

Transport Protocols for Networked Games

TCP or UDP ?

Why use TCP?

- TCP provides reliable, in-order delivery
- TCP goes through most firewalls, UDP does not
- TCP manages connection for us

Why not to use TCP?

- TCP incurs higher latency
- Don't always need reliability and in-order delivery
- High header overhead

position = 10 →
position = 13 → X
position = 15 →

Updated position not delivered to
application until (outdated) lost packet
is received

A's position = 10 \longrightarrow
B's position = 13 \longrightarrow X
C's position = 15 \longrightarrow

Some messages need not be delivered in
sequence.



**Gestures from someone far away need
not be received reliably.**

A study on ShenZhou
Online shows that
46% of the bandwidth
is occupied by TCP
header

enet.cubik.org

**A library that provides
reliability, sequencing,
connection managements
over UDP**

Delivery can be stream-oriented (like TCP) or message-oriented (like UDP)

Supports partial reliability

```
enet_packet_create (“abc”,  
4, ENET_PACKET_FLAG_RELIABLE)
```

**Retransmission triggered by
timeout based on RTT**

**Data in queue are bundled into
one packet if there is space**

enet.cubik.org

**Portable, easy to use,
but still, most firewalls
block UDP traffic**

- MMORPG that uses **TCP**:WoW, Lineage I/II, Guild Wars, Ragnarok Online, Anarchy Online, Mabinogi
- MMORPG that uses **UDP**: EverQuest, SW Galaxies, City of Heroes, Ultima Online, Asherons Call, FFXI

Need to study the use
of TCP for networked
games


How to provide
reliability over UDP?

How slow is TCP, really?

**Which part of TCP is the
root of slowness?**

Can we fix TCP?

A Quick Review of TCP




TCP Delayed ACK

TCP Spec: max **500ms** delay
Most implementation: **200ms**


Why delay ACK?


- reduce num of ACKs
- in case receiver wants to send data within 200ms (in which case it can piggyback the ACK with data)
- give sender time to buffer more data for sending (avoid silly window syndrome)



TCP Fast Retransmission

**Definition of Dup ACKs in
4.4BSD and Stevens:
“pure ACK, cannot
piggyback with data”**





Spurious Retransmission

RTO estimation

$$E_i = 7E_{i-1}/8 + RTT/8$$

$$V_i = 3V_{i-1}/4 + |RTT - E_{i-1}|/4$$

$$RTO = \max(E_i + 4V_i, 1s)$$

Linux's RTO estimation

$$E_i = 7E_{i-1}/8 + RTT/8$$

$$V_i = 3V_{i-1}/4 + |RTT - E_{i-1}|/4$$

$$W_i = \min(V_i, 50\text{ms})$$

$$RTO = \max(200\text{ms}, E_i + W_i)$$

Note:


