

# TCP Goes to Hollywood

Stephen McQuistin and Colin Perkins  
University of Glasgow

Marwan Fayed  
University of Stirling

# Multimedia Applications and HTTP

- HTTP-based multimedia delivery protocols (e.g., MPEG-DASH, HLS) are popular
- They allow applications to make use of the existing HTTP infrastructure (e.g., CDNs)
- These protocols can be configured for low latency applications: smaller segment sizes and buffers, for example
- ... but latency is added at the transport layer

# TCP adds latency

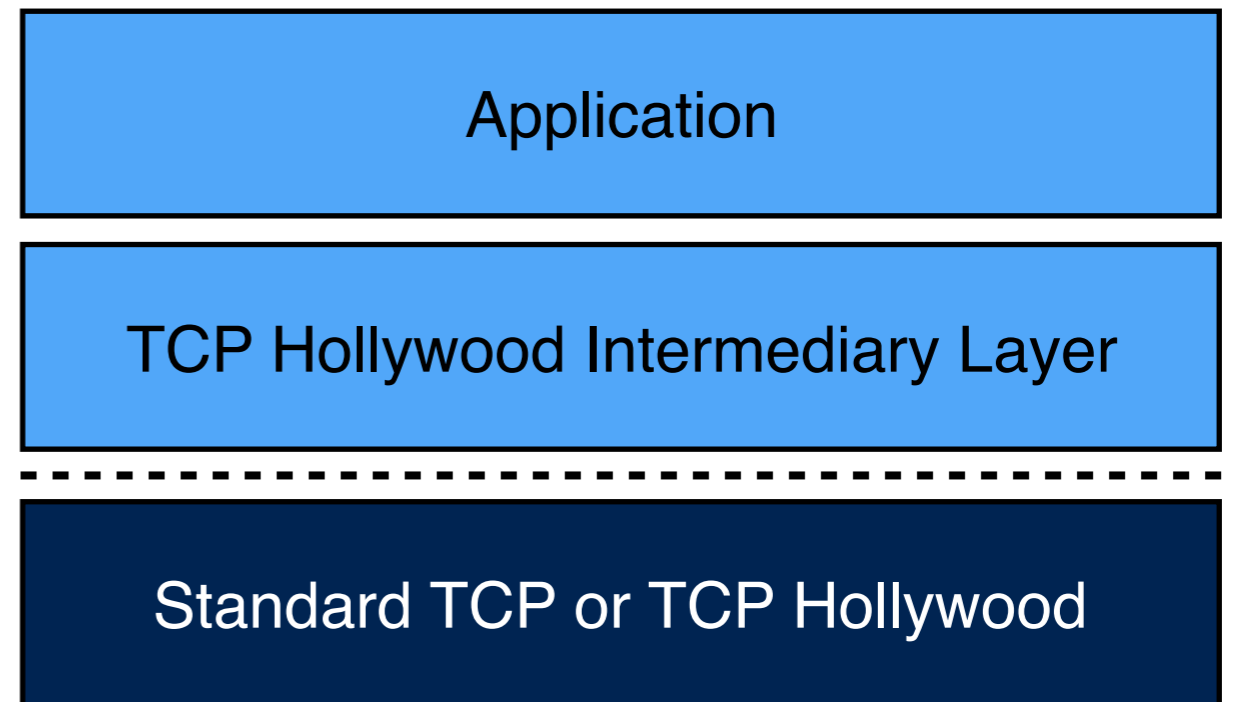
- In-order delivery means buffering out-of-order segments, waiting for the delivery of earlier data
- Reliability involves detecting that a segment has been lost, and retransmitting it
- Both of these mechanisms add latency, making TCP a poor choice for real-time multimedia applications

# Introducing TCP Hollywood

- Uses TCP as a substrate, to overcome ossification, but modified to reduce latency
- Message-oriented to allow application data units to be sent
- Unordered delivery of messages, given independent utility
- Partially reliable based on time and dependency information

# Architecture

- Functionality split between user-space intermediary layer, and kernel extensions
- Intermediary layer works over either standard TCP or TCP Hollywood
- Supports partial deployments




# Unordered message delivery

- Builds on standard TCP's byte stream, so need to frame messages
- At sender, applications pass messages to intermediary layer, for encoding (to escape framing bytes), and framing
- Nagle's algorithm disabled — it would add latency
- At receiver, incoming segments delivered as they arrive, decoded, and passed to application — no latency added by buffering
- ACKs generated as with standard TCP

# Partial reliability

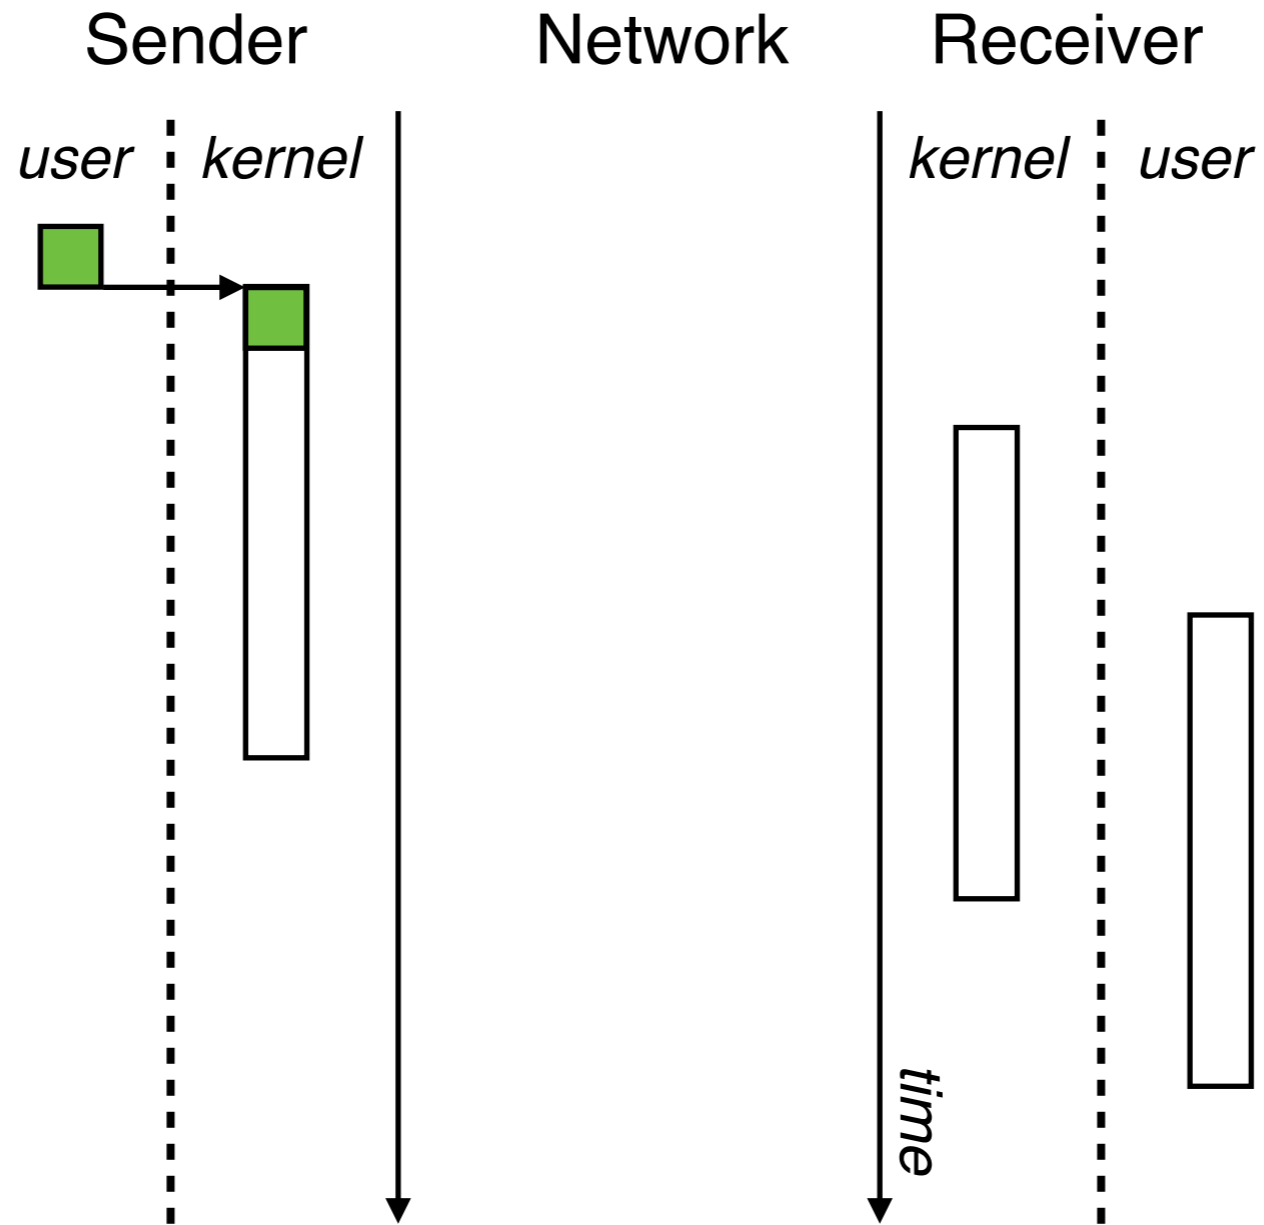
- Applications pass metadata with messages: deadline and dependency information
- Messages might expire — either they won't arrive on time to be played out, or they depend on an undelivered message
- Expired messages aren't retransmitted — a live message is sent instead, as an *inconsistent retransmission*
- Payload is different, but with the same TCP sequence number and length
- Recovers utility lost by retransmitting expired data

