

# TCP Gearup II

---

# Overview

- How to think about send/recv
- About buffers
- How to debug/test in wireshark
- Implementation notes
- Any questions you have

# Roadmap

## Milestone II

- Basic sending and receiving using your sliding window/send receive buffers
- Plan for the remaining features

# Key resources

- Lecture 15: **How sliding window works** (sending side)
- Lecture 16: **Sliding window receiver**, retransmissions, zero-window probing
- HW3: Do it sooner rather than later—it will help!
- Testing and tools stuff: “TCP getting started” and “Testing with Wireshark” in the docs

## VWrite ("s" command in REPL)

- Input: some normal socket, data you want to send
  - => You need to define your send/rcv buffer, what variables/state etc you need to represent them
- Load data into your send buffer
- Block if send buffer is full, otherwise return number of bytes send

## VRead ("r" command)

- Input: normal socket, buffer for received data
- Read from rcv buffer, write that data to whatever buffer was passed in
- If rcv buffer is empty, block
- Return: number of bytes read\*\*\*

## Your goals:

- Defining data structures (buffers, etc), variables for how you keep track of things in the buffer

# Sending and receiving: API

More info: "Socket API example" in docs

## VWrite ("s" command)

- Input: normal socket, data to send
- Loads data into send buffer
- Block if send buffer is full

## VWrite ("r" command)

- Input: normal socket, buffer for received data
- Read from recv buffer, write to app buffer
- Block if recv buffer empty
- Return: number of bytes read

Demo!

# Your buffers

- Should use a [circular buffer](#)
- You get to decide on mechanics
  - How to keep track of read/write pointers
  - How to translate between sequence numbers => buffer indices
- You MAY use a library, but you should decide if this is what you want
  - You can also try generating a *generic* buffer library, but make sure you test it...

For detailed info


=> **RFC9293 Sec 3.3**: what all the variables mean

=> **Lecture 15-16**: detailed breakdown of how to use buffers

*You should decide, as a team, how you want to proceed here!*


Did your team use an existing library to represent your circular buffers?

65 responses




- Yes, we used an existing library
- Yes, we used an existing library, but we modified its source code
- No, we built our own circular buffer
- We didn't use a circular buffer

# SENDING SIDE (SND)




# RECEIVING SIDE (RCV)



Want to see a better version of this? See the notes from lectures 15-16.

For more explanations, see RFC9293, Sec 3.3.

For more info on this part, we recommend doing HW3—it is designed to help here!



# What happens in the TCP stack?

Your TCP stack will have some threads—you decide what they do

When you get a new packet...

=> Look up 4-tuple in socket table => find socket struct

=> Socket struct => all your per-connection TCP state  
(buffers, sequence numbers, etc....)

What to do with each segment? **RFC9293 Sec 3.7.10 is your friend**

=> + our modifications in "TCP notes" docs

# How does all of this fit into your work from before????

After Milestone I, most of your logic will be part of how you represent "normal-type" sockets  
For any packet received/API call => map to a socket  
=> Based on that socket's state (buffers, state machine, etc), what should happen?

REPL

TCP STACK

API CALLS

SOCKET API (VCONNECT, VLISTEN, ...) (LIKE GO/C/ETC) SOCKET API

SOCKETS: TWO TYPES

"Normal" sockets  
- One per active TCP connection  
- Has TCB (buffers, TCP state, etc.)

Listen sockets  
- One per open listen port  
- Has no TCB (can't send/recv)

TCP LOGIC  
STATE MACHINE,  
SLIDING WINDOW..

Socket table  
Maps packets => sockets based on header info

PACKET EVENTS

DECIDE WHAT/WHEN TO SEND

NEW HANDLER (PROTO=6)

USE SEND FROM IP!  
SendIP(destAddr, protocol, bytes)

TCP STACK

IP

IP LAYER

"Normal" sockets  
- One per active TCP connection  
- Has TCB (buffers, TCP state, etc.)

