INSTITUTE OF COMMUNICATIONS AND CONNECTED SYSTEMS
ADVANCING THE FUTURE TECHNOLOGIES OF A CONNECTED WORLD

The inter-dependence of network transport and application behaviour

Stuart Clayman
Institute of Communications and Connected Systems
University College London
s.clayman@ucl.ac.uk

Overview

- We present how services and the network are often intrinsically linked, and how there is an inter-dependence between the network transport and application behaviour.
- The purpose is to discover how new network protocols such as BPP can be used in the future, but how services need to be adapted to be effective.
- To highlight this we cover some high-level basics of digital video and network behaviour, and show the effects of BPP.
- We will consider:
  - a high-level view of video encoding and decoding,
  - how the behaviour of the network has an impact on what the viewer sees
  - the impact of BPP and on BPP
  - where next .....

BPP - Big Packet Protocol

- BPP is based on the idea of injecting meta-information into packets in order to provide programmability for network nodes on how to process those packets.
- The most important aspect to consider is that a BPP aware router, given specific commands, can drop parts of the payload. These are called chunks in BPP.
- With BPP, we don't see whole packets being dropped, but parts of the packet payload - the chunks - being dropped.
- The concept is to reduce the load on the network somewhat, with reduced bandwidth, but to keep the flow of packets arriving at the receiver, so there is no stalling.
- It is designed for high bandwidth applications where retransmission times make resending too slow. So it is more at the UDP end of the spectrum than TCP like.

High Level View

- Encoder: video in / get digital stream out.
- Decoder: digital stream in / get video out.

Encoder

digital video stream

Decoder

Network transmission

- For transmission, the task is to send the video stream from a sender to a receiver over the network.
- Depending on the socket type used and how busy the network is, the receiver will get different behaviour.
- The network looks like a pipe with funny attributes.

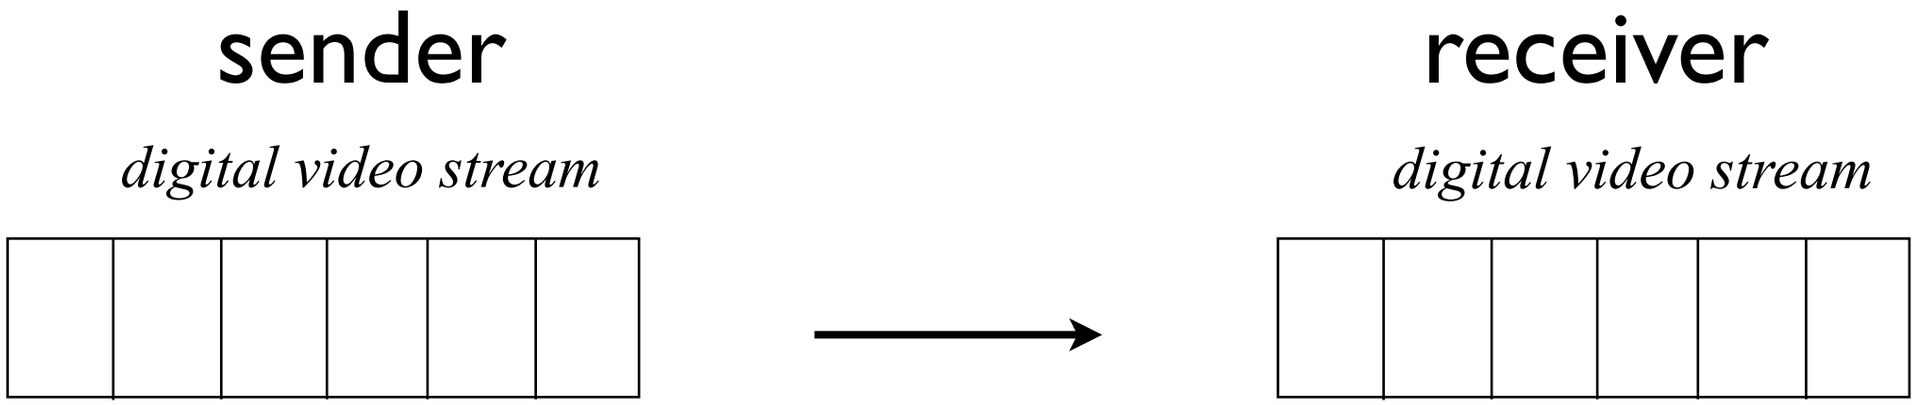

Frames and Blocks and Packets

- Each stream of video is a representation of a collection of digital video frames.
- Each frame is split into smaller blocks called Macro Blocks.

