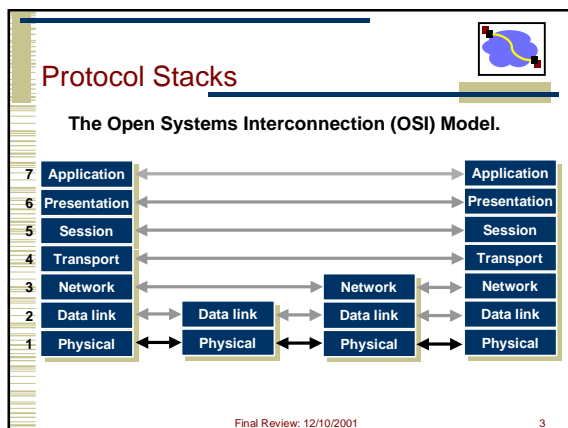


Course Review

Outline

- Networks: A top down view (for a change).
- Other topics.
 - Security
 - QoS
 - Multicast
- Questions?

Final Review: 12/10/2001 2



HTTP Request Example

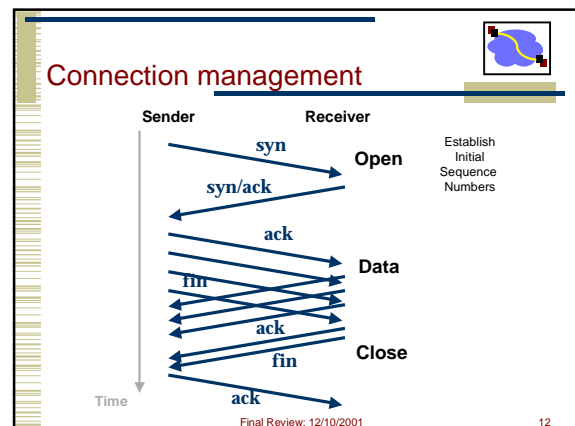
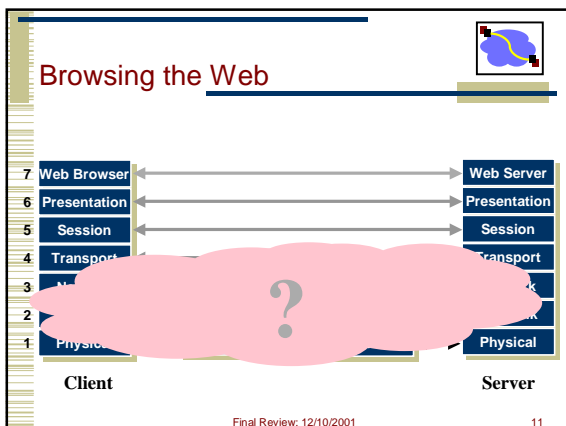
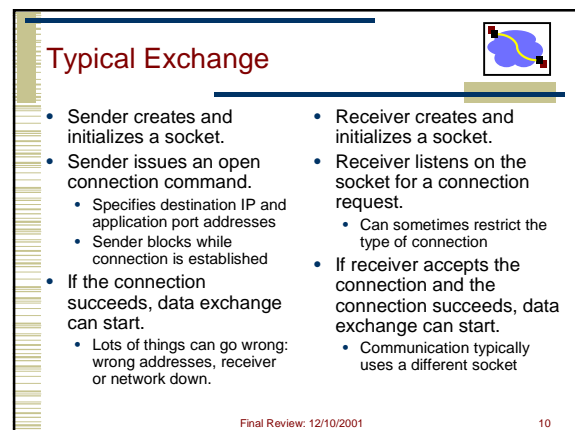
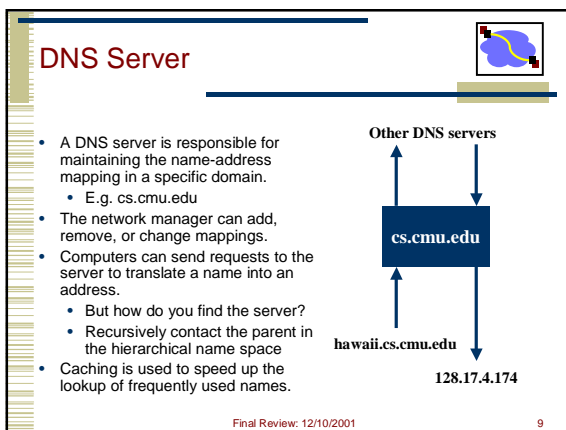
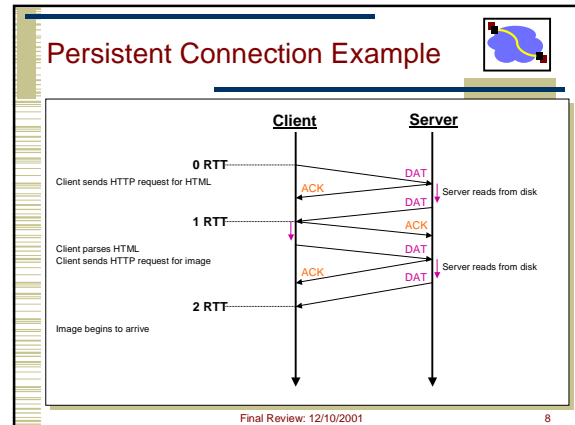
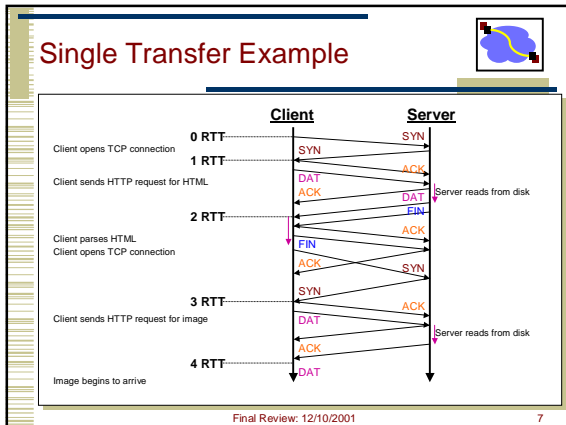
```
GET / HTTP/1.1
Accept: */*
Accept-Language: en-us
Accept-Encoding: gzip, deflate
User-Agent: Mozilla/4.0 (compatible; MSIE 5.5;
Windows NT 5.0)
Host: www.seshan.org
Connection: Keep-Alive
```

