Issues In Transporting MPEG over IP Networks

Transporting MPEG over a Network

- MPEG-specific issues to be dealt with:
  - Bandwidth requirements.
  - Burstiness of traffic.
  - Effects of packet loss.
  - Latency.
  - Delay Jitter.
- Application requirements:
  - Multicasting

Bandwidth Requirements

- Unlike traditional data traffic, MPEG video uses a relatively high amount of bandwidth for extended periods of time.
- For variable bit-rate MPEG, the bandwidth is not even known and is a function of the encoded material
- The network ***must*** somehow guarantee the bandwidth or the system simply does not work, or the video quality is unacceptable.

Burstiness of Traffic

- MPEG video is inherently bursty:
  - Video enters the system at 30 frames/second
  - Encoder transmits 30 frames/second
  - Frames are encoded into I-, P-, and B-frames
  - I-Frames are larger than P-Frames, which are larger than B-frames.
- Result: ***bursty traffic!***
- Note that CBR MPEG is still bursty but the ***average*** bit rate is exactly what was set in the encoder.

Example

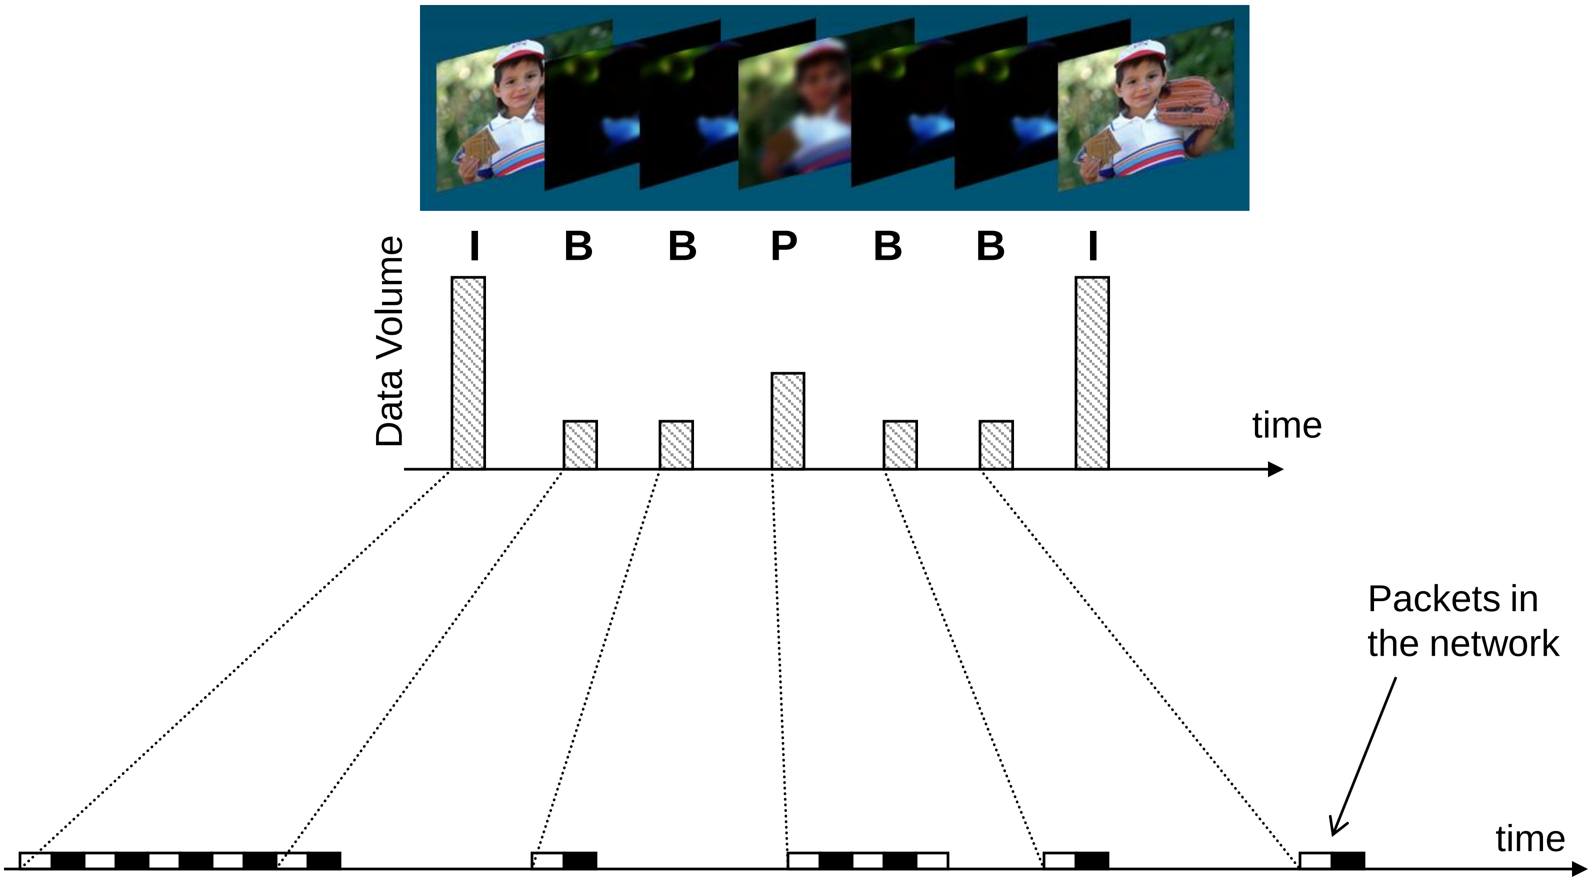