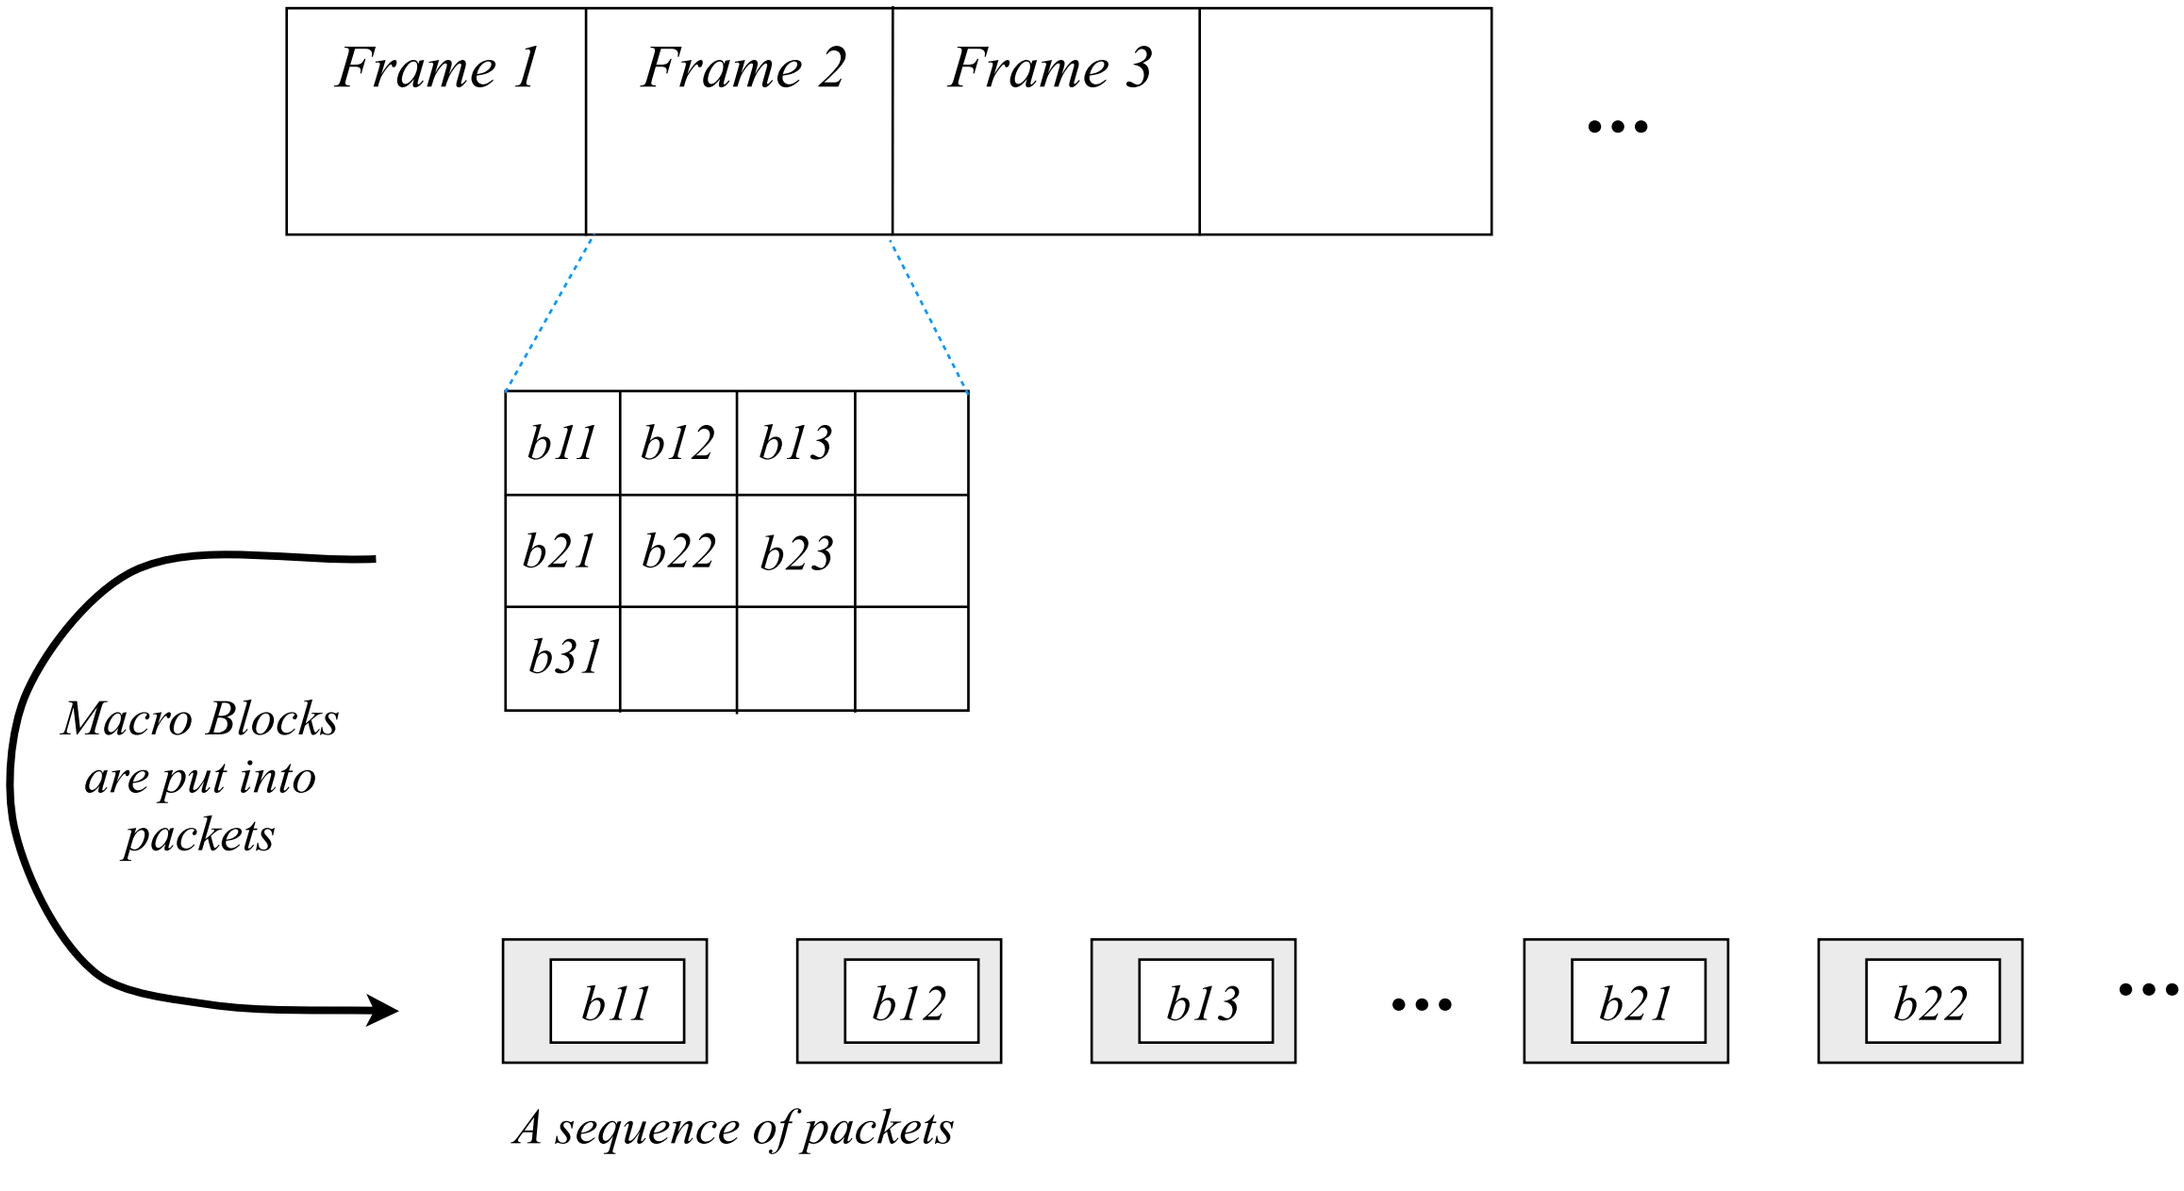

Frames and Blocks and Packets

- The packets are sent over the network to the receiver and the blocks are collected and the frames are rebuilt.
- Each frame is then sent to the decoder.

