

Lecture 10

**Transport Protocol  
for  
Networked Games**

29 March 2010

**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application until (outdated) lost packet  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.**

# 46%

of ShenZhou Online 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**

- **Uses TCP:** WoW, Lineage I/II, Guild Wars, Ragnarok Online, Anarchy Online, Mabinogi
- **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

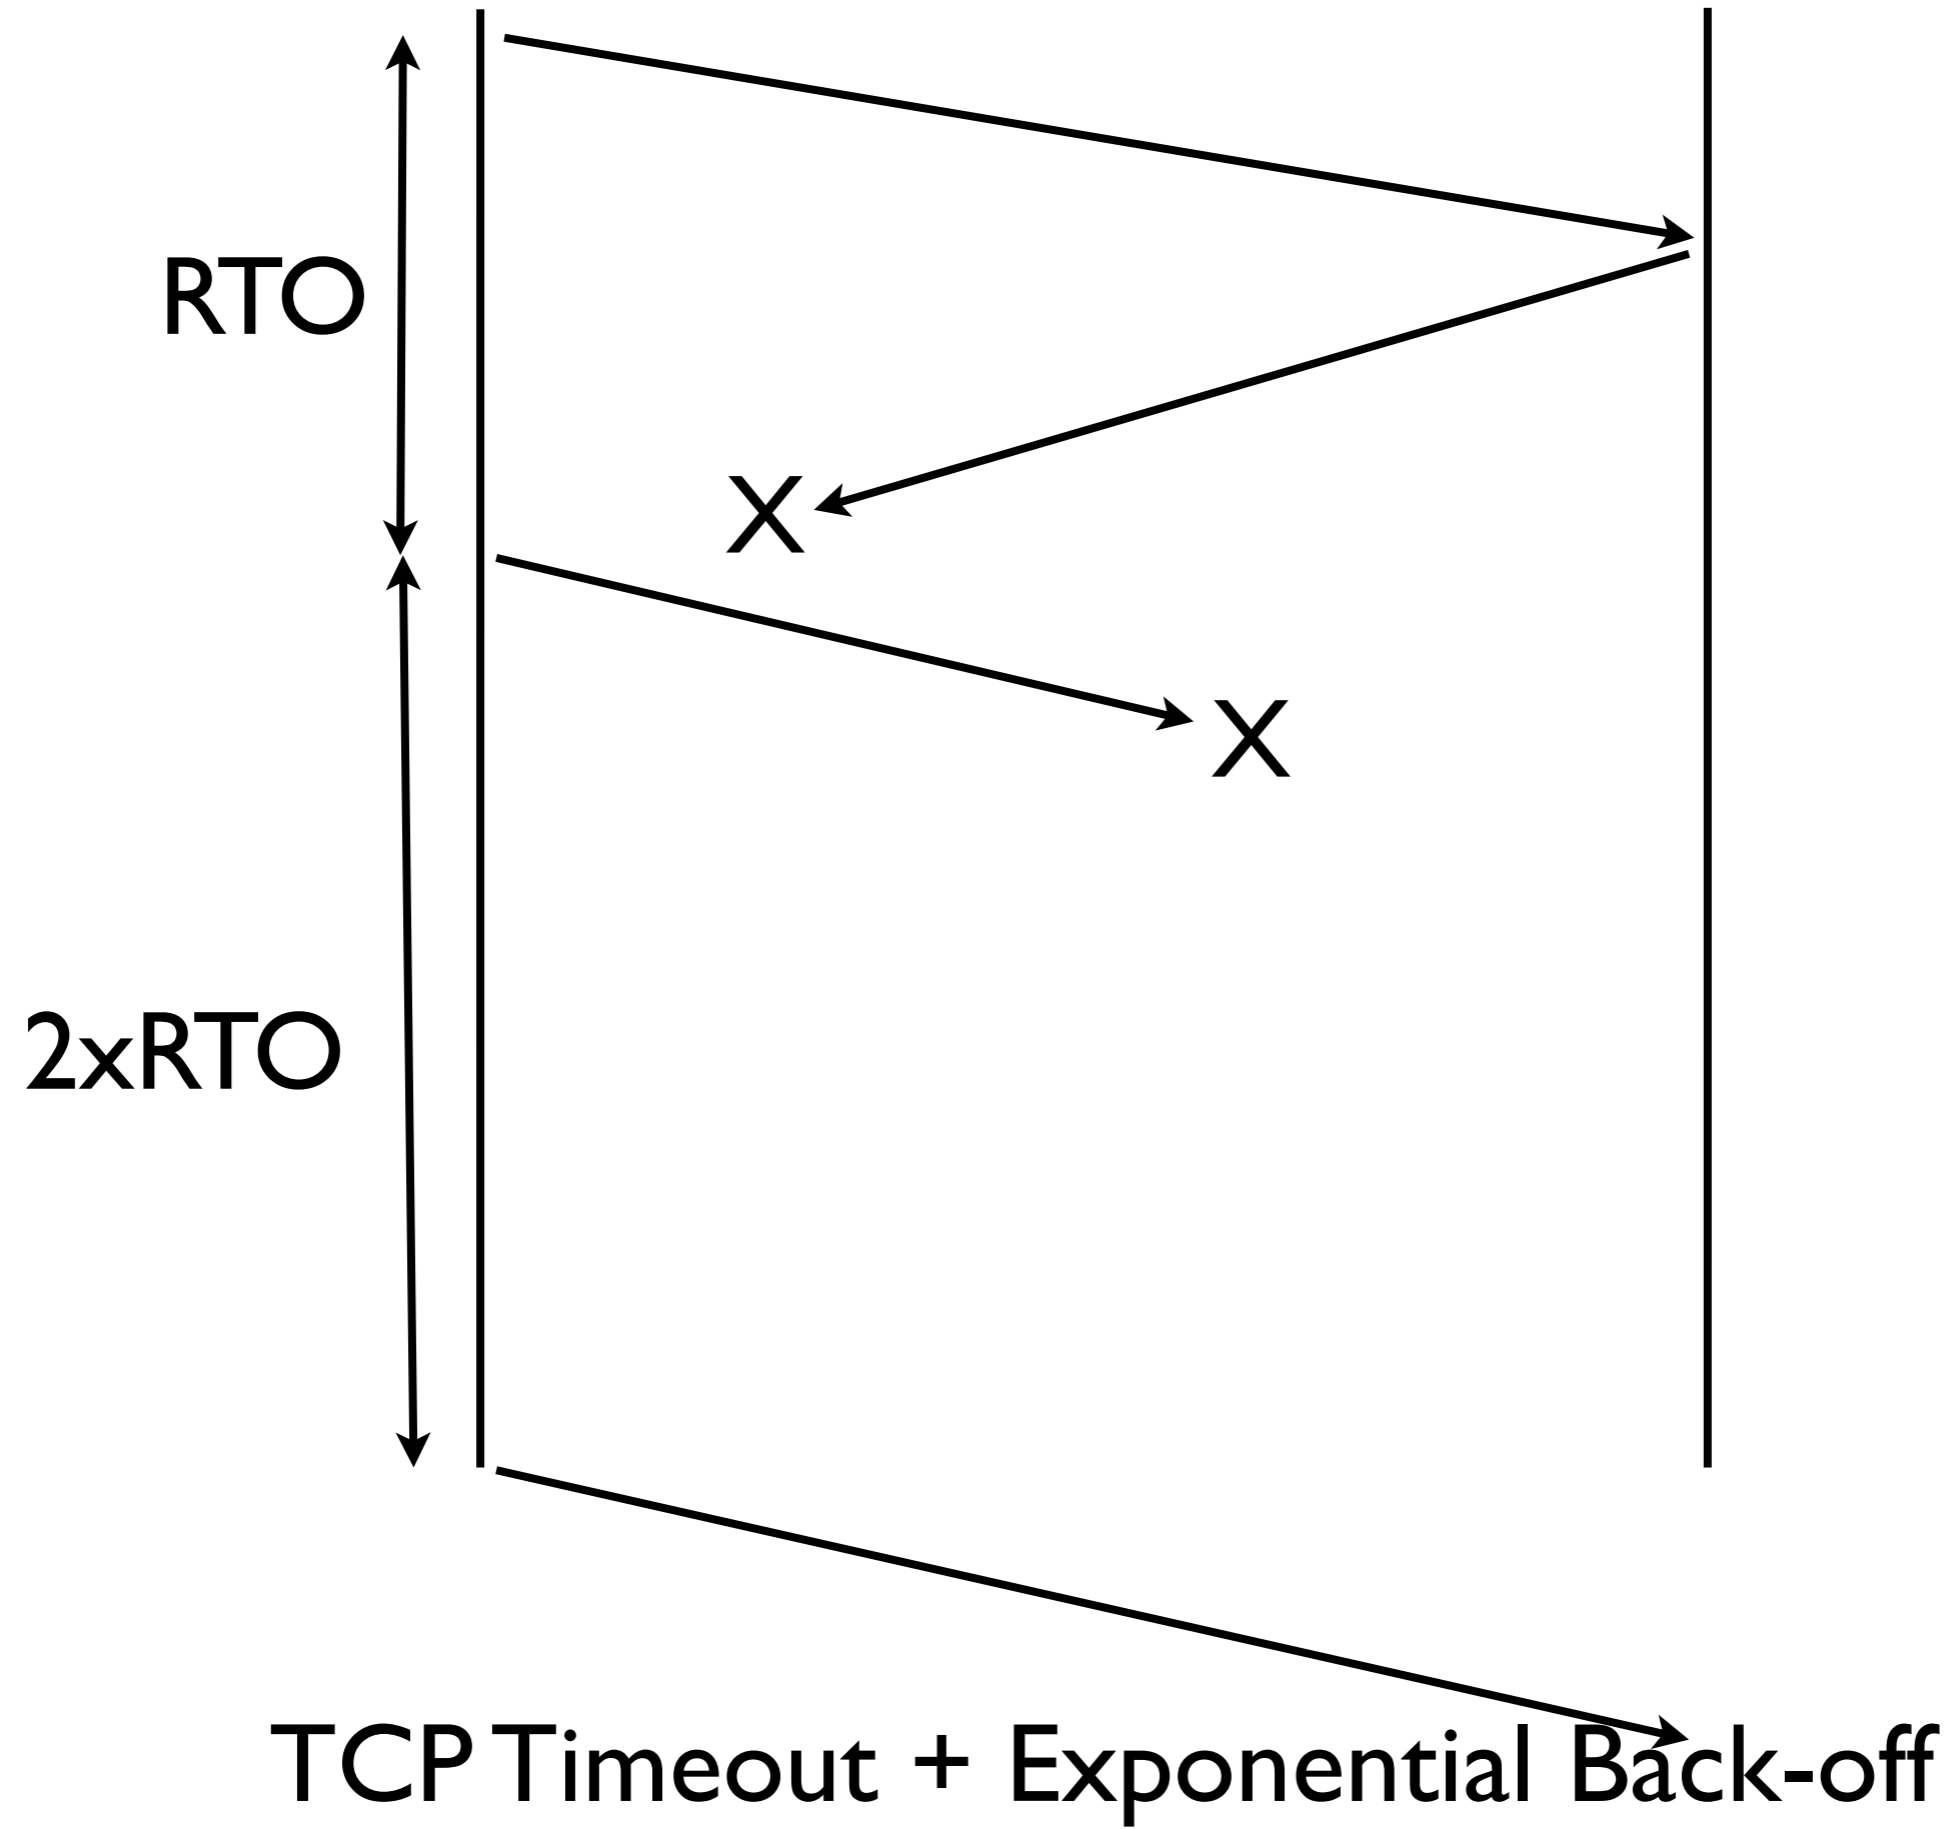
TCP Spec: max **500ms** delay  
Most implementation: **200ms**