Final Review: 12/10/2001 5

HTTP Response Example

```
HTTP/1.1 200 OK
Date: Tue, 27 Mar 2001 03:49:38 GMT
Server: Apache/1.3.14 (Unix) (Red-Hat/Linux) mod_ssl/2.7.1
OpenSSL/0.9.5a DAV/1.0.2 PHP/4.0.1pl2 mod_perl/1.24
Last-Modified: Mon, 29 Jan 2001 17:54:18 GMT
ETag: "7a11f-10ed-3a75ae4a"
Accept-Ranges: bytes
Content-Length: 4333
Keep-Alive: timeout=15, max=100
Connection: Keep-Alive
Content-Type: text/html
.....
```

Final Review: 12/10/2001 6



Reliability

- Checksum guarantees end-end data integrity.
- Sequence numbers detect packet sequencing problems:
 - duplicate: ignore
 - reordered: reorder or drop
 - lost: retransmit
- Lost packets detected by sender.
 - uses time out to detect lack of acknowledgment
 - requires reliable roundtrip time estimate
- Retransmission requires that sender keeps copy of the data until ACK is received.
 - performance issue

Final Review: 12/10/2001

13

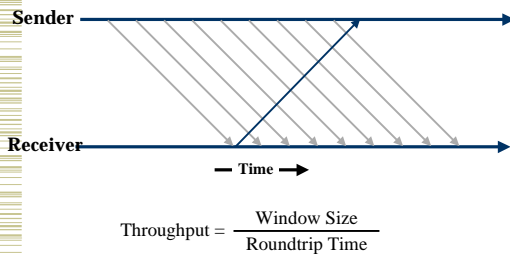
When to Send a Packet?