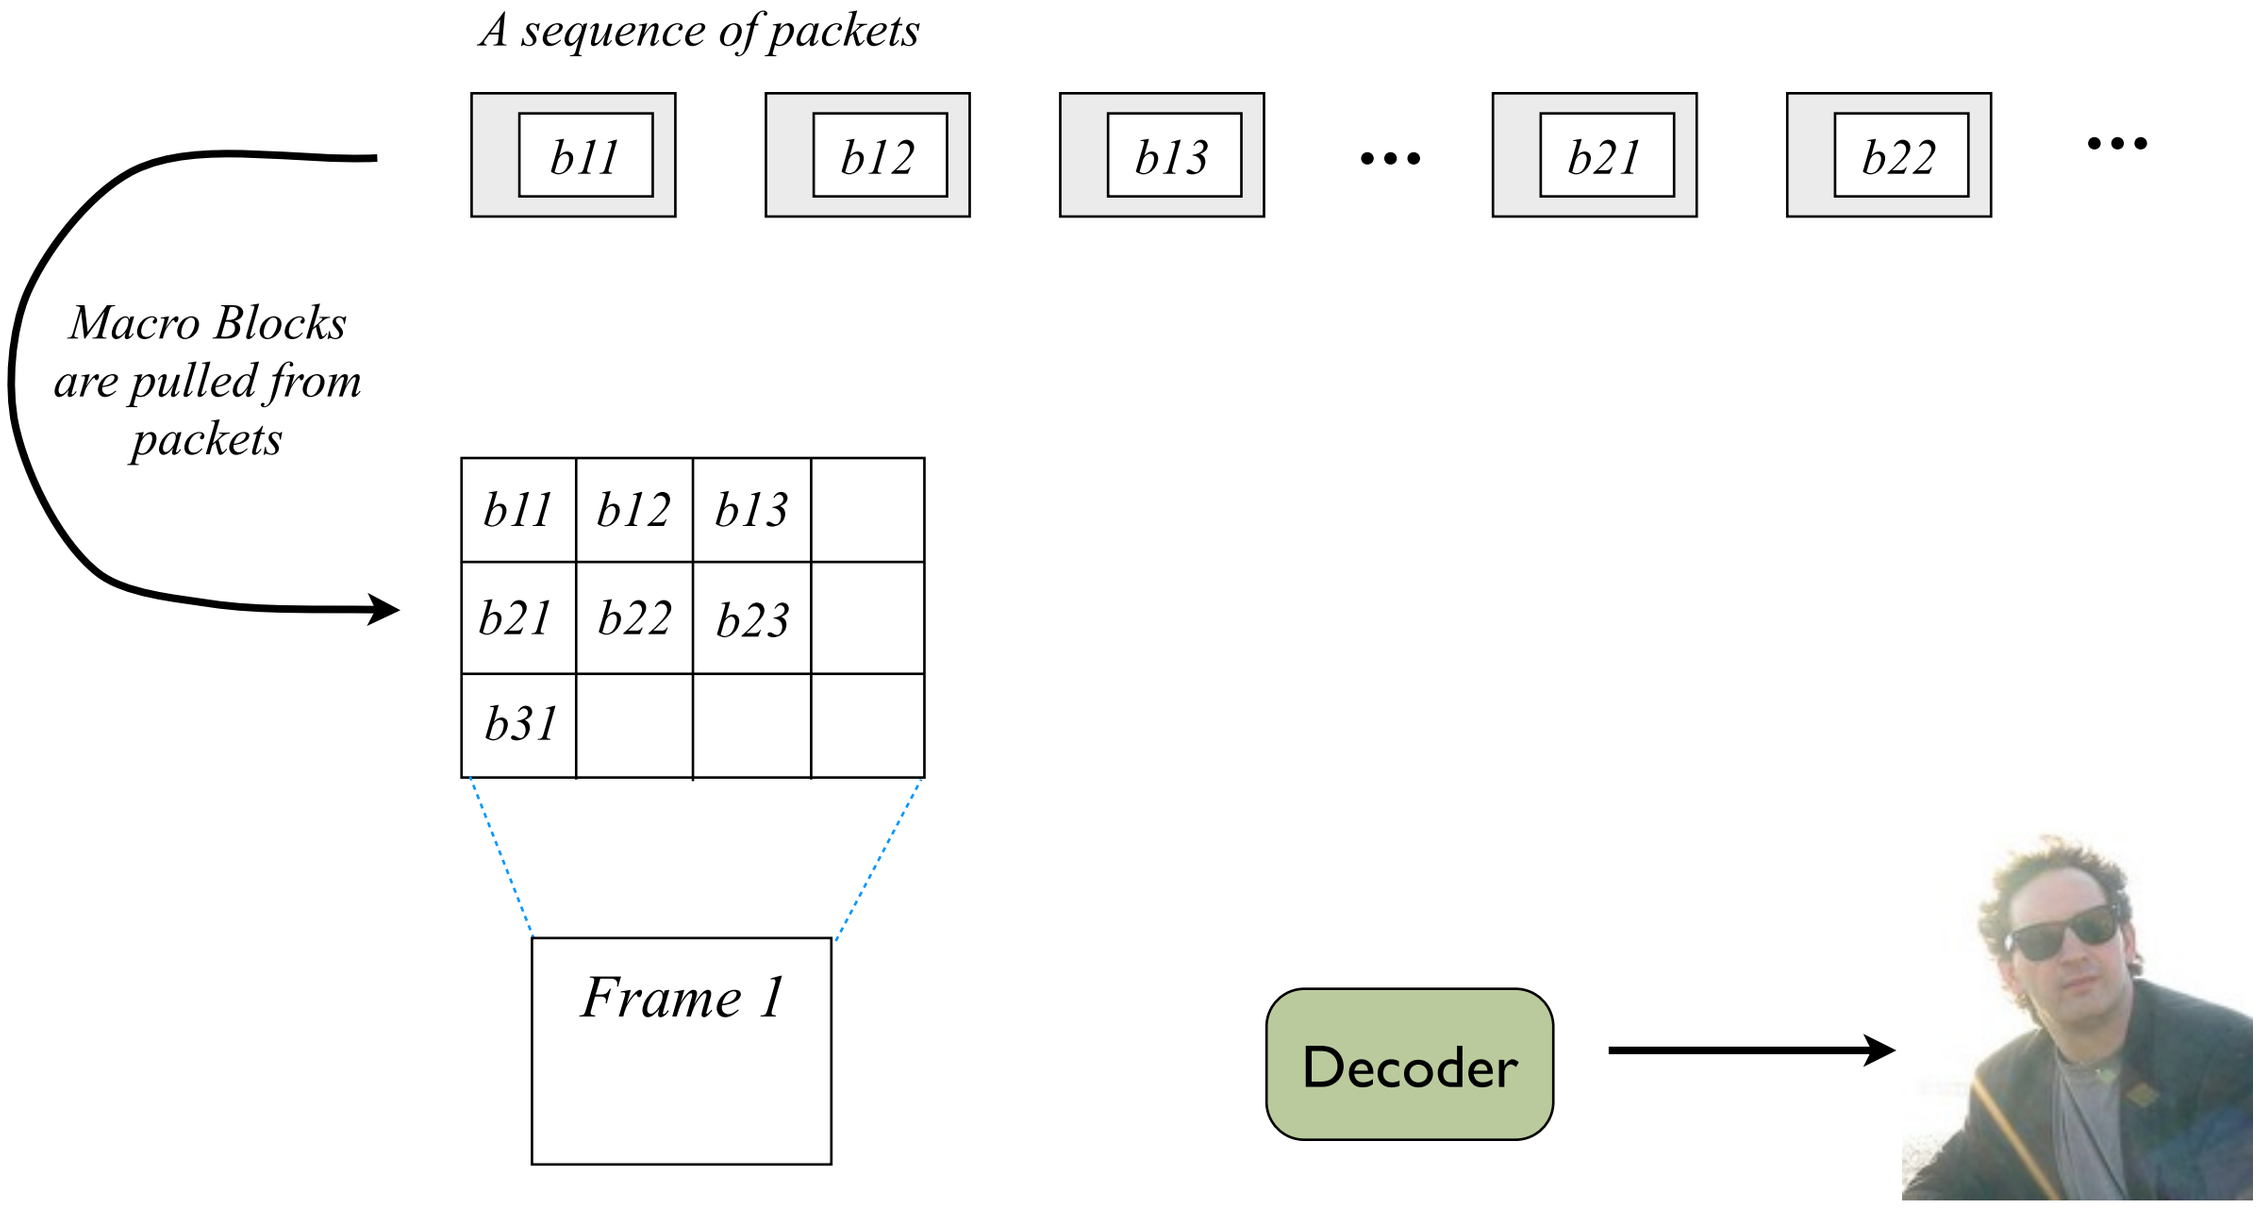

Frames and Blocks and Packets

- Consider what happens when there is loss.
- There are missing blocks, so the frame will be partially reconstructed and then decoded incorrectly.

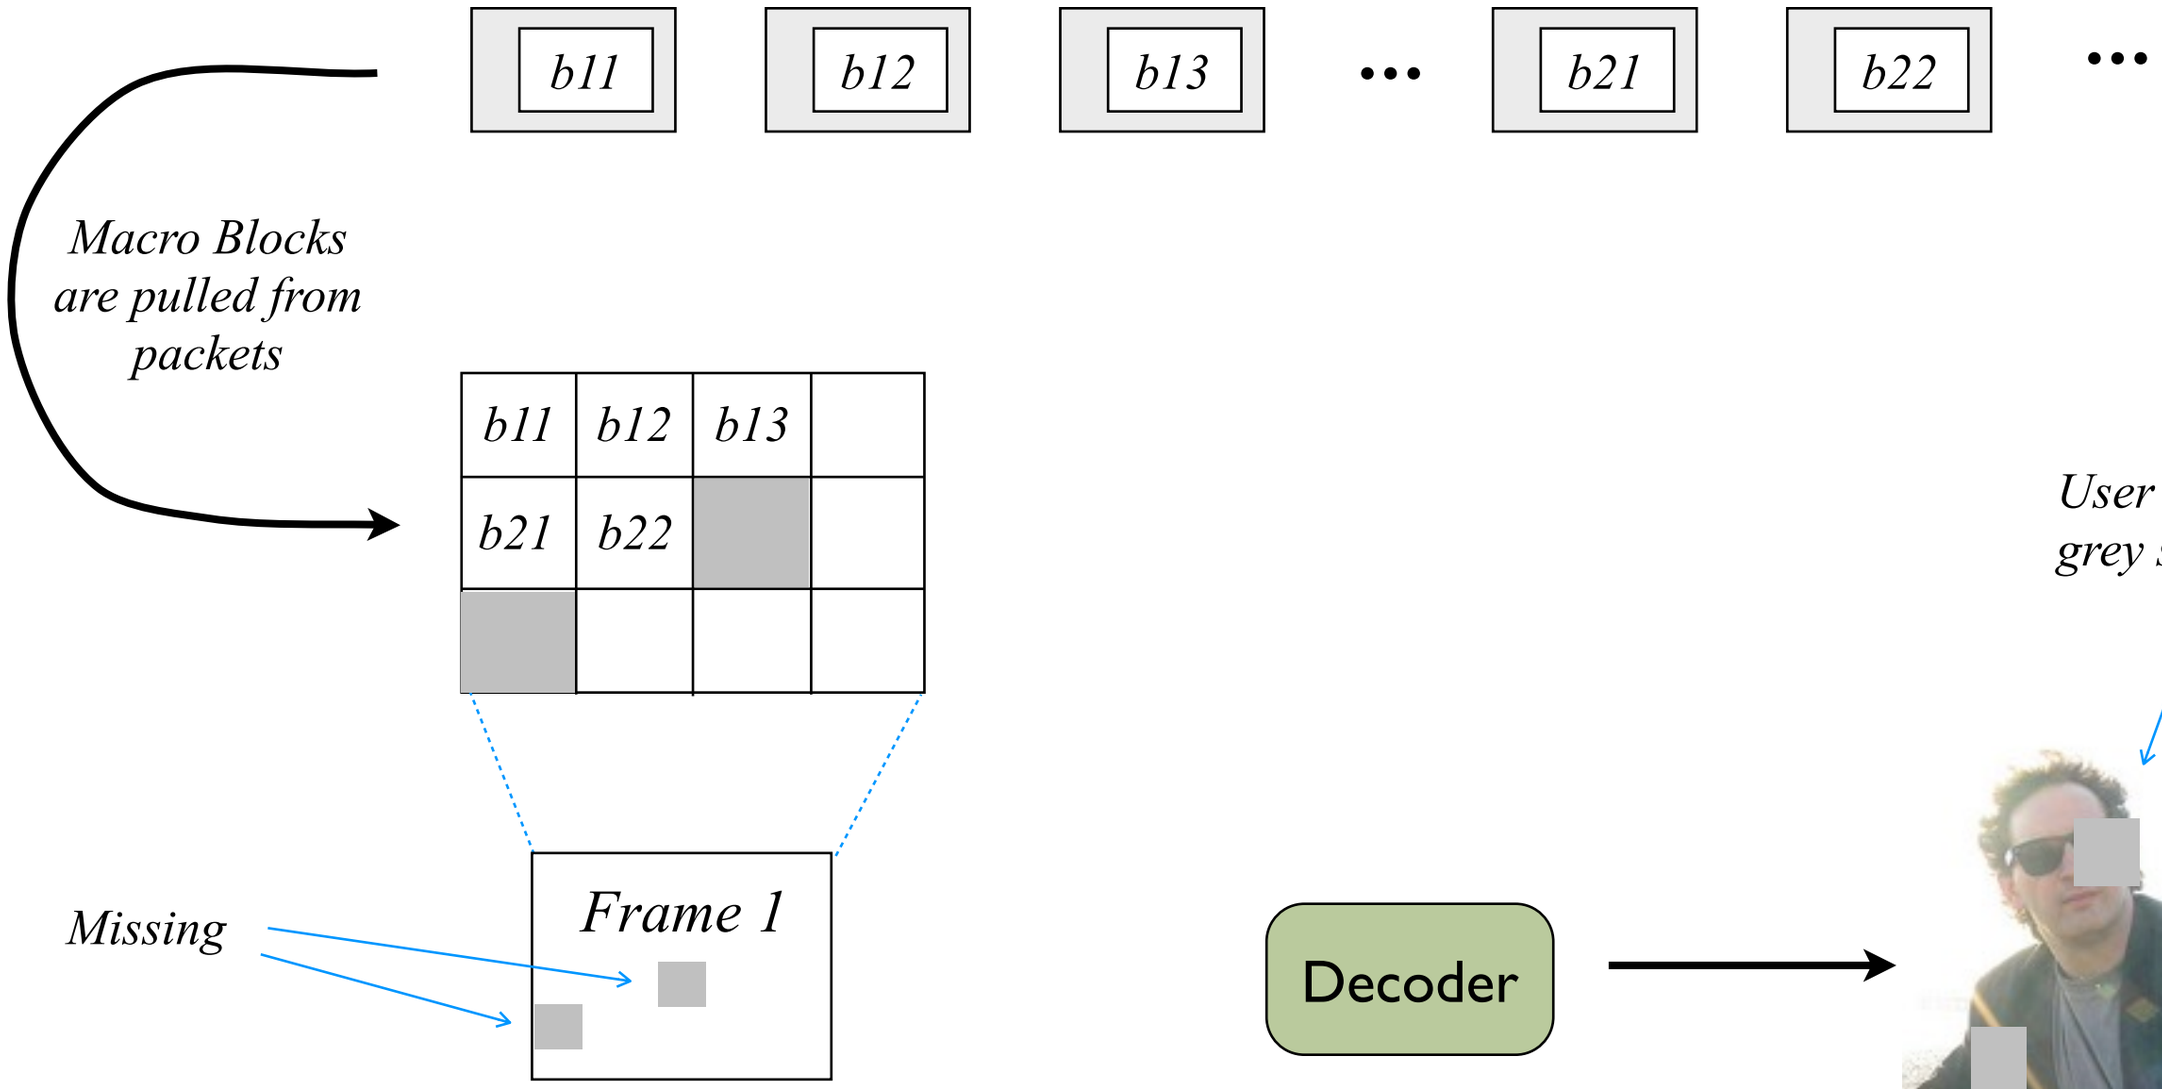

UDP Network transmission

- If a UDP socket is used then there is a view of a network pipe that displays loss at the receiver.
- The receiver application has to deal with the loss. Usually by showing missing data as grey squares in the video. The application has control over requests for resends, if needed.
- This works well for interactive video, with low latency.