Delayed ACK \Rightarrow

increase RTT \Rightarrow

increase RTO

Congestion Control

Window Size



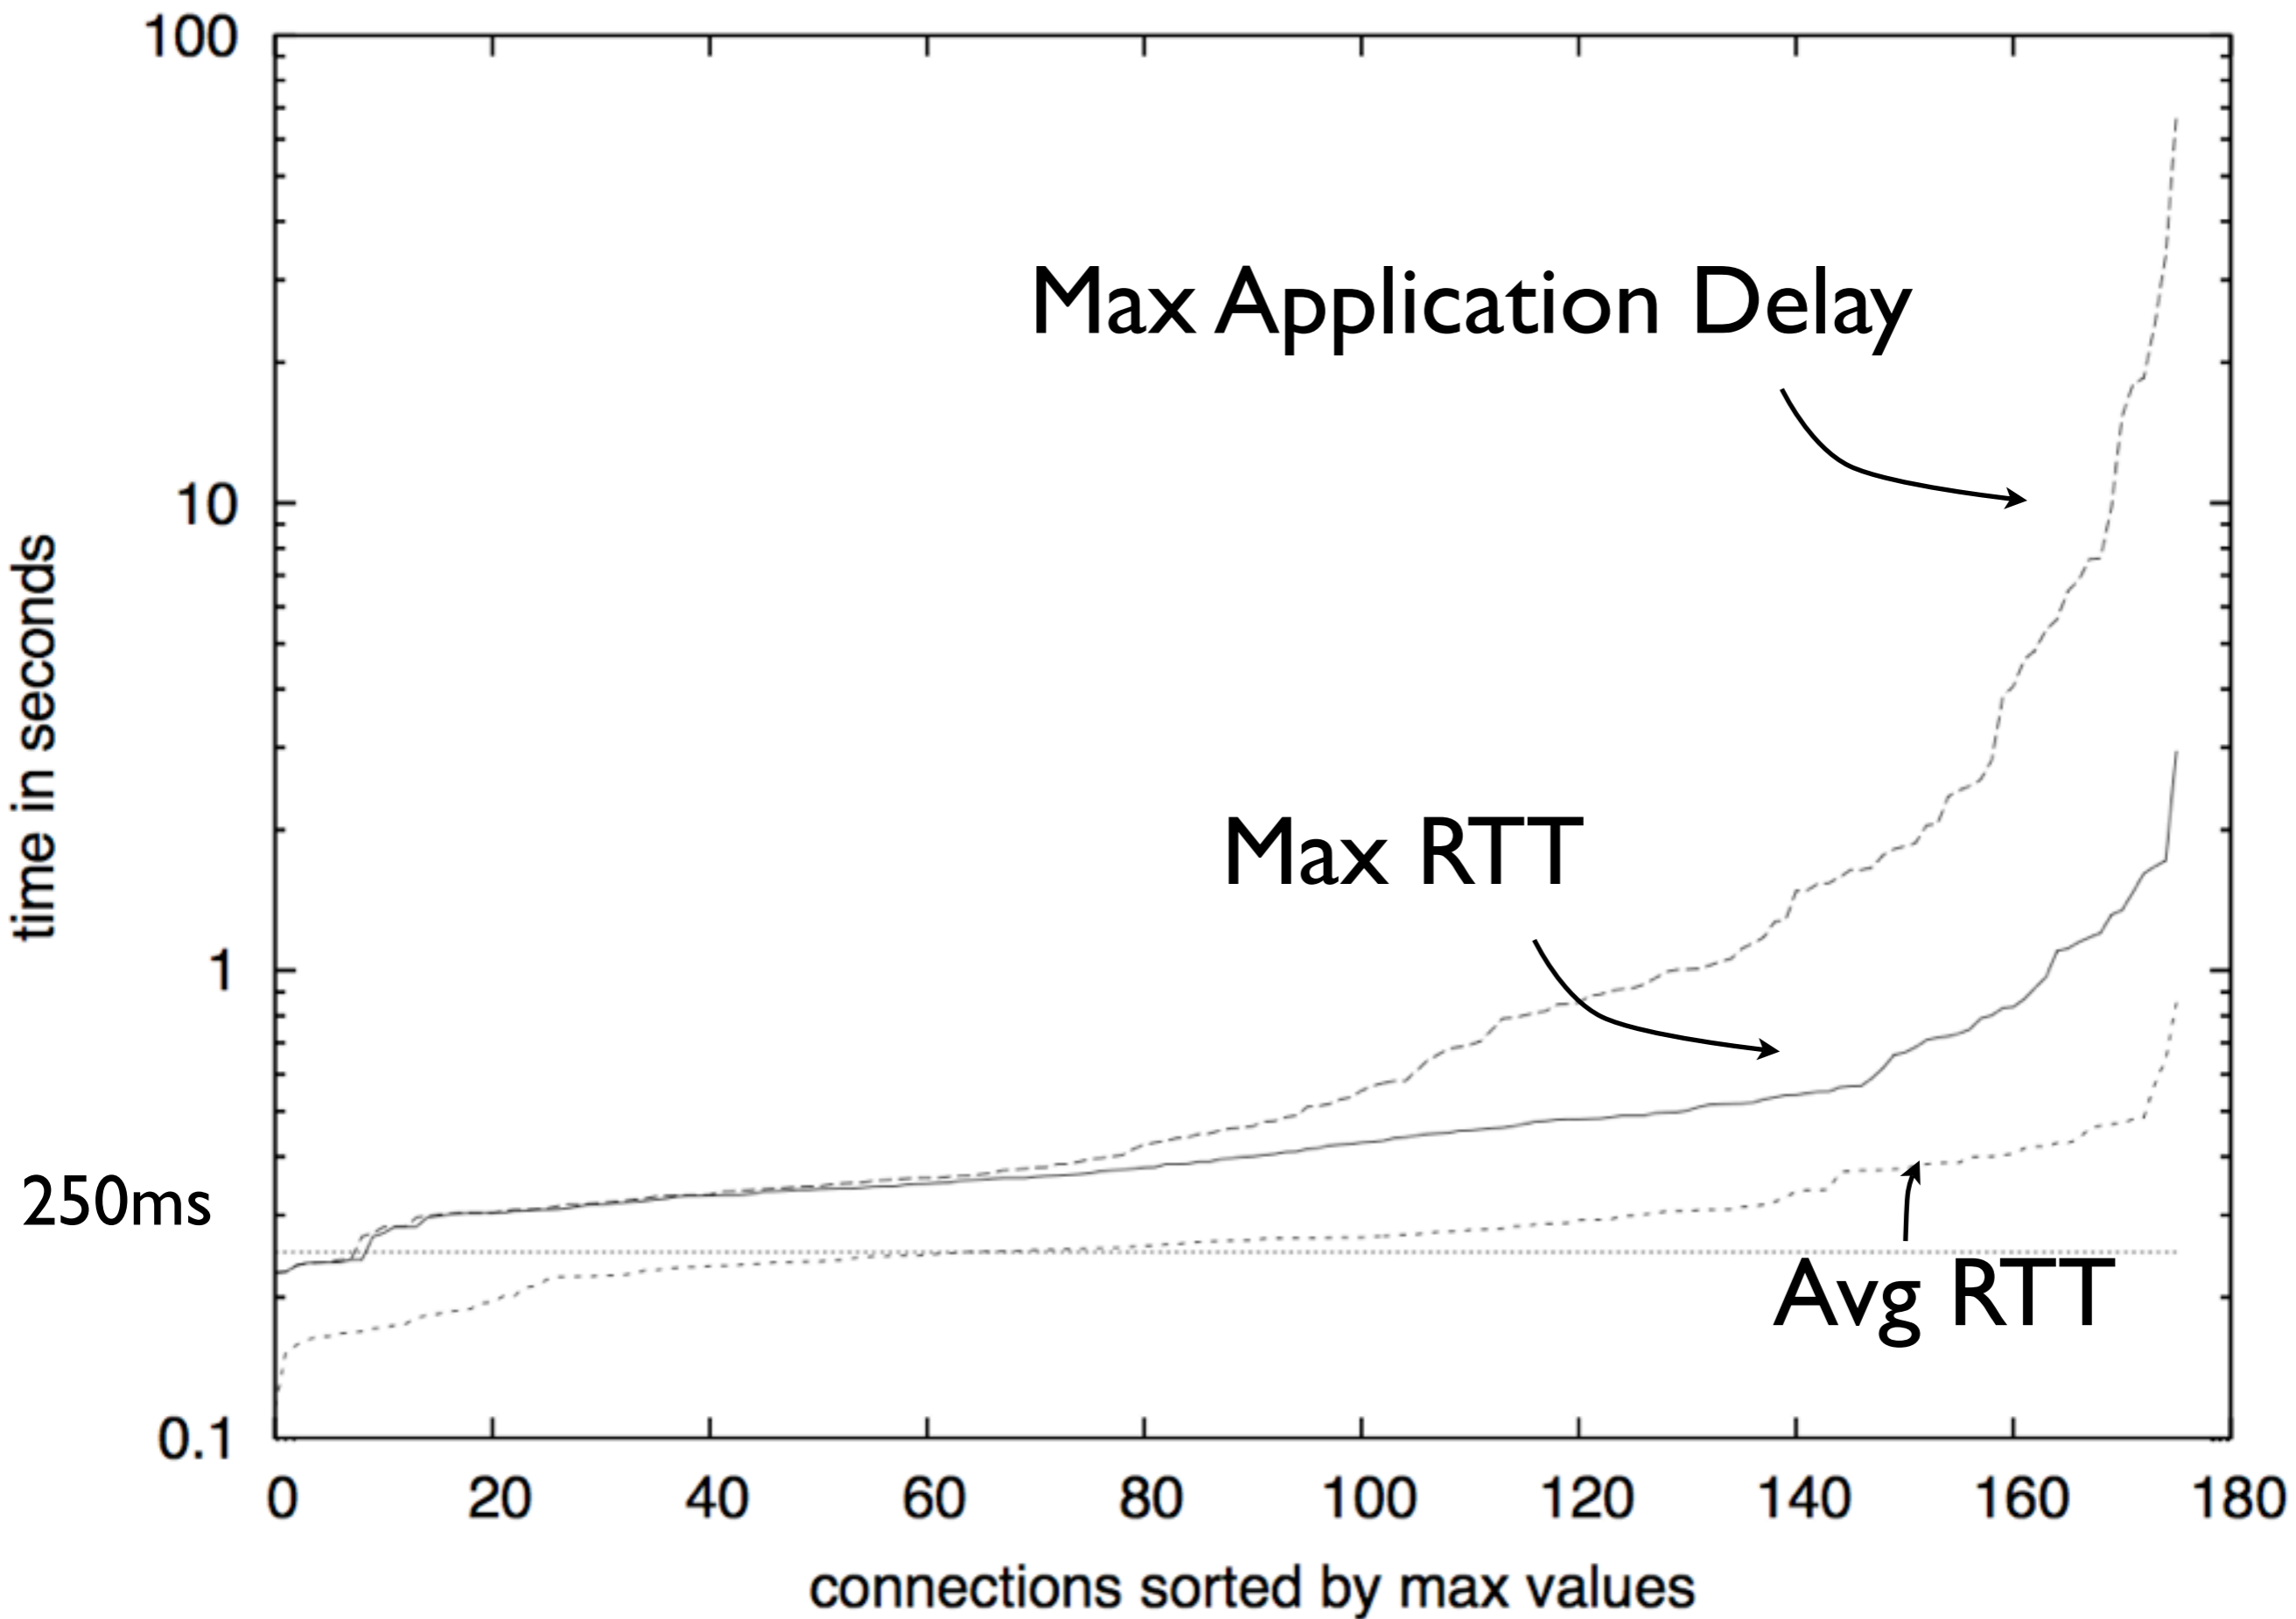
Time