# TCP Hollywood in action






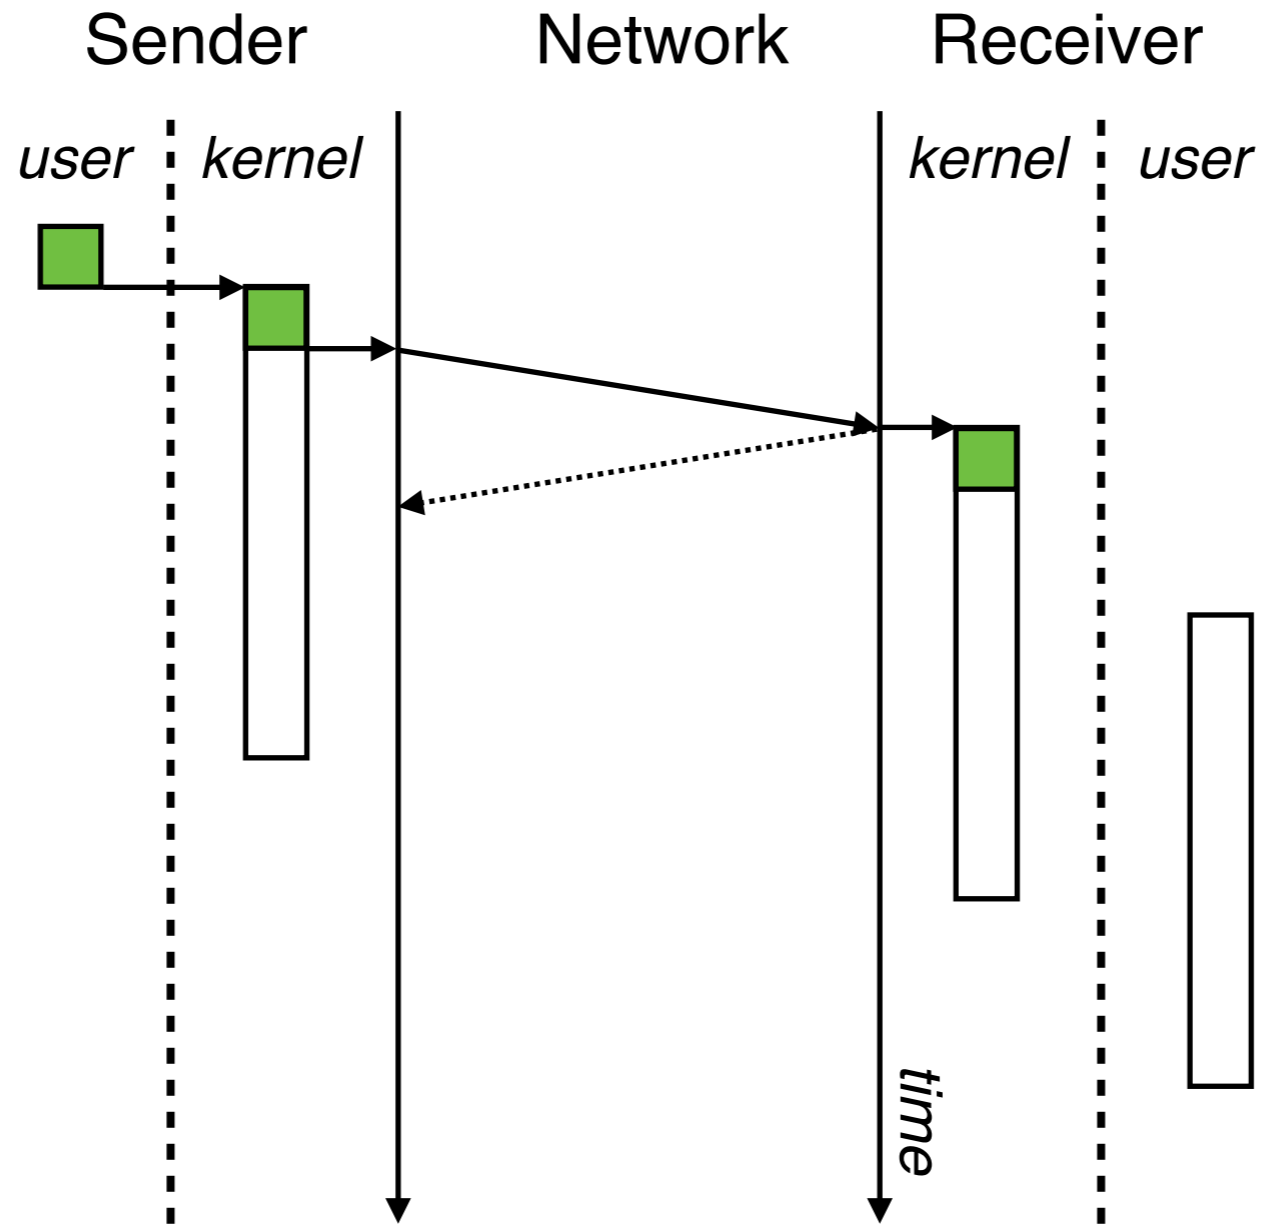
# TCP Hollywood in action



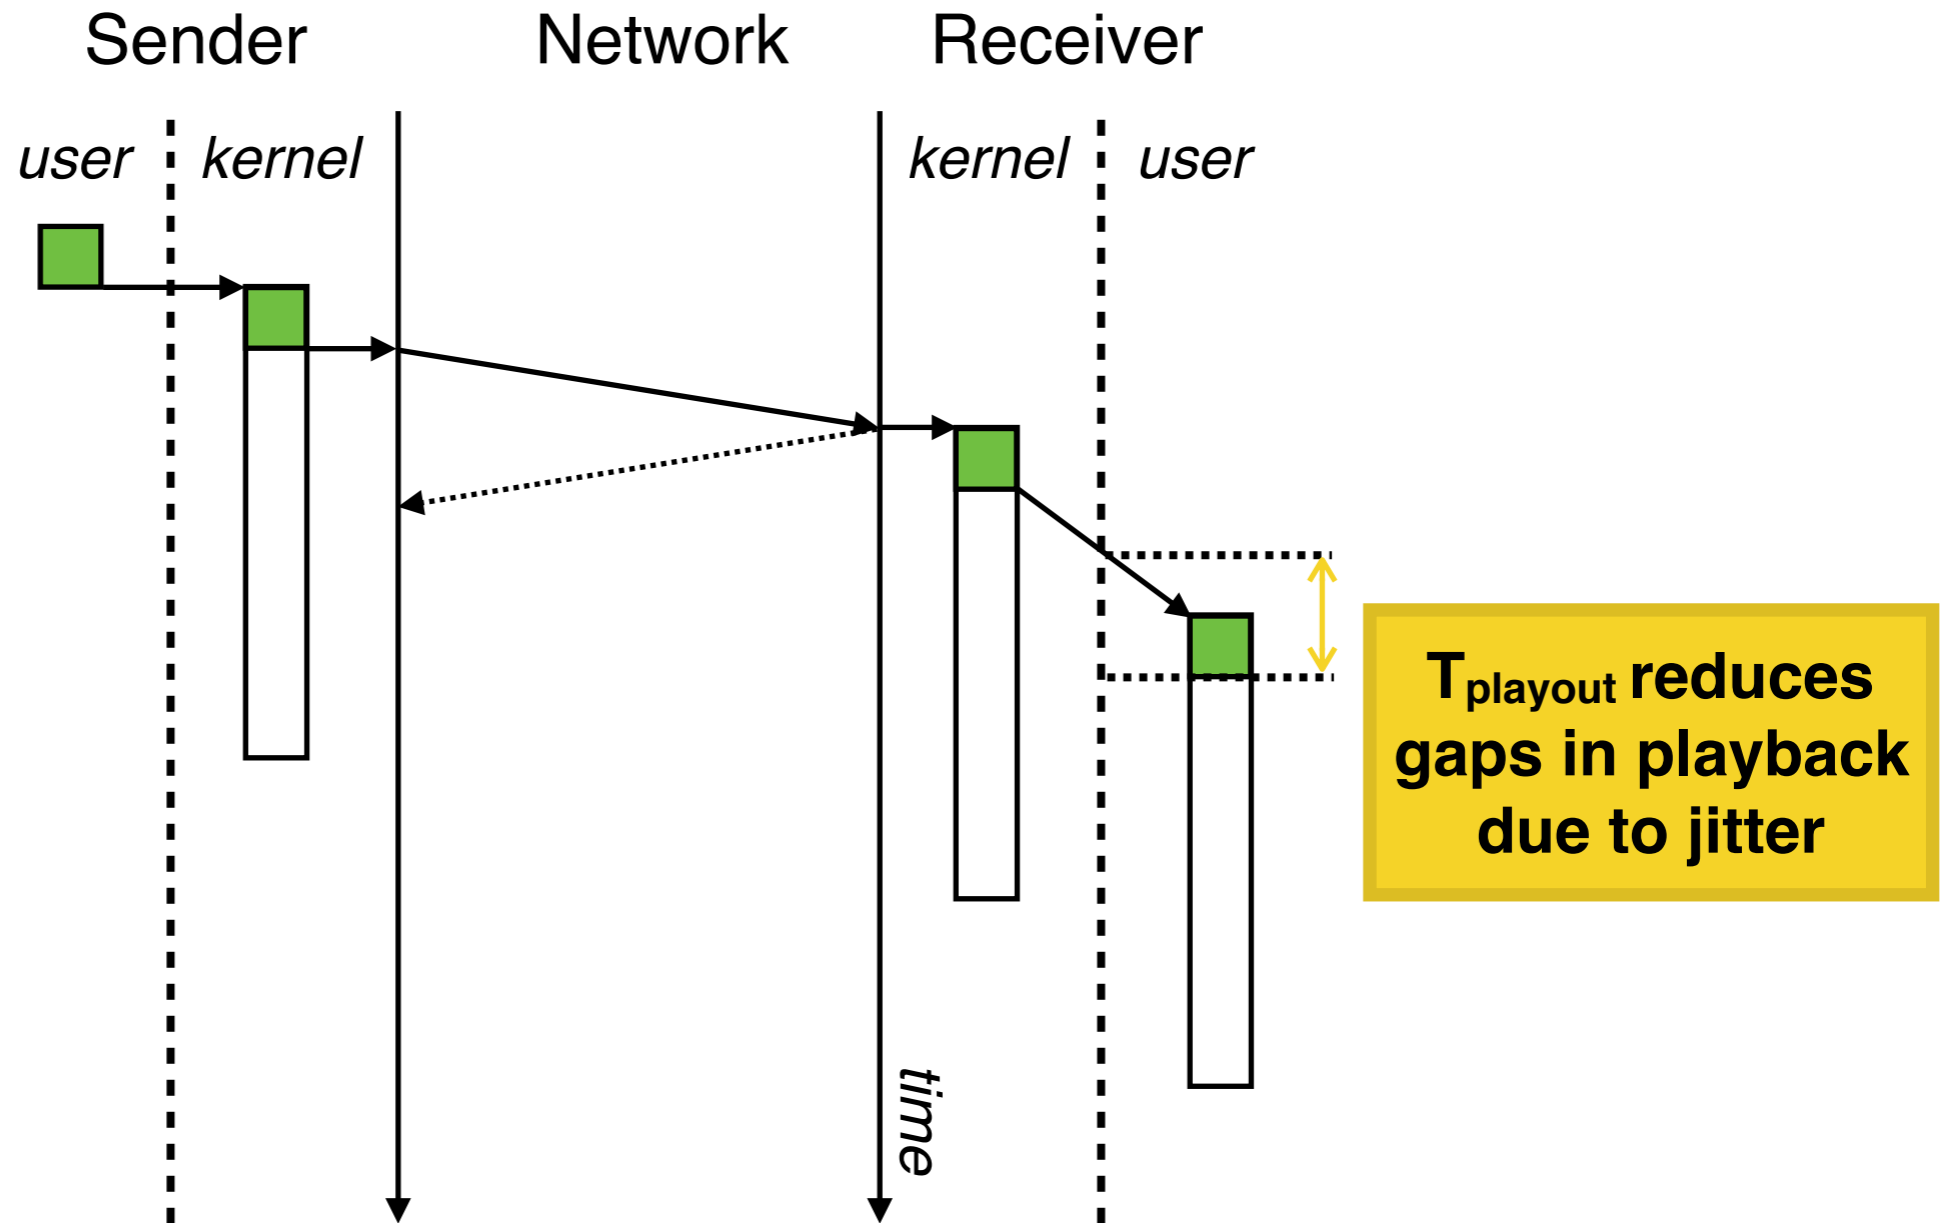
# TCP Hollywood in action



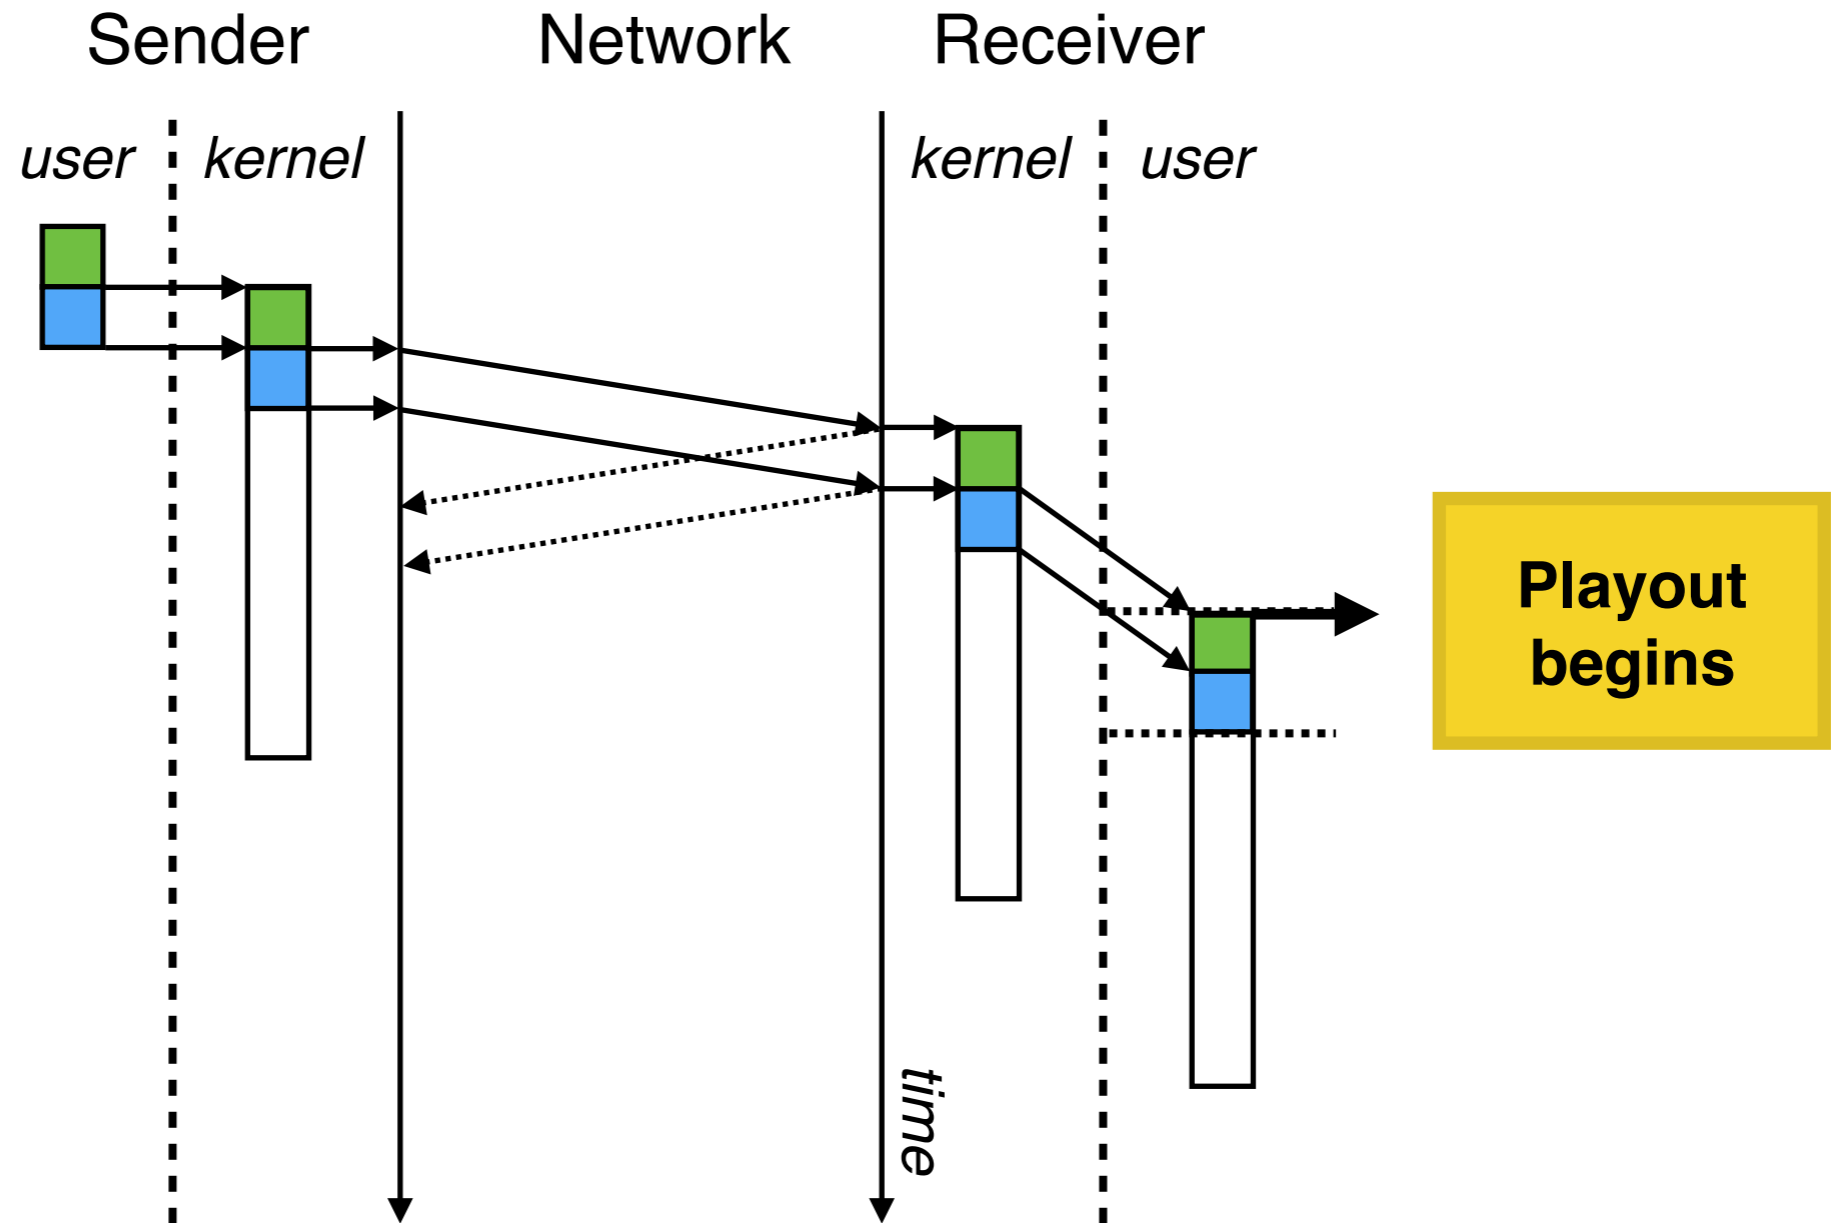
# TCP Hollywood in action



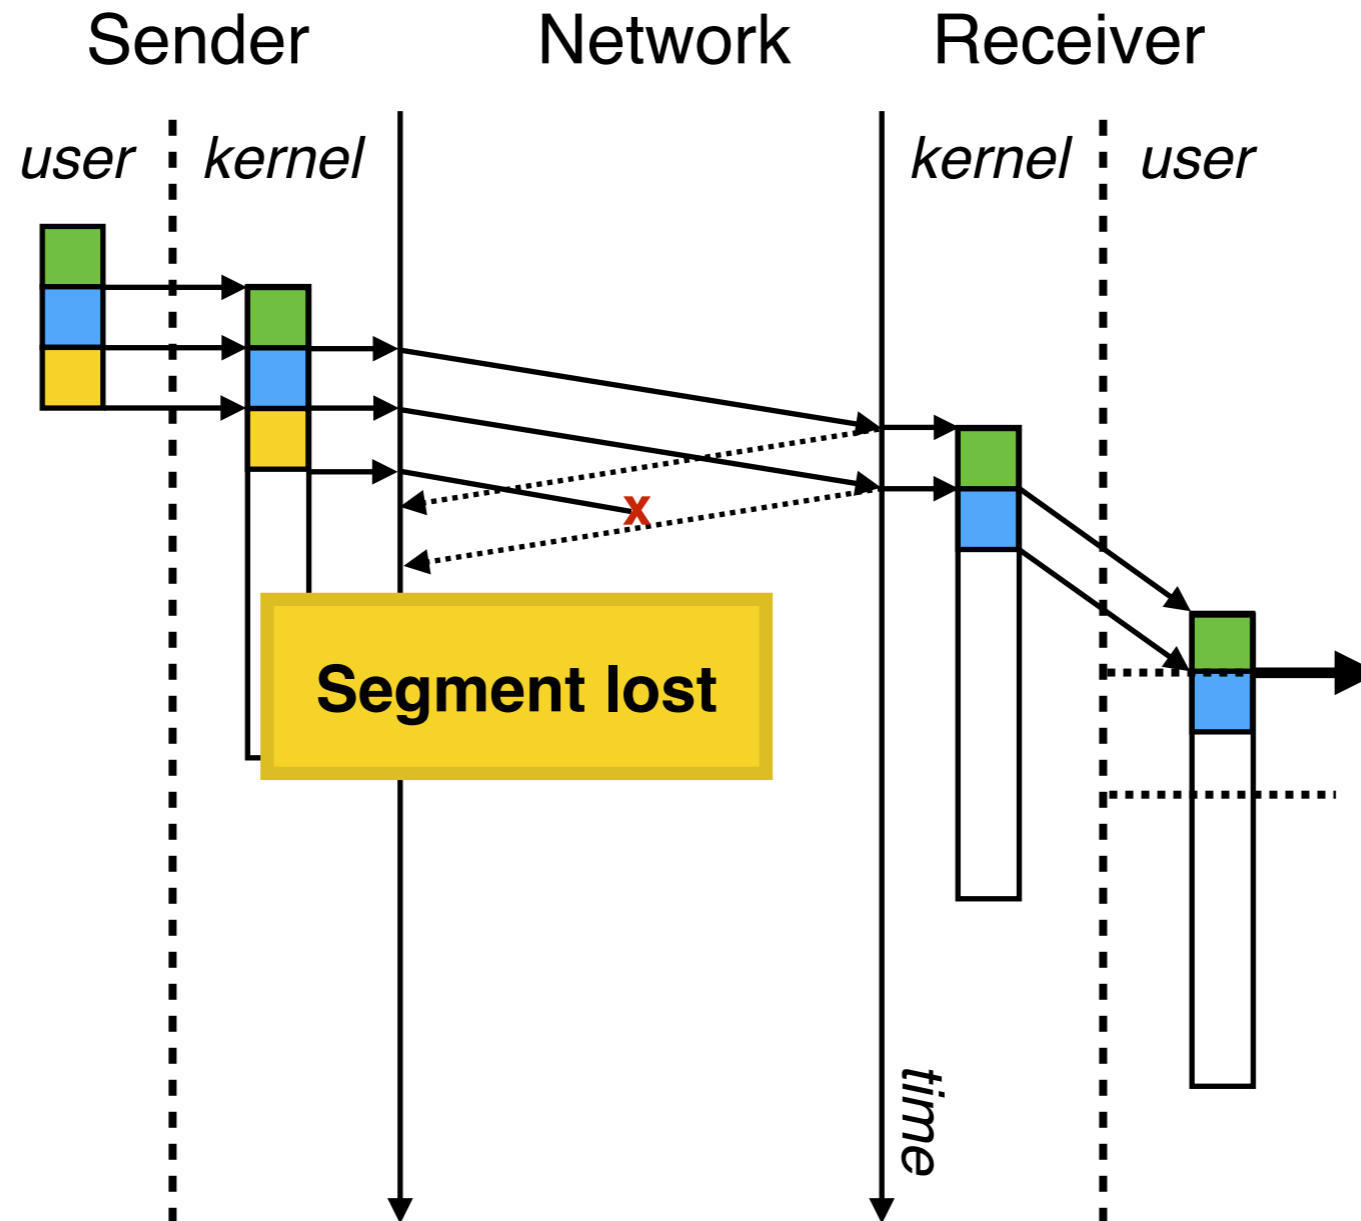
# TCP Hollywood in action




# TCP Hollywood in