- End-to-end flow control.
 - avoid buffer overflow on receiver
 - receiver advertizes a window size
- Congestion control.
 - estimates amount of data that can be in network
 - implemented using the congestion window, slow start, and fast retransmit/recovery mechanisms
- Efficiency considerations.
 - try to send large packets (if possible)
 - more efficient in the network and on end points
 - piggybacking of acks

Final Review: 12/10/2001

14

Window Size versus Throughput



Final Review: 12/10/2001

15

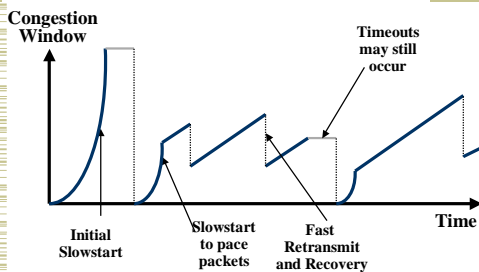
TCP Congestion Avoidance

- Congestion avoidance limits how fast TCP can send data.
 - Implemented using a congestion window that limits how much data can be in the network
 - independent from flow control window
 - transmission is limited by minimum of the two windows
 - window grows in response to acknowledgement
- Packet loss is seen as sign of congestion.
 - multiplicative decrease of the congestion window
 - have to cut back fast since cost of congestion is high
- How do you detect when more bandwidth becomes available?
 - gradually increment congestion window (probing)
 - results in oscillation around congestion window size!

Final Review: 12/10/2001

16

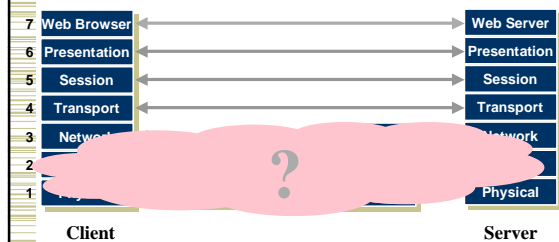
TCP Saw Tooth Behavior



Final Review: 12/10/2001

17

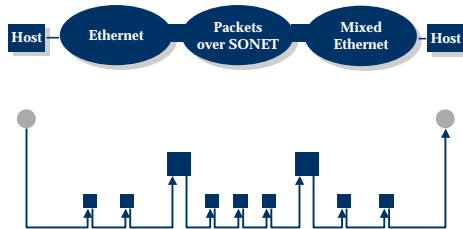
Browsing the Web



Final Review: 12/10/2001

18

Hop-by-Hop Packet Forwarding in the Internet



Final Review: 12/10/2001

19

Addressing in IP v4 (Basic)

- Each host has an Internet address.
- Addresses are hierarchical.
 - address contains hint about location
- Address space is divided in three classes of point-to-point addresses, multicast addresses, and some special addresses.

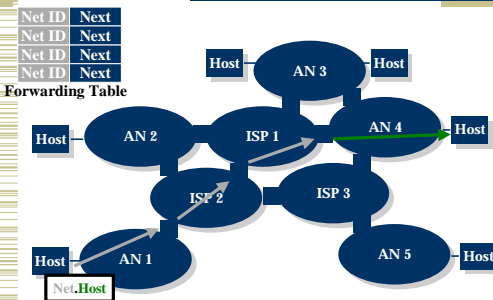
	type	network	host
A	1	7	24
B	2	14	16
C	3	21	8
D	4	(multicast)	28