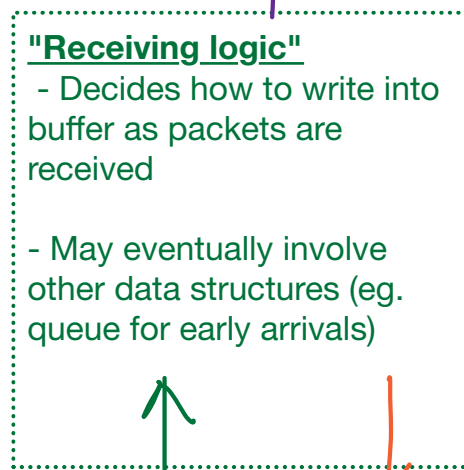
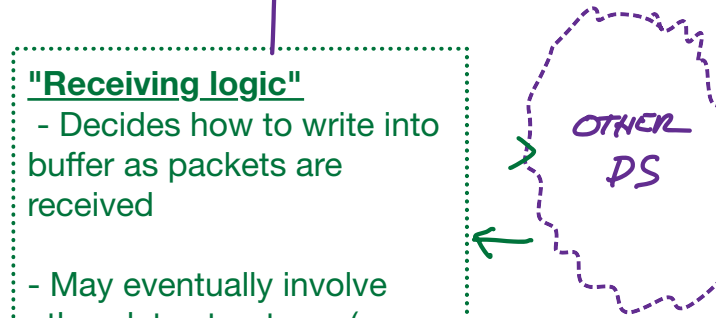
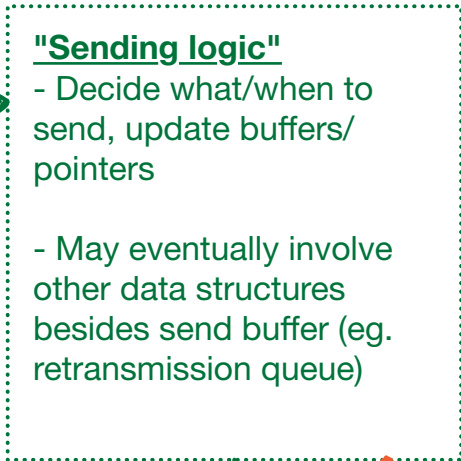

TCP LOGIC  
STATE MACHINE,  
SLIDING WINDOW...

# What happens inside a socket?

API

WRITE()


READ()



Sends segments via your IP stack (eg. SendIP)



Packet arrivals (when len(TCP payload) > 0)



# Implementing VRead/VWrite

Performance requirement: send/recv process **MUST** be event driven

- No `time.Sleep`
- No busy-waiting

Where does this apply?

- REPL: `s`, `r`, `sf`, `rf`
- VRead/VWrite
- Deciding when to send, or check for new data

=> Channels, condition variables, etc. are your friends


***Channels?***

***Condition variables?***

**=> See code demo in video (REALLY)**

Also "channels demo" in docs and resources

# How to test TCP



## Useful wireshark mechanics

- SEQ/ACK analysis
- Follow TCP stream
- Validating the checksum

Note: watching traffic in wireshark works differently in this project!  
=> See "TCP getting started" guide for details