# 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)



**Definition of Dup ACKs  
in 4.4BSD and Stevens:  
“pure ACK with no data”**





Spurious Retransmission

# RTO estimation

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

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

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

# Linux's RTO estimation


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

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


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

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

**delayed ACK**




**increase RTT**



**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?**

**low** packet rate  
**small** packet size

**“Thin Streams”**







**About 4 packets / sec**


Average Payload:  
**100 Bytes**

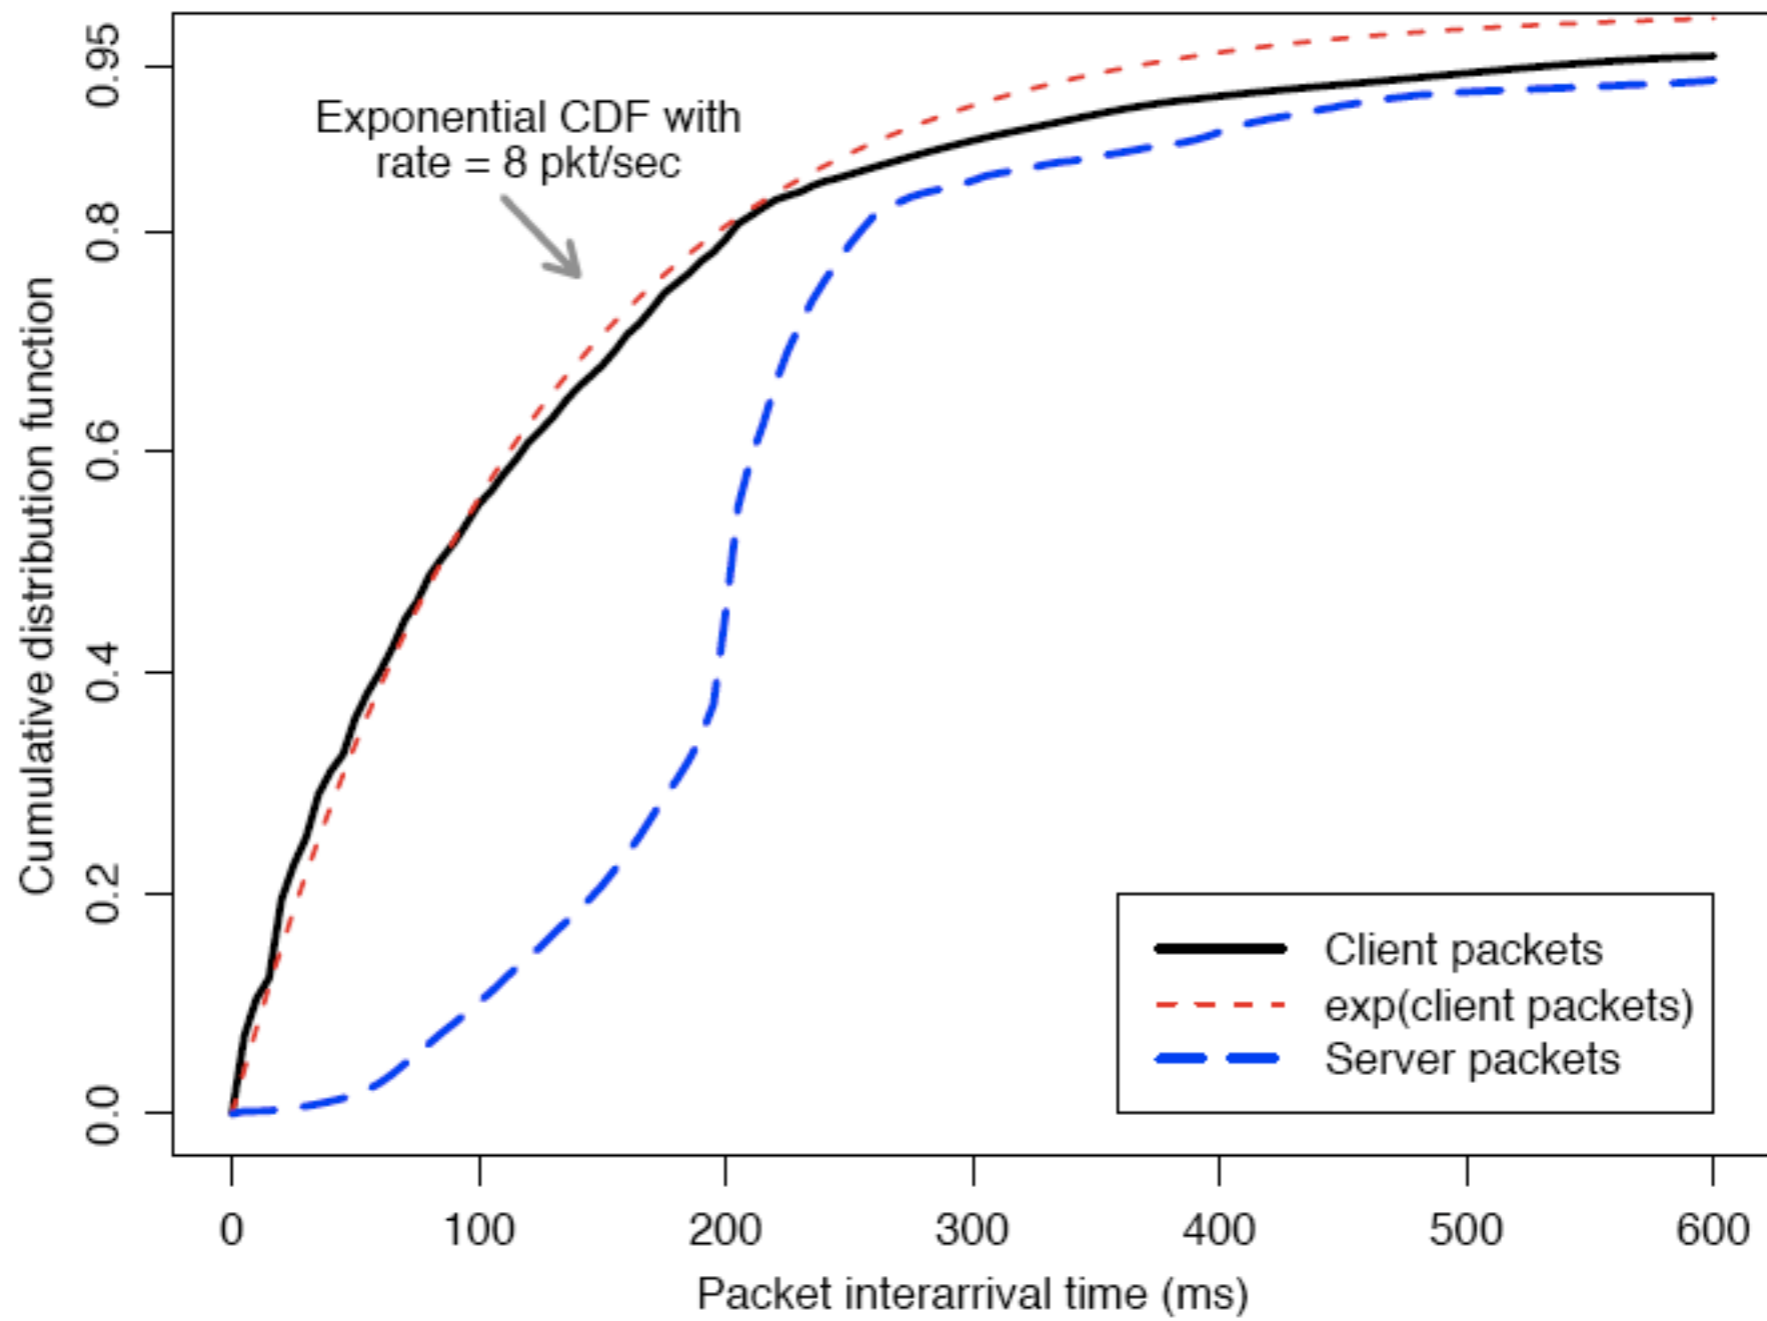
**Loss Rate 1%**


**But some experience 6  
retransmissions**