action




# TCP Hollywood in action



# TCP Hollywood in action

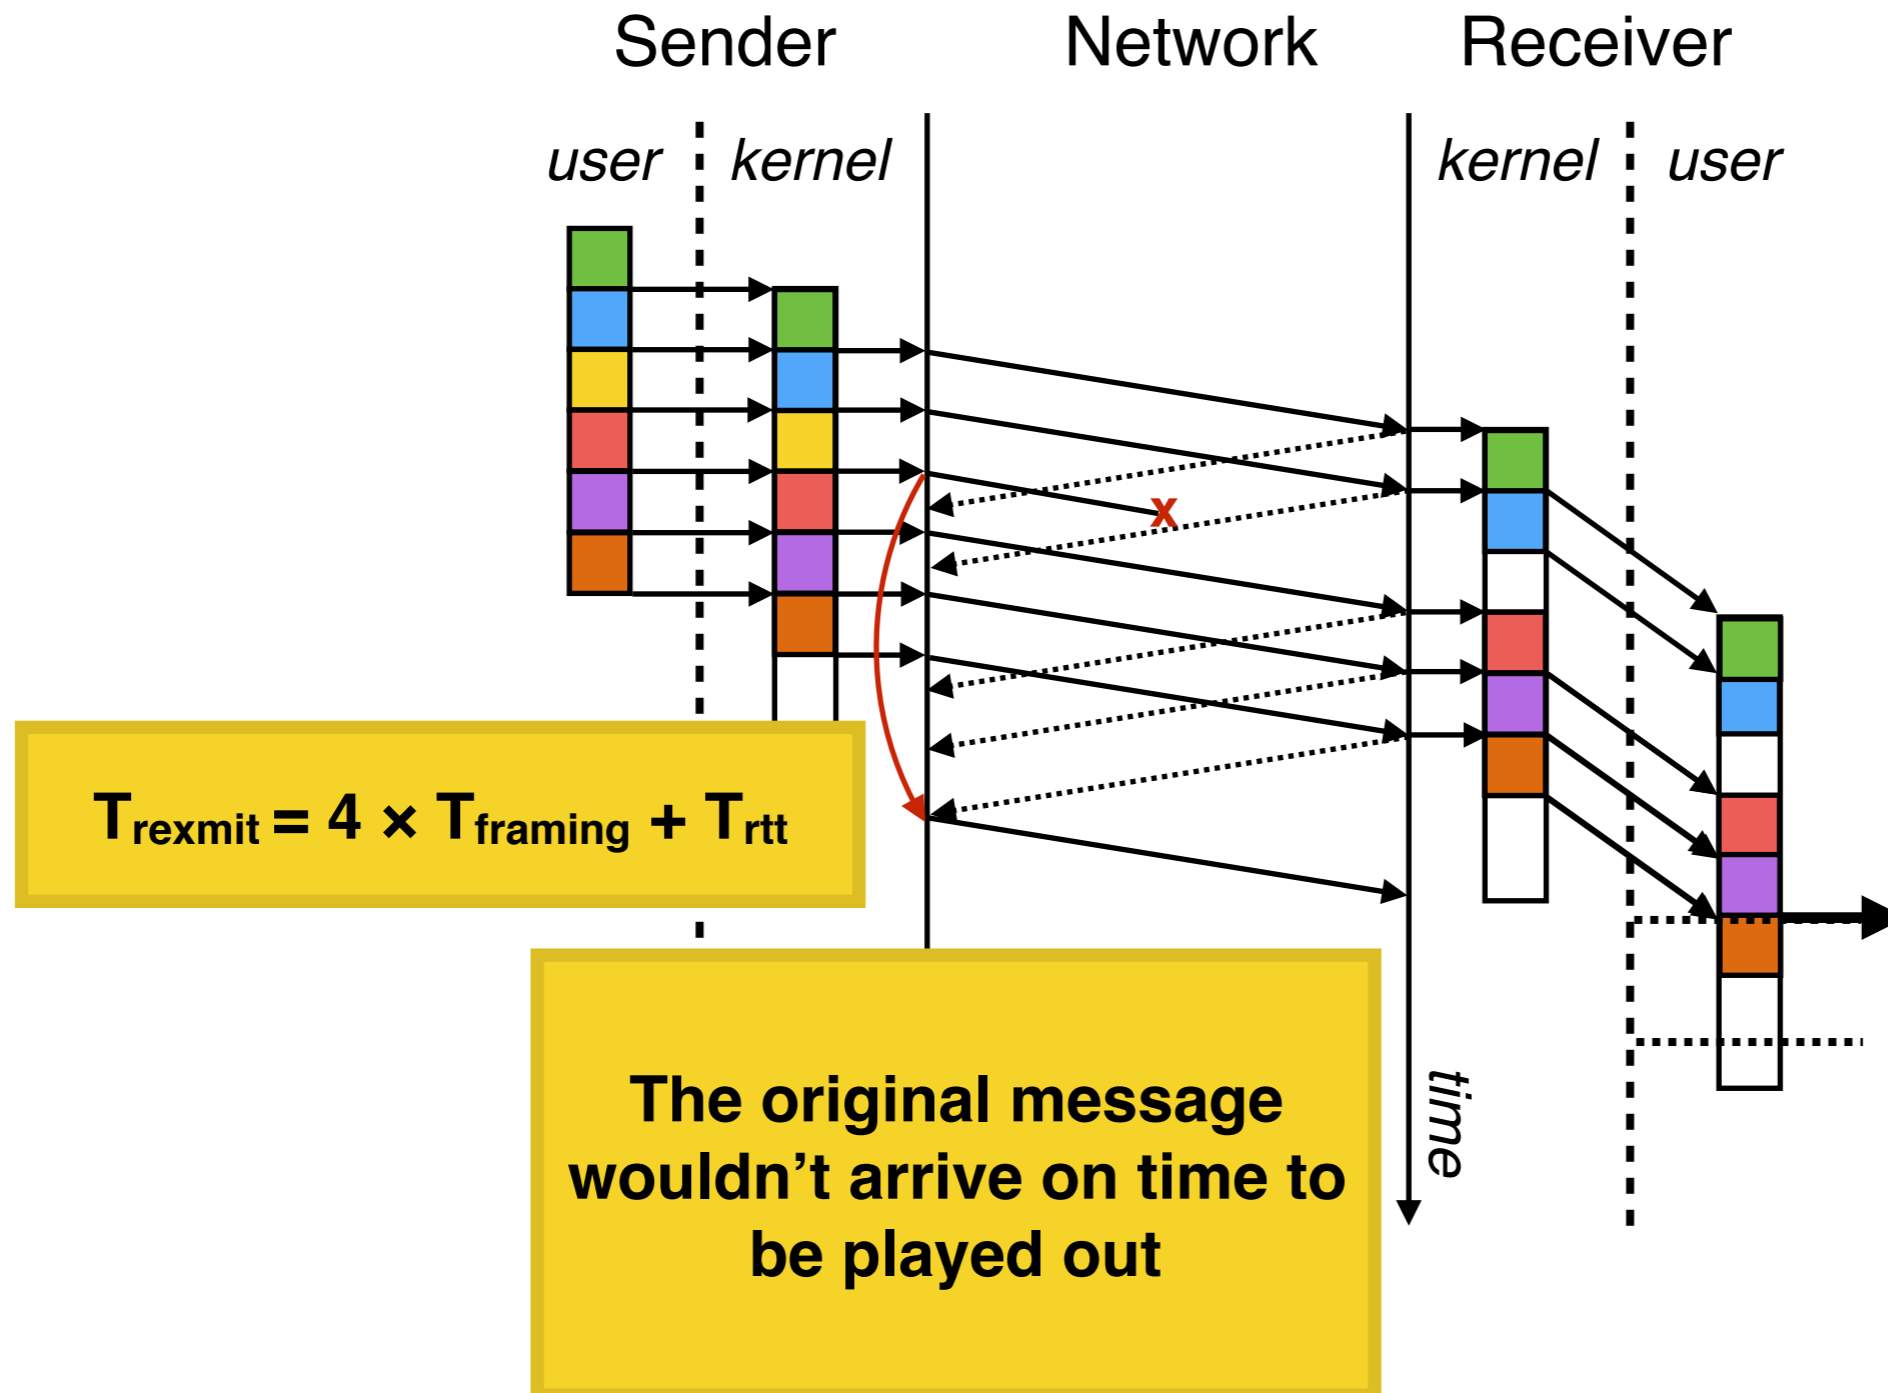


# TCP Hollywood in action

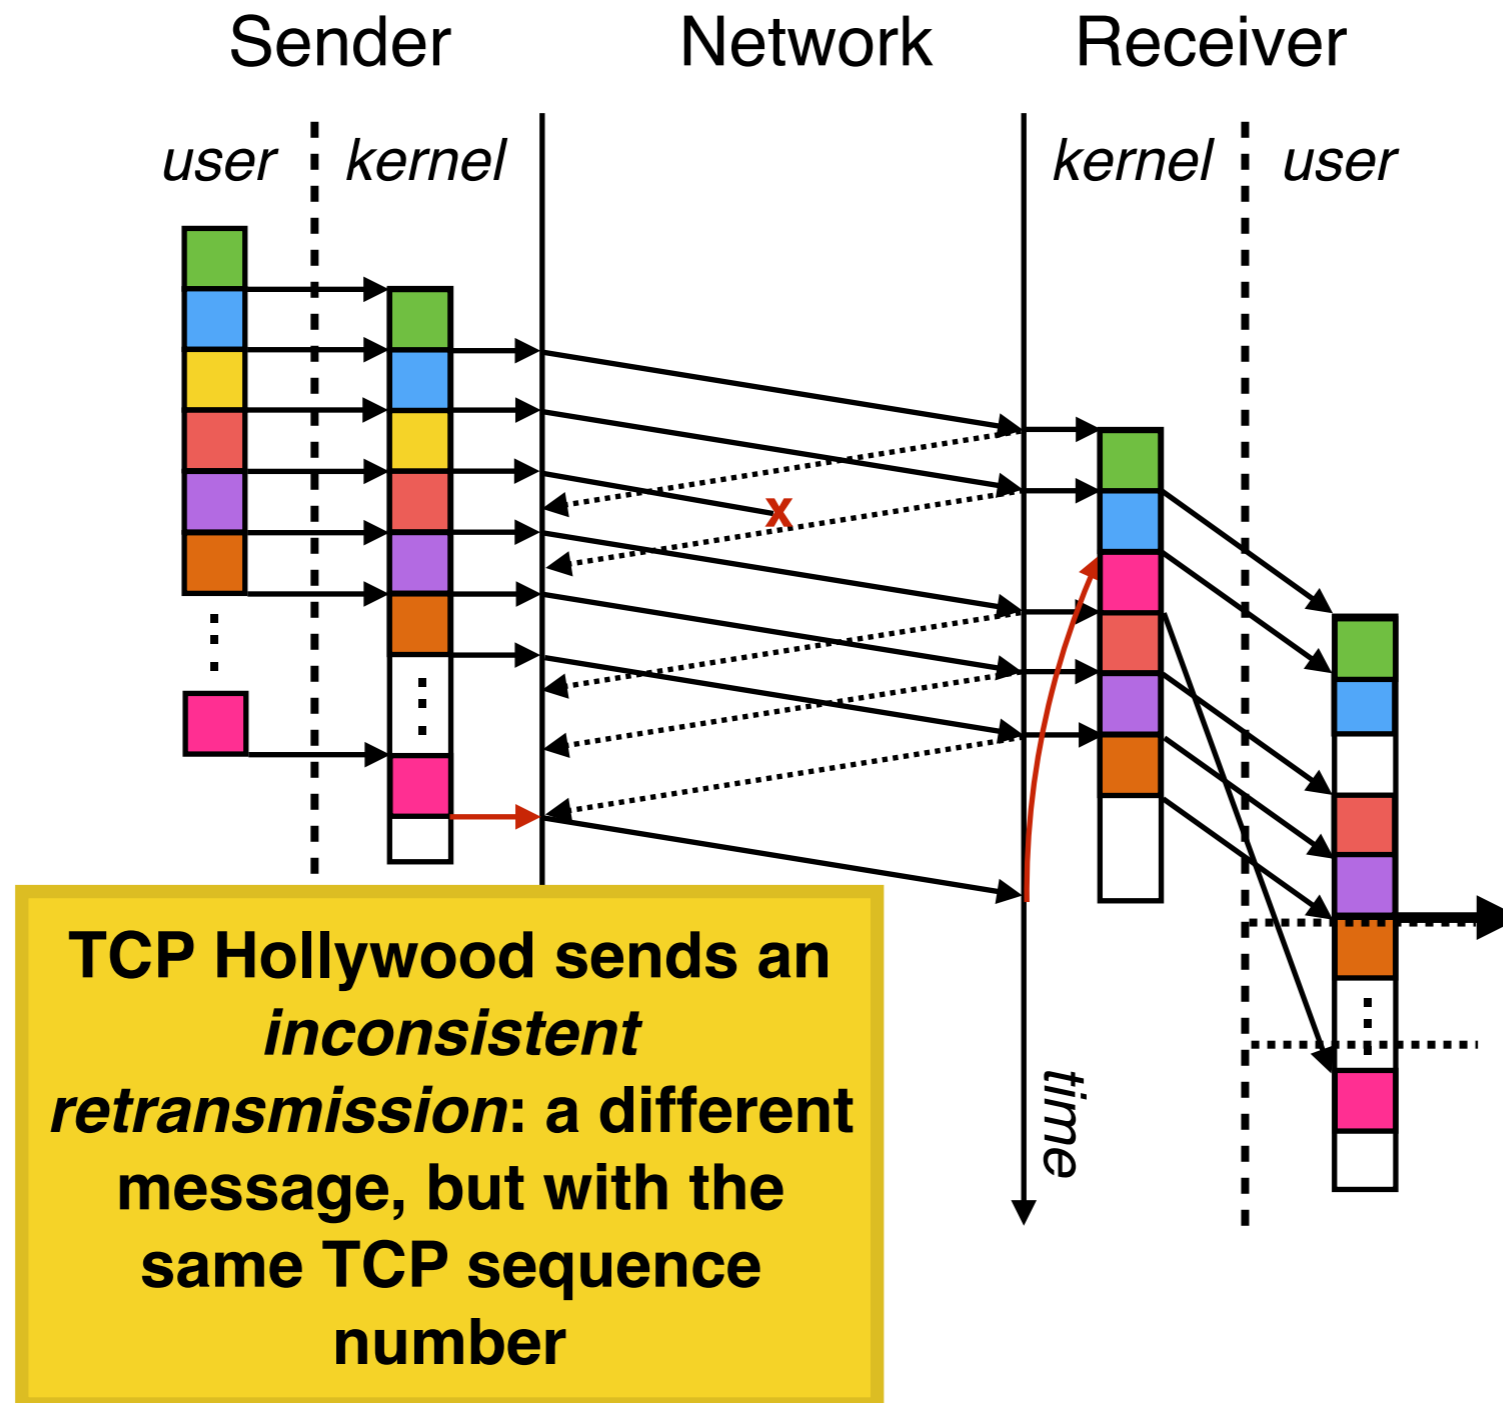




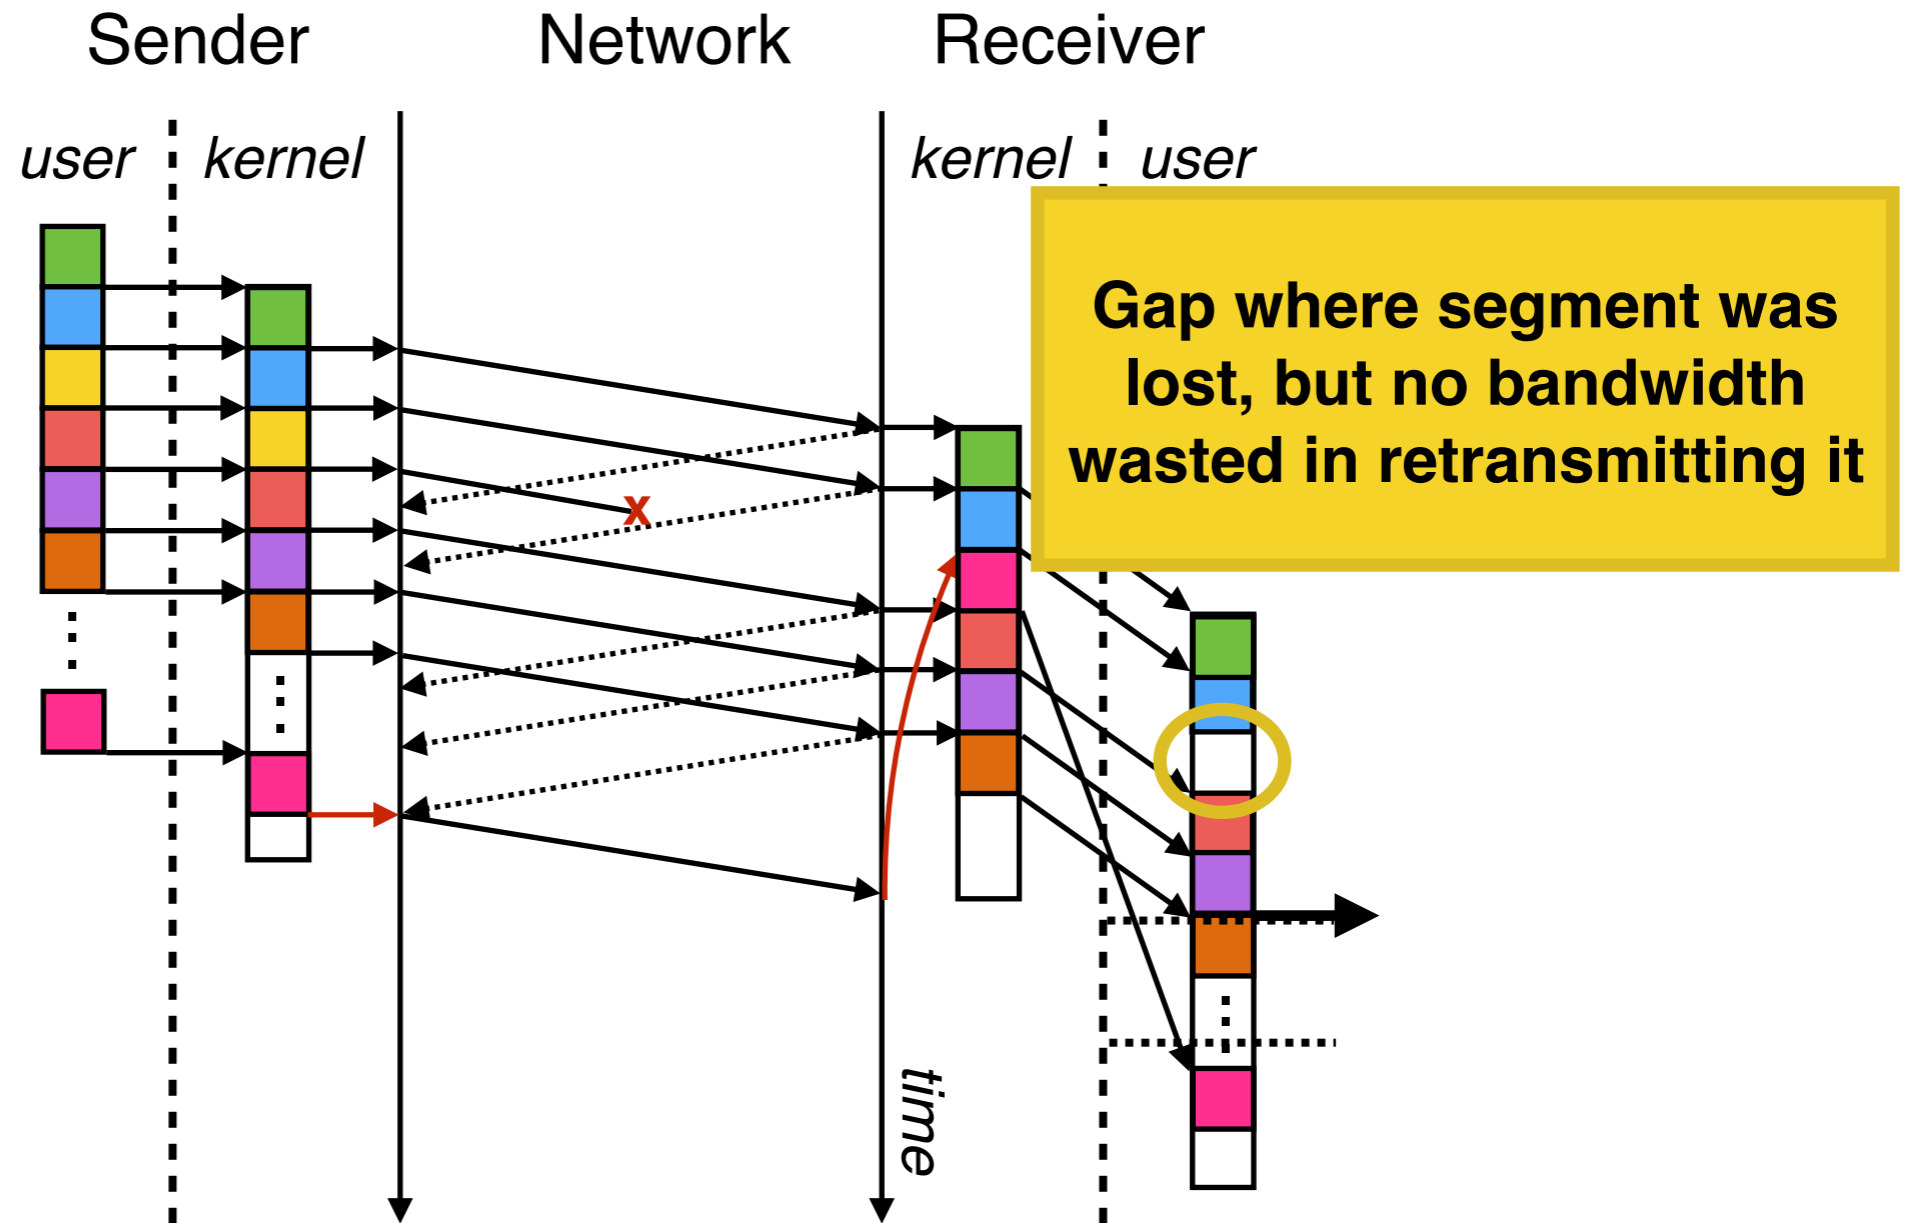
# TCP Hollywood in action




# TCP Hollywood in action



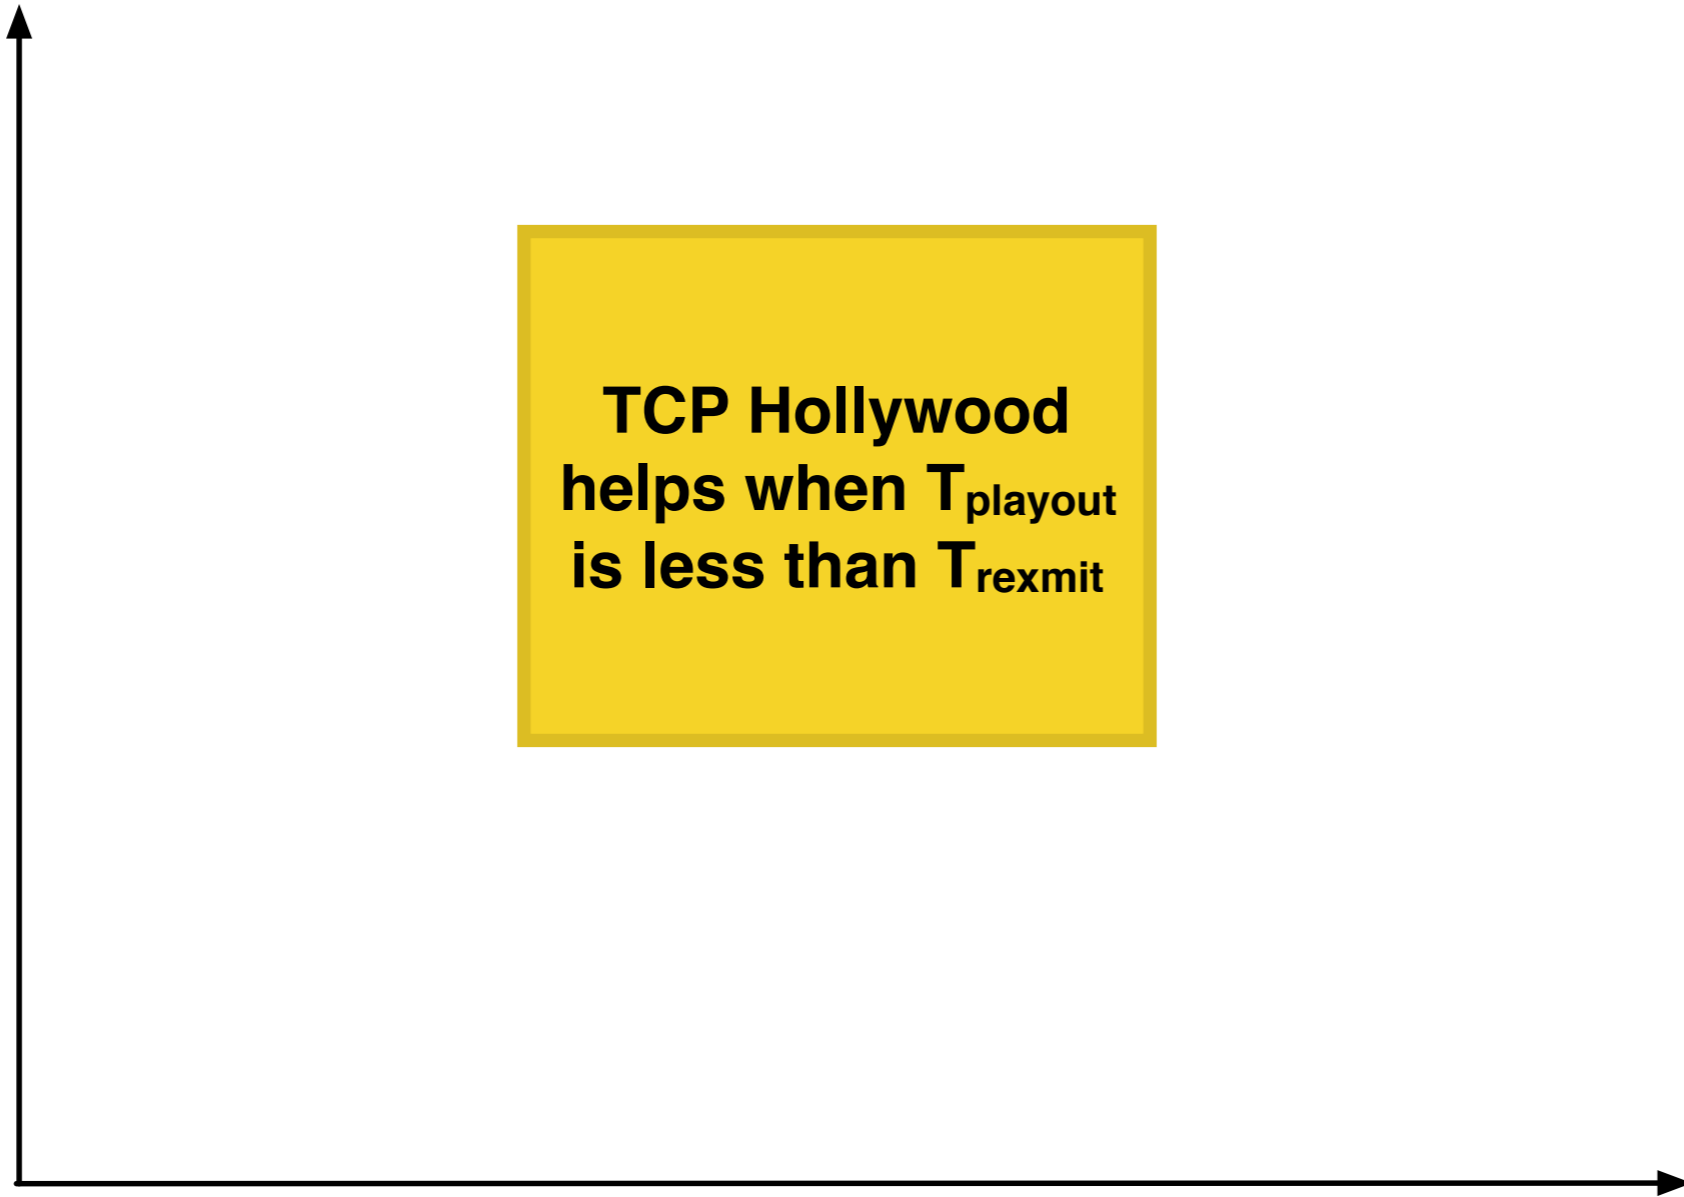