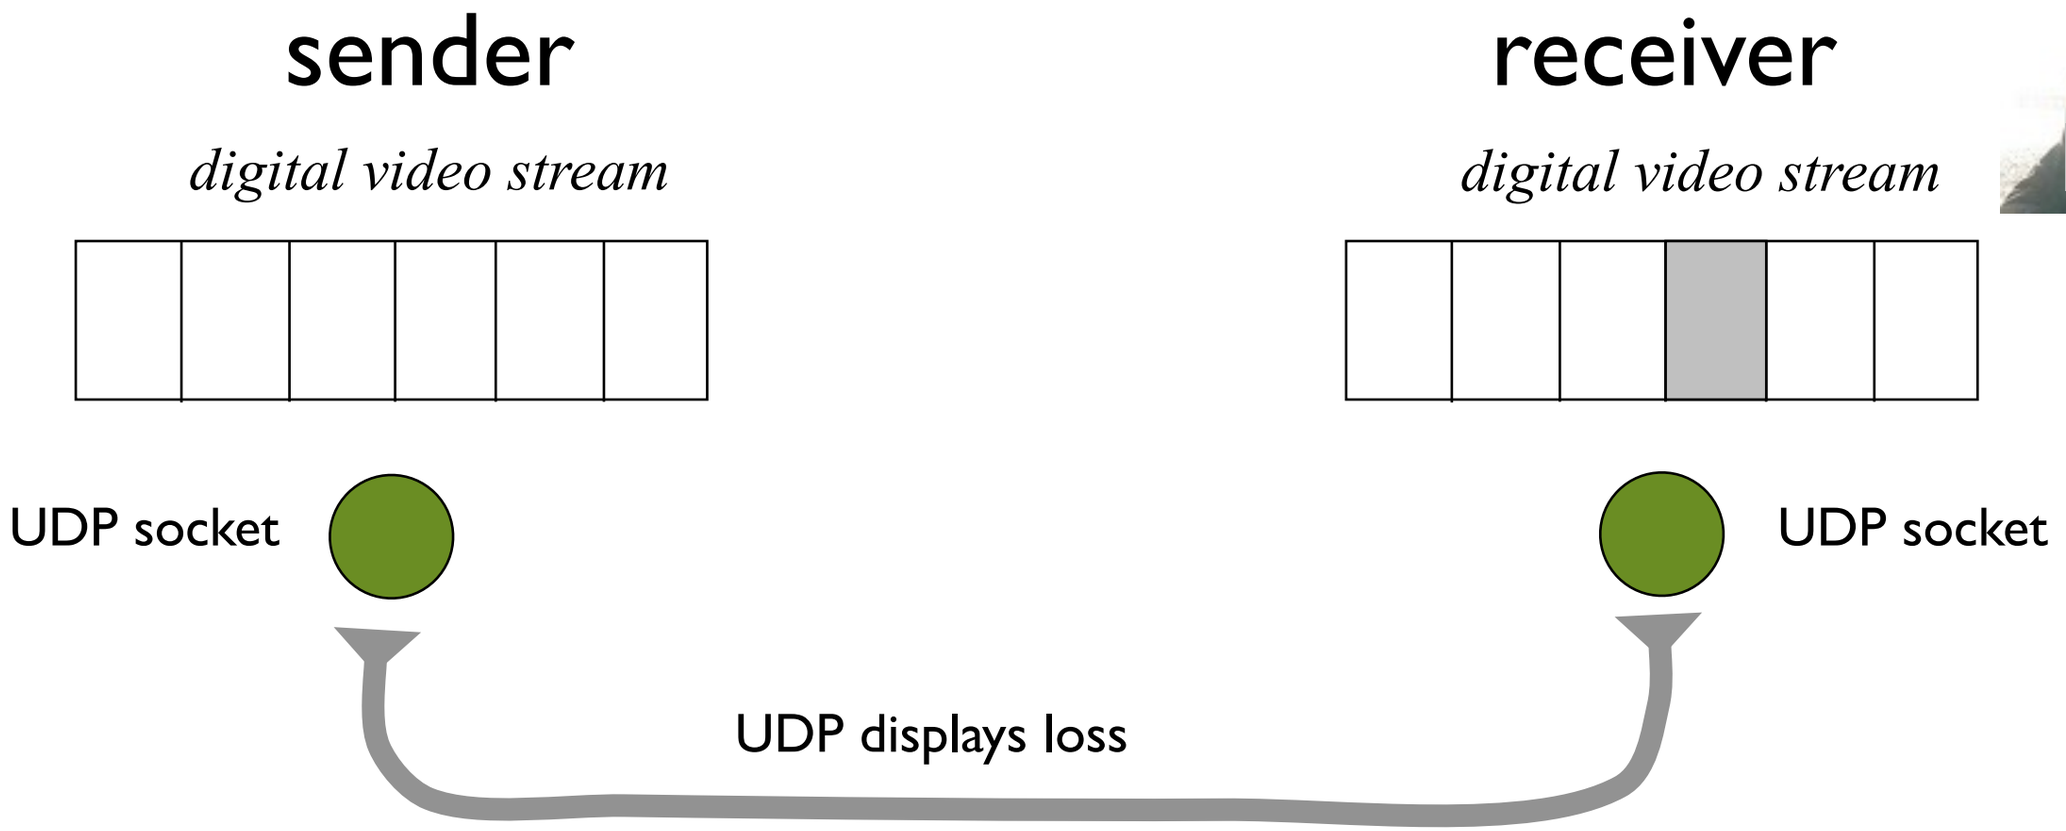

TCP Network transmission

- If a TCP socket is used then there is a view of a network pipe that has no loss but has delay / latency.
- The receiver application has to deal with the delay. Usually done by buffering a few seconds of video. The application has no control over the resend mechanism, it is done in TCP.
- This works well for video playback. No interactivity.