# ShenZhou Online





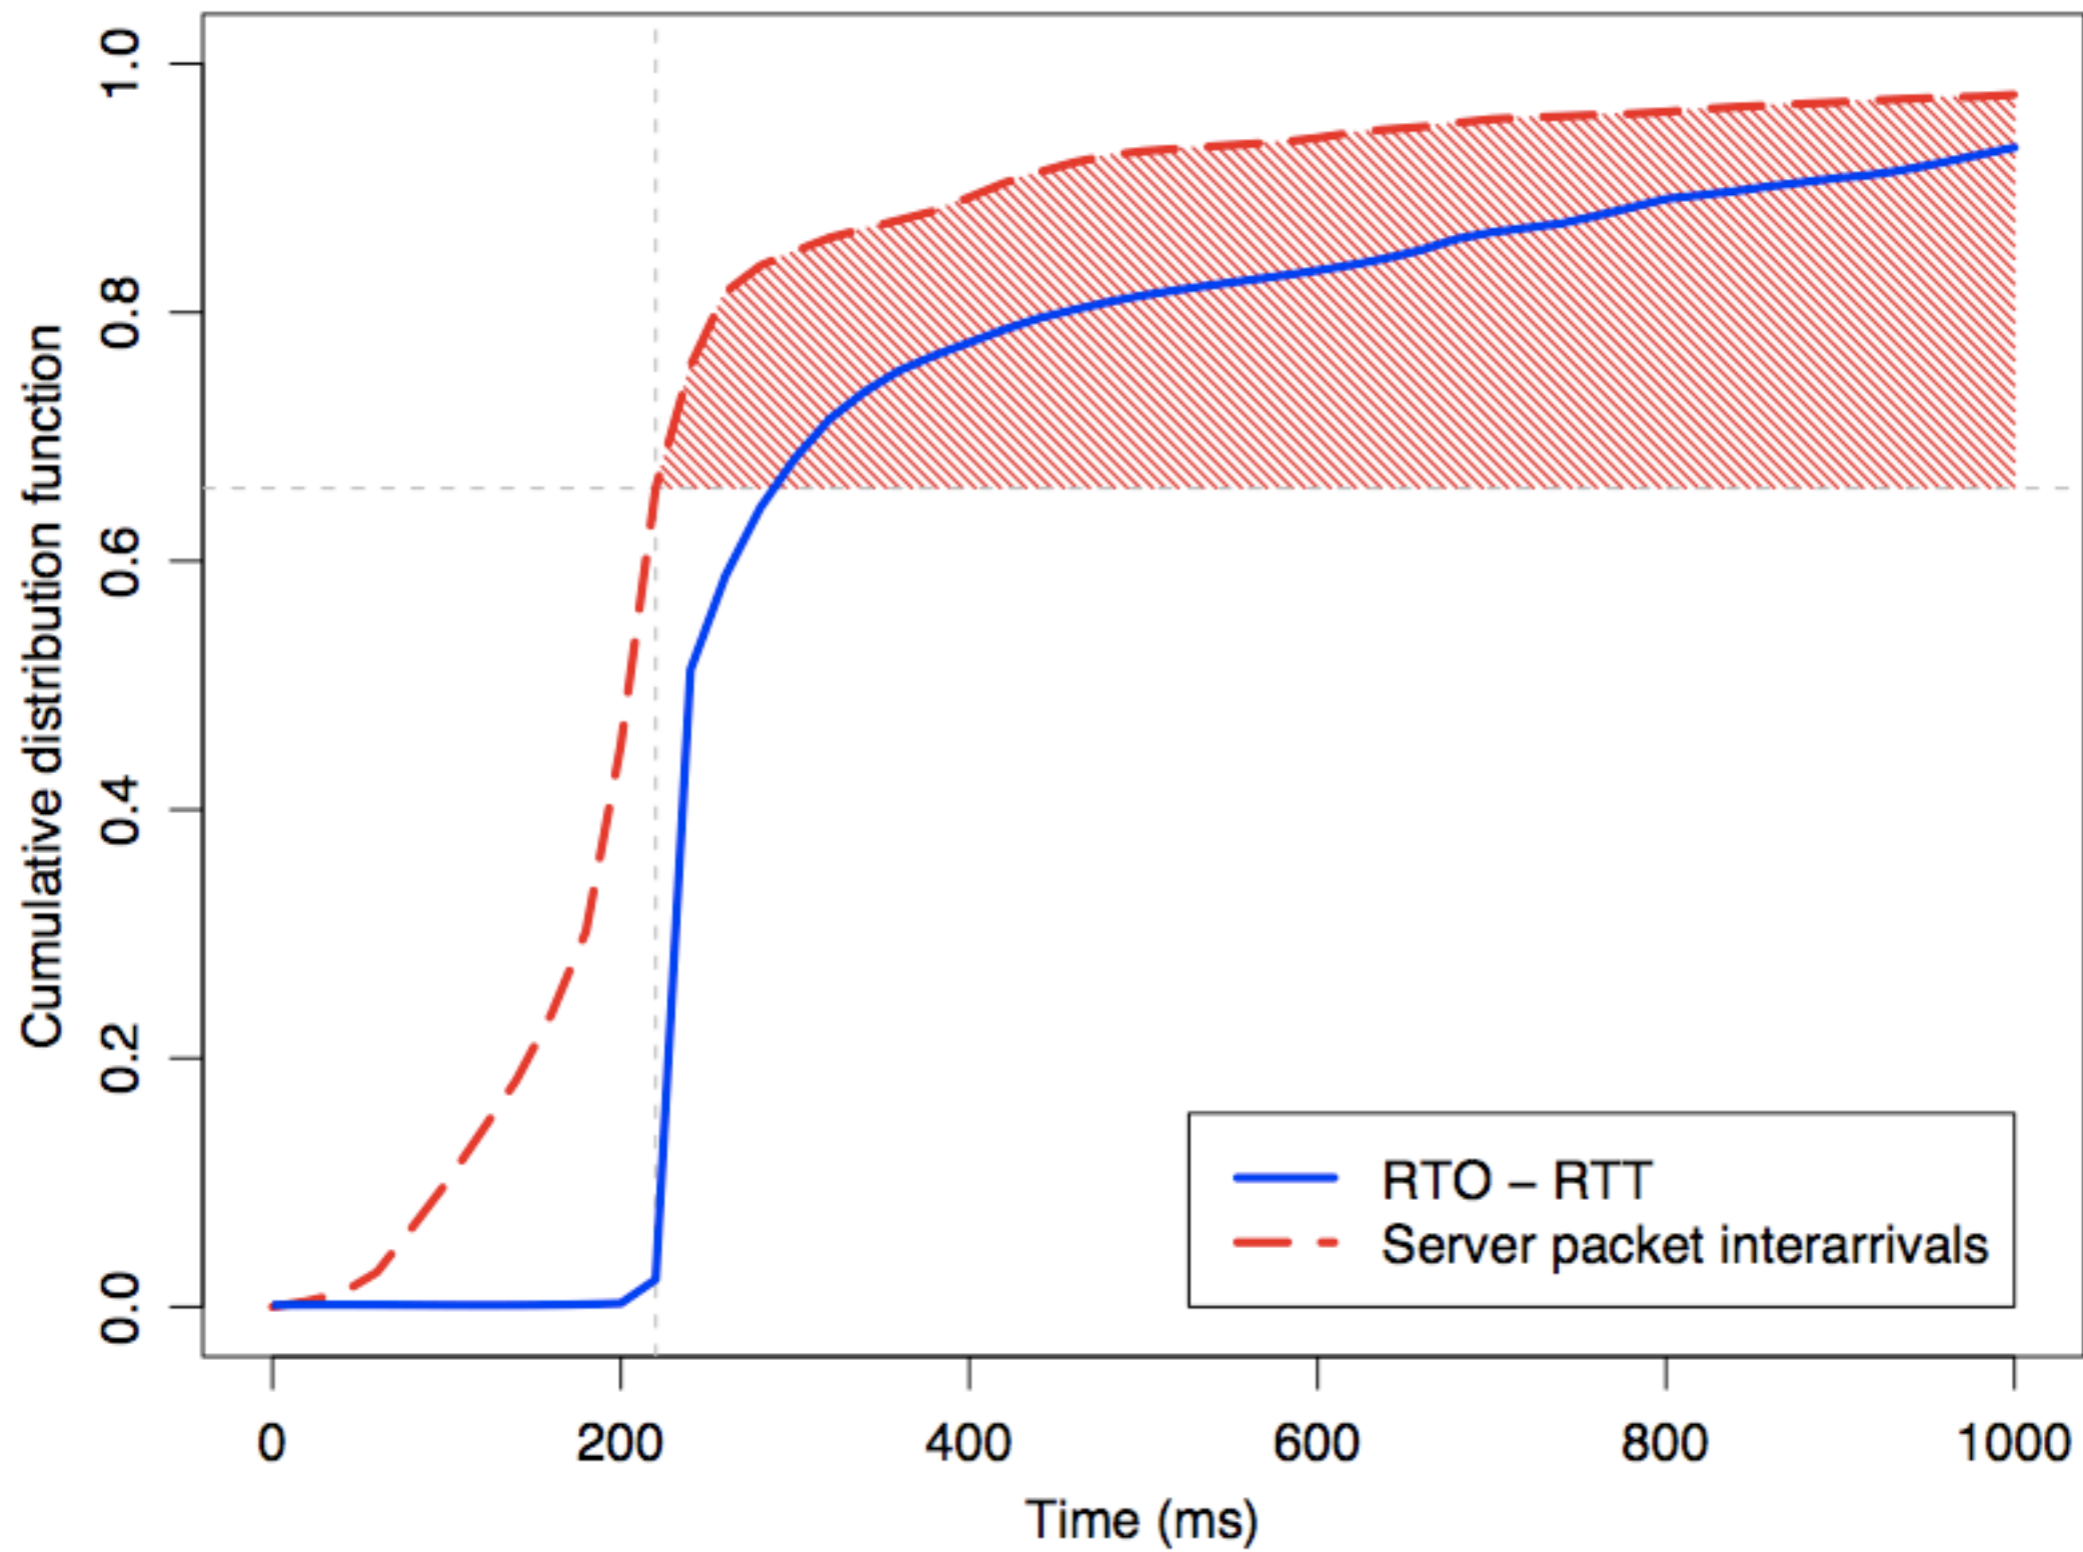





**“Thin Streams”**

**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

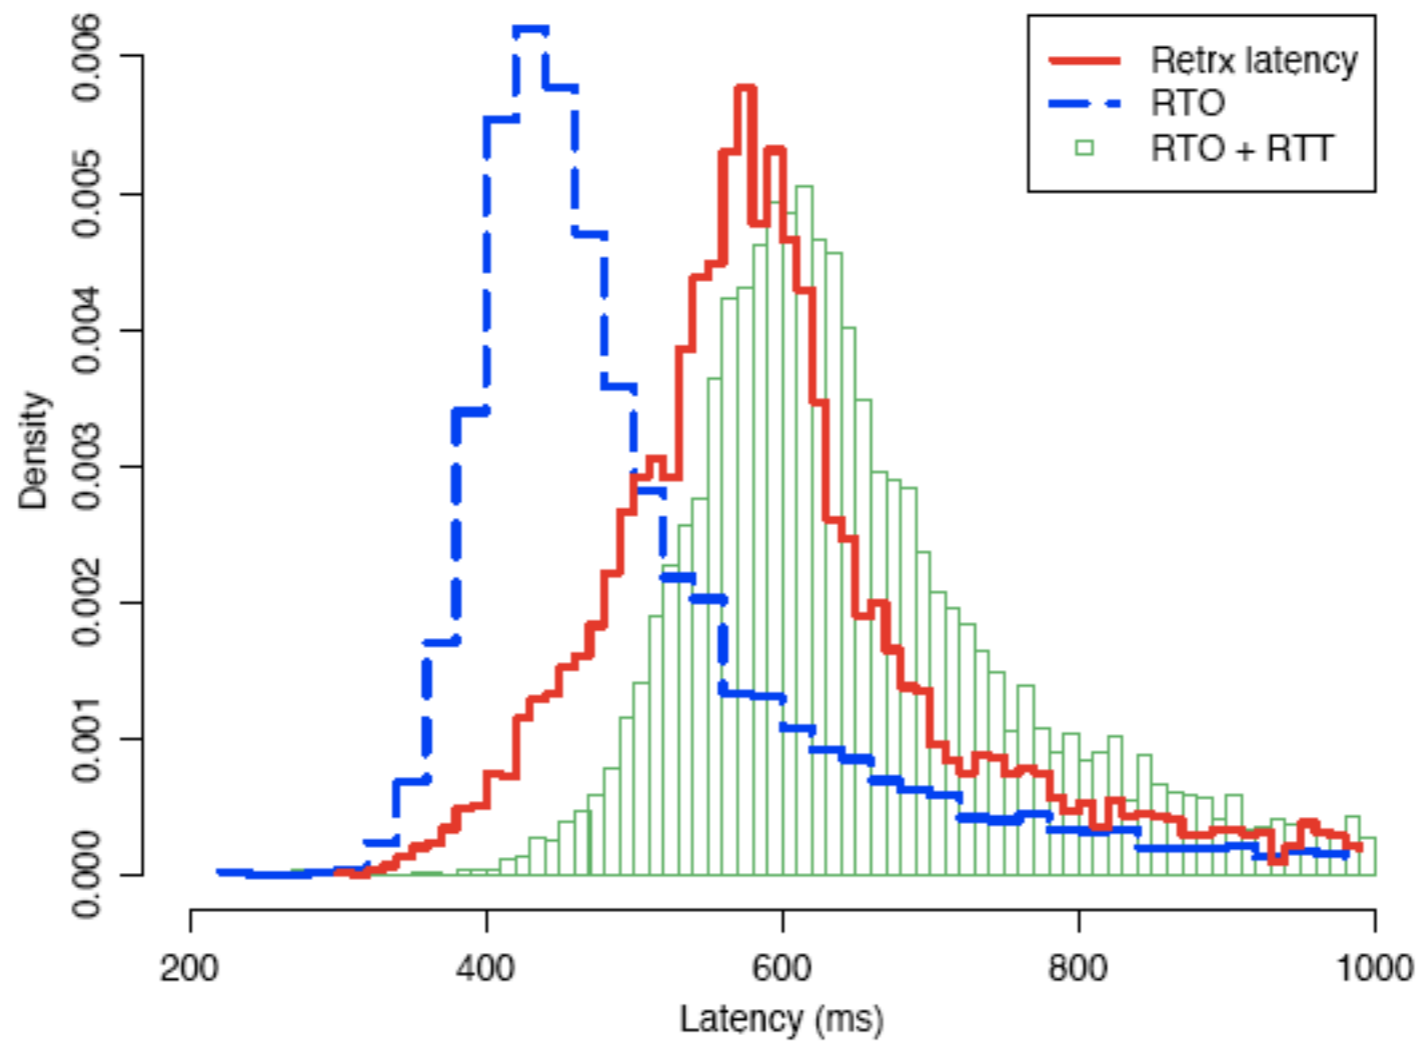



Figure 9: Average latency of dropped packets