# TCP Hollywood in action



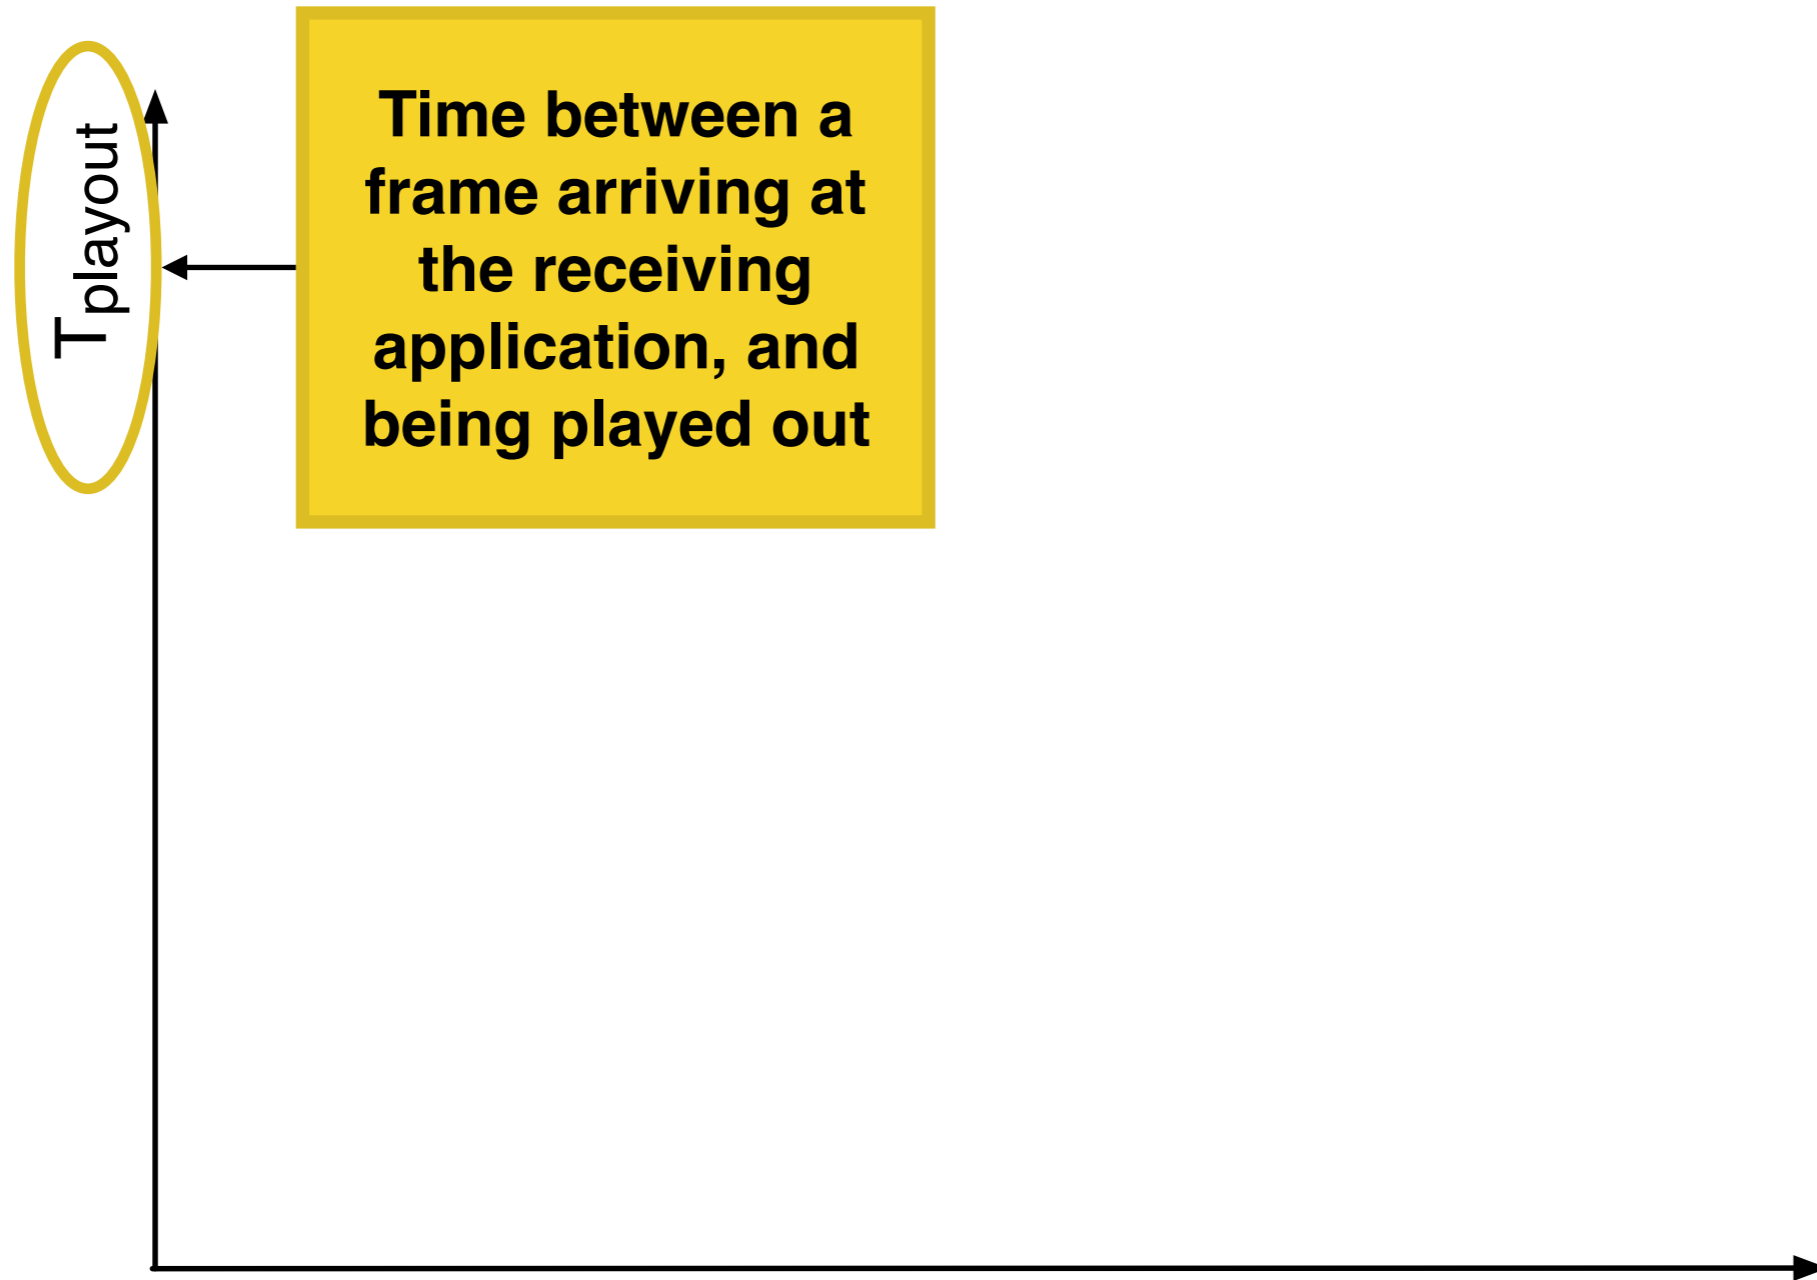
# TCP Hollywood in action



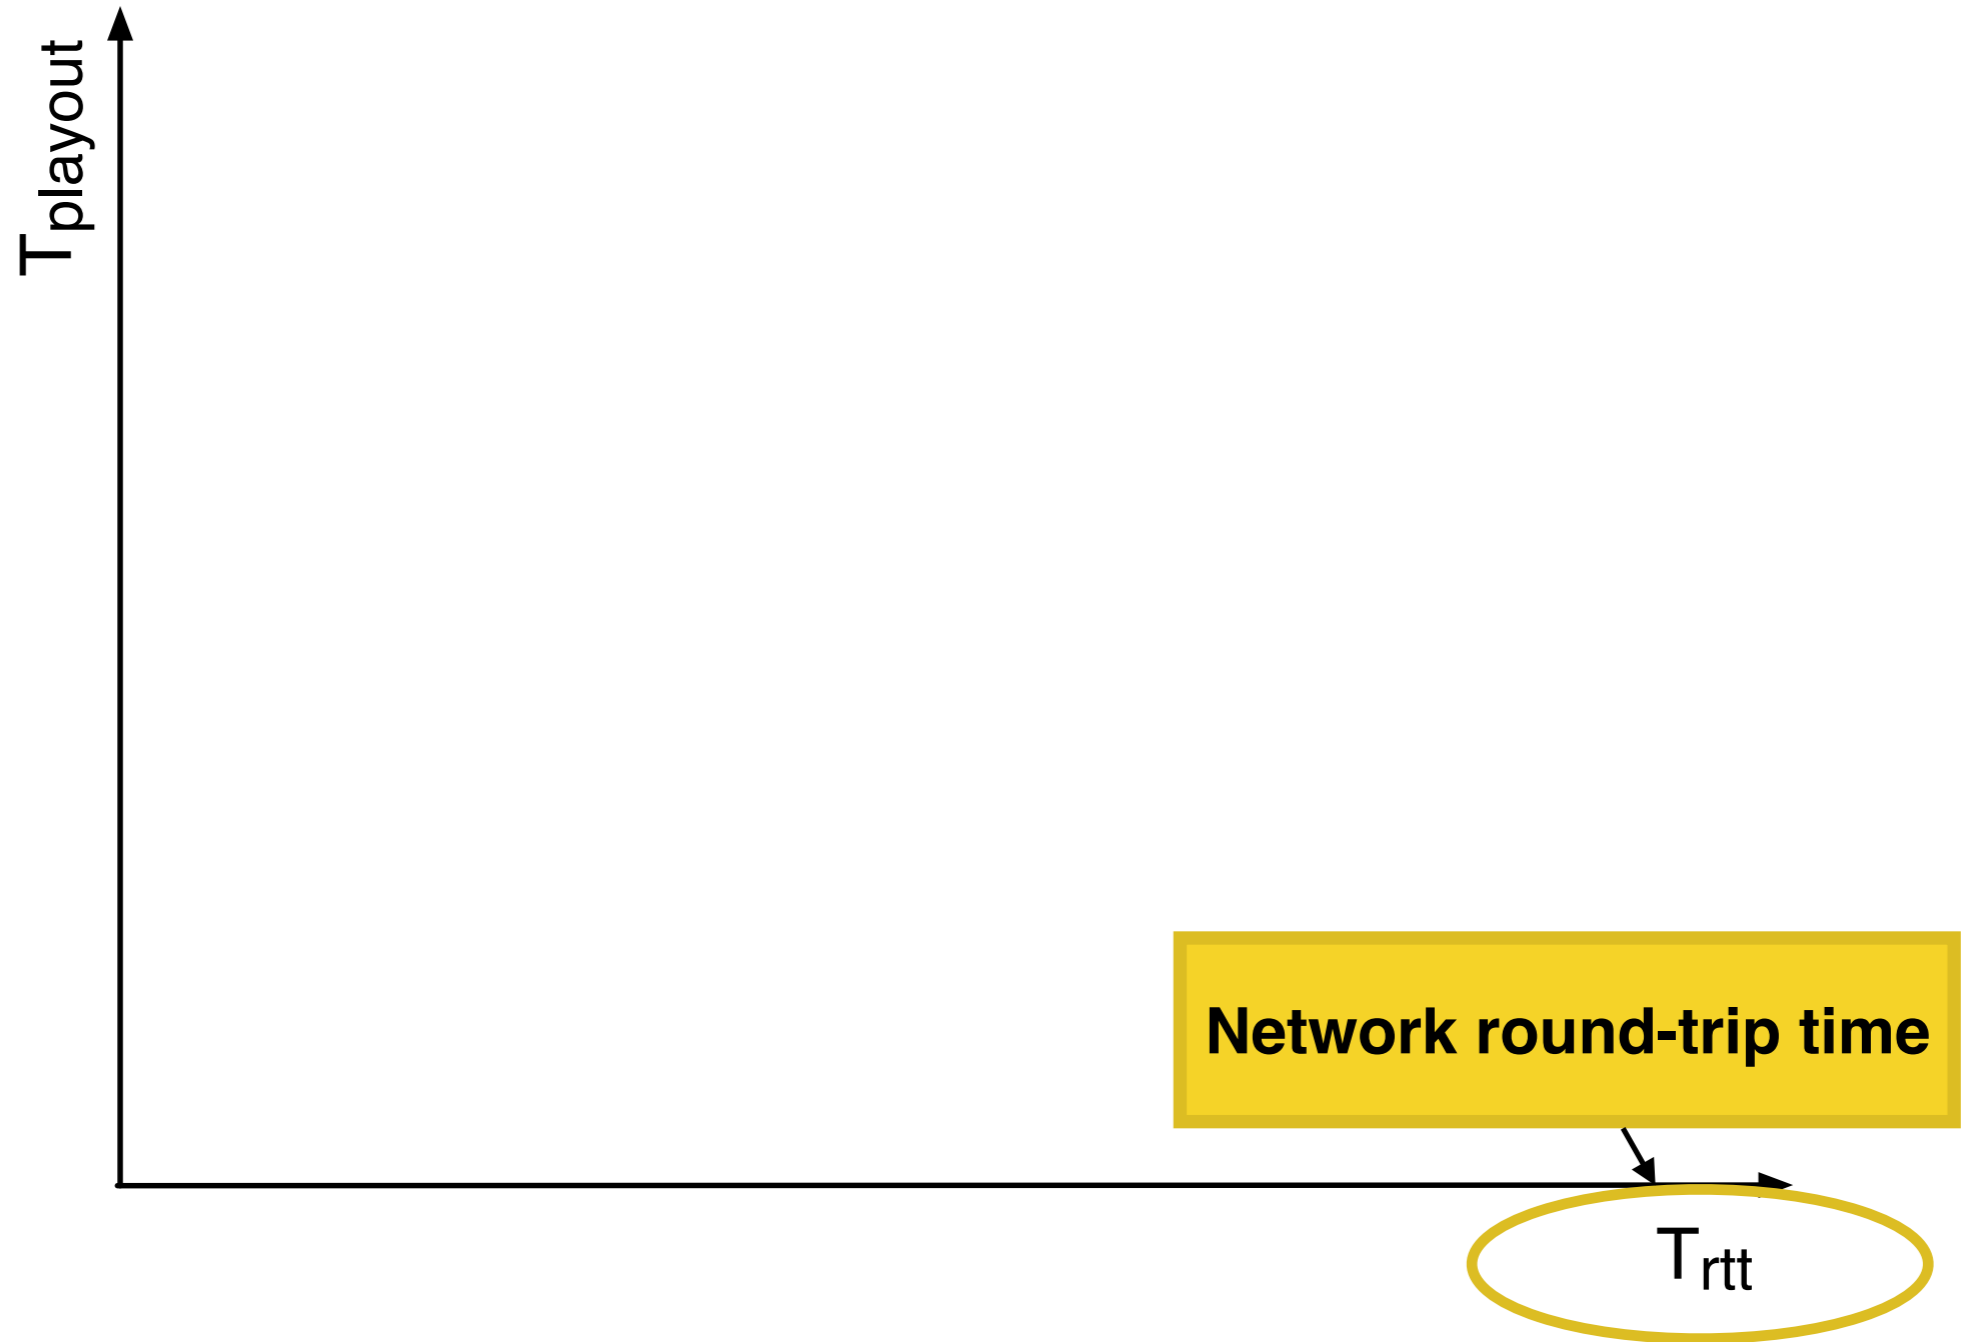
# Feasibility Region



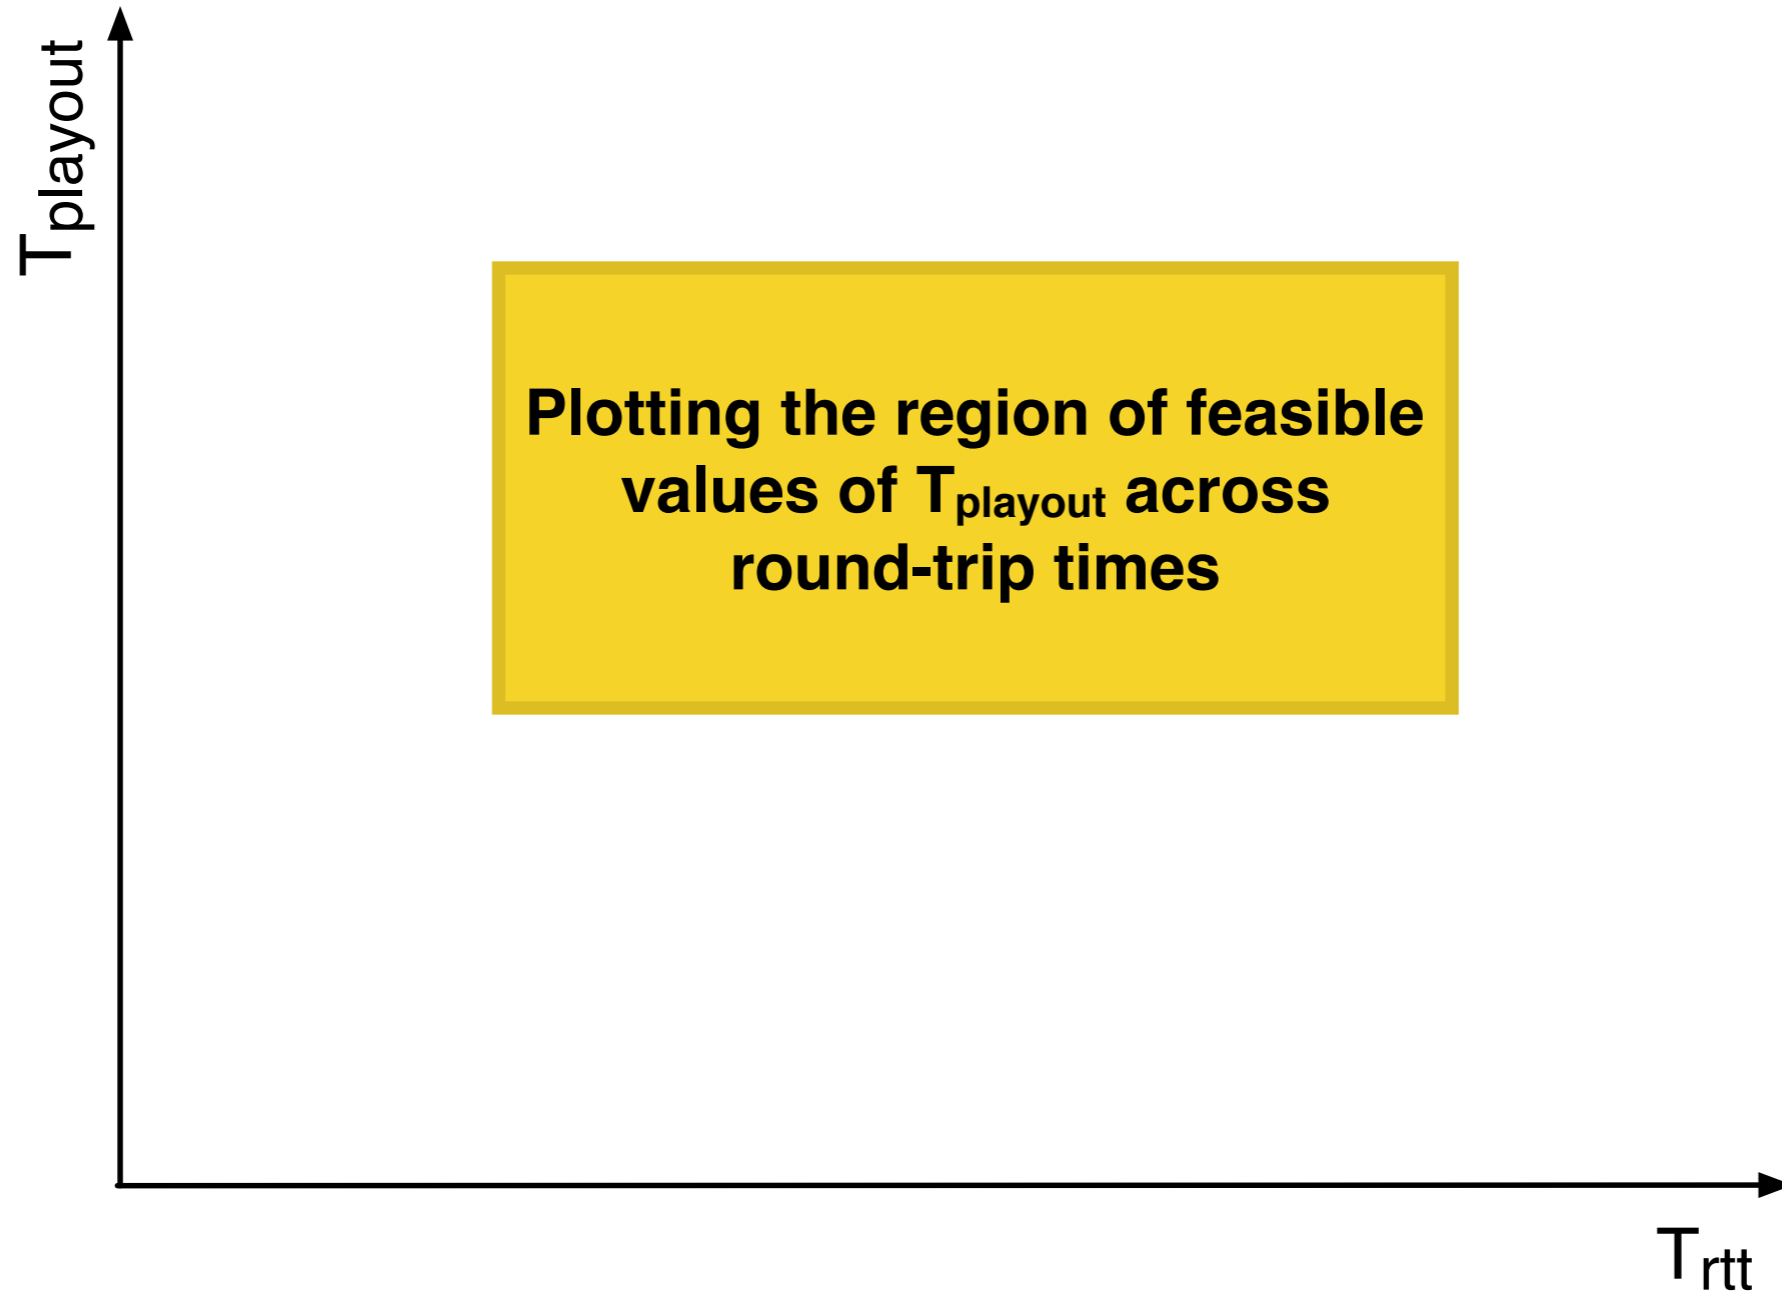
# Feasibility Region



# Feasibility Region

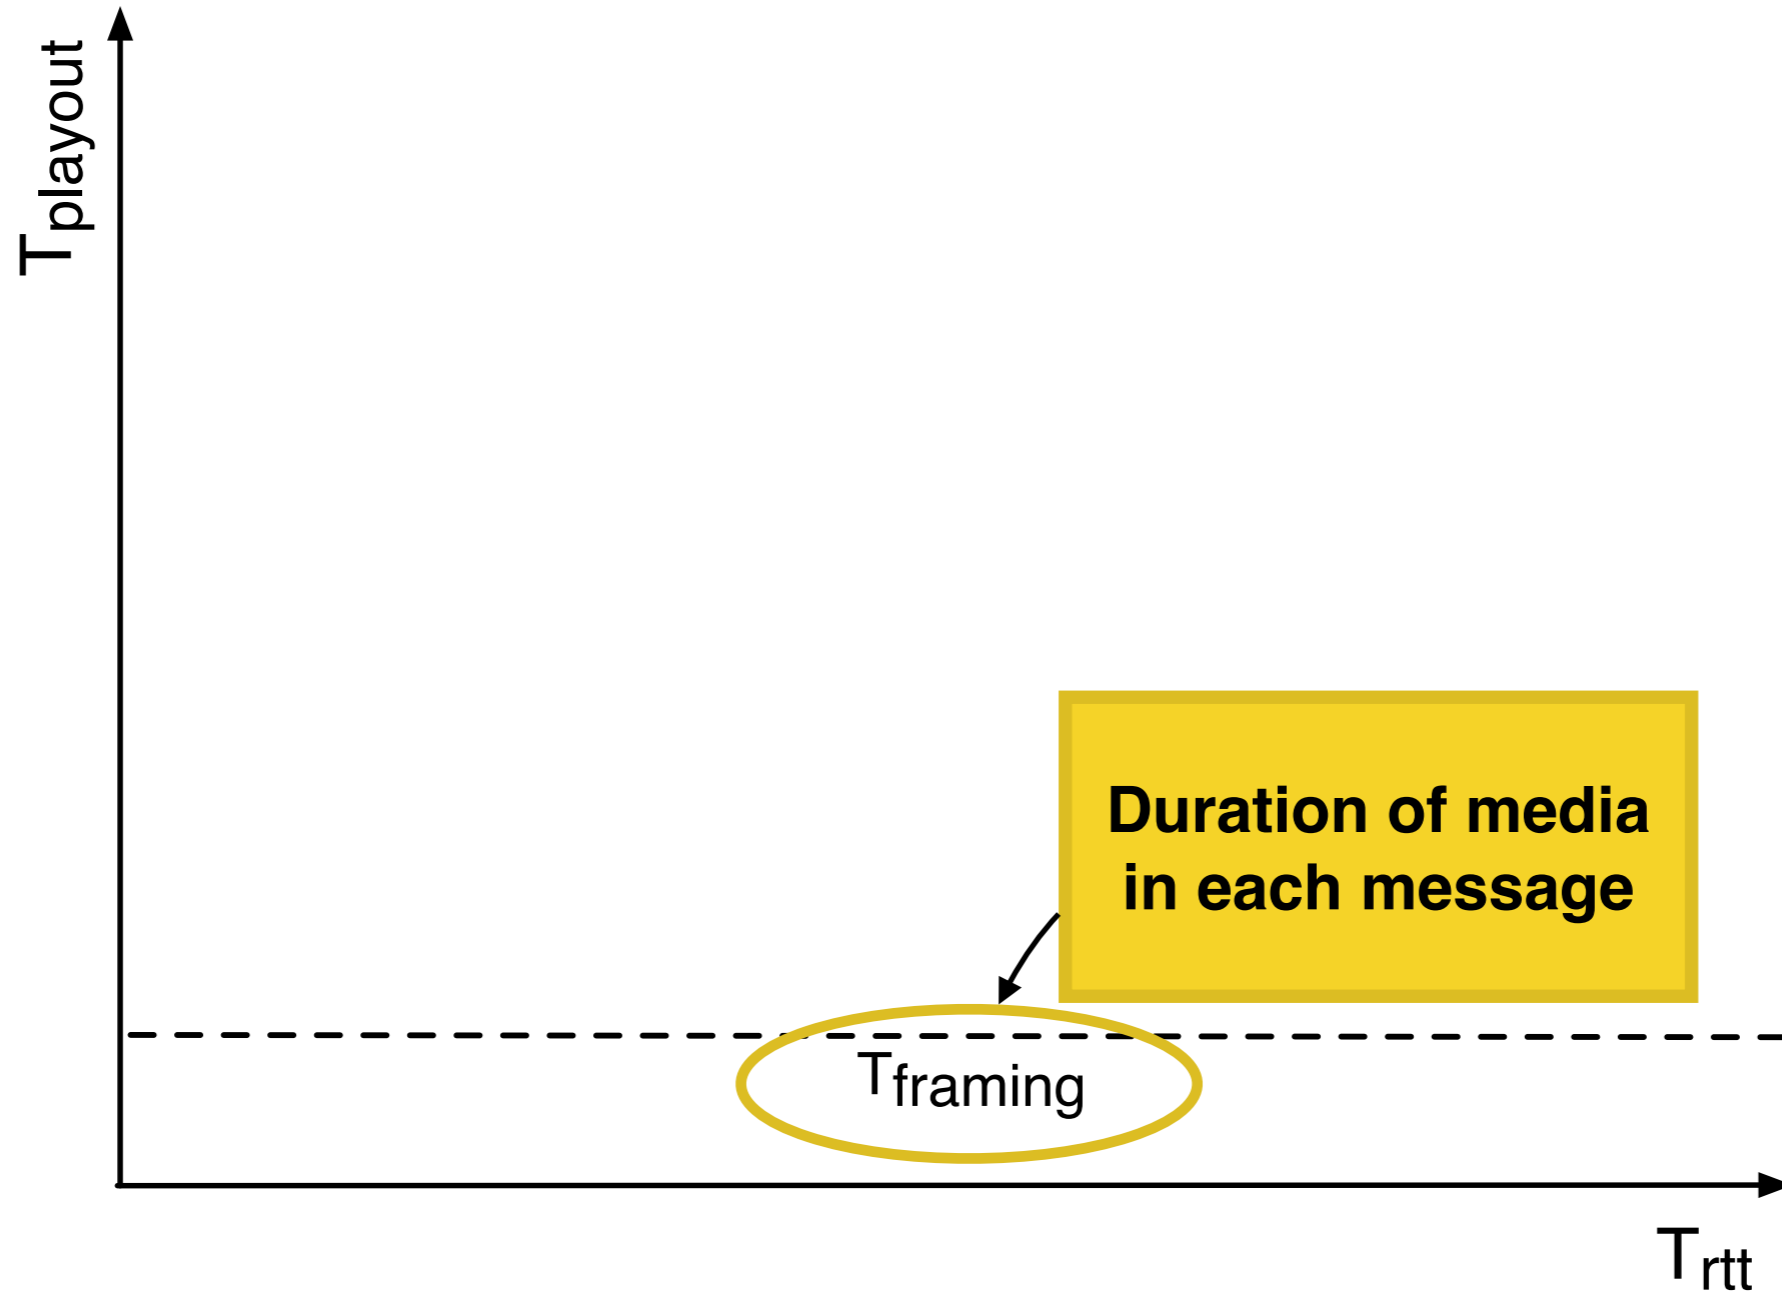


# Feasibility Region

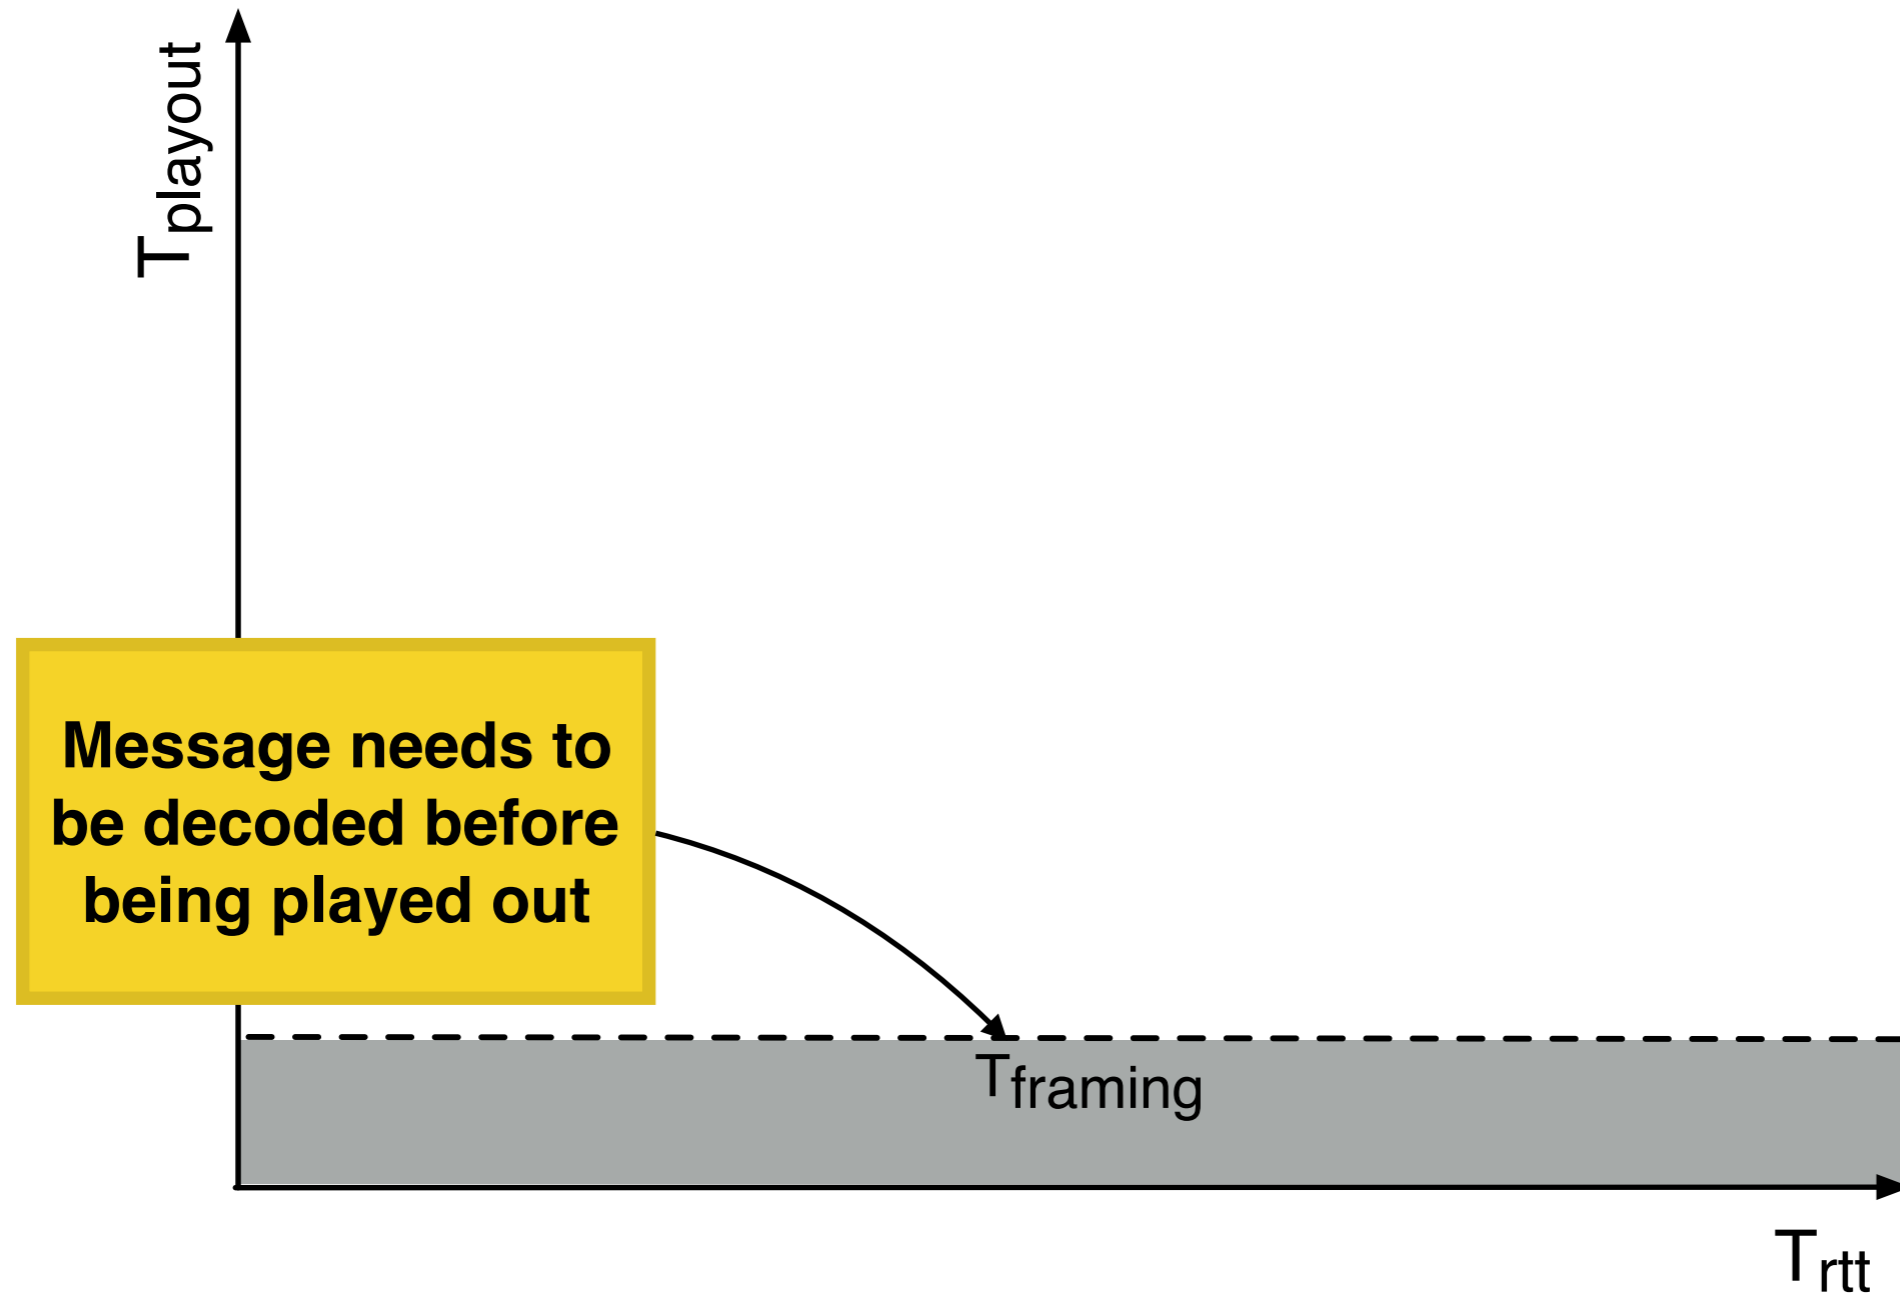