Example: Packet Interdepartures

Two-minute run
1.8 Mb/s
1400-byte packets
Audio/Video muxed
100 Mb/s Ethernet

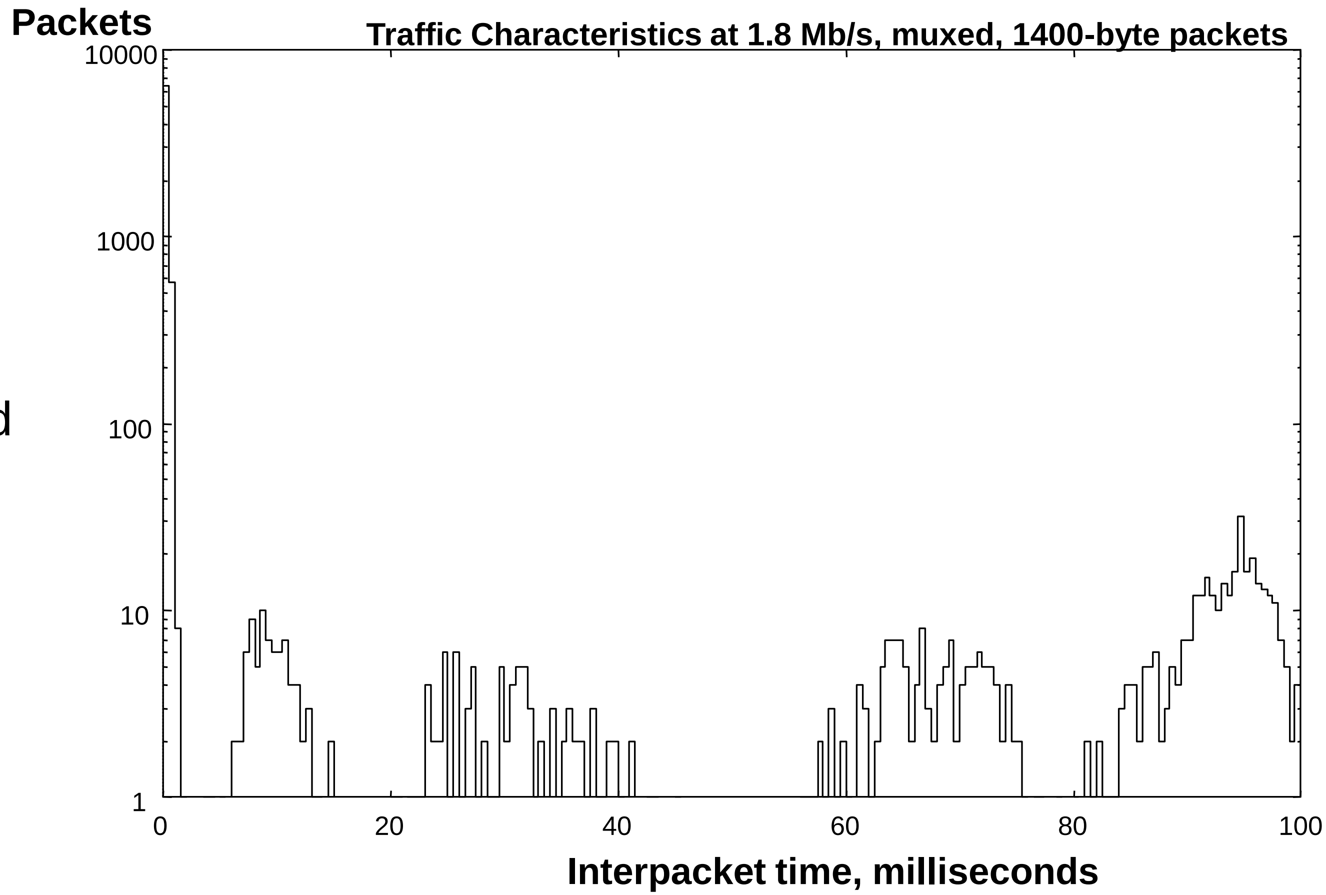

Constraining the Bit Rate

- It is possible to “smooth out” the traffic and constrain the MPEG stream to use only the average rate (i.e., make the peak rate = average rate).
- Price paid: increased latency.

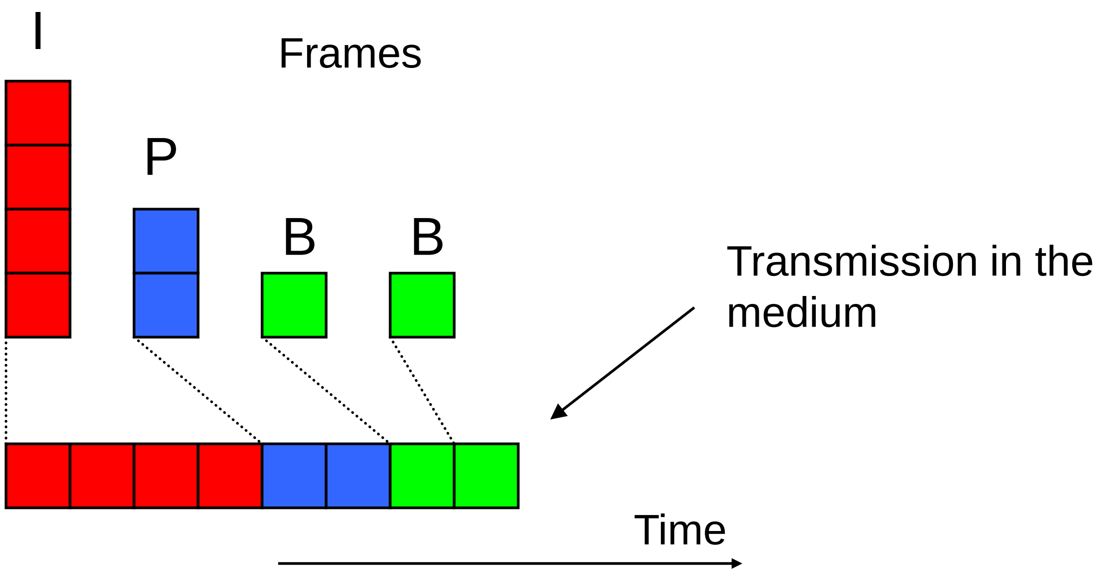

Constrained Interdepartures

Two-minute run
1.8 Mb/s
1400-byte packets
Audio/Video muxed
Peak Rate: 2.1 Mb/s
(1/0.85 of avg.)
100 Mb/s Ethernet

