

# Implementing Real-Time Transport Services over an Ossified Network

Stephen McQuistin and Colin Perkins  
University of Glasgow

Marwan Fayed  
University of Stirling

# Talk Overview

- Multimedia Applications and the Transport Layer
- Ossification and Innovation
- Transport Services
- ... for Real-Time Multimedia Applications
- Realising Transport Services
- Example: TCP Hollywood

# Multimedia Applications

- 64% of consumer Internet traffic in 2014 → 80% by 2019 (Cisco VNI)
- Difficult to develop and standardise
- WebRTC and DASH standardisation work highlights challenges



**WebRTC**



# Transport Layer

- Neither TCP or UDP provides all the features we require
- UDP adds minimal features beyond those of IP
- TCP adds many desired services (e.g., congestion control), but includes others we don't want (e.g., reliability)
- Can build the features we need within UDP's payload — large amount of effort, lacks reusability
- In principle, we could build a new protocol that provides the features we need

# Ossification

- Middleboxes expect packets that look like either TCP or UDP: rejecting everything else is a common security policy
- New protocols (e.g., DCCP, SCTP) see little deployment on the public Internet
- TCP and UDP can be used as substrates for new protocols
- Need to ensure that middlebox compatibility is maintained

# Innovation at the transport layer

- Two broad architectural approaches
- Develop a new, monolithic protocol that uses TCP or UDP as a substrate — e.g., QUIC
- Add a layer of indirection, and develop reusable building blocks — transport services

# Transport Services

- *“an end-to-end facility provided by the transport layer”*
- Need to define the set of services required by applications
- Determine how these services can be realised by transport protocols
- Map the set of services on to an appropriate transport protocol (TCP, UDP, and others where available)
- Results in a set of reusable services that help application developers, and improve performance

# Real-Time Multimedia Applications

- Maximum delay, depending on interactivity
- Interactive applications: low hundreds of milliseconds (for VoIP) — depends on human perceptibility
- Non-interactive: tens of seconds (for VoD) — depends on desired experience
- Services need to respect timeliness constraint, and add minimal latency

# Timing and Deadlines

- Data has set time that it needs to have arrived by, otherwise it is skipped, and not useful
- If the transport layer doesn't know about this deadline, useless data might be sent
- With the deadline, likelihood of data arriving on time can be estimated
- Requires network delay estimate, receive buffer size
- Fundamental service: others follow from this

# Partial Reliability

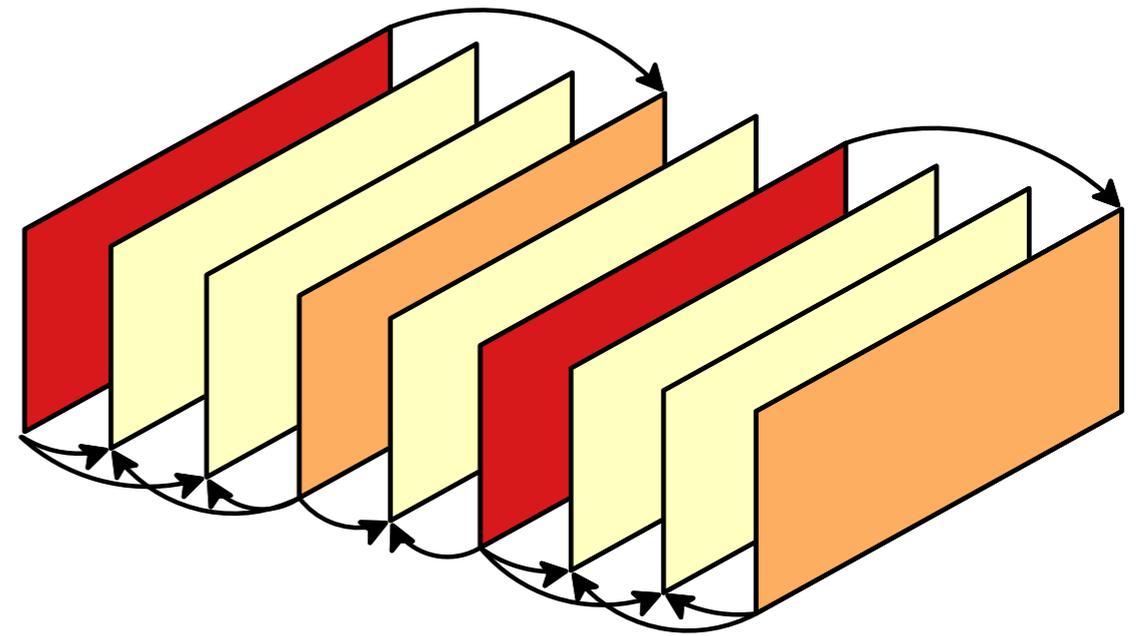
- IP provides best-effort packet delivery, so some packets will be lost
- Timeliness constraint means that data is useless after its deadline
- Guaranteed reliability would result in useless data being sent, deadlines not being met
- Need *partial* reliability: retransmit lost packets, but only if they will arrive within their deadline

# Message-oriented

- Partial reliability means that some packets may not be delivered
- The packets that do arrive need to be independently useful
- Implies application-level framing, with application data units (ADUs) being sent
- Given independent utility, and need to reduce latency, ADUs can be delivered in the order they arrive
- Support for multiple sub-streams

# Dependencies

- Partial reliability means that not all data will arrive successfully
- Interdependencies exist within data
- Data shouldn't be sent if it relies on a previous transmission that was not received
- Utility difficult to define for some applications