TCP Congestion Control

**Congestion window
resets to 2 after an idle
period ($> \text{RTO}$)**

**What does real game
traffic look like?**







About 4 packets / sec


Average Payload:
100 Bytes


Loss Rate 1%

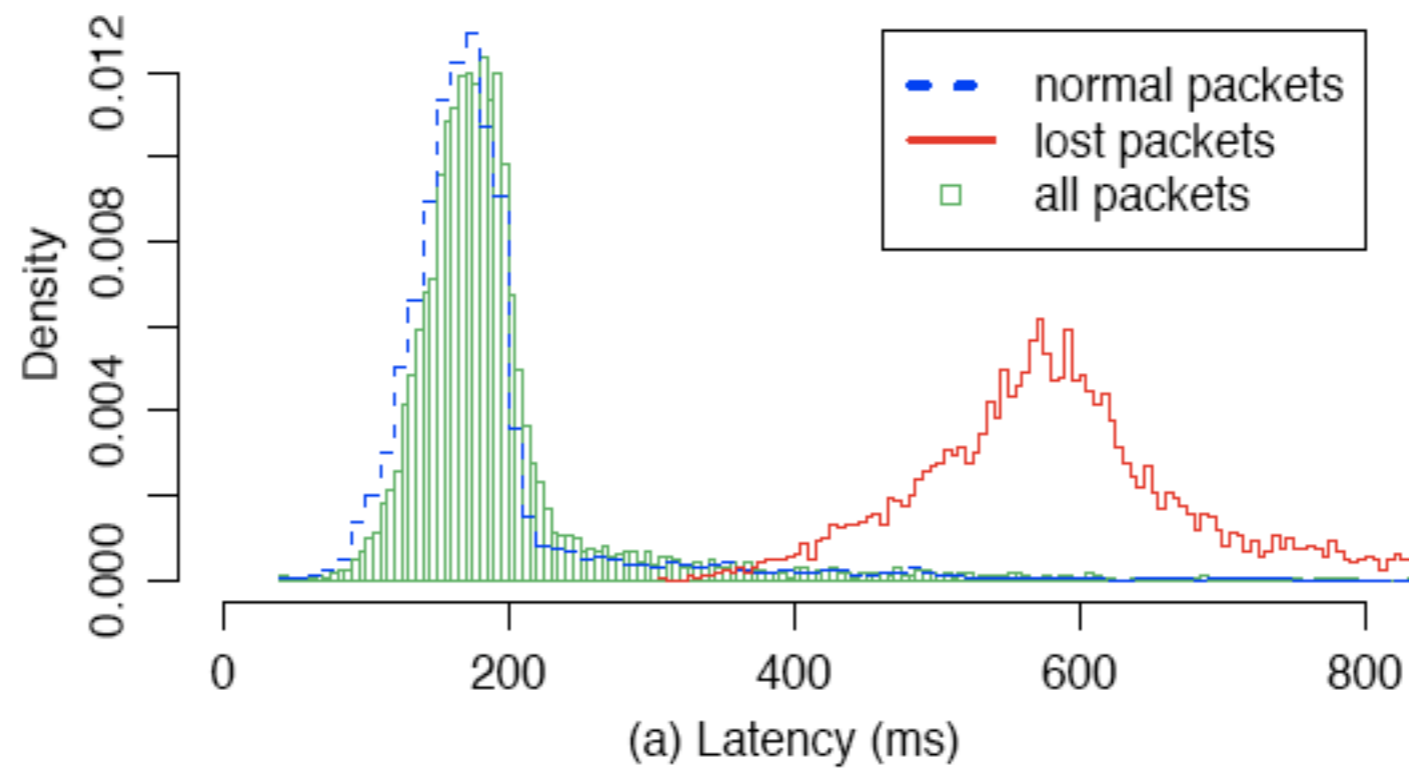
**But some experience 6
retransmissions**