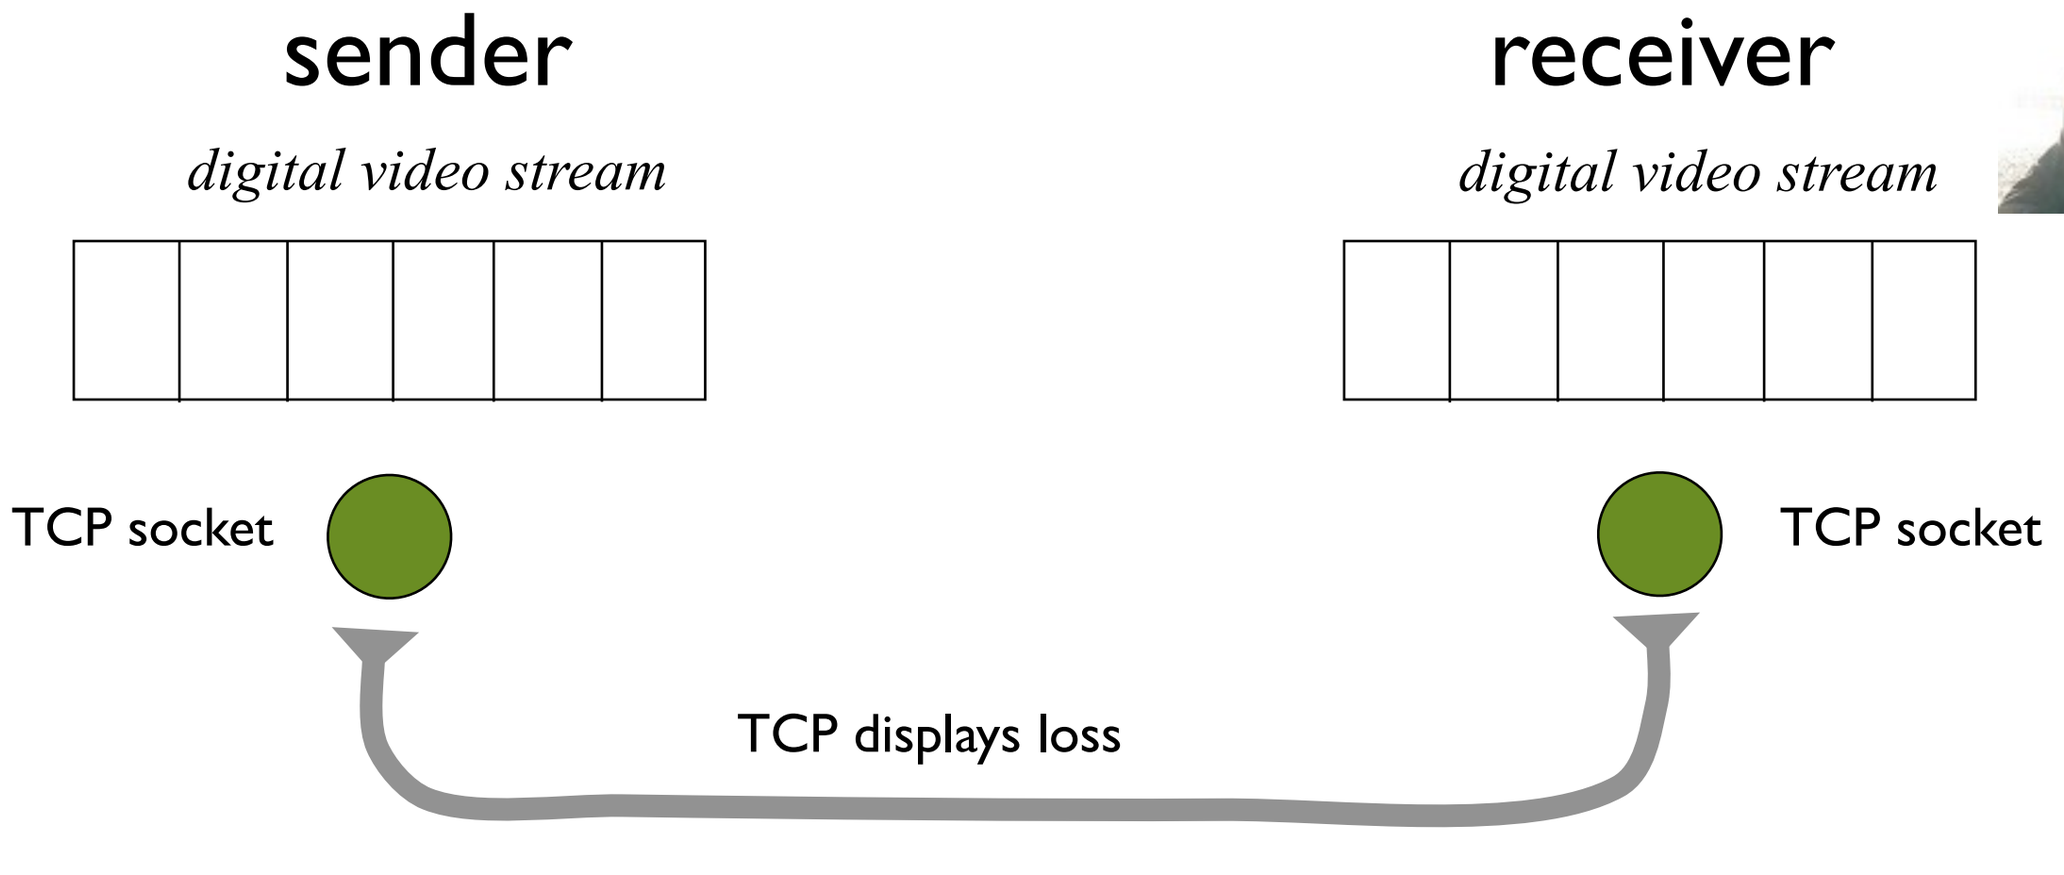

Bitrate differences

- To deal with different bandwidth availabilities digital video is often encoded multiple times with different visual qualities.
- The higher the quality, the higher the data rate, the bigger the file, and the more bandwidth that will be used.

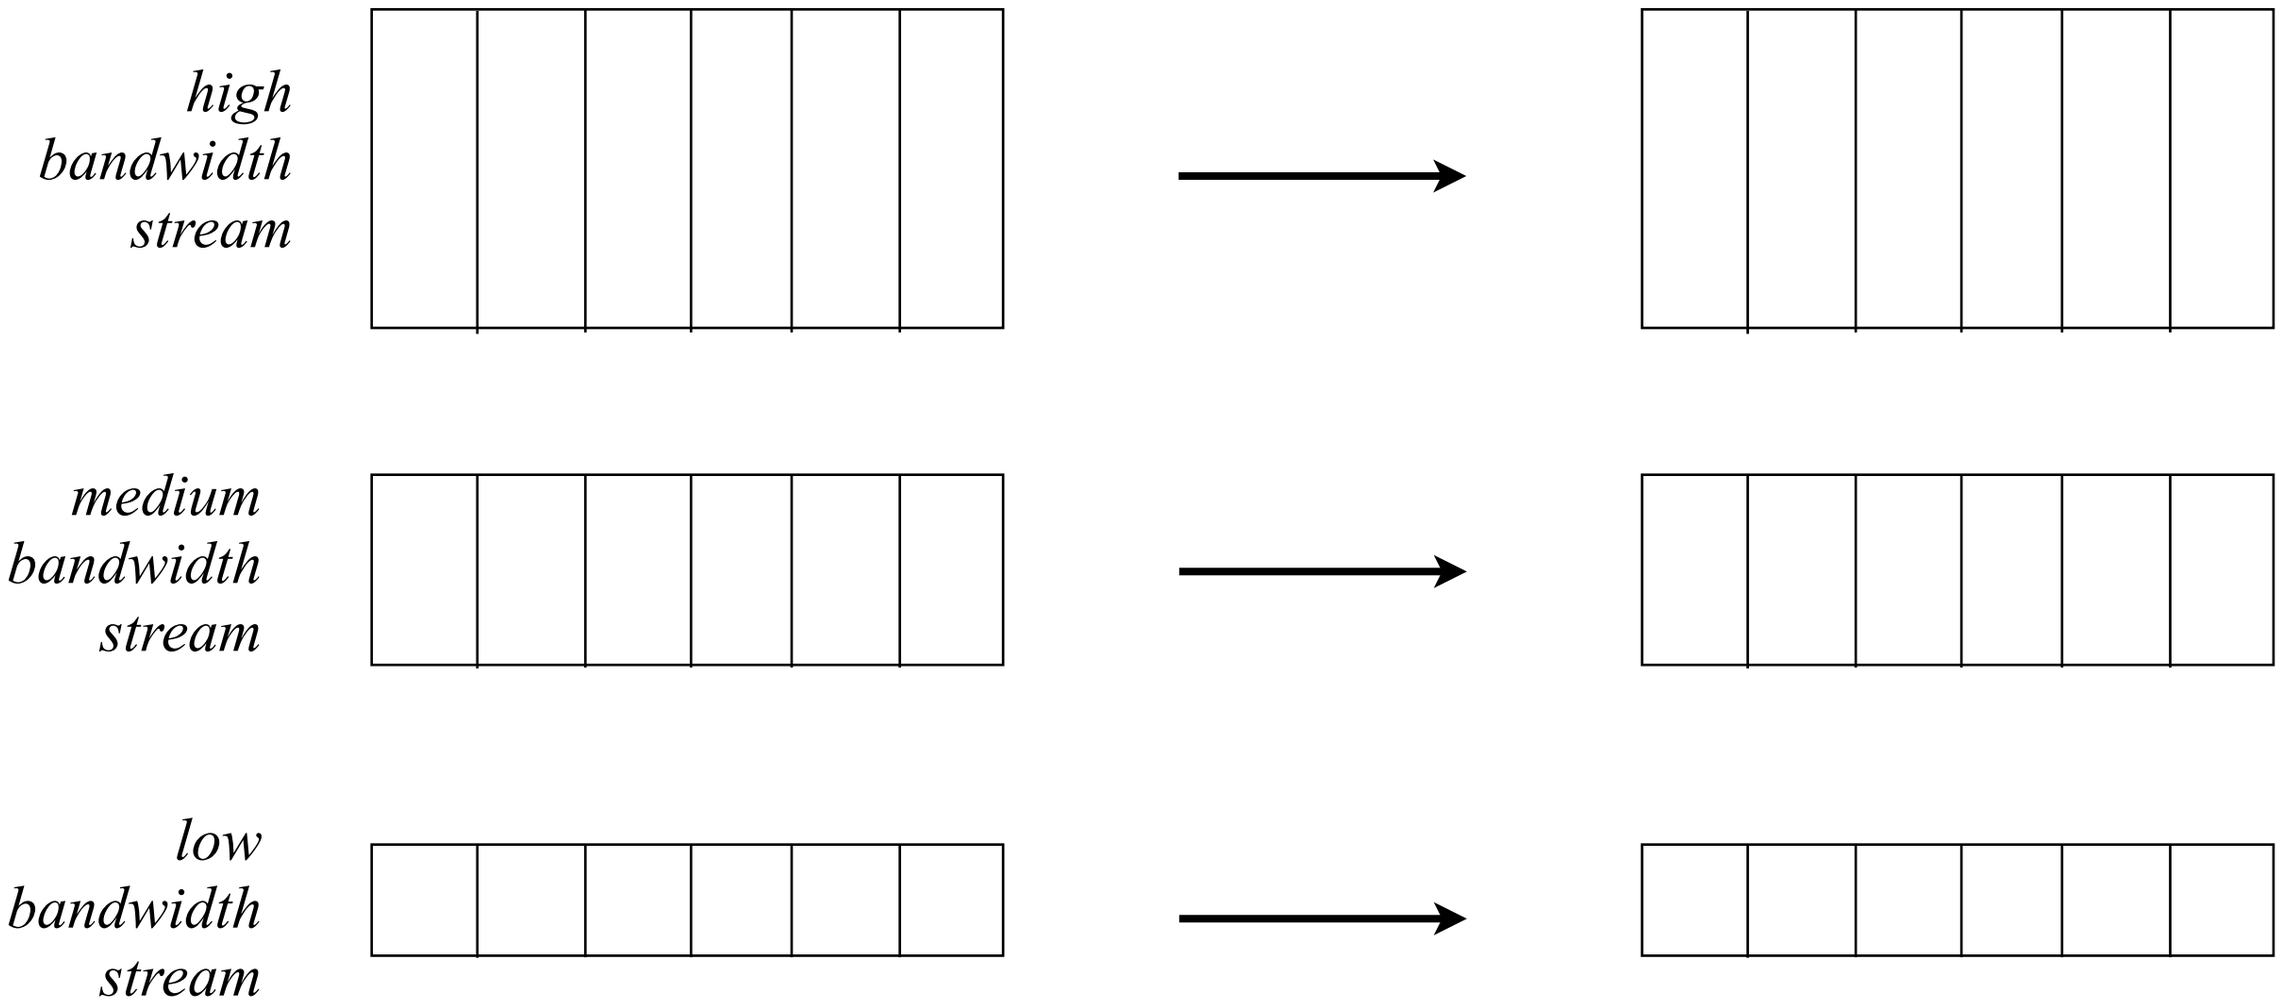

Adaptability

- To deal with varying network behaviour and loss changes over time, the MPEG-DASH standard (Dynamic Adaptive Streaming over HTTP) has been devised. As it runs over HTTP, it is inherently TCP.
- The system is designed to send segments, of a few seconds length, of video at the requested quality.
- It starts with the lowest quality, and progressively goes to higher quality if the receiver observes that congestion is low and the network bandwidth is available.
- The receiver can dynamically adapt the quality of the video being sent by requesting segments of a different quality.
- The receiver puts the sequence of segments into the decoder for viewing.