MPEG-1: I, P, and B frames

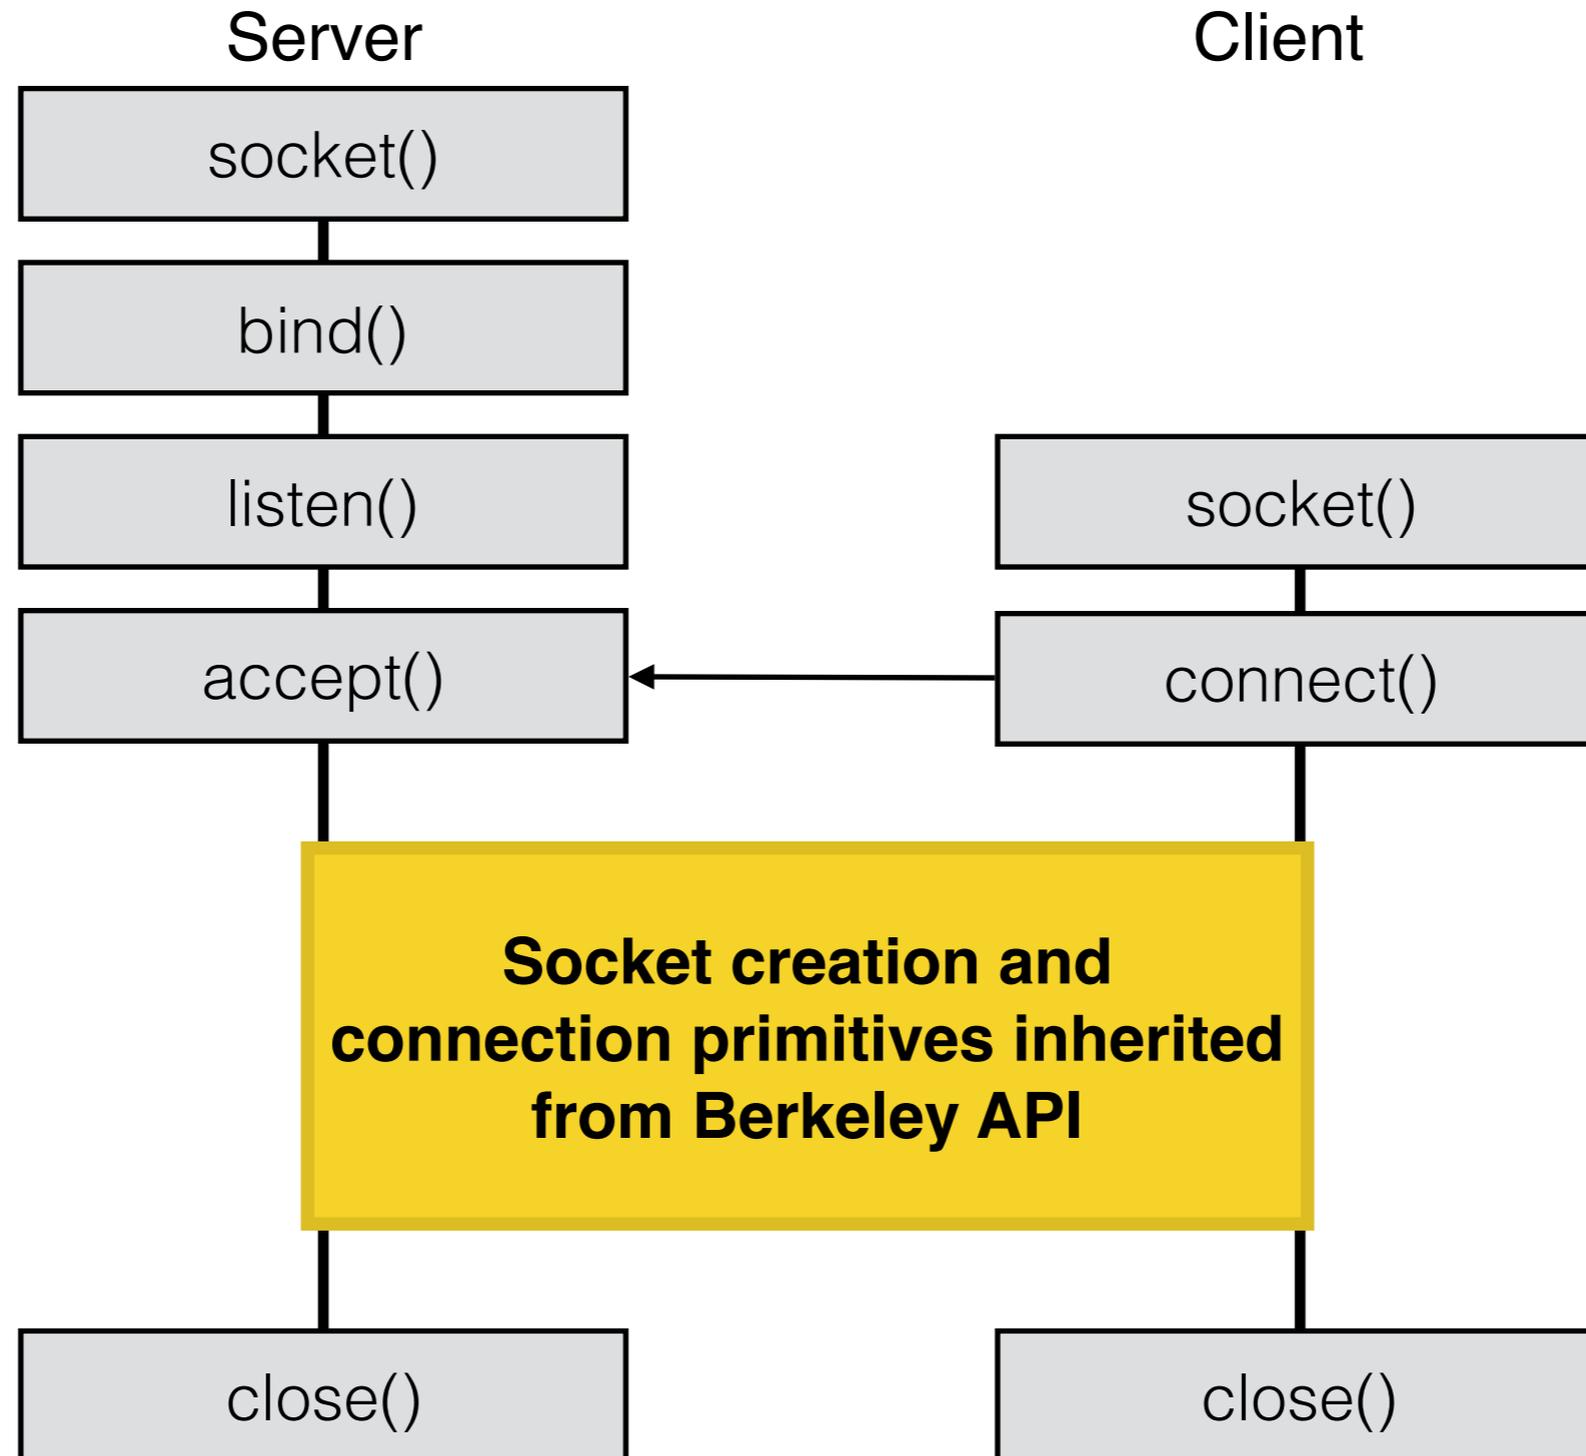
# Connections & Congestion Control

- Congestion control important to protect the network and other applications
- Need to select algorithm appropriate to application
- Connection-oriented service is useful in some scenarios
- Enables explicit setup and teardown of in-network state (e.g., for NAT mappings)