**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

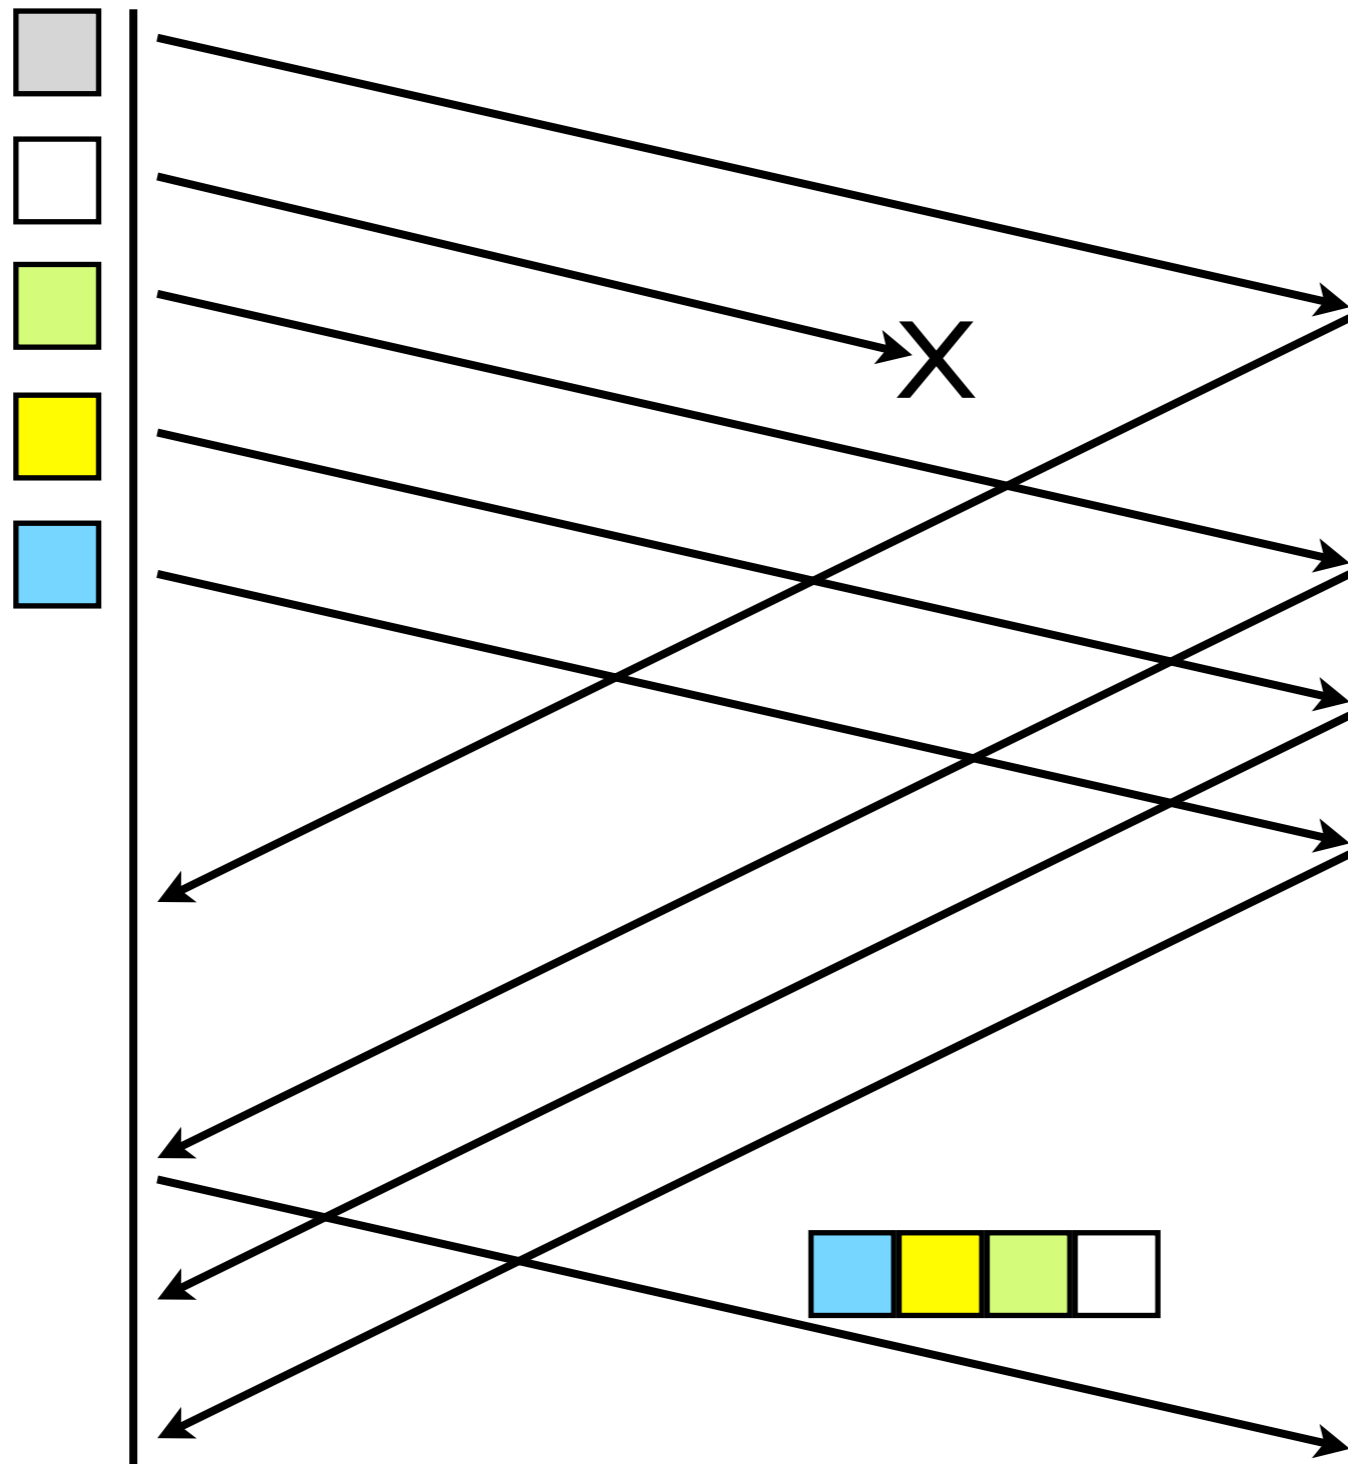
**spurious retransmission  
is not disastrous**

# 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

- not natively available in Windows
- available in many UNIXes