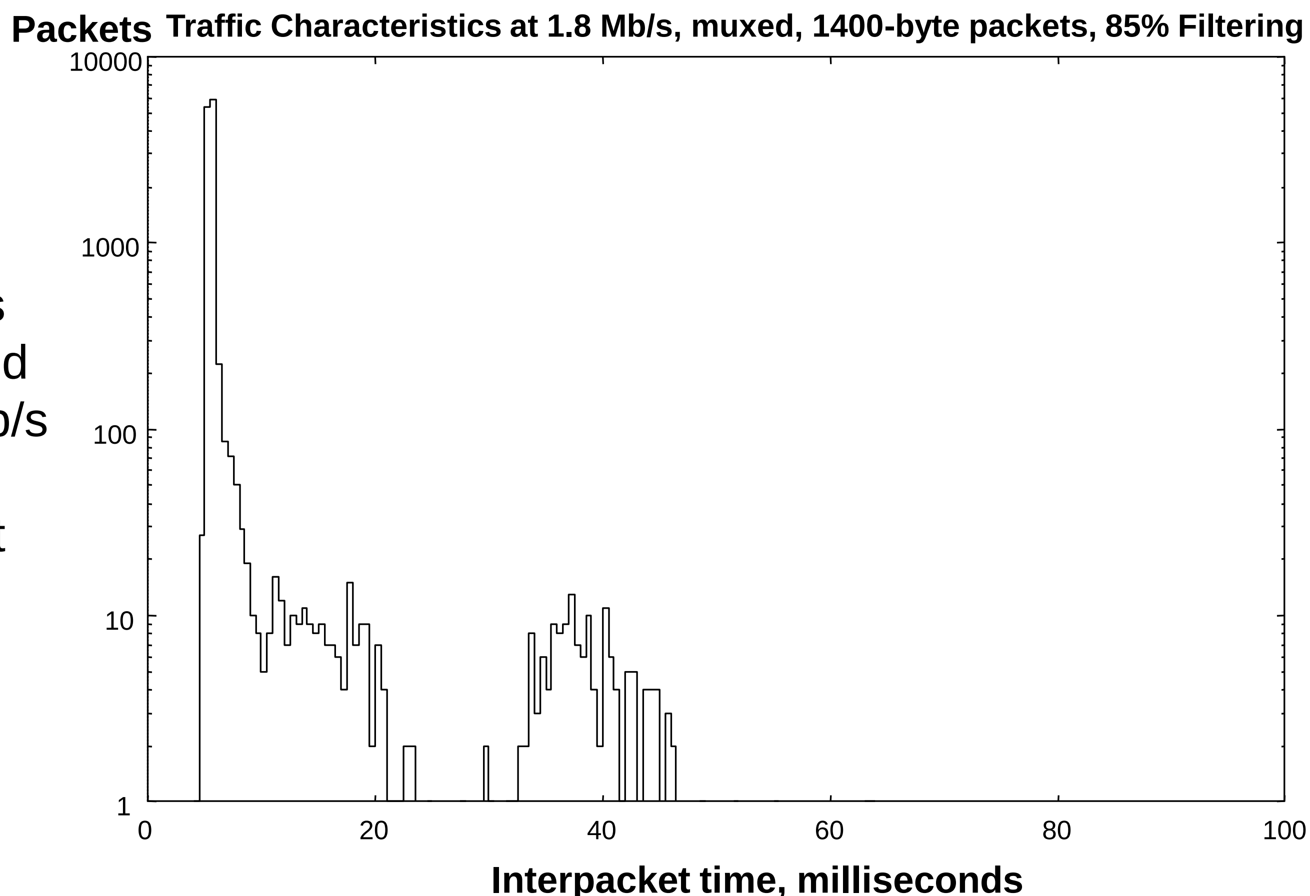

Packet Loss

- "Conventional wisdom"says that with real-time media, some data loss is acceptable.
  - This is certainly true with uncompressed video, since a packet loss will provoke a glitch of 1/30 of a second.
- Because of the frame dependencies, with MPEG, packet loss causes very noticeable artifacts that can persist for up to 15 frames (0.5 second).
- Packet loss ***must be VERY LOW!***

  (or the viewers get upset!)

Acceptable Packet Loss

- Example: packet loss rate 10^{-6} (one packet in a million), bit rate 4.0 Mb/s, packet size 1400 bytes:

$$\text{Glitch Interval} = \frac{8 \times \text{Packet Size}}{\text{Bit Rate} \times \text{Packet Loss}} = \frac{8 \times 1400}{4{,}000{,}000 \times 10^{-6}} = 2800 \text{ seconds}$$

Loosing one packet in	**Gives you a glitch every**
10 million	7 hours 47 minutes
1 million	47 minutes
100 thousand	4 minutes 40 seconds
10 thousand	28 seconds
1 thousand	2.8 seconds

Latency

- Interactive applications such as videoconferencing require latencies on the order of 250 ms or less.
- Minimum latency for MPEG with B-frames:
  - 2 frames latency in the encoder
  - 2 frames latency in the decoder
  - Total: 4 frames @ 33 ms/frame = 133 milliseconds ONE WAY!
- Actual MPEG encoder chips may add up to 3 more frames to the encode latency (due to pipelining)
- Decoder must fill to the VBV before starting
  - Bursty networks have an advantage here!
  - Cut down on startup time if the encoder can burst!