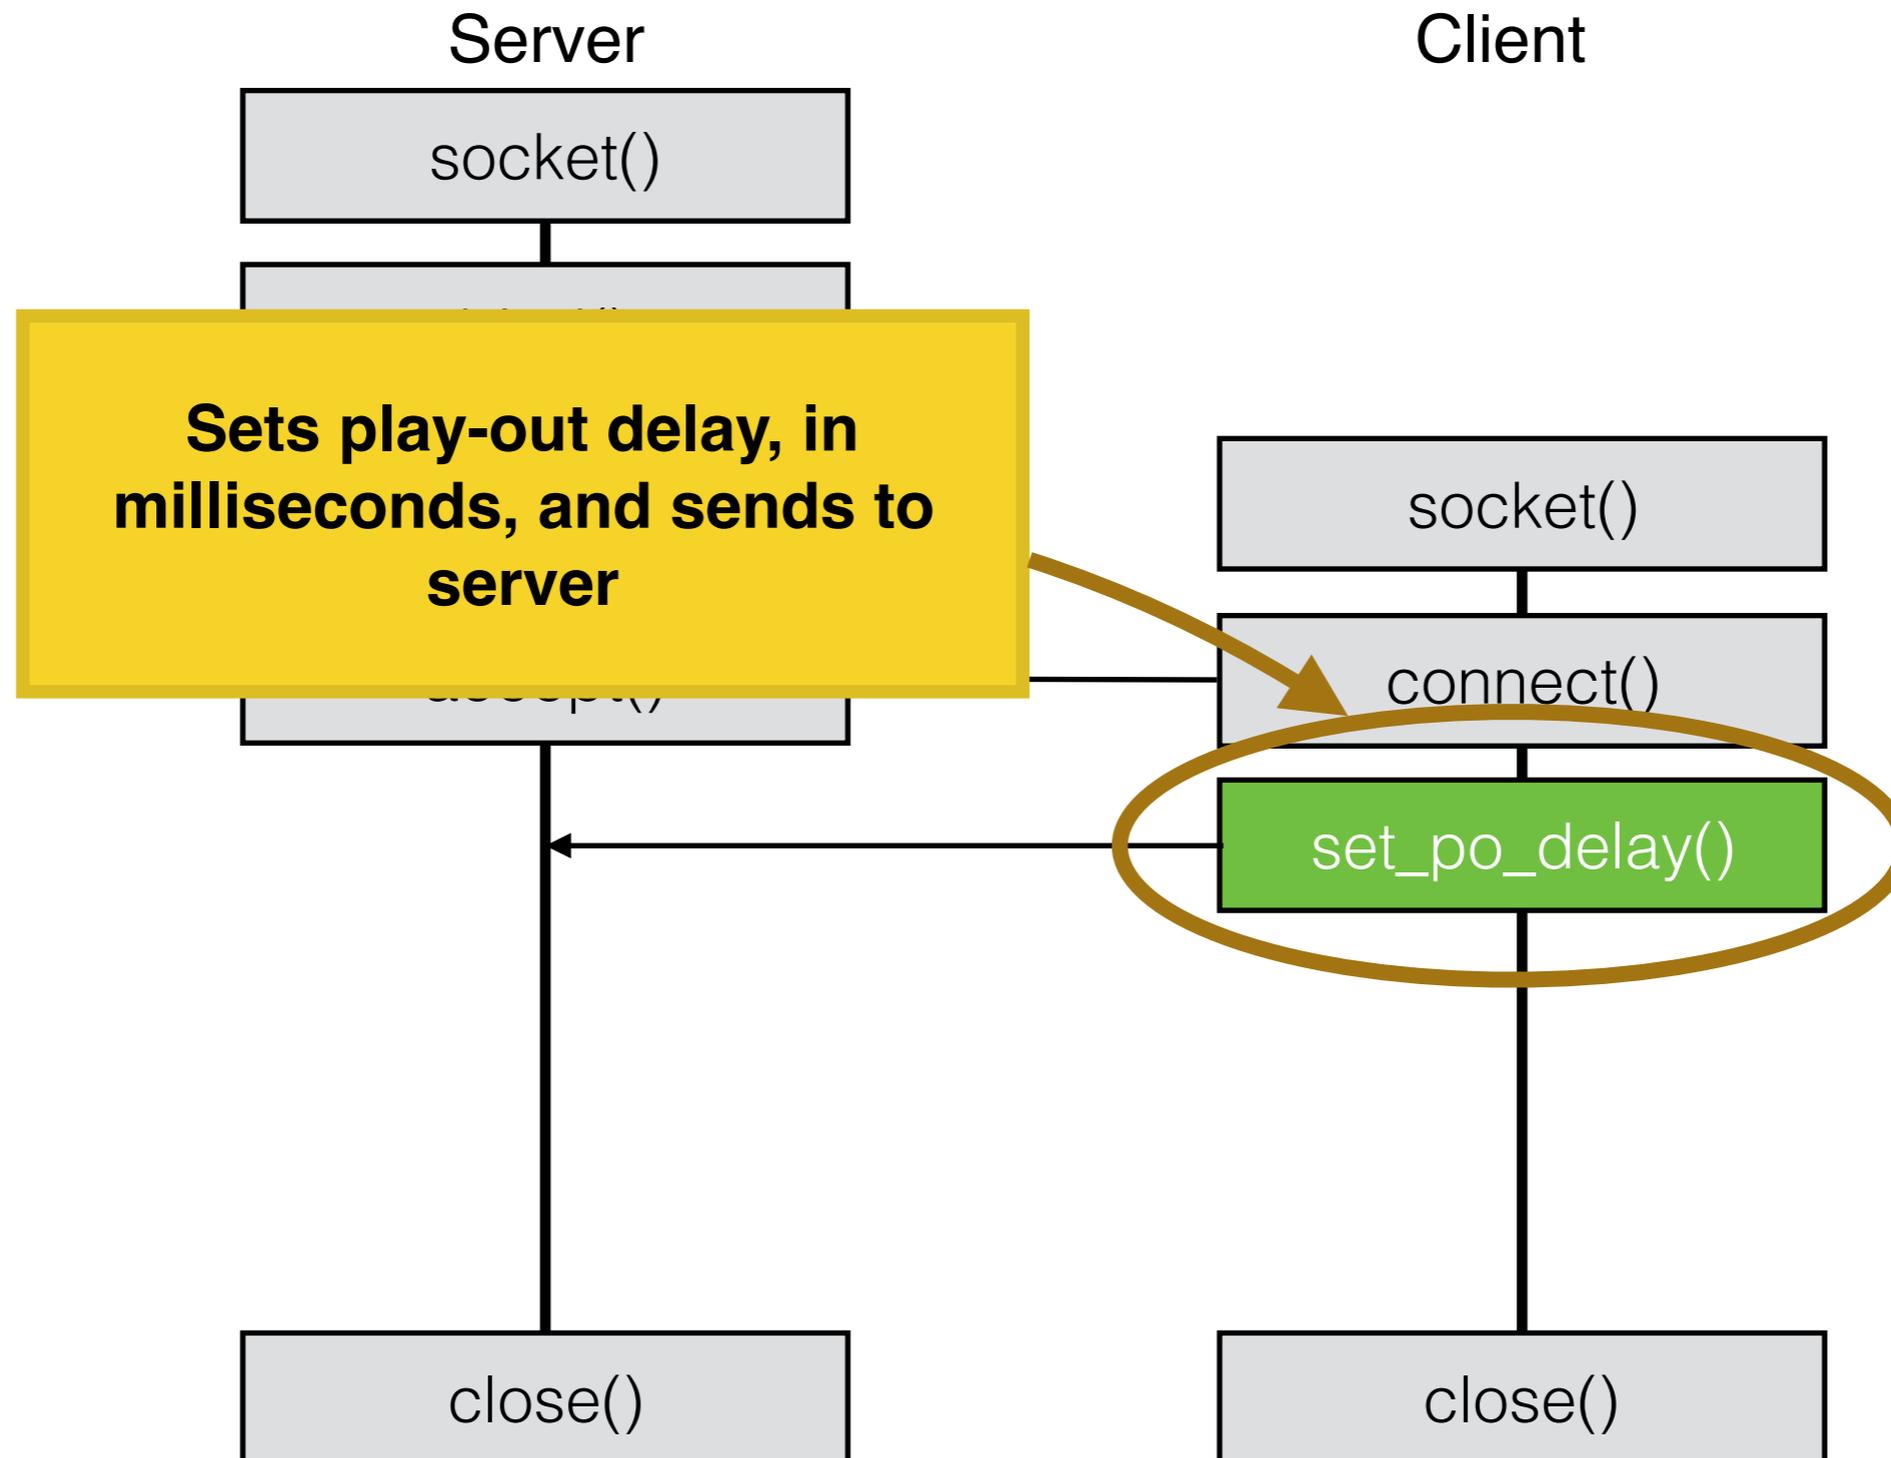
# Real-Time Transport Services

Transport Service	Requirement
Deadlines	Core
Partial reliability	Core
Dependencies	Core
Message-oriented	Core
Sub-streams	Core
Congestion controlled	Core
Connection oriented	Subsidiary
Keep-alive	Subsidiary