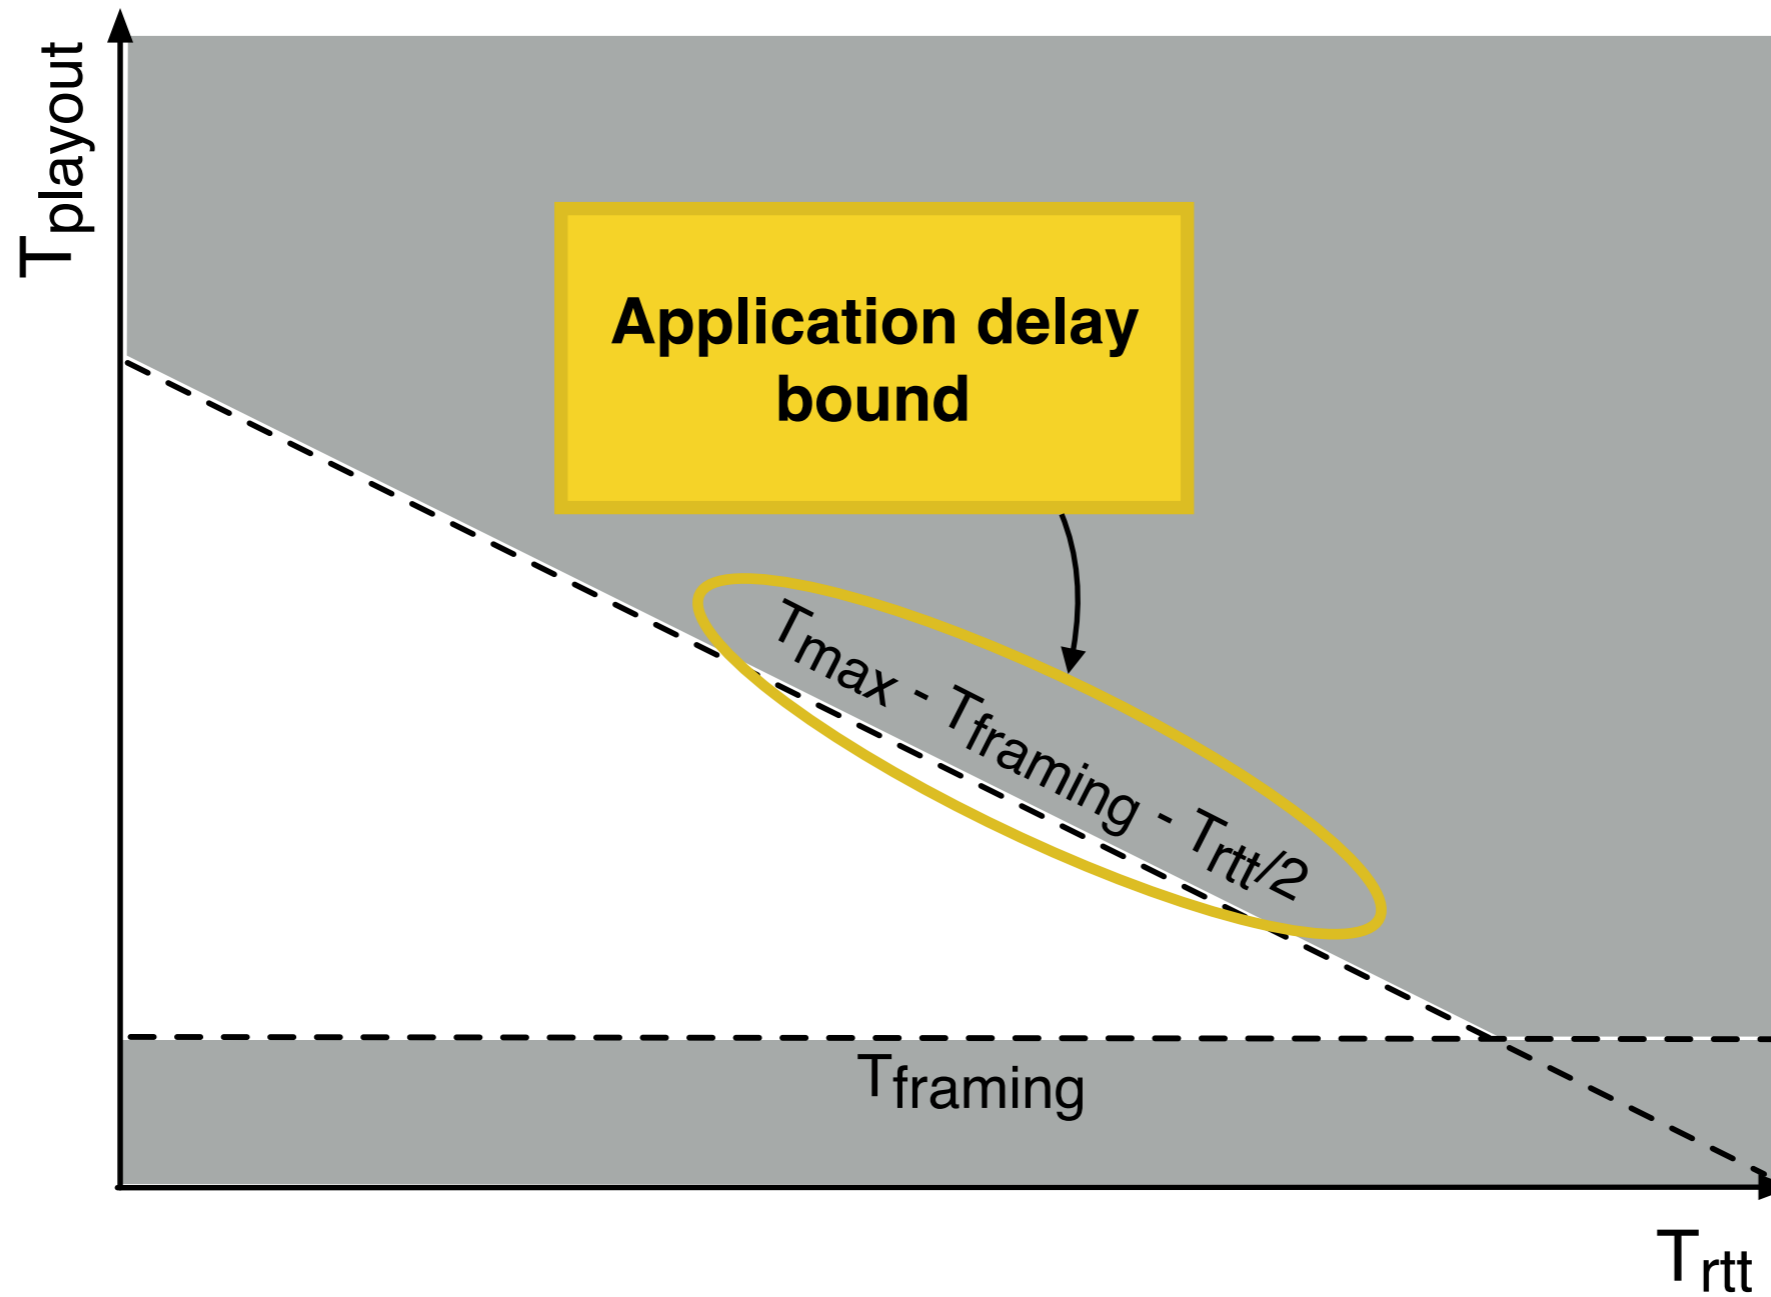
# Feasibility Region



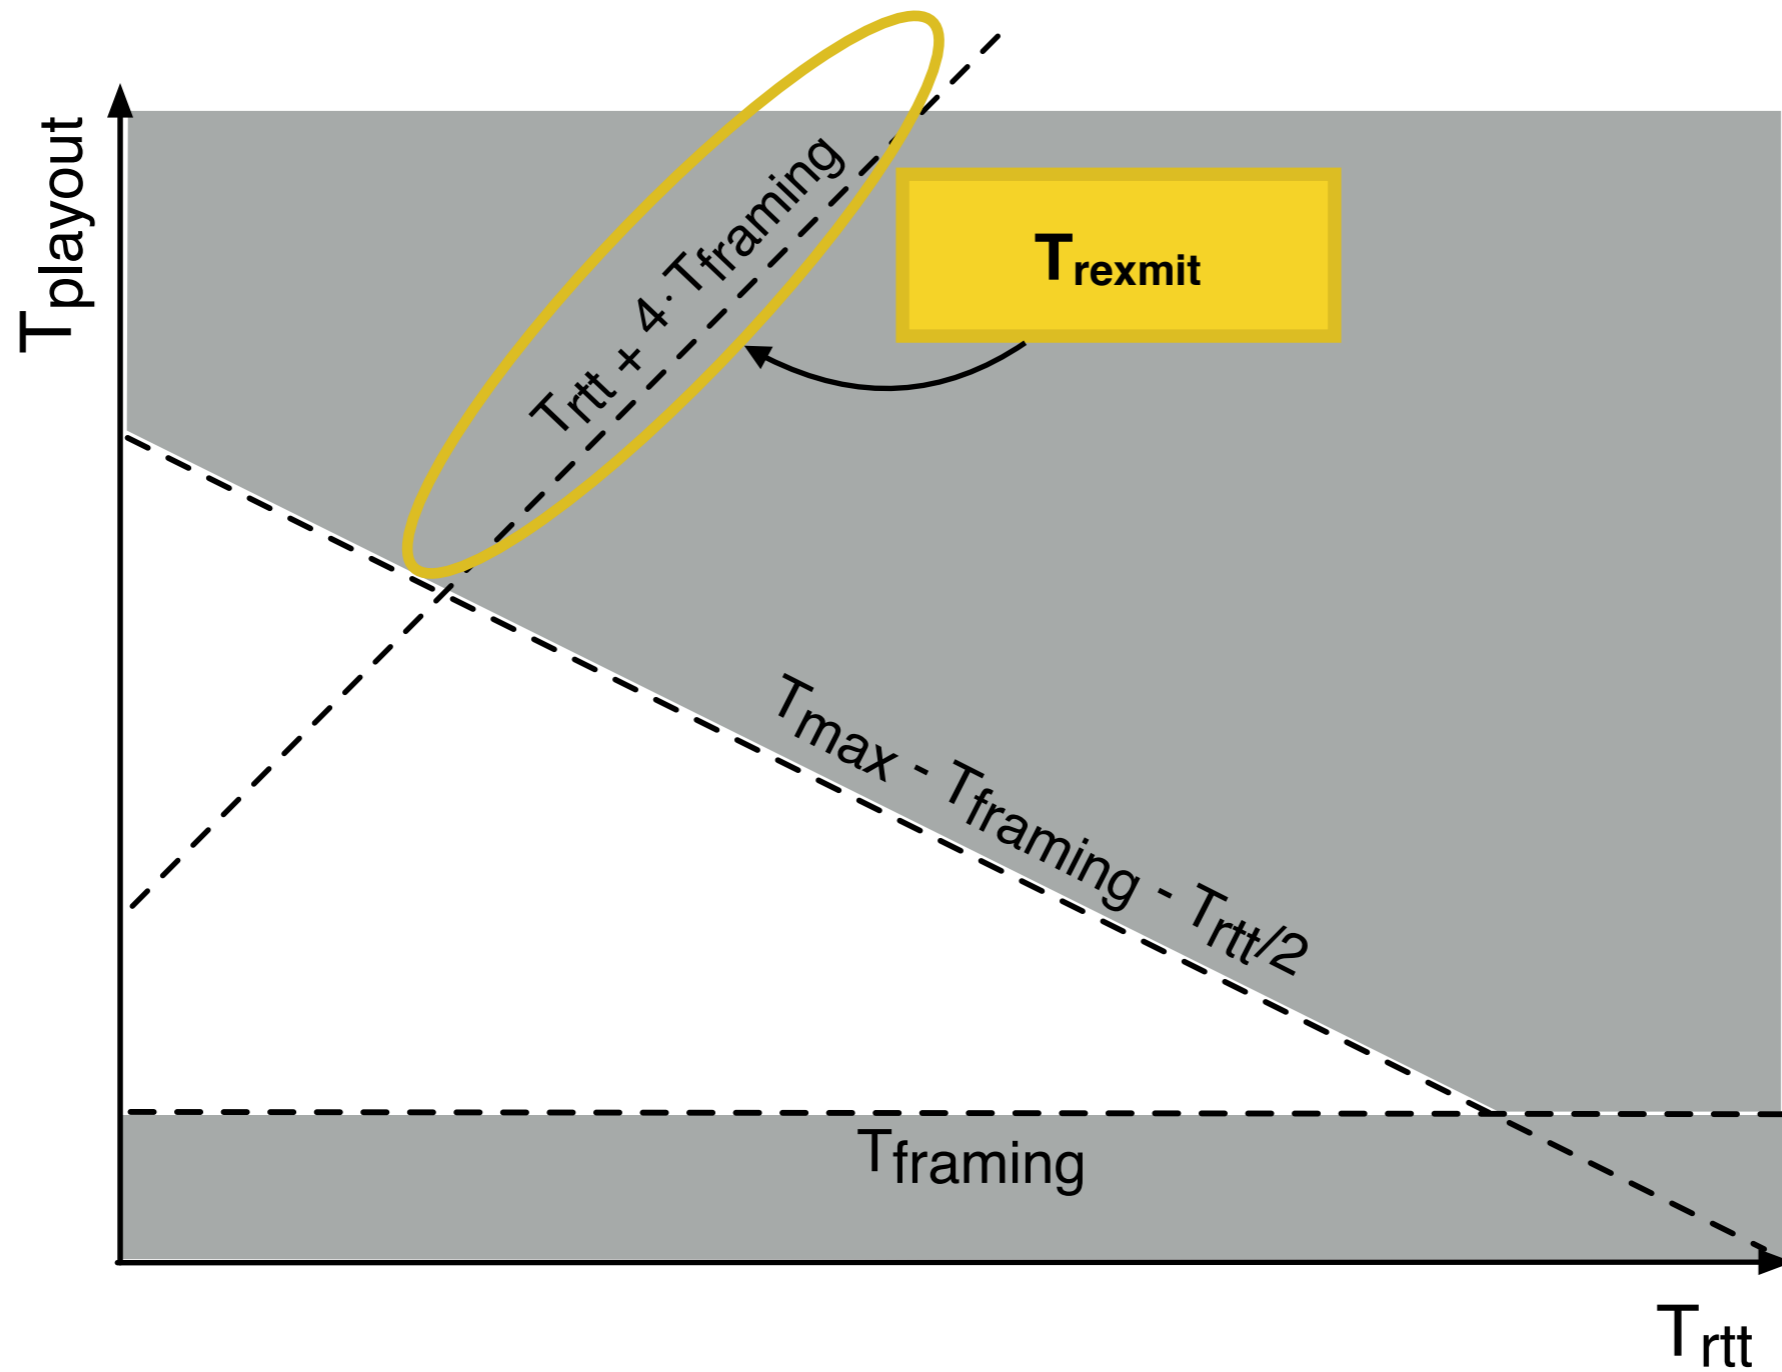
# Feasibility Region




# Feasibility Region



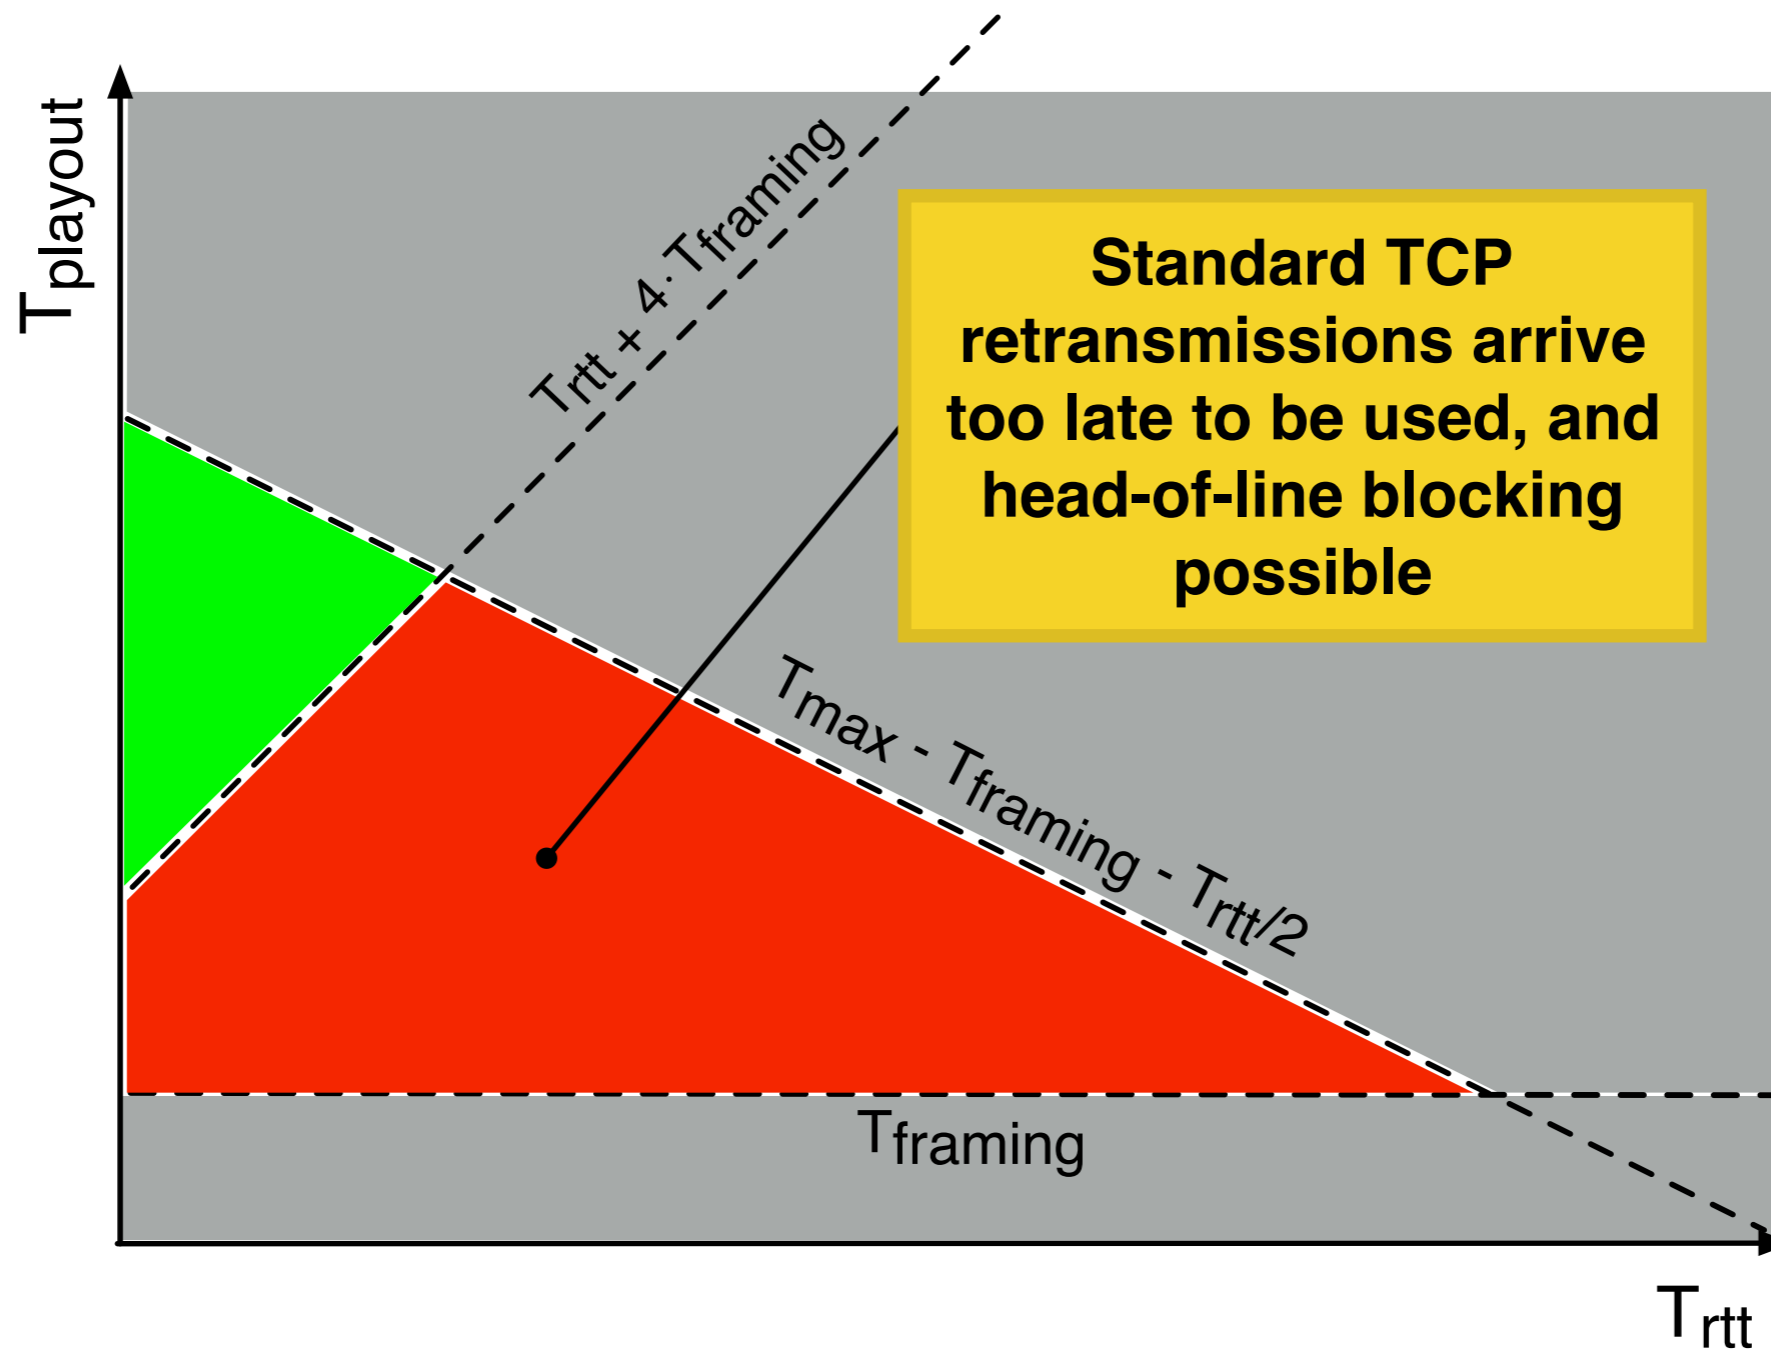
# Feasibility Region




# Feasibility Region



# Feasibility Region



# Feasibility Region

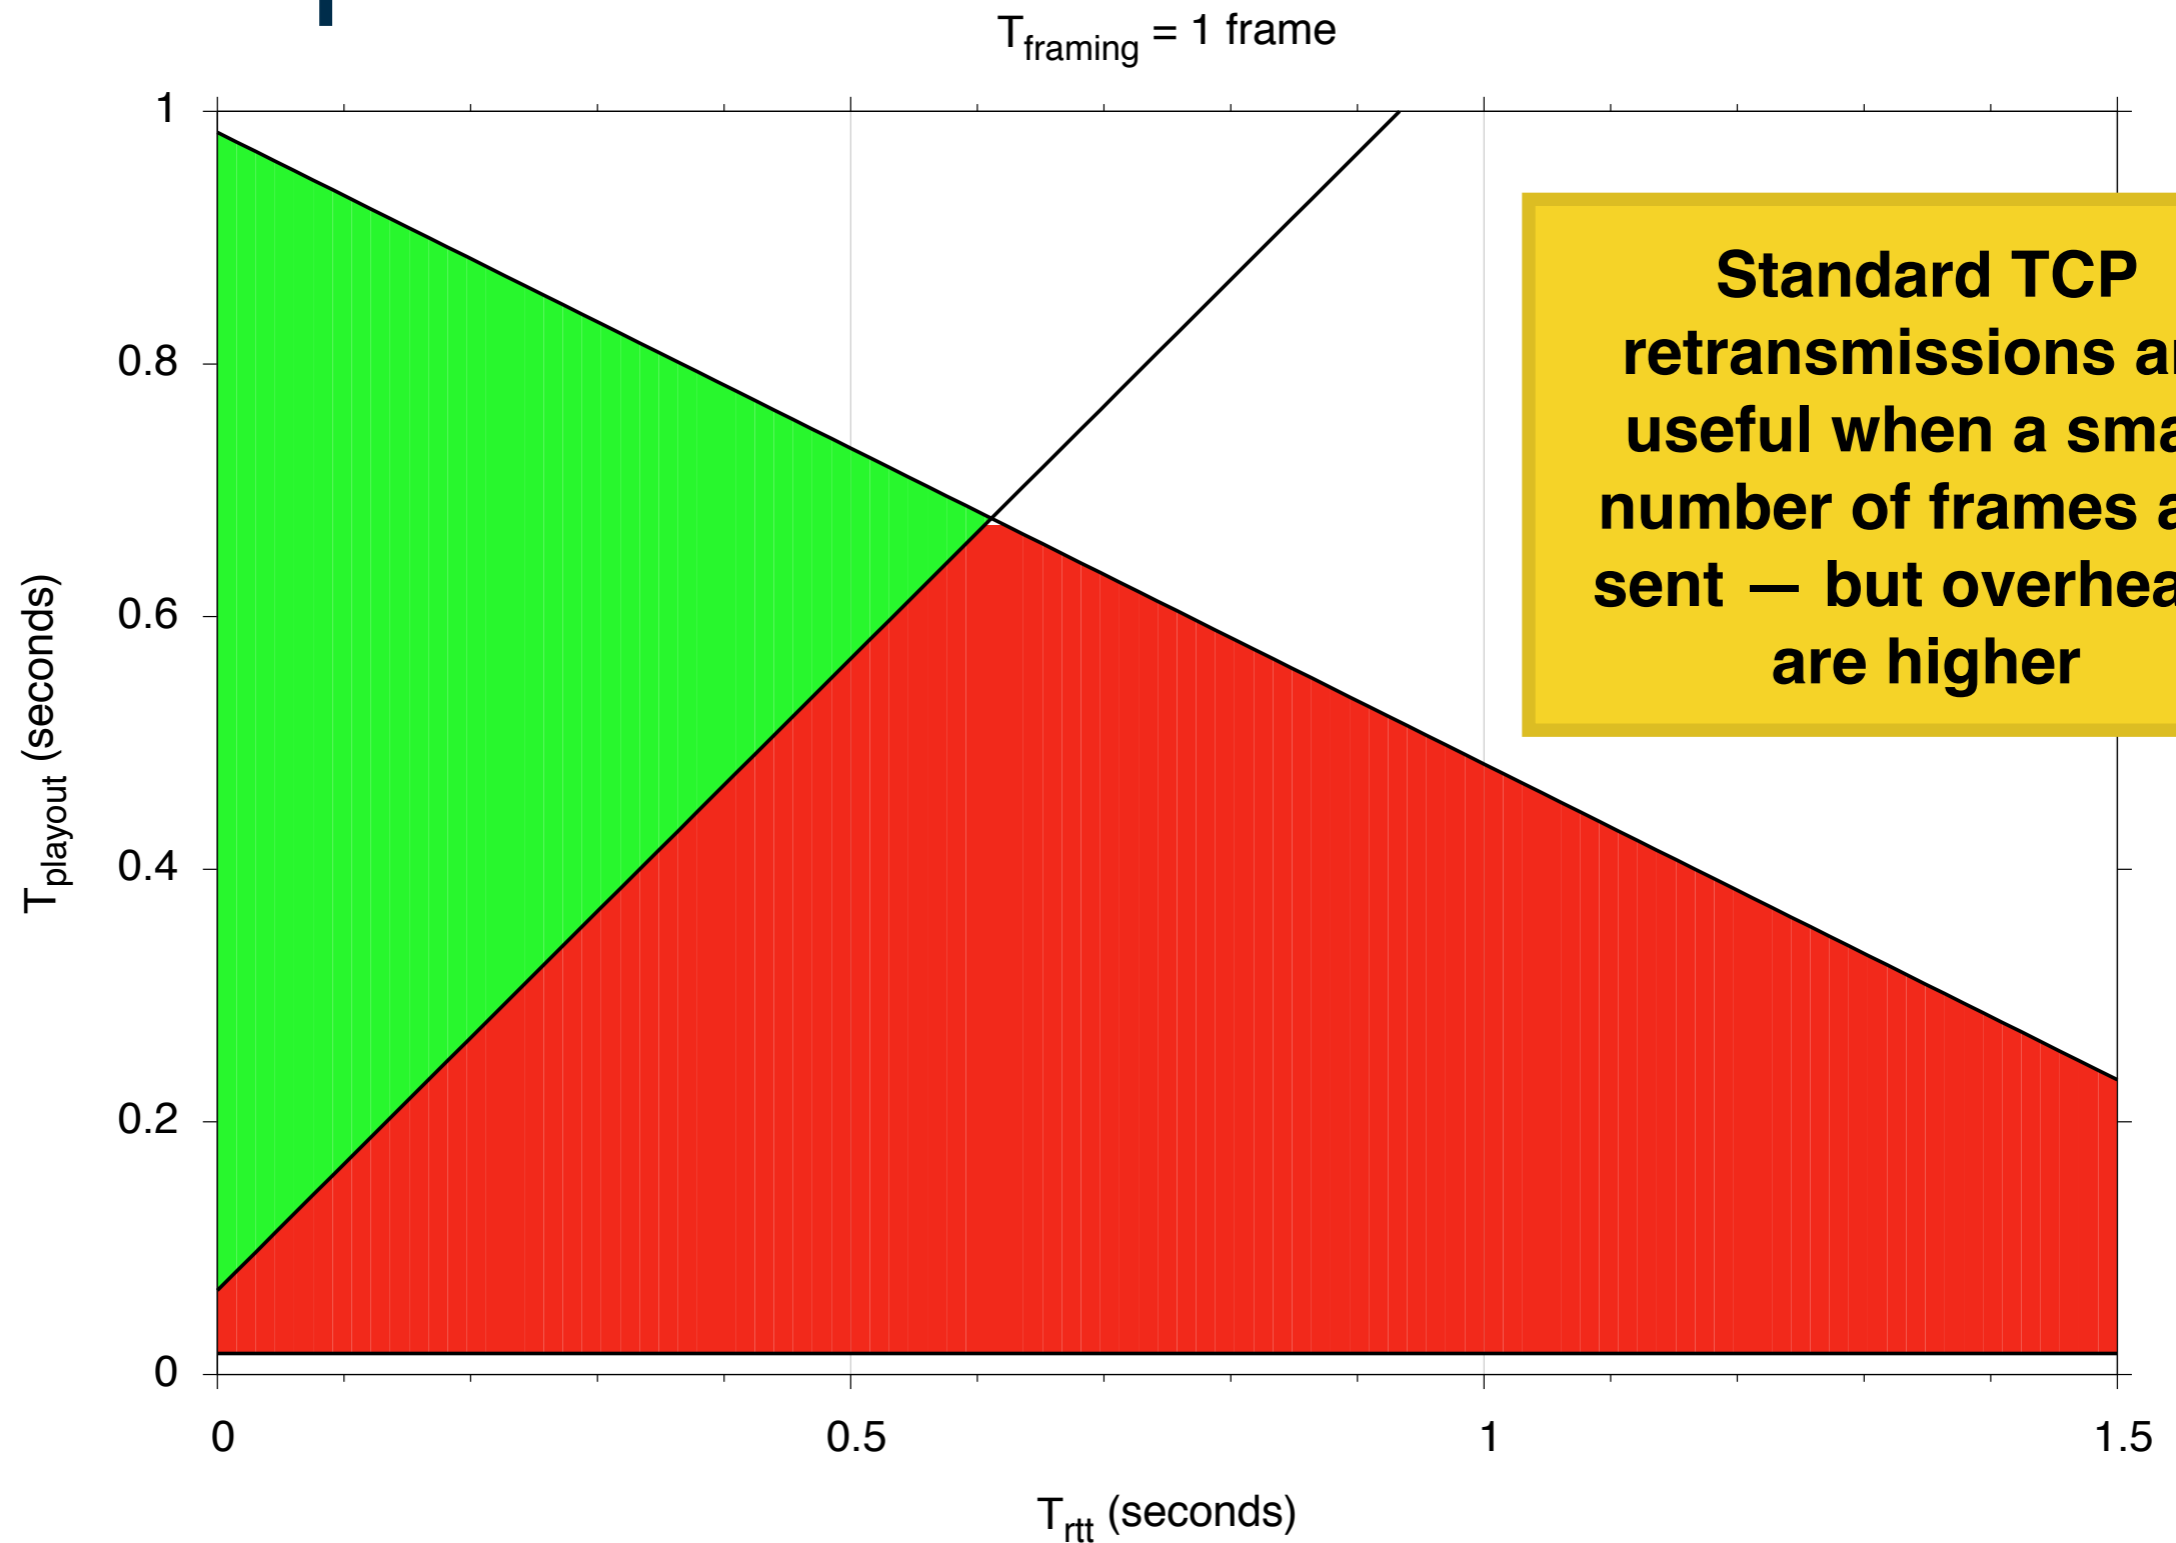


# Example Application


- IPTV application, using MPEG-DASH configured for low-latency delivery
- $T_{\max} = 1$  second, within zap time recommendations
- $T_{\text{framing}}$  determined by number of frames in message
- Trade-off between size of  $T_{\text{framing}}$ , and utility of standard TCP retransmissions