ShenZhou Online









Other games have
low packet rate and
small payload size too

application (platform)	payload size (bytes)			average
	average	min	max	
Anarchy Online (PC) [‡]	98	8	1333	
World of Warcraft (PC)	26	6	1228	
Counter Strike (PC) [‡]	36	25	1342	
Halo 3 (Xbox 360) ^{†‡}	247	32	1264	
Halo 3 (Xbox 360) ^{†‡}	270	32	280	
Gears of War (Xbox 360) [‡]	66	32	705	
Tony Hawk's Project 8 (Xbox 360) [‡]	90	32	576	
Test Drive Unlimited (Xbox 360) [‡]	80	34	104	

[†] For Halo 3 (beta version), we also show differences between intens

[‡] The presented values are average values over all players (sending

Table 1: Examples of game stream packet st

Table 1: Examples of game stream packet st

packet interarrival time (ms)				percentiles		avg. bandwidth requirement	
average	median	min	max	1%	99%	(pps)	(bps)
632	449	7	17032	83	4195	1.582	2168
314	133	0	14855	0	3785	3.185	2046
124	65	0	66354	34	575	8.064	19604
36	33	0	1403	32	182	27.778	60223
67	66	32	716	64	69	14.925	35888
457	113	3	10155	14	8953	2.188	10264
308	163	0	4070	53	2332	3.247	5812
40	33	0	298	0	158	25.000	22912

intensive (the upper row) and moderate (the lower row) action.
 (including minimum 1000 packets) within the period of the trace.

Packet statistics per stream based on packet traces

Packet statistics per stream based on packet traces

“Thin Streams”

Findings I:
Fast retransmission
rarely triggered

**In ShenZhou Online traces, fail
to trigger fast retransmission
because
insufficient dup ACK (50%)
interrupted by data (50%)**

Findings 2:
Delay due mostly to
timeout




Figure 9: Average latency of dropped packets

Figure 9: Average latency of dropped packets



Findings 3:
Congestion window
reset is frequent

**12% - 18% of packets
faces window reset**

think..

think..

think..

click (tank attack here)



click (missile launch there)




click (charge soldiers)



The last command is delayed as congestion
window = 2

How to make TCP (or, transport protocol) go faster in these games?

I. Remove exponential backoff




TCP Timeout

2. Make RTO Smaller

make sure minimum
RTO is not 1s


**spurious retransmission
is not disastrous**

3. Make Fast Retransmit Faster




Retransmit after one duplicate ACK

4. Retransmission Bundling



Retransmit all unacknowledge data in queue

5. Redundant Data Bundling




Send any unacknowledged segment in queue as long as there is space. Lost data gets recovered in the next transmission before retransmission.

6. Turn off or reduce Delayed ACKs

Packet interarrival time
on average $> 200\text{ms}$
(can't combine two ACKs into one)

7. Combine Thin Streams into Thicker Stream



TCP for Games

- remove exponential backoff
- reduce RTO
- make fast retransmit faster
- retransmit aggressively
- don't delay ACK
- combine into thick streams

Beyond TCP and UDP?


Ideal Transport Protocol for Games

- optional reliability
- optional order-of-delivery
- flexibility in organizing messages into different classes with different requirements

SCTP

Stream Control Transport Protocol

Multi-streaming: multiple independent streams



**A stream can be either reliable or
non-reliable**

**Data from multiple streams can be
bundled into one packets**

Message-oriented (like UDP)

**Message can be flagged for
unordered delivery**

Ideal Transport Protocol for Games

- optional reliability
- optional order-of-delivery
- flexibility in organizing messages into different classes with different requirements

SCTP for Games?

States of SCTP

- rumored to be in Vista
- available in FreeBSD 7 and Solaris 10
- lksctp project for Linux