Example: 128.2.209.19

Final Review: 12/10/2001

20

Routing based on Network Identifier



Final Review: 12/10/2001

21

Problems with Simple Address Structure

- Running out of addresses.
 - Especially true for mid-sized networks
- Routing tables are becoming too big.
 - 100 of thousands of entries
- Temporary solution: classless inter-domain routing.
 - Use address space more efficiently by relaxing the strict address structure,
 - length of network address is variable
 - generalization of subnetting idea
 - have internet service providers hand out blocks of addresses to their customers

Final Review: 12/10/2001

22

Route Lookup with CIDR

- Problem: with CIDR there can be multiple matches when looking up an address.
 - Can for example happen when a customer switches ISPs but keeps addresses
- Solution: lookup is based on longest prefix match.
 - If there are multiple matches in the lookup, the longest match (longest netmask) wins

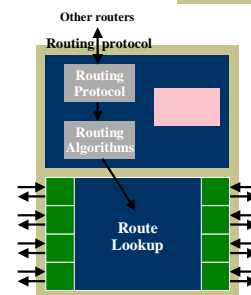
10110110	hosts
10110110 010	hosts
10110110 010 0100011	

Final Review: 12/10/2001

23

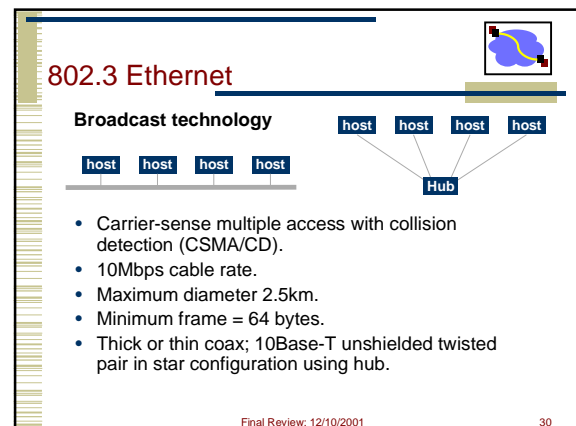
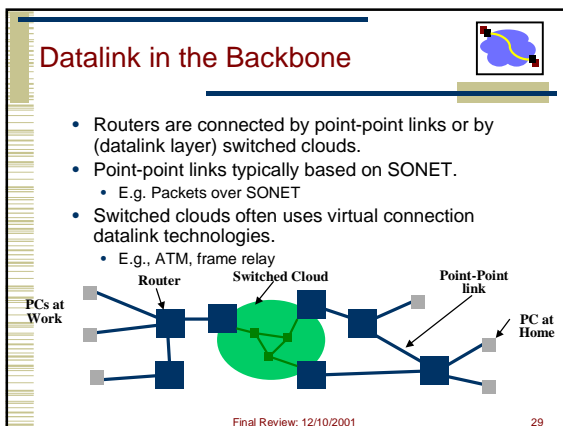
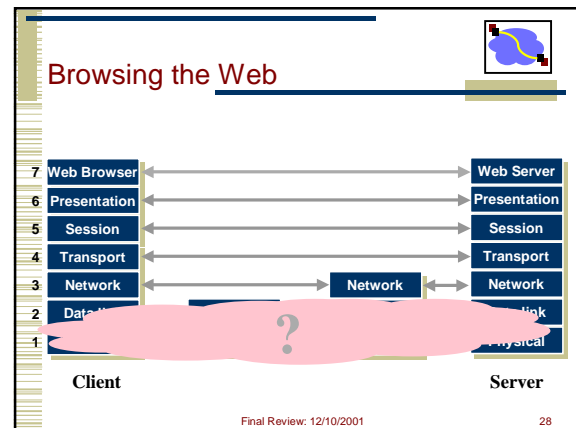
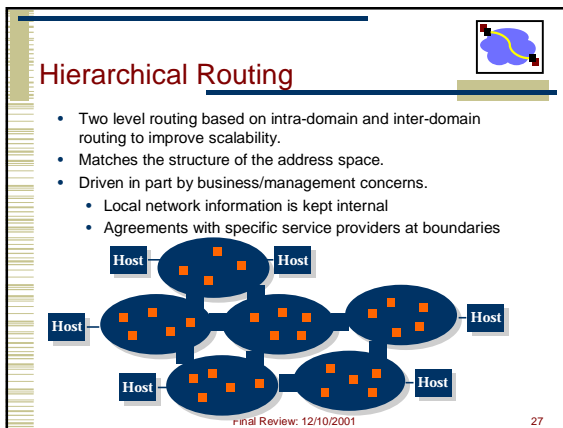
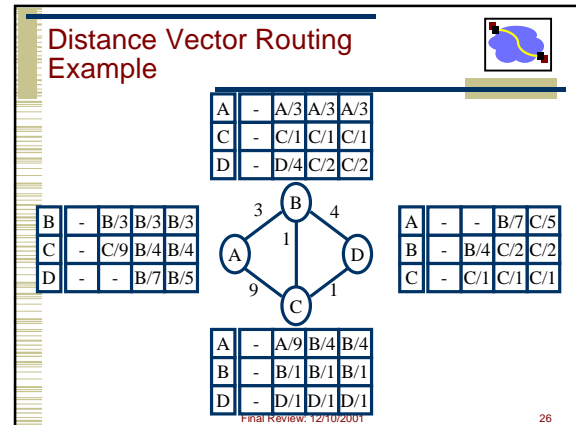
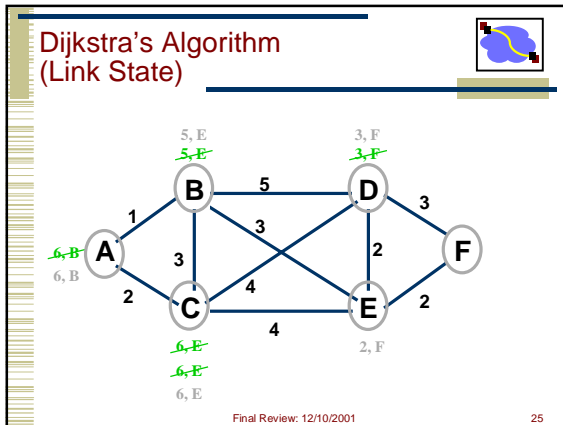
What Does Routing Do?