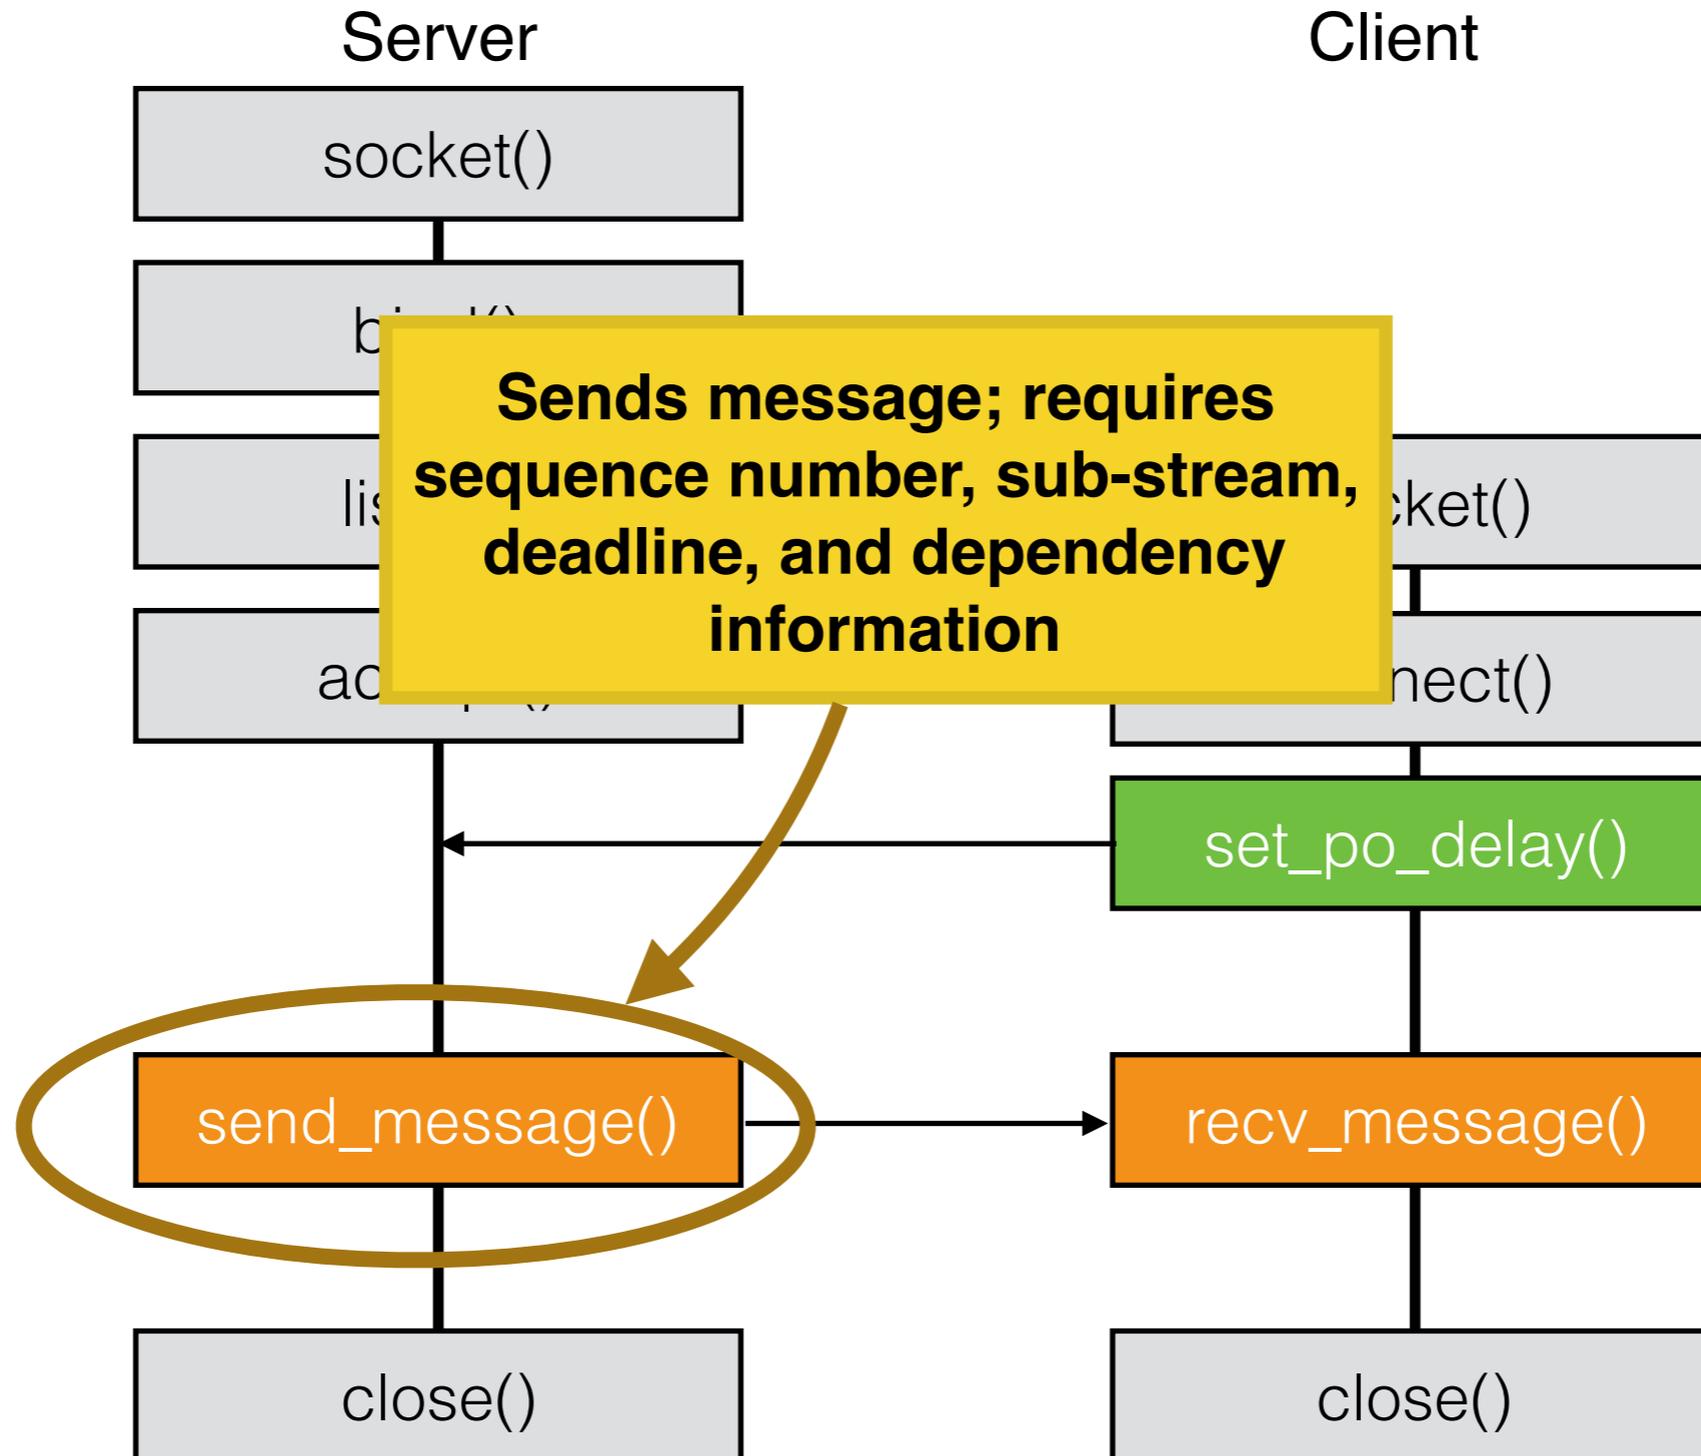
# Abstract API Lifecycle



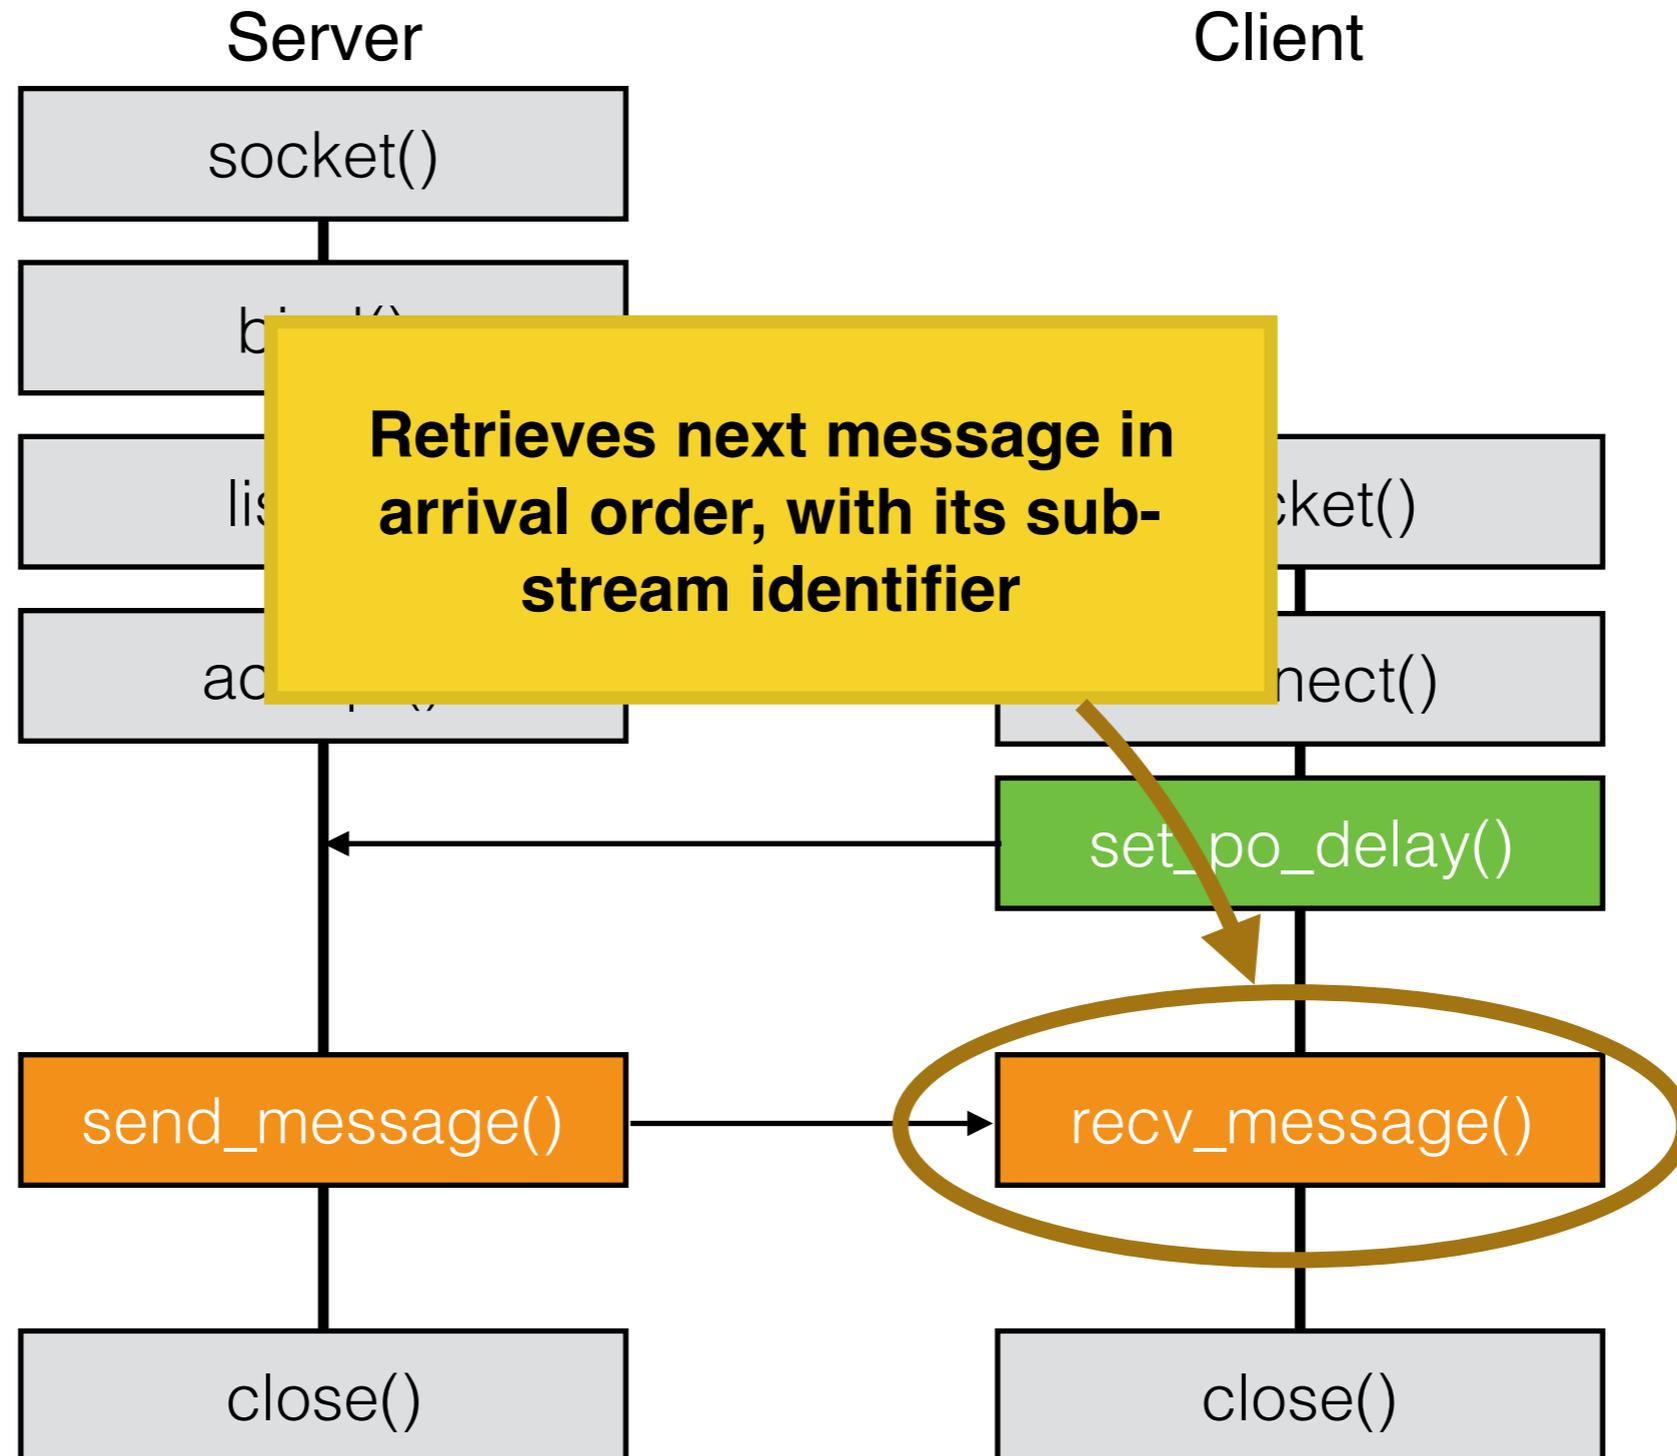
# Abstract API Lifecycle



# Abstract API Lifecycle



# Abstract API Lifecycle



# Realising transport services

- Need to support this combination of transport services
- Ossification restricts us to using either TCP or UDP — might change over time
- UDP first → fallback to TCP
- UDP not always available (5-10% - Google, 1-5% MAMI)