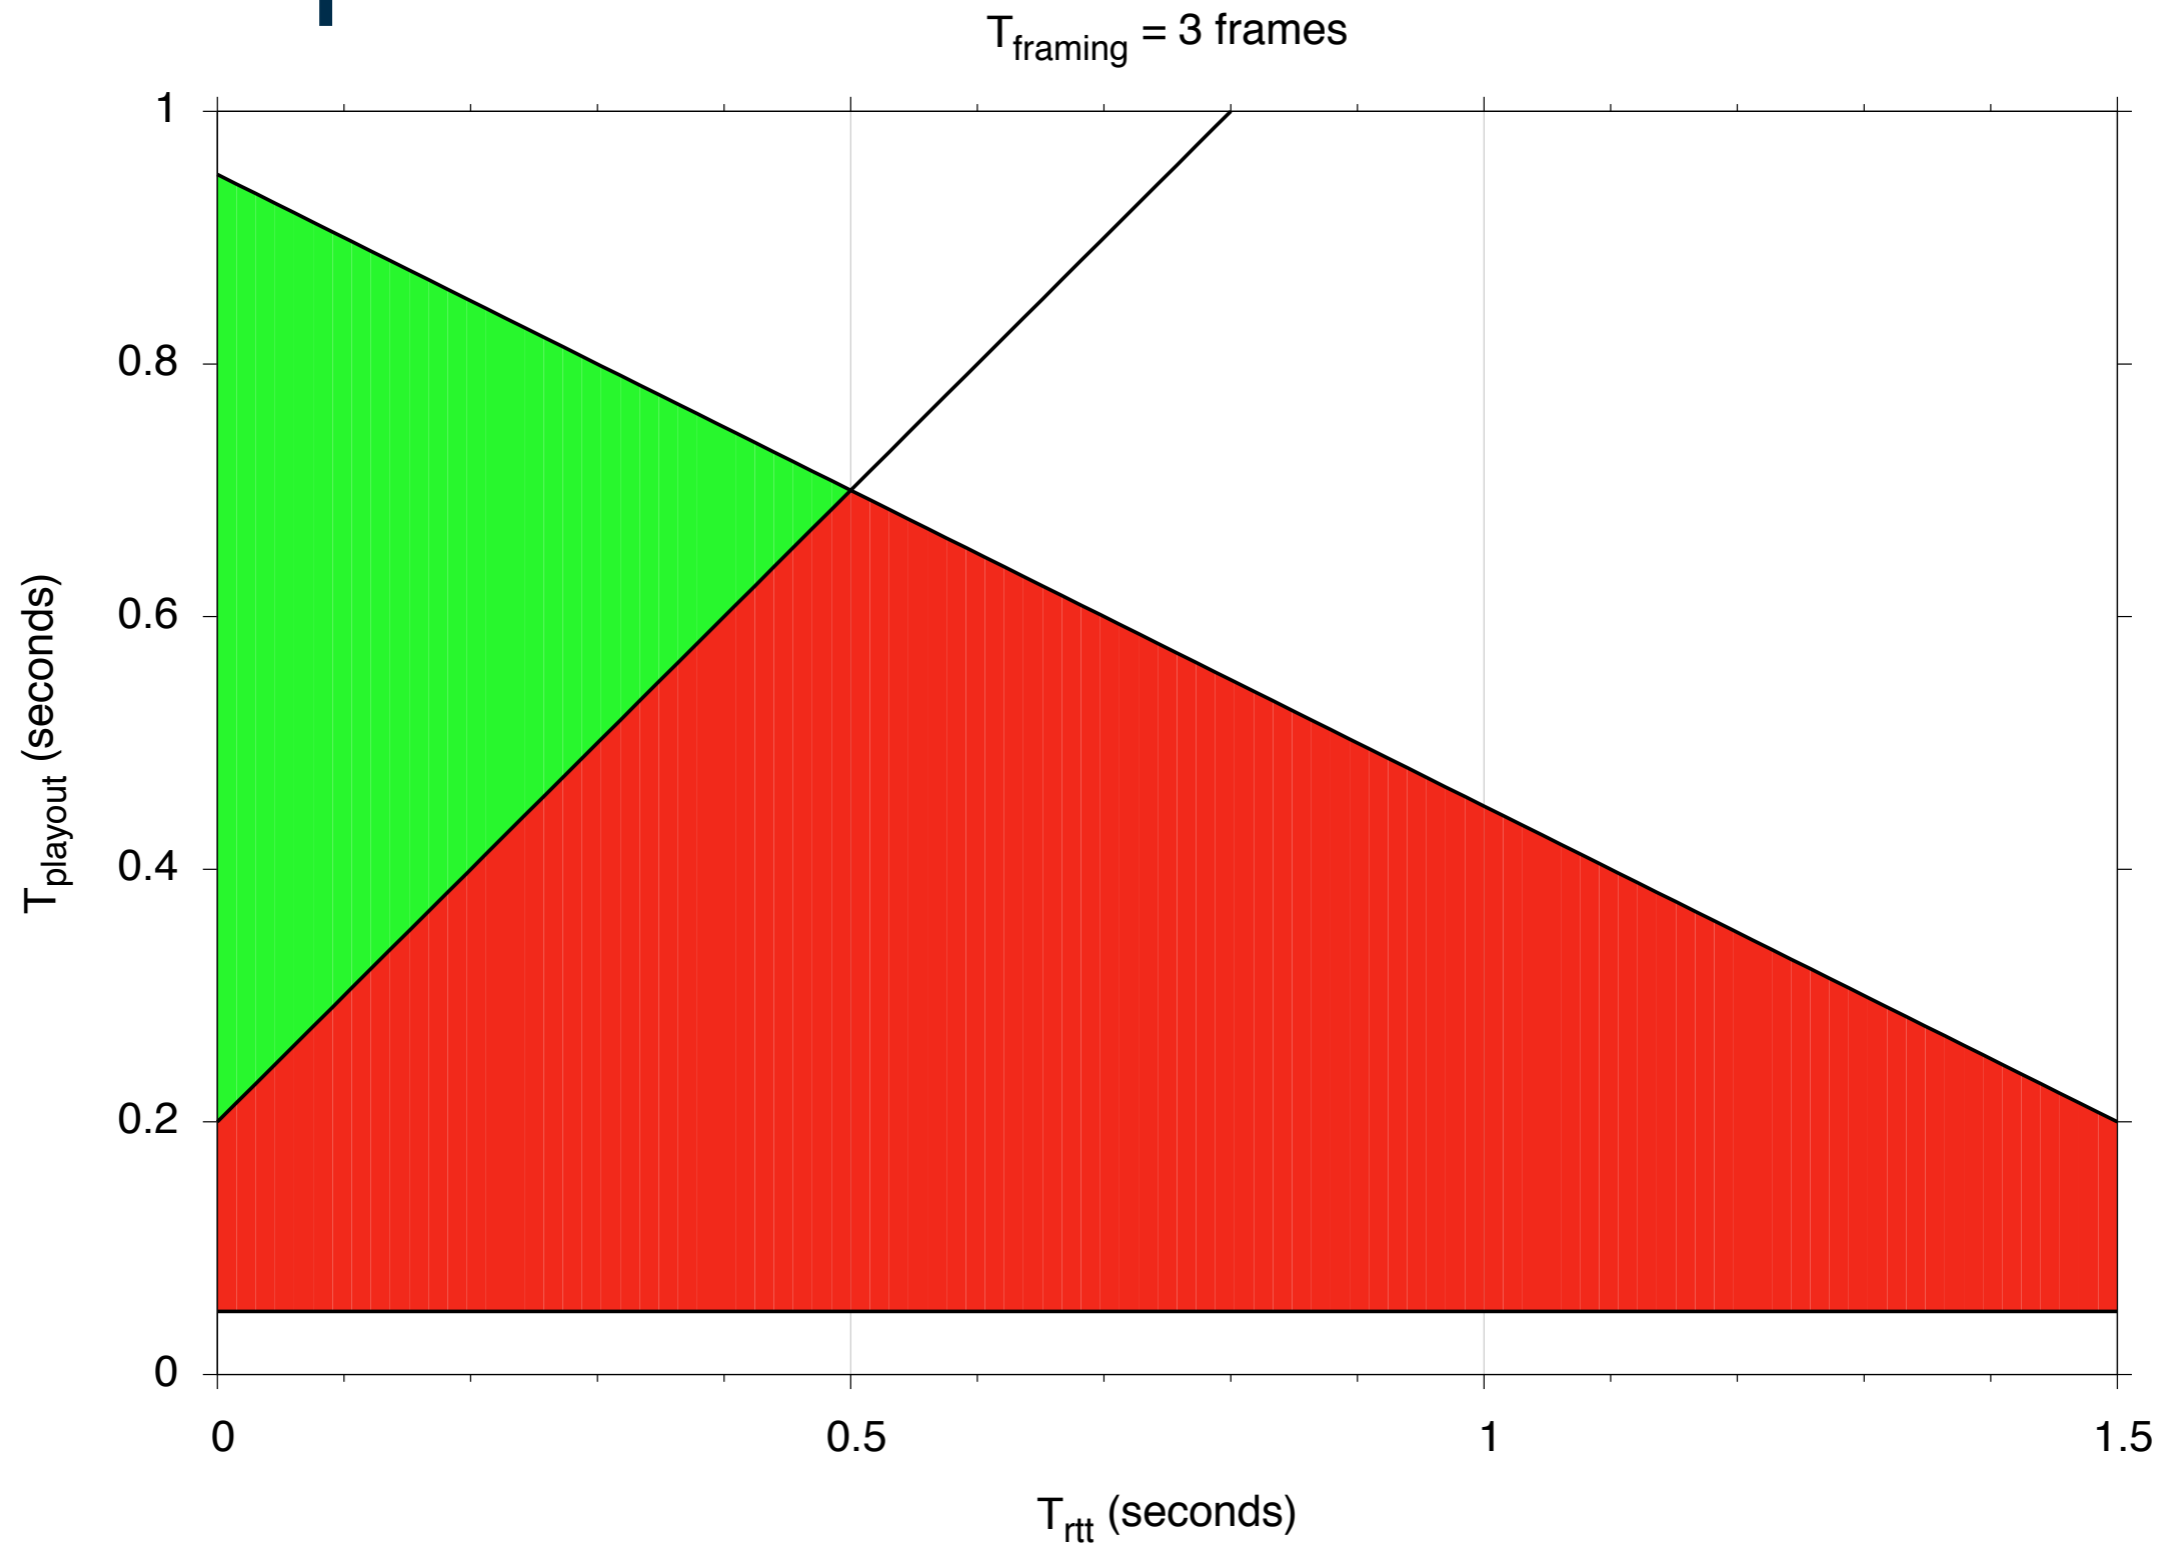
# Example




# Example




# Example




# Example



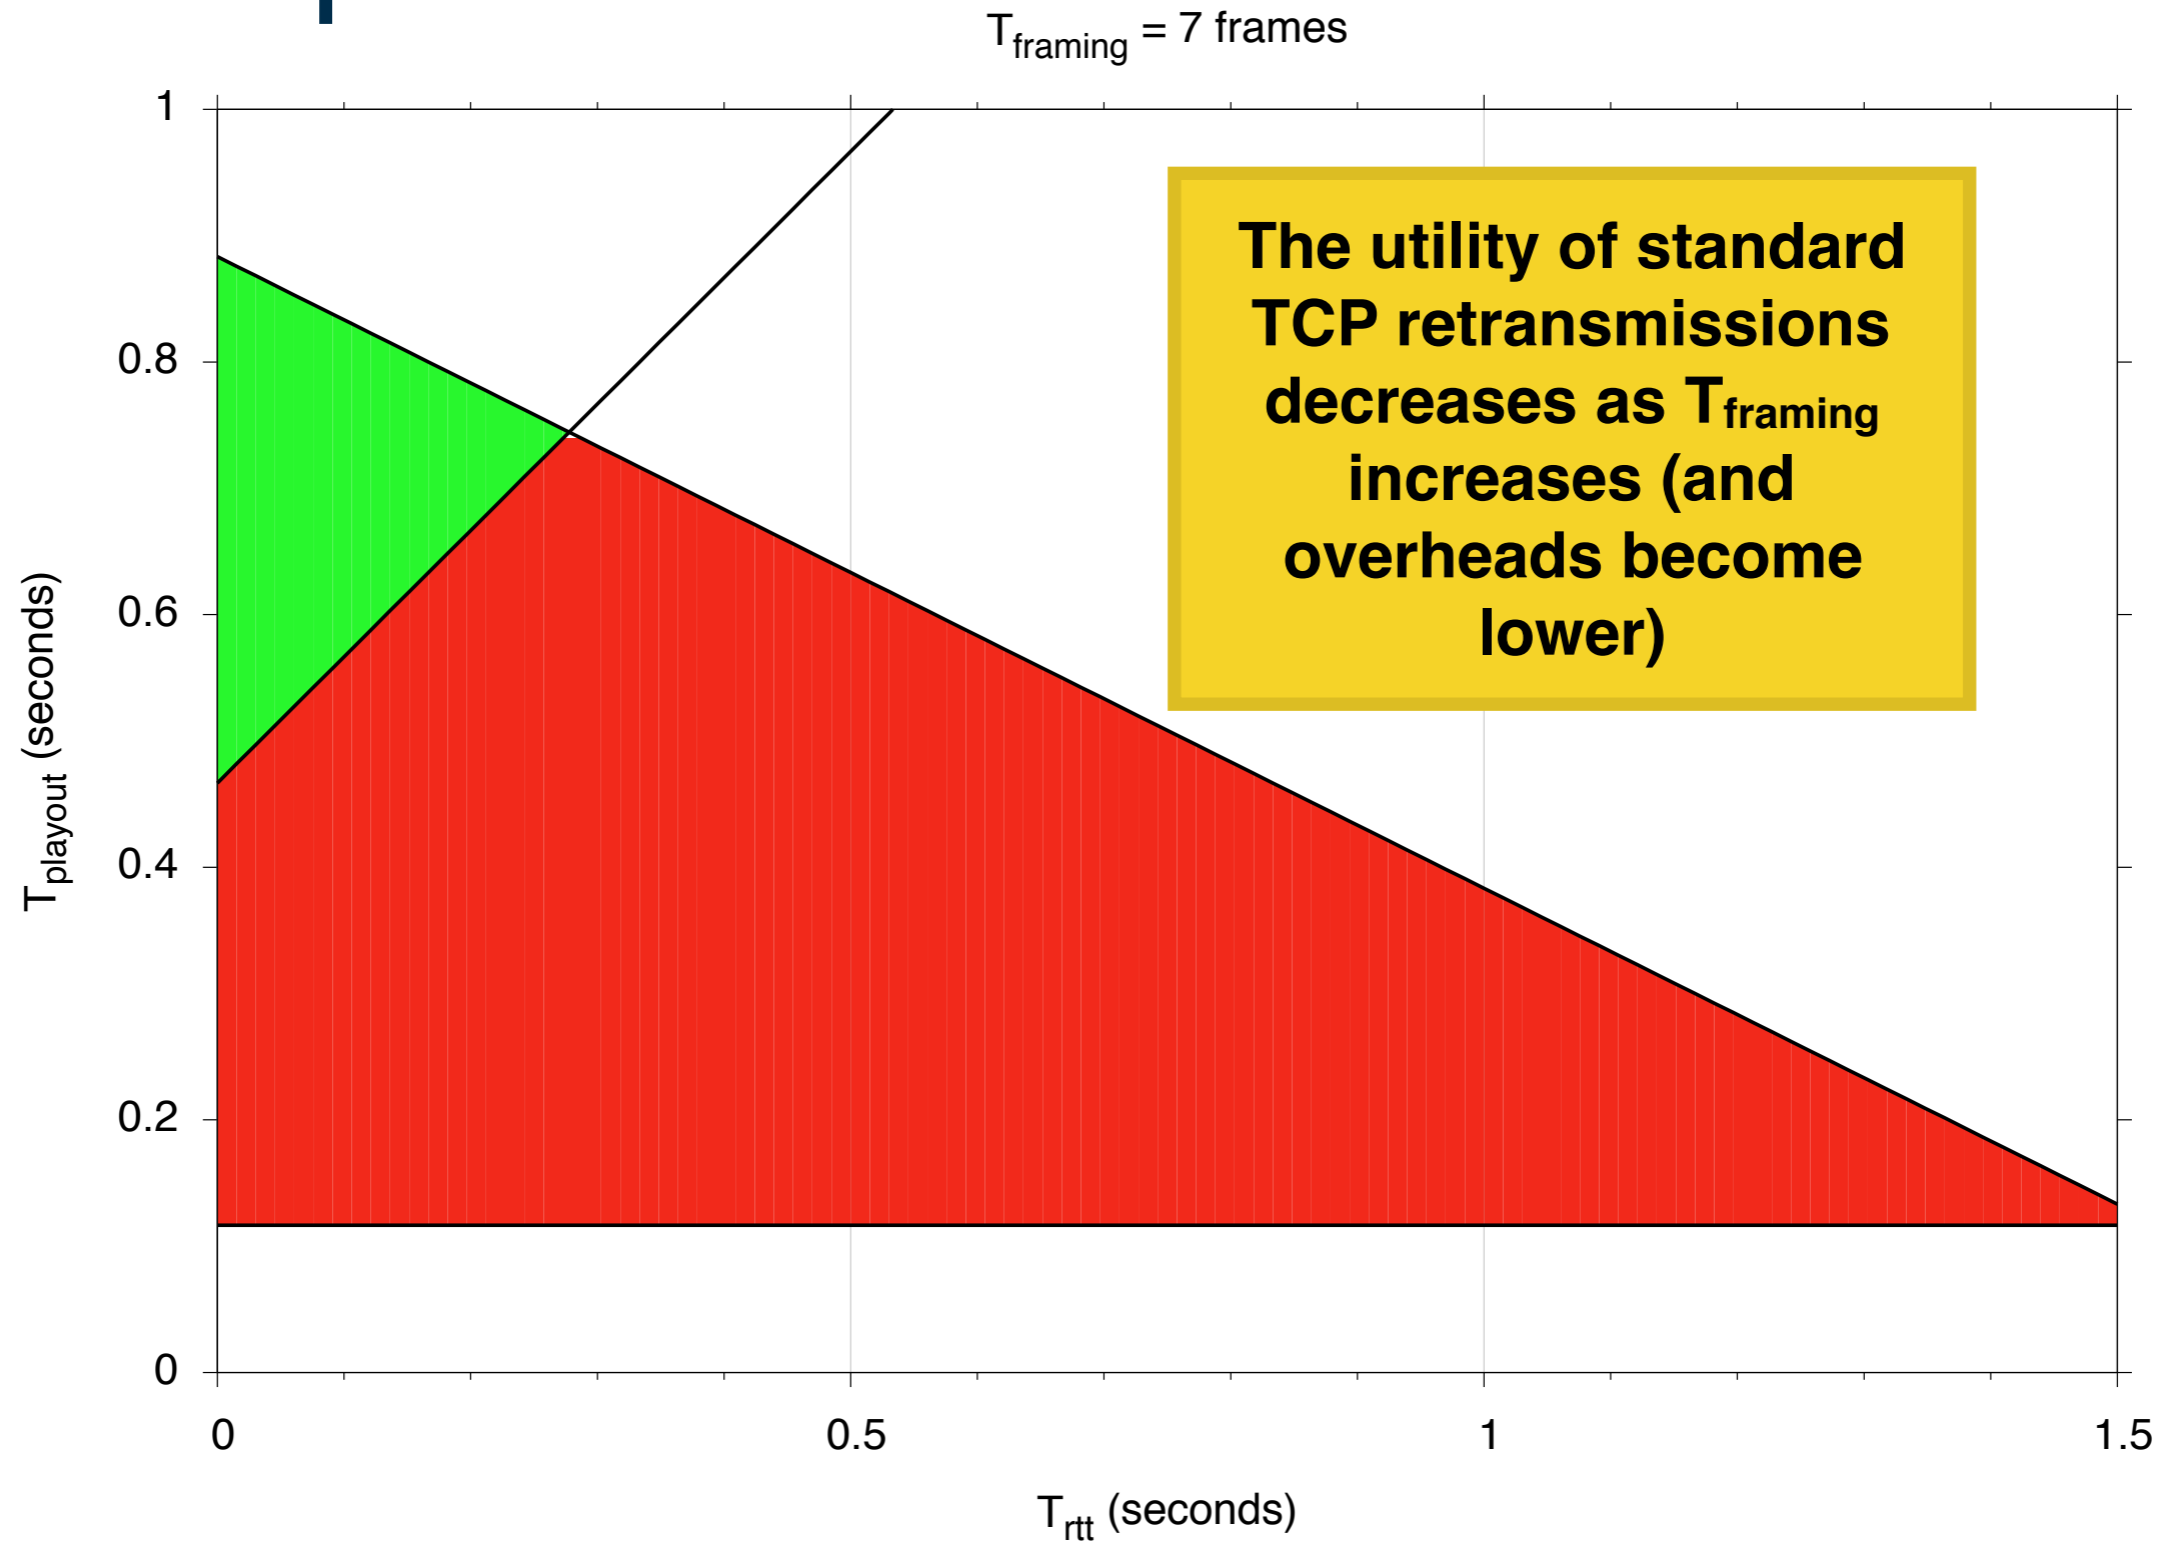
# Example




# Example



# Example

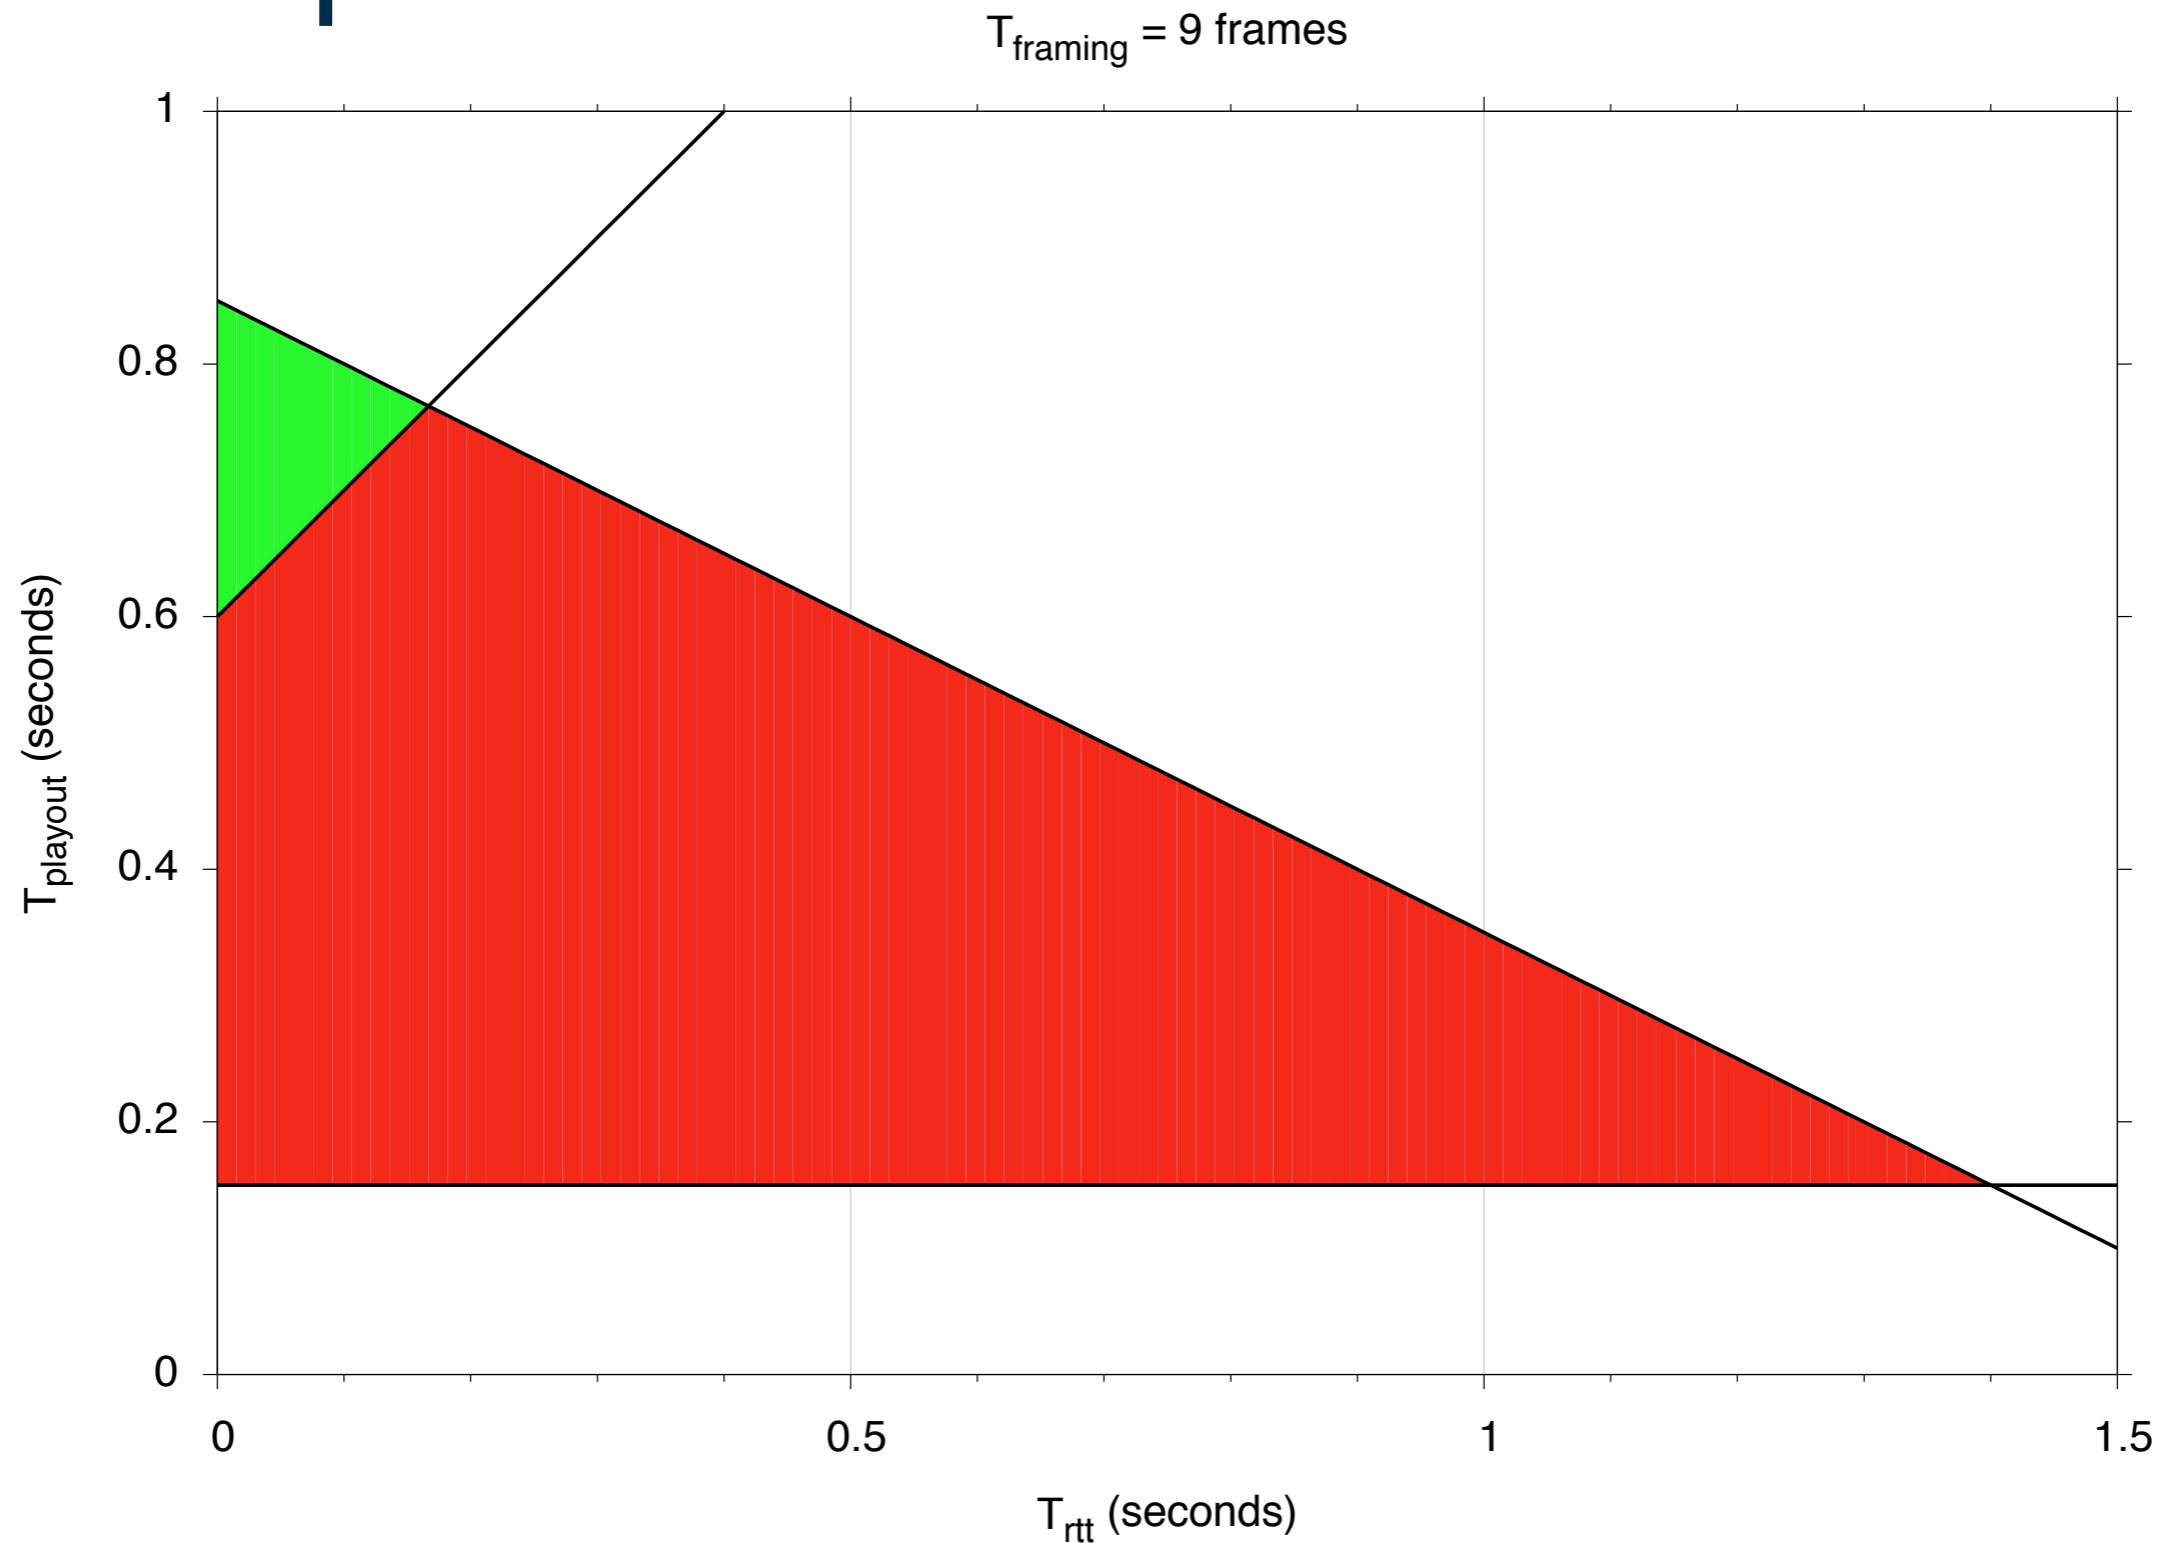


# Example






# Example




# Example

$T_{\text{framing}} = 10 \text{ frames}$



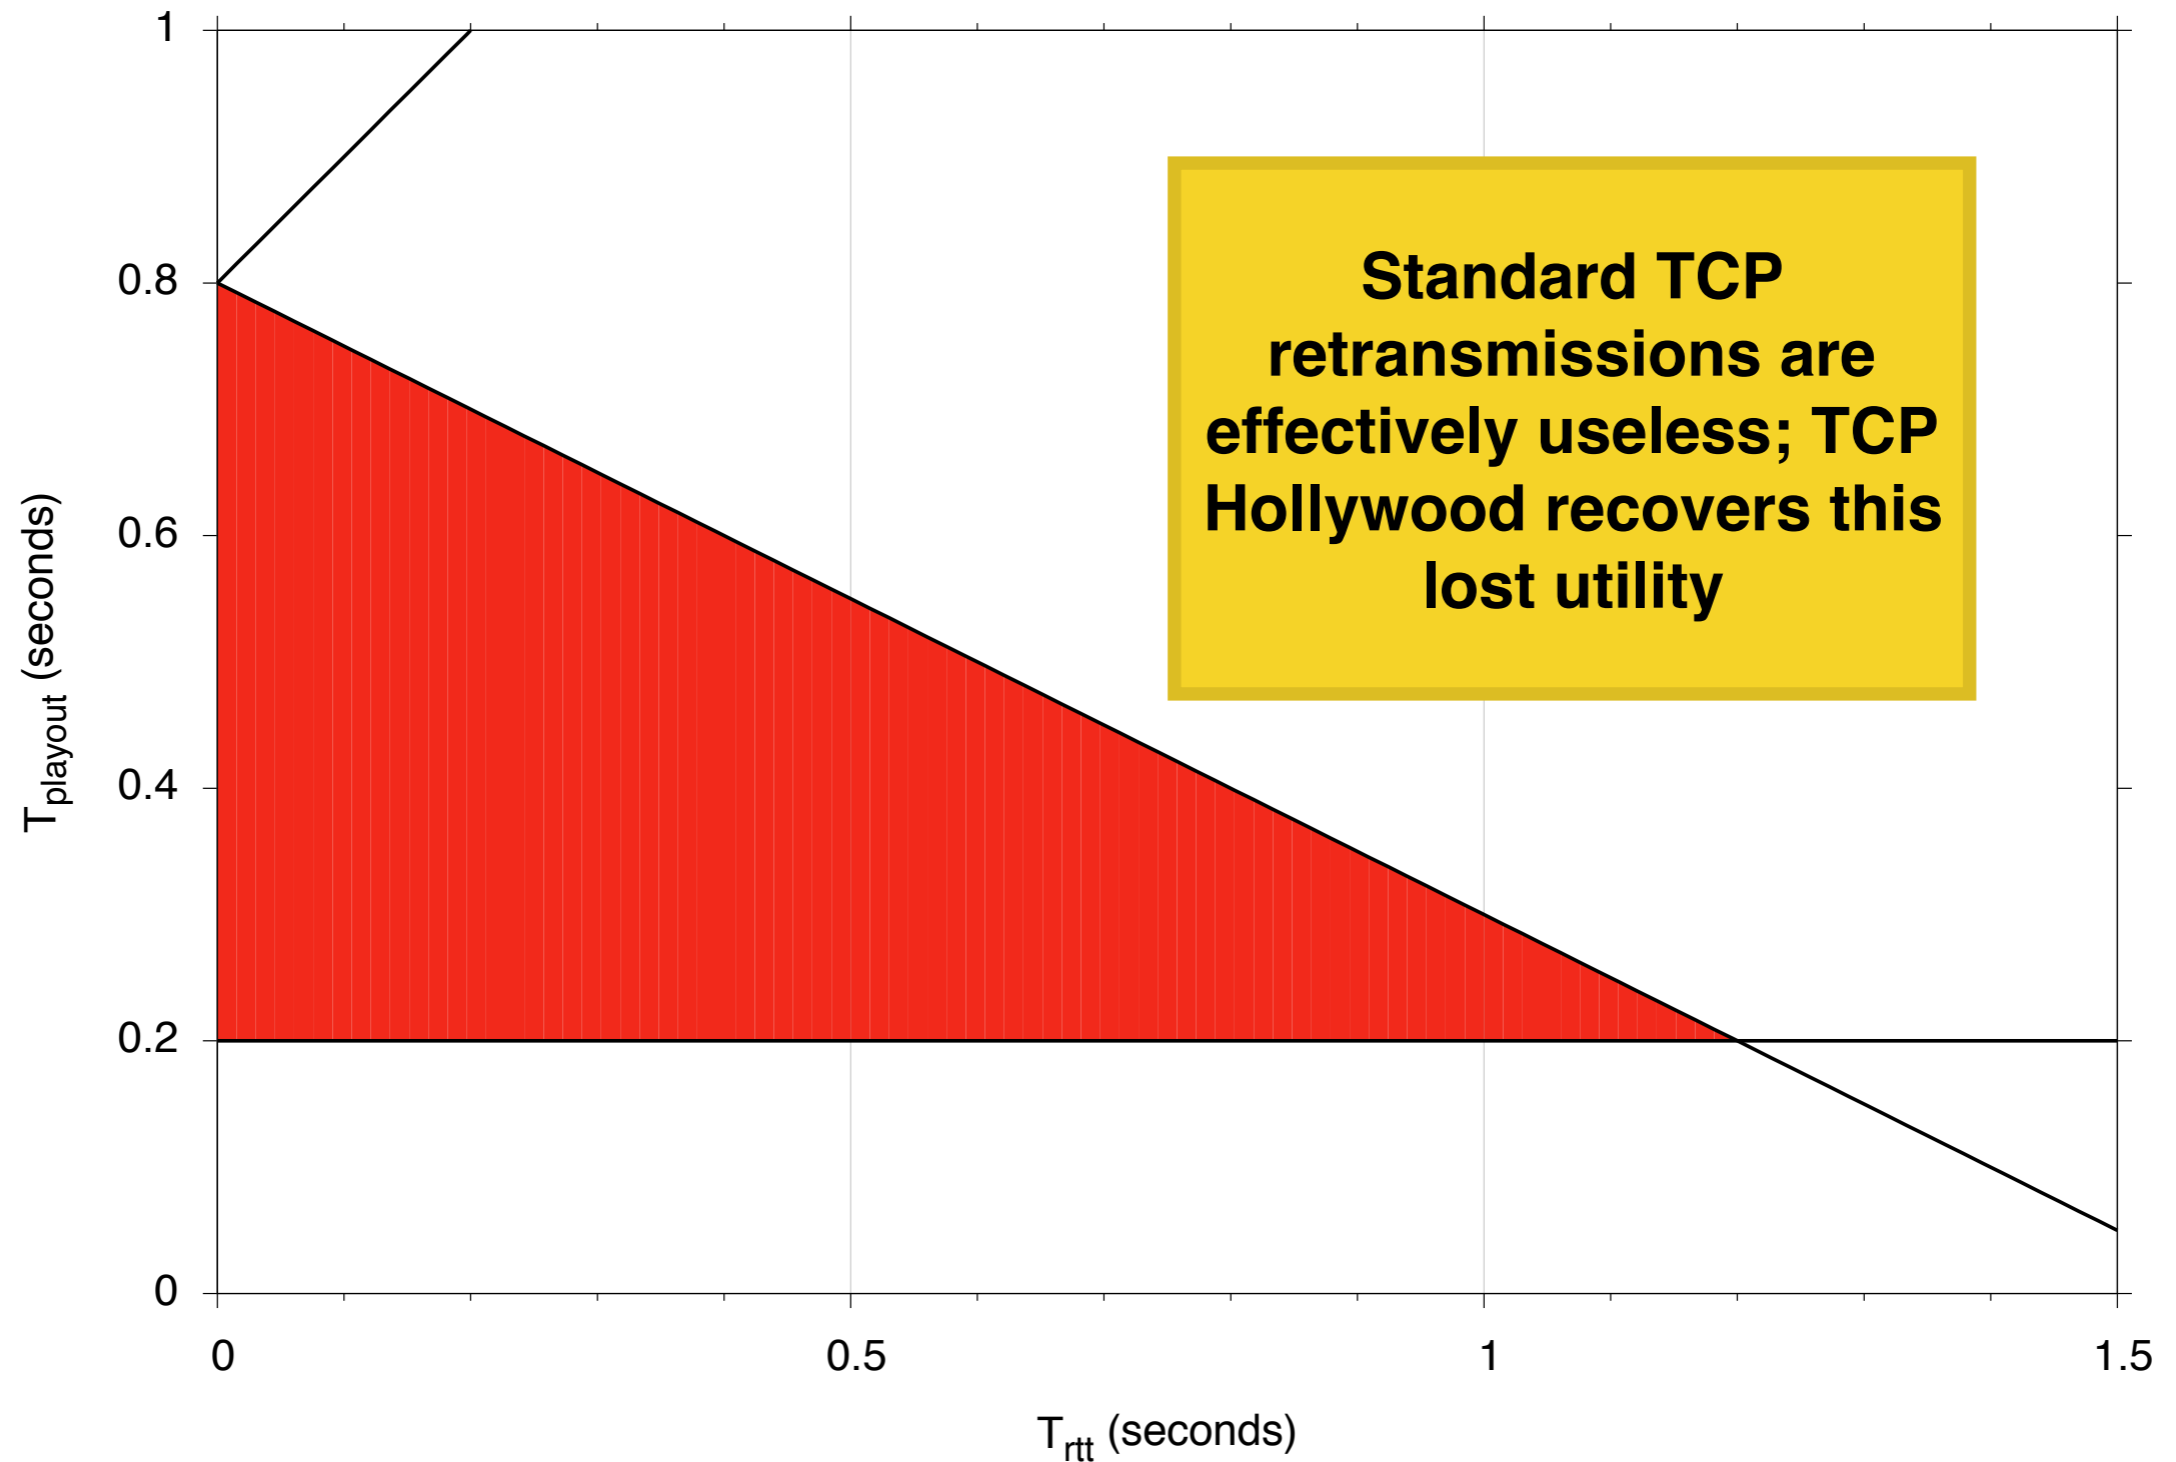
# Example

$T_{\text{framing}} = 11 \text{ frames}$



# Example



$T_{\text{framing}} = 12 \text{ frames}$




# Deployability

- Inconsistent retransmissions are the only wire-visible modification vs. standard TCP — same TCP sequence number, different payload
- Middleboxes performing payload inspection may interpret the behaviour as an attack — man on the side
- Experiments between TCP Hollywood server, and 14 UK clients
- 8 fixed-line ISPs, 4 cellular operators - all major UK ISPs

# TCP Hollywood is deployable

- Tested on two ports to check if HTTP traffic treated differently
-  inconsistent retransmission delivered successfully
-  original segment delivered instead
- Intermediary layer handles cases where original delivered - performance no worse than standard TCP



# TCP Hollywood

- Unordered, partially reliable message-oriented TCP-based transport protocol
- Eliminates head-of-line blocking, reducing latency
- Prevents retransmission of expired data, increasing utility
- Deployable across all major UK ISPs



“Hollywood Sign”, Gnaphron - CC BY-SA 2.0  
[flickr.com/photos/gnaphron/8485145044](https://www.flickr.com/photos/gnaphron/8485145044/)