- Routing protocol specifies how routers jointly collect information about the network.
 - Routing protocols must be standardized
- Routing algorithm uses network information to select appropriate routes and to set up the routing table.
- The data forwarding engine performs route lookup in the routing table.
 - through which interface should a packet be forwarded?



Final Review: 12/10/2001

24



Ethernet Switches

- Bridges make it possible to increase LAN capacity.
 - Packets are no longer broadcasted - they are only forwarded on selected links
 - Adds a switching flavor to the broadcast LAN
- Ethernet switch is a special case of a bridge: each bridge port is connected to a single host.
 - Simplifies the protocol and hardware used (only two stations on the link)
 - Can make the link full duplex (really simple protocol!)
 - Can have different port speeds

Final Review: 12/10/2001

31

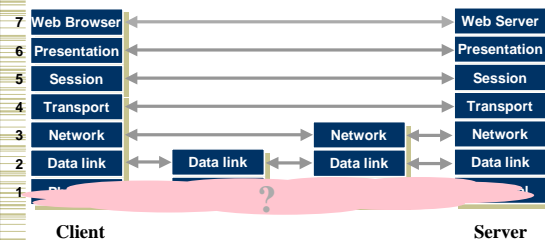
Framing

- A link layer function, defining which bits have which function.
- Minimal functionality: mark off units of transmission.
- Some techniques:
 - frame delimiter characters with character stuffing
 - frame delimiter codes with bit stuffing
 - out of band delimiters (e.g. FDDI control symbols)
 - synchronous transmission (e.g. SONET)