What to do?

- To reduce latency:
  - Get rid of the B-frames.
  - Reduce the VBV.
- The steps above represent a tradeoff between latency and quality:
  - For the same bit rate, reducing the latency reduces the quality.
  - For the same quality, one has to increase the bit rate when the latency is lowered.
- Measured latency numbers (unloaded 100 Mb/s Ethernet):
  - 4 Mb/s, VBV=224k, IPB frame structure: **350 ms**
  - 4 Mb/s, VBV=224k, IP frame structure: **290 ms**
  - 4 Mb/s, VBV=60k, IP frame structure: **260 ms**
- New MPEG-2 technology is bringing latency to under 120 ms, using techniques similar to H.261.

Latency Drift

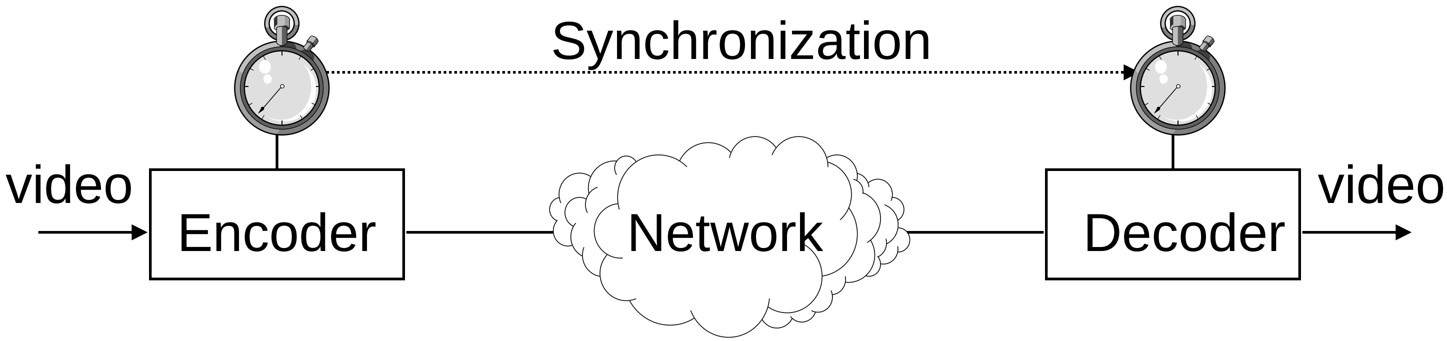

- Encoder clock determines data generation rate
- Decoder clock determines data consumption rate
- If clocks are not synchronized, decoder will eventually underflow or overflow
  - Latency in the system changes with time!