Adaptability

- The sender has multiple streams, split into segments of a few seconds, which can be sent using different bandwidths.
- Each segment is requested and chosen dynamically by the receiver, based on loss patterns. It gets a mix of bit-rates.

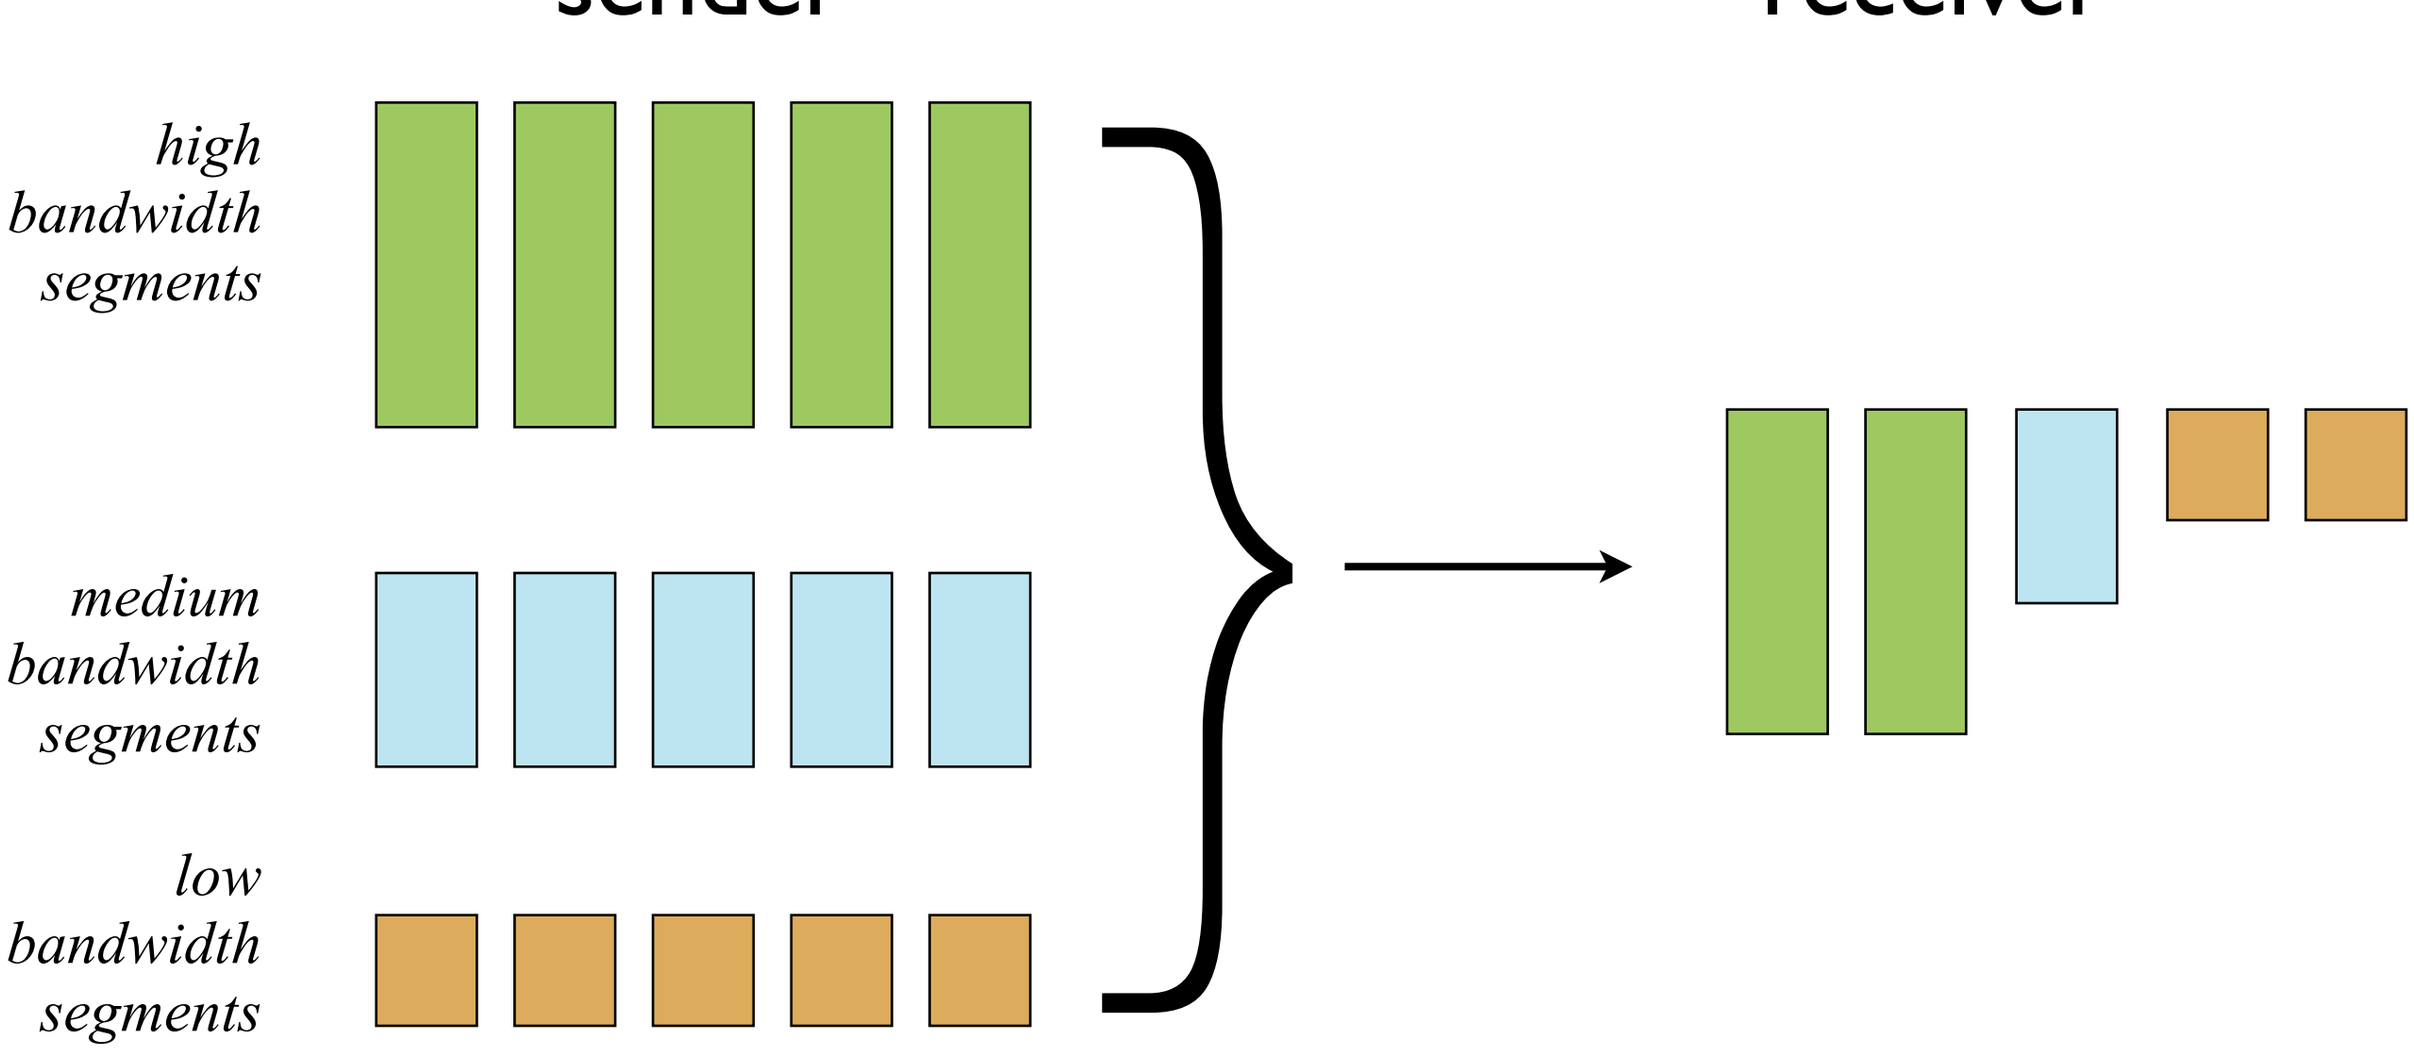

Using BPP

- What happens if we use BPP. Such an approach has a big impact on applications at the sender and at the receiver.
- If we use BPP for video, then some of the chunks in a packet will be dropped, rather than the whole packet.
- If we put Macro Blocks into chunks, then we still see loss at the receiver, and these will be decoded as grey squares.
- So not really a benefit.