## Transport Service

Deadlines

Partial reliability

Dependencies

Message-oriented

Sub-streams

Congestion controlled

Connection oriented

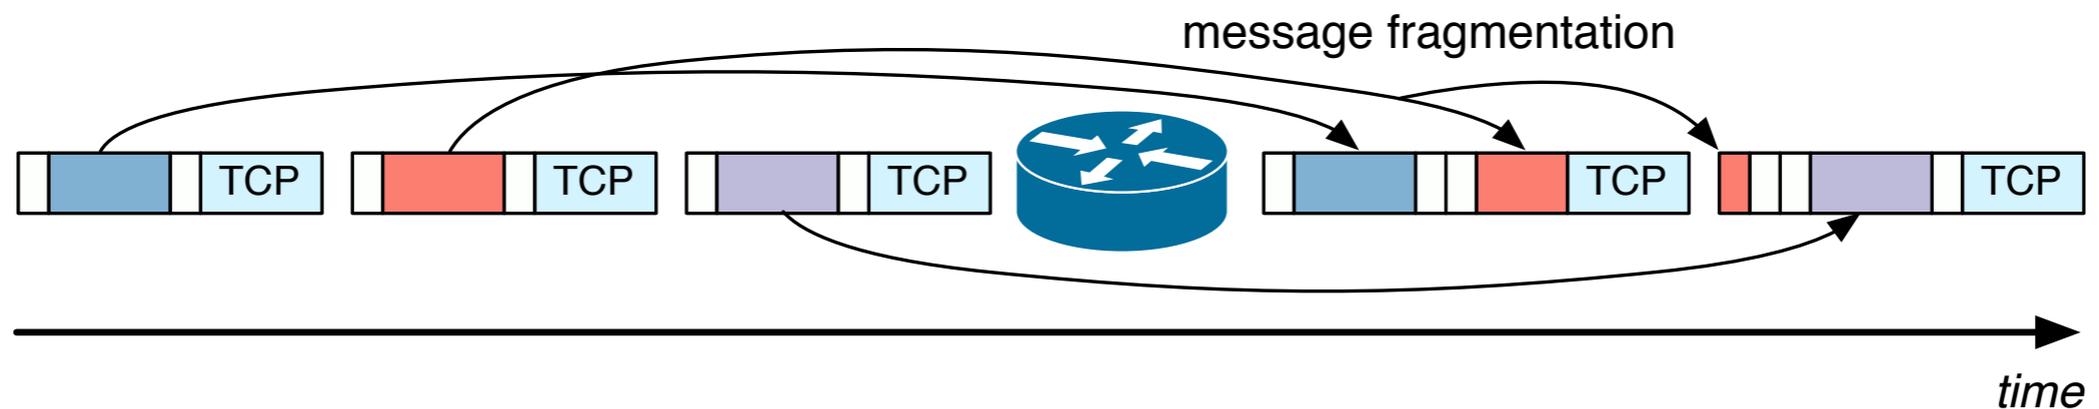
Keep-alive

# UDP as a substrate

- Already supports the sending of datagrams/messages
- Support for partial reliability requires detecting loss, retransmitting if message will arrive before deadline
- Need an estimate of one-way network delay
- Sub-stream support requires small header in each message
- Connections and congestion control can be added

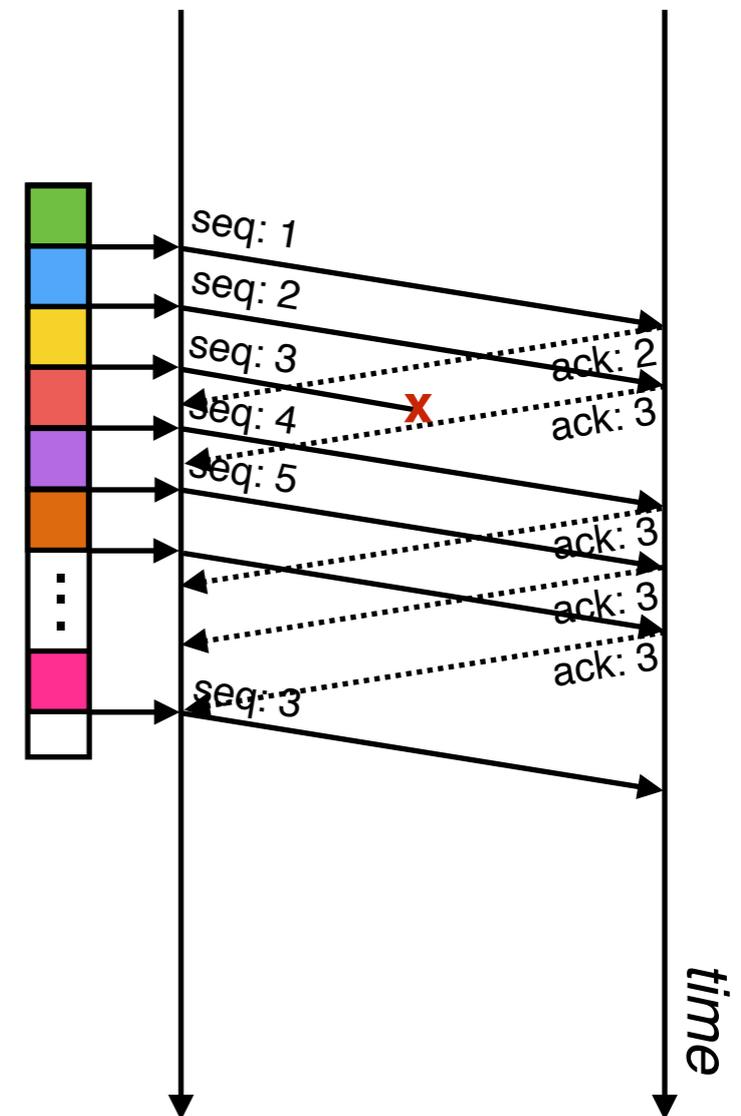
# TCP as a substrate

- Messaging requires a framing mechanism, to support resegmenting middleboxes — e.g., COBS, as in Minion/uTCP
- Sub-stream support requires small header in each message
- Already supports connections
- Congestion control supported, but algorithm fixed: support for other algorithms as in DCCP