Final Review: 12/10/2001

32

Browsing the Web

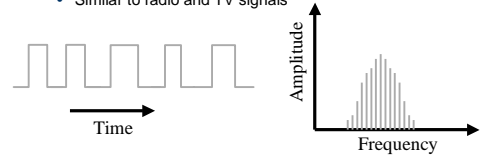


Final Review: 12/10/2001

33

The Frequency Domain

- A (periodic) signal can be viewed as a sum of sine waves of different strengths.
- Every signal has an equivalent representation in the frequency domain.
 - What frequencies are present and what is their strength
 - Similar to radio and TV signals



Final Review: 12/10/2001

34

Wireless: Good News Bad News

- Great technology: no wires to install, convenient mobility, ..
- High attenuation limits distances.
 - Wave propagates out as a sphere
 - Signal strength reduces quickly ($1/\text{distance}^2$)
- High noise due to interference from other transmitters.
 - Use MAC and other rules to limit interference
 - Aggressive encoding techniques to make signal less sensitive to noise
- Other effects: multipath fading, security, ..
- Ether has limited bandwidth.
 - Try to maximize its use

Final Review: 12/10/2001

35

TCP Problems Over Noisy Links

- Wireless links are inherently error-prone
 - Fades, interference, attenuation
 - Errors often happen in bursts
- TCP cannot distinguish between corruption and congestion
 - TCP unnecessarily reduces window, resulting in low throughput and high latency
- Burst losses often result in timeouts
- Sender retransmission is the only option
 - Inefficient use of bandwidth

Final Review: 12/10/2001

36

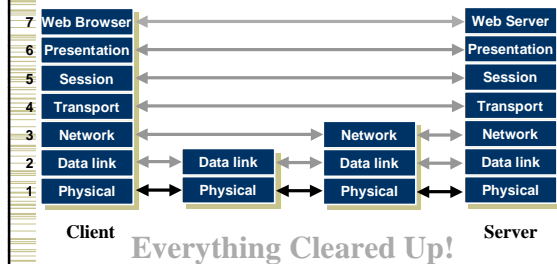
Proposed Solutions

- End-to-end protocols
 - Selective ACKs, Explicit loss notification
- Split-connection protocols
 - Separate connections for wired path and wireless hop
- Reliable link-layer protocols
 - Error-correcting codes
 - Local retransmission

Final Review: 12/10/2001

37

Browsing the Web



Final Review: 12/10/2001

38

Security Threats

- Impersonation.
 - Pretend to be another user with the intent of getting access to information or services
- Secrecy.
 - Get access to the contents of packets
- Message integrity.
 - Change a message unbeknownst to the sender or receiver
- Repudiation
 - Denying to have sent a message
- Denial of service.
 - Flooding the system so users with legitimate needs cannot get service
- Range of other threats: password guessing, exploiting programming bugs, ...

Final Review: 12/10/2001

39

Encryption

$$\text{ciphertext} = E(\text{plaintext}, k)$$

$$\text{plaintext} = D(\text{ciphertext}, k')$$

- Private key (symmetric, e.g. DES)
 - the two parties share a common private key k
- Public key (asymmetric, e.g. RSA)
 - derive two keys, k_{private} and k_{public}
 - k_{private} is kept private by its owner
 - k_{public} is published
- Tradeoffs between private and public key cryptography.
 - Key management, speed
- Challenge: key management.

Final Review: 12/10/2001

40

Example Applications

- Kerberos.
 - Support security in corporate environment
 - Based on key distribution center that knows all the entities
 - Know = share secret
- Secure socket layer (SSH).
 - Support secure channels in open internet environment
 - Based on certificates and certification authorities
 - Provides privacy, but trust is limited
- Pretty good privacy (PGP).
 - Provides privacy, authentication, repudiation in internet environment
 - Key management based on a "web of trust"

Final Review: 12/10/2001

41

How to Provide QoS?