Avoiding Latency Drift

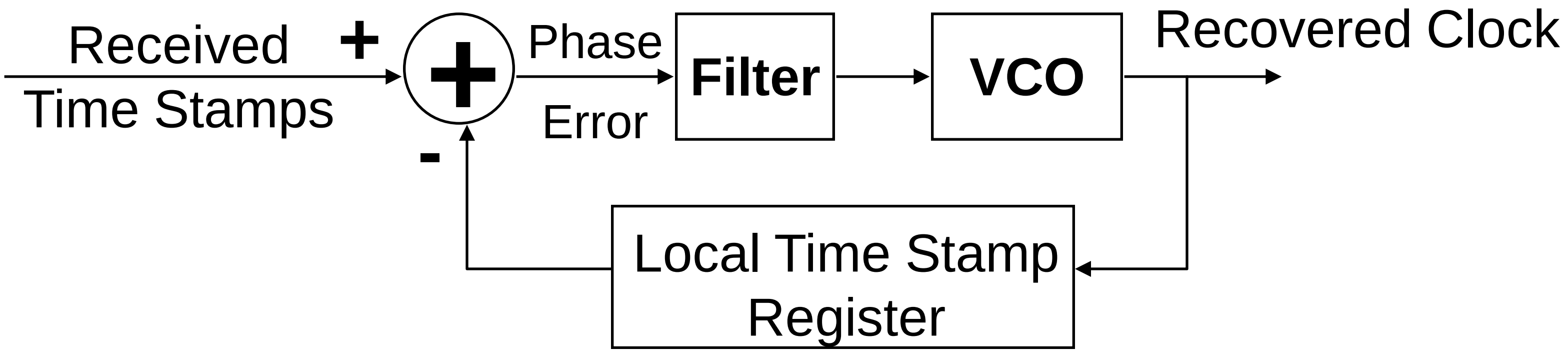

- Clock can be recovered at the decoder by a PLL
- Incoming signal has periodic time stamps
- The decoder clock is adjusted based on the received time stamps

Delay Jitter

- Variations in delay from packet to packet.
- Early MPEG set-top boxes required very small delay jitter because:
  - Clock recovery (PLL) required time stamps to be regular; not much filtering.
  - Minimizing decoder memory.
- Intrinsic MPEG latencies are two orders of magnitude or more higher than typical network delay jitter.
- Current systems have plenty of memory.
- Bottomline: ***no real requirements on delay jitter!***

Multicasting

- Video applications very often have multicasting (one-to-many distribution) requirements.
- Requirements on the network infrastructure:
  - IP Multicast support.
  - IGMP V2 support - important for leaving groups.
  - Layer 2 multicast support.
- Typical problem areas:
  - ATM LANE 1.0 backbones - poor multicast support.
  - Lack of IGMP V2 support.
  - Lack of layer 2 multicast support.

Resource Reservation

- Video ***must have*** appropriate bandwidth set aside, or it will not work!
- Approaches currently used in industry:
  - Isolated or dedicated network.
  - Traffic engineering.
  - Router configuration to prioritize streams based on IP address or UDP port.
  - Use of the IP Precedence bits.
- Possible future evolution:
  - Use a reservation protocol such as RSVP.

MPEG-2 Transport Over ATM

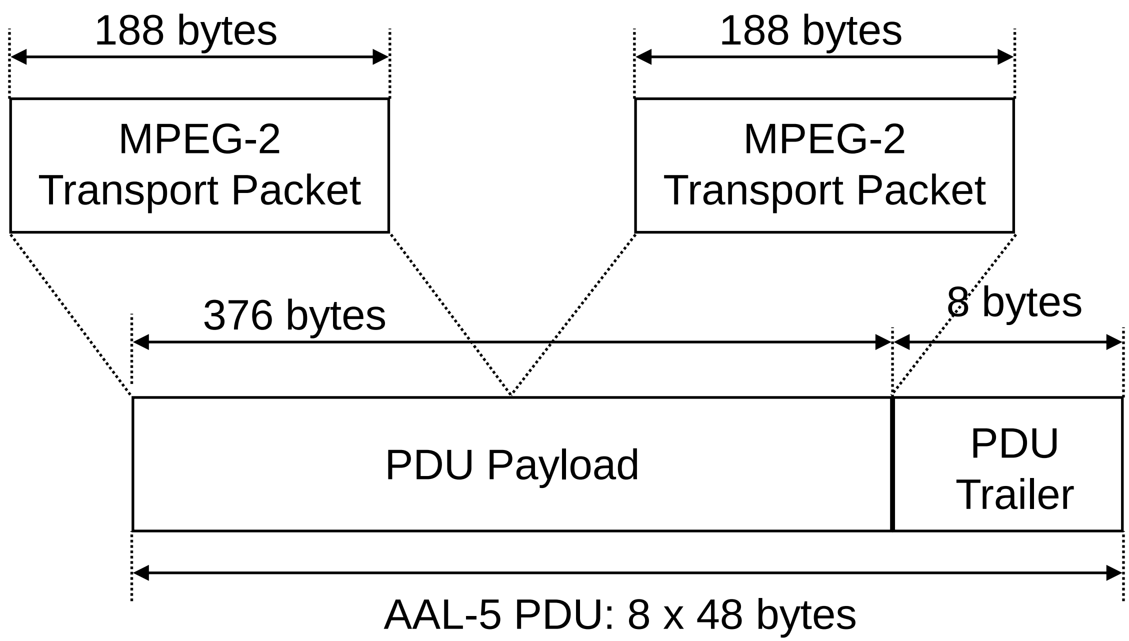

- Standardized by the ATM Forum.
- Basic mode: two transport packets per AAL-5 PDU.