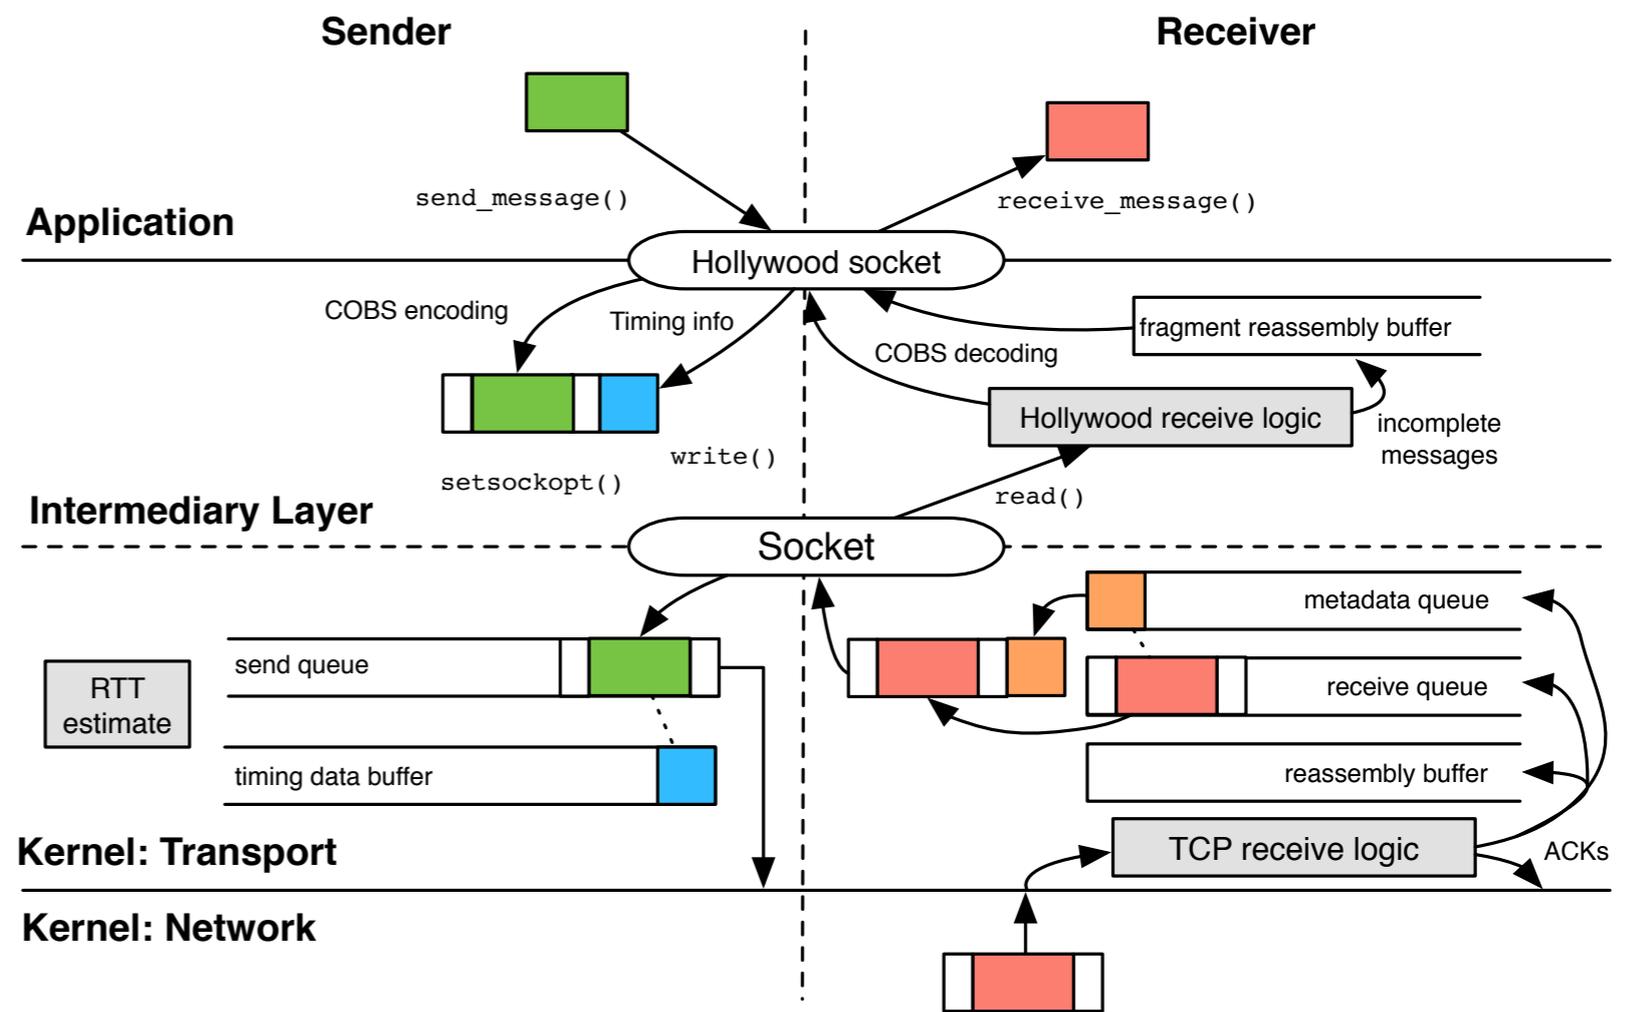
# Relaxing reliability in TCP

- Middleboxes ossified around TCP do not expect gaps in the TCP sequence space
- Need to “retransmit” missing TCP sequence numbers, without retransmitting payloads — *inconsistent retransmissions*
- Mapping between data and TCP sequence number is no longer constant



# TCP Hollywood

- Unordered, partially reliable message-oriented delivery
- Intermediary layer: COBS encoding to maintain message boundaries
- Kernel: unordered delivery of incoming segments



# TCP Hollywood

- Uses inconsistent retransmissions to support partial reliability
- Evaluation between TCP Hollywood server and 14 clients around the UK
- Evaluations also conducted by Honda et al.
- Small scale — more evaluations needed

ISP	Port	
	80	4001
<b>Fixed-line</b>		
Andrews & Arnold	●	●
BT	●	●
Demon	●	●
EE	●	●
Eclipse	●	●
Sky	●	●
TalkTalk	●	●
Virgin Media	●	●
<b>Mobile</b>		
EE	▲	▲
O2	▲	▲
Three	●	●
Vodafone	▲	●

# Summary

- Services can be implemented on TCP and UDP
- TAPS WG formulating list of services by breaking down existing protocols
- Here, top down: start with application, define services without constraints of existing protocols

## Transport Service

Deadlines

Partial reliability

Dependencies

Message-oriented

Sub-streams

Congestion controlled

Connection oriented

Keep-alive