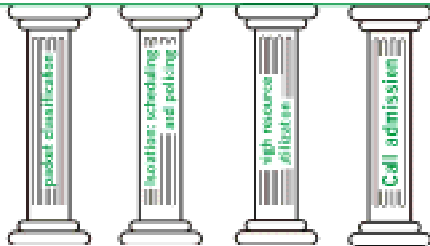
- Admission control limits number of users.
 - You cannot provide guarantees if there are too many users sharing the same set of resources (bandwidth)
 - For example, telephone networks - busy tone
 - This implies that your request for service can be rejected
- Traffic enforcement limits how much traffic users can inject based on predefined limits.
 - Make sure user respects the traffic contract
 - Data outside of contract can be dropped (before entering the network!) or can be sent at a lower priority
- Scheduling support in the routers guarantee that users get their share of the bandwidth.
 - Again based on pre-negotiated bounds
- Signaling protocol gives routers the information they need to provide QoS.
 - E.g. RSVP

Final Review: 12/10/2001

42

Qos Summary

QoS for networked applications



Final Review: 12/10/2001

43

IETF QoS Models

- Integrated services: diverse QoS at the micro-flow level.
 - Range of QoS: best effort, controlled load, guaranteed
 - Specific end-to-end service defined for each class
 - Requires end-to-end support, e.g. edge and core routers
 - Concern about complexity, cost, marketing/charging
- Differentiated services: QoS at the aggregate flow level.
 - Defines range of "forwarding behaviors", but services are defined by the providers
 - Pushes most complexity to the edge of the network – fast core routers work only with small number of traffic classes
- Based on the same building blocks.

Final Review: 12/10/2001

44

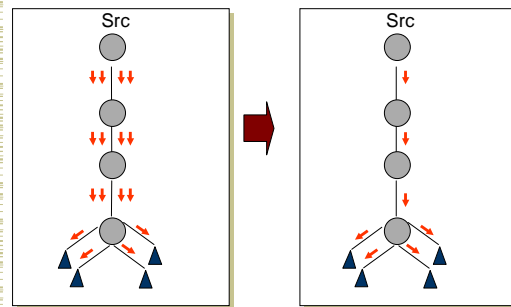
Multimedia Challenges

- TCP/UDP/IP suite provides best-effort, no guarantees on expectation or variance of packet delay
- Streaming applications delay of 5 to 10 seconds is typical and has been acceptable, but performance deteriorate if links are congested (transoceanic)
- Real-Time Interactive requirements on delay and its jitter have been satisfied by over-provisioning (providing plenty of bandwidth), what will happen when the load increases?...

Final Review: 12/10/2001

45

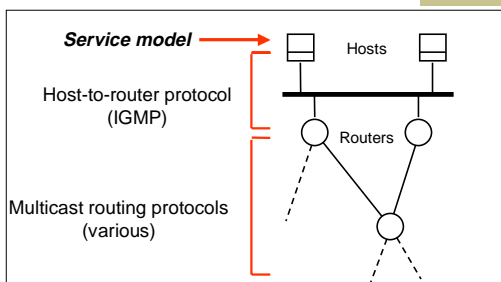
Multicast – Efficient Data Distribution



Final Review: 12/10/2001

46

IP Multicast Architecture



Final Review: 12/10/2001

47

Multicast Routing

- Basic objective – build distribution tree for multicast packets
- Core based protocols
 - Examples: CBT, PIM-SM
- Flood and prune
 - Examples: DVMRP, PIM-DM
- Link-state multicast protocols
 - Example: MOSPF

Final Review: 12/10/2001

48

Shared vs. Source-based Trees



- Source-based trees
 - Separate shortest path tree for each sender
 - DVMRP, MOSPF, PIM-DM, PIM-SM
- Shared trees
 - Single tree shared by all members
 - Data flows on same tree regardless of sender
 - CBT, PIM-SM

Final Review: 12/10/2001

49

Questions?



Final Review: 12/10/2001

50