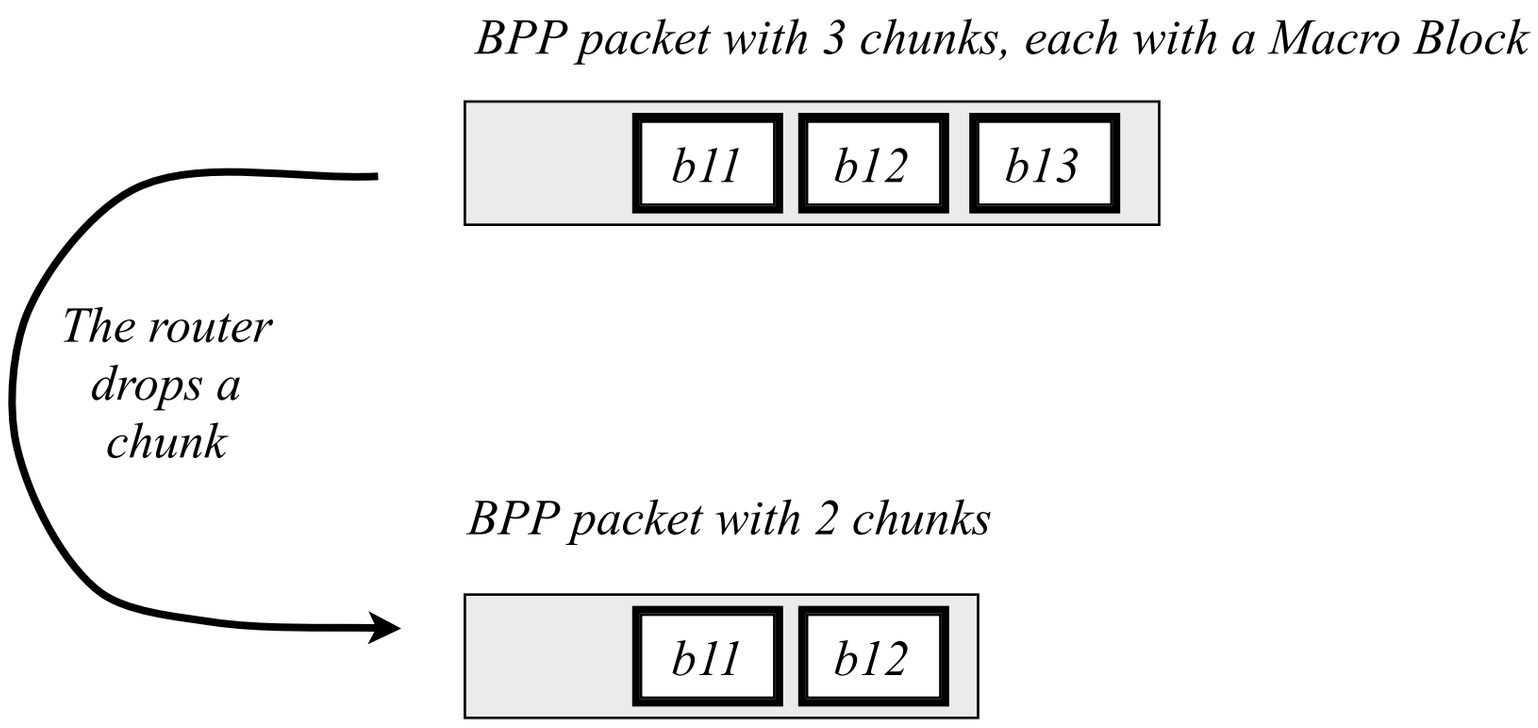

Using BPP

- What if we borrow some ideas from MPEG-DASH and have packets which contain multiple bitrate segments.
- The trouble is these segments are generally 2 seconds long and are far too big for a packet.
- Even if we choose much smaller segments, the video decoder is not designed for different bit-rate data per frame. They are designed for this over a longer time scale.

BPP packet with 3 chunks, each with a segment

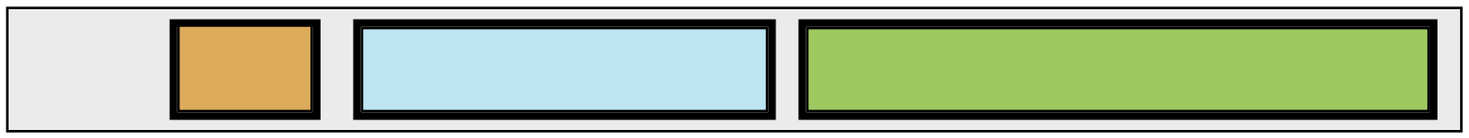

Way bigger than a packet

BPP packet with 3 chunks, each with a massively cut down segment, to fit in a packet

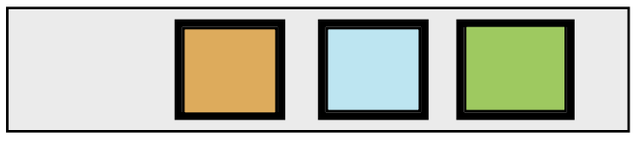

Decoder cannot deal with this

Enhancements for BPP

- To make BPP effective for continuous media, and for future schemes such as holographic transmissions, BPP will need to be coupled with an encoder and decoder that can do multiple encodings for the same region.
- These encodings can be put into a packet, and if any chunks are dropped, then the receiver will still have data.
- For video we would need a scaling multi-layer encoder that can embed multiple encodings into the same Macro Block.

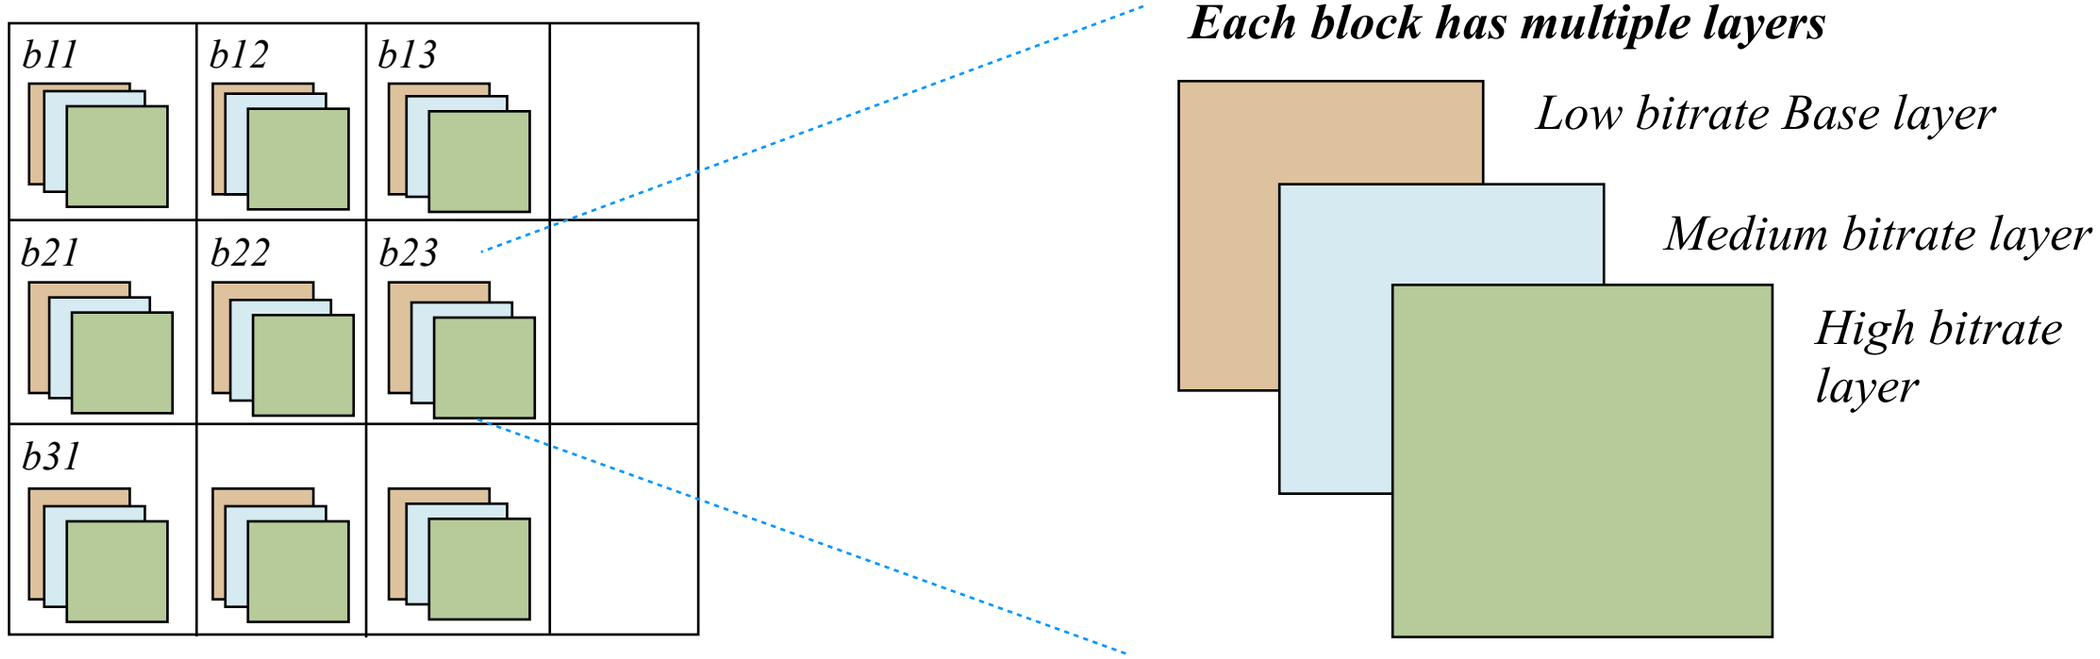

Enhancements for BPP

- With a multi-layer encoder, each packet could contain all of the encodings for that Macro Block, from Base, upwards.
- The router can drop chunks, but packets will always have the Base layer encoding, ensuring continuous video delivery.
- The decoder applies higher encodings onto the Base layer.
- The receiver always gets a packet with some video data.