MPEG over ATM (cont.)

- The ATM Forum specifies the use of a CBR channel of rate equivalent to the MPEG bit rate.
- Advantage of ATM: resource reservation
  - Encoder needs tell the ATM switch the target bit rate.
  - ATM switch will reserve the resources.
- Disadvantages of ATM:
  - Less efficient, much more expensive than Ethernet.
  - Multicast is rooted: encoder needs to create the tree and add the receivers.

Transporting Scalable Video (MPEG-4)

- MPEG-4 defines a “base layer” and “enhancement layers”.
- Enhancement layers add quality to the base layer.
- DCT-based enhancement data is continuous, i.e., any received enhancement data is usable.
- The number of enhancement layers does not need to be defined at encode time
  - Encoder generates “base” and “enhancements”
- Usage of this feature is a function of the application and services provided by the network infrastructure.

Application: Variable Available Bandwidth

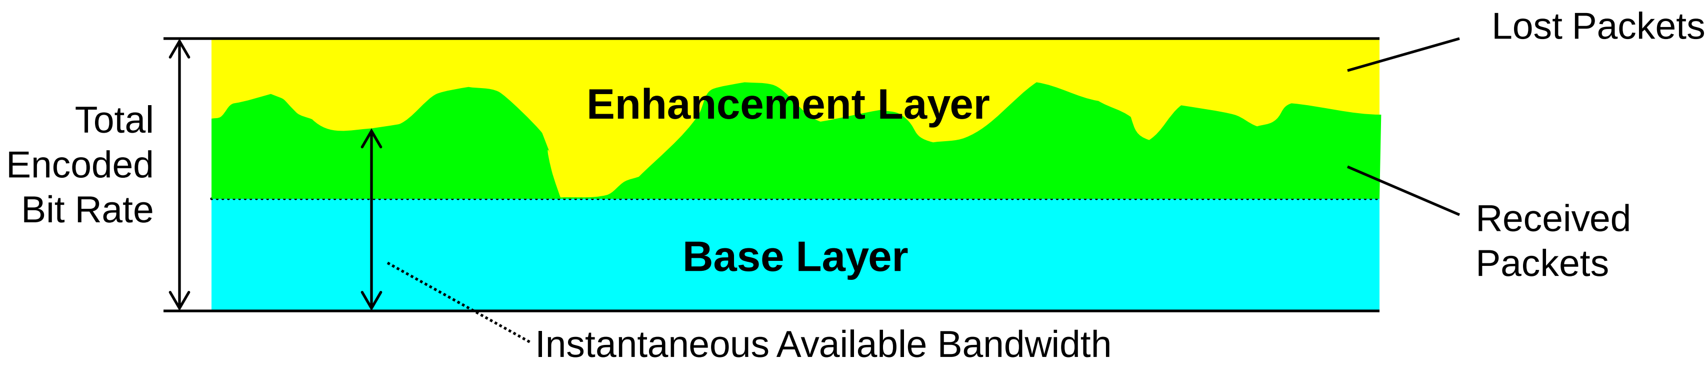

- Server (or real-time encoder) sends base plus all the enhancement.
- A random portion of the enhancement packets is lost due to traffic or congestion.
- Base layer needs some form of protection.
- Any enhancement packet received is usable.

Possible Enhancements

- Feedback from the client to throttle the server, based on measured packet loss
  - Do not send enhancements if they are going to be lost.
- Client knows the connection bandwidth and asks only for the suitable enhancement data.
- In a VOD application with multiple clients, the server may stop sending enhancement data if it runs out of resources
  - Little overall quality degradation allows server to accommodate more clients.

Adding Network Features

- Multicasting
  - Divide the enhancement data into a number of layers.
  - Assign a different multicast address to the base layer and each enhancement layer.
  - Client joins the base layer and as many enhancement layers as it can get.
- IP Precedence
  - Send each layer with a different IP Precedence value.
  - Base layer gets highest precedence; enhancement layers get decreasing precedence.
- Resource Reservation