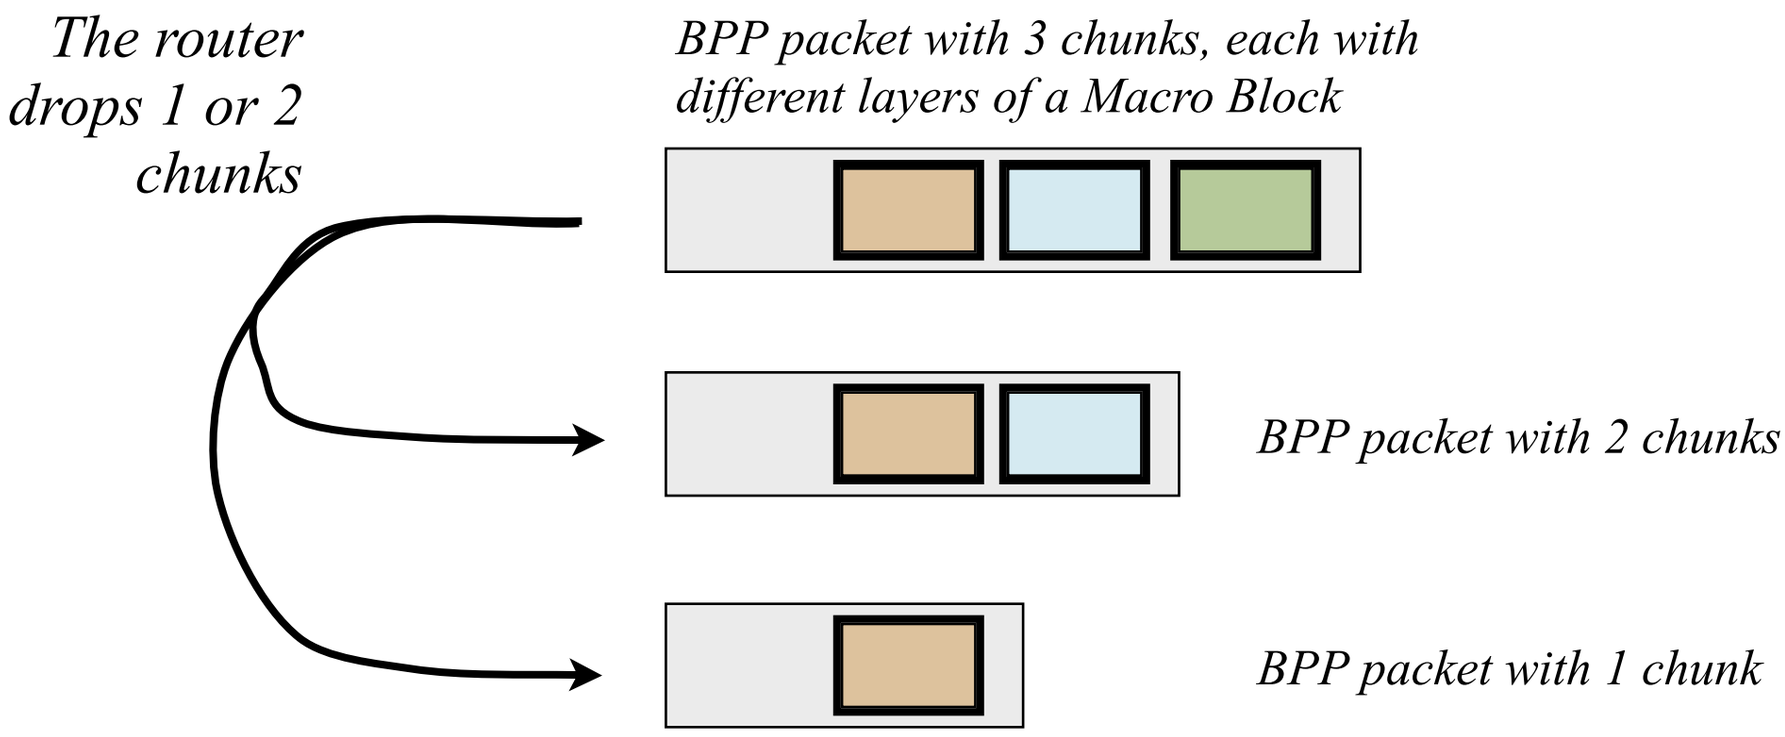

Combining Techniques

- We can add dynamicness and adaptability by having a feedback loop from the receiver to the sender so that it can inform the number of ***or*** %age of lost chunks.
- The sender can create packets with a suitable amount of chunking based on the network congestion.
- This approach combines the idea of live feedback with BPP.

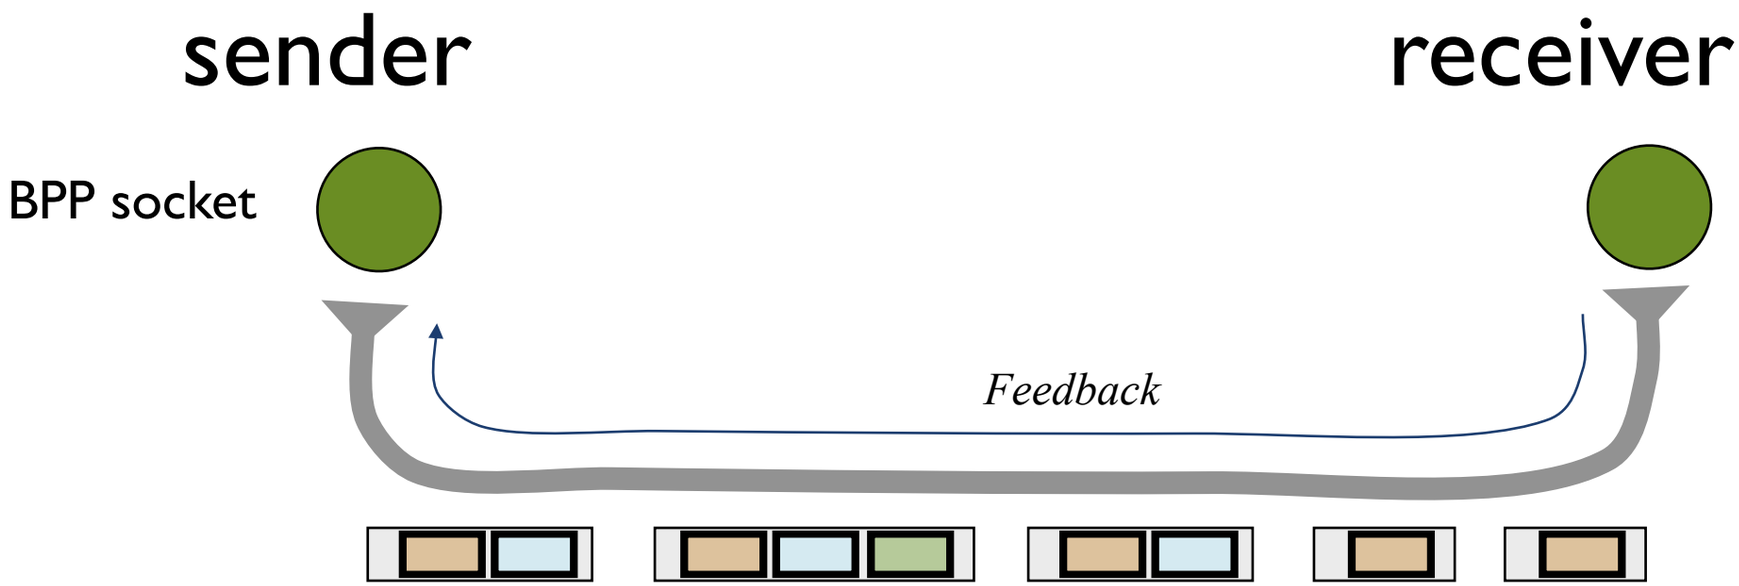

BPP packets with variable number of chunks

Conclusions

- Services and the network are intrinsically linked. If you try to add special features to one and overlook the behaviour of the other, the outcome can be unexpected.
- For the future use of BPP, we will need applications that are aware of BPP's behaviour.
- To access BPP from the application space, we will need a new kind of Socket that accesses the BPP programmable functionality and send BPP commands.
- The applications will need to BPP aware, and construct packets with the correct structure.
- We will need a feedback loop from receiver to sender so that it can create suitable amount of chunking based on the network congestion, providing live dynamics.

Spares

Frames and Blocks

- Digital Video Stream
- Unlike film, each frame does not need to cover the whole area. There are partial frames and update frames that are deltas to the previous frame.
- Talk about encoding / compression
- Mention I, P, B frames