# Reference implementation

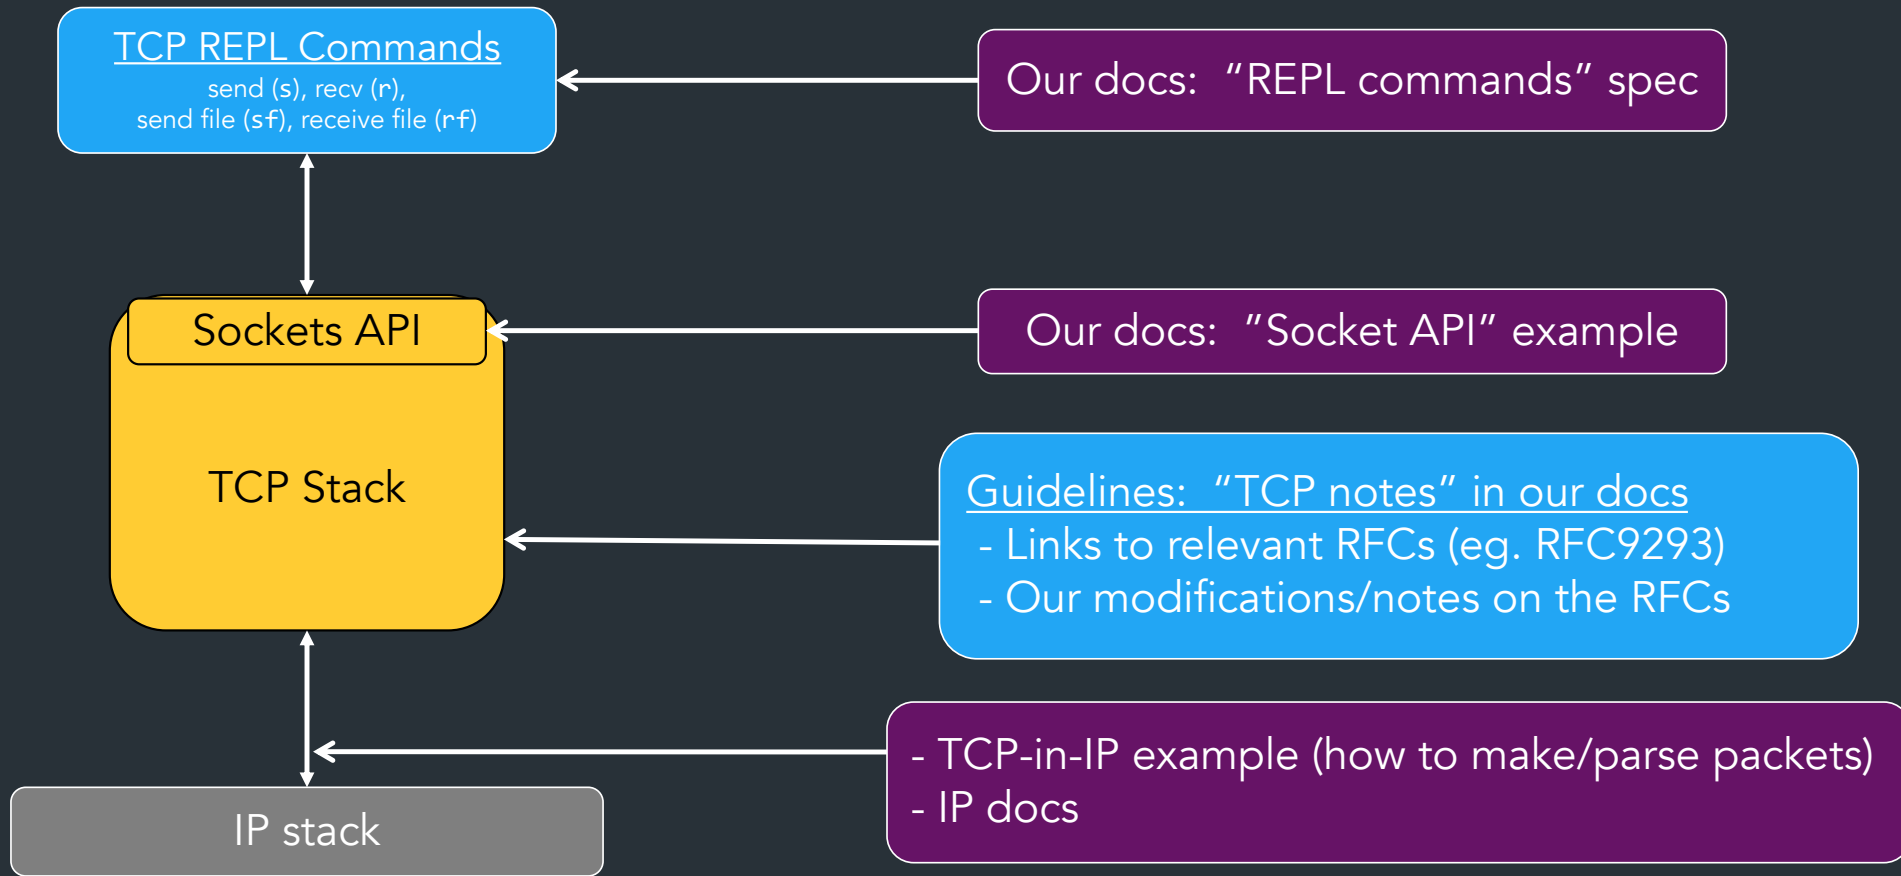
- Our implementation of TCP
- Try it and compare with your version!

Note: we're using a new reference this year!!

- We've tested as best we can, but there may be bugs
- See Ed FAQ, docs FAQ for list of known bugs
- Let us know if you have issues!

⇒ If the spec disagrees with the reference implementation,  
the spec wins—**don't propagate buggy behavior**  
(please help us find any discrepancies!)

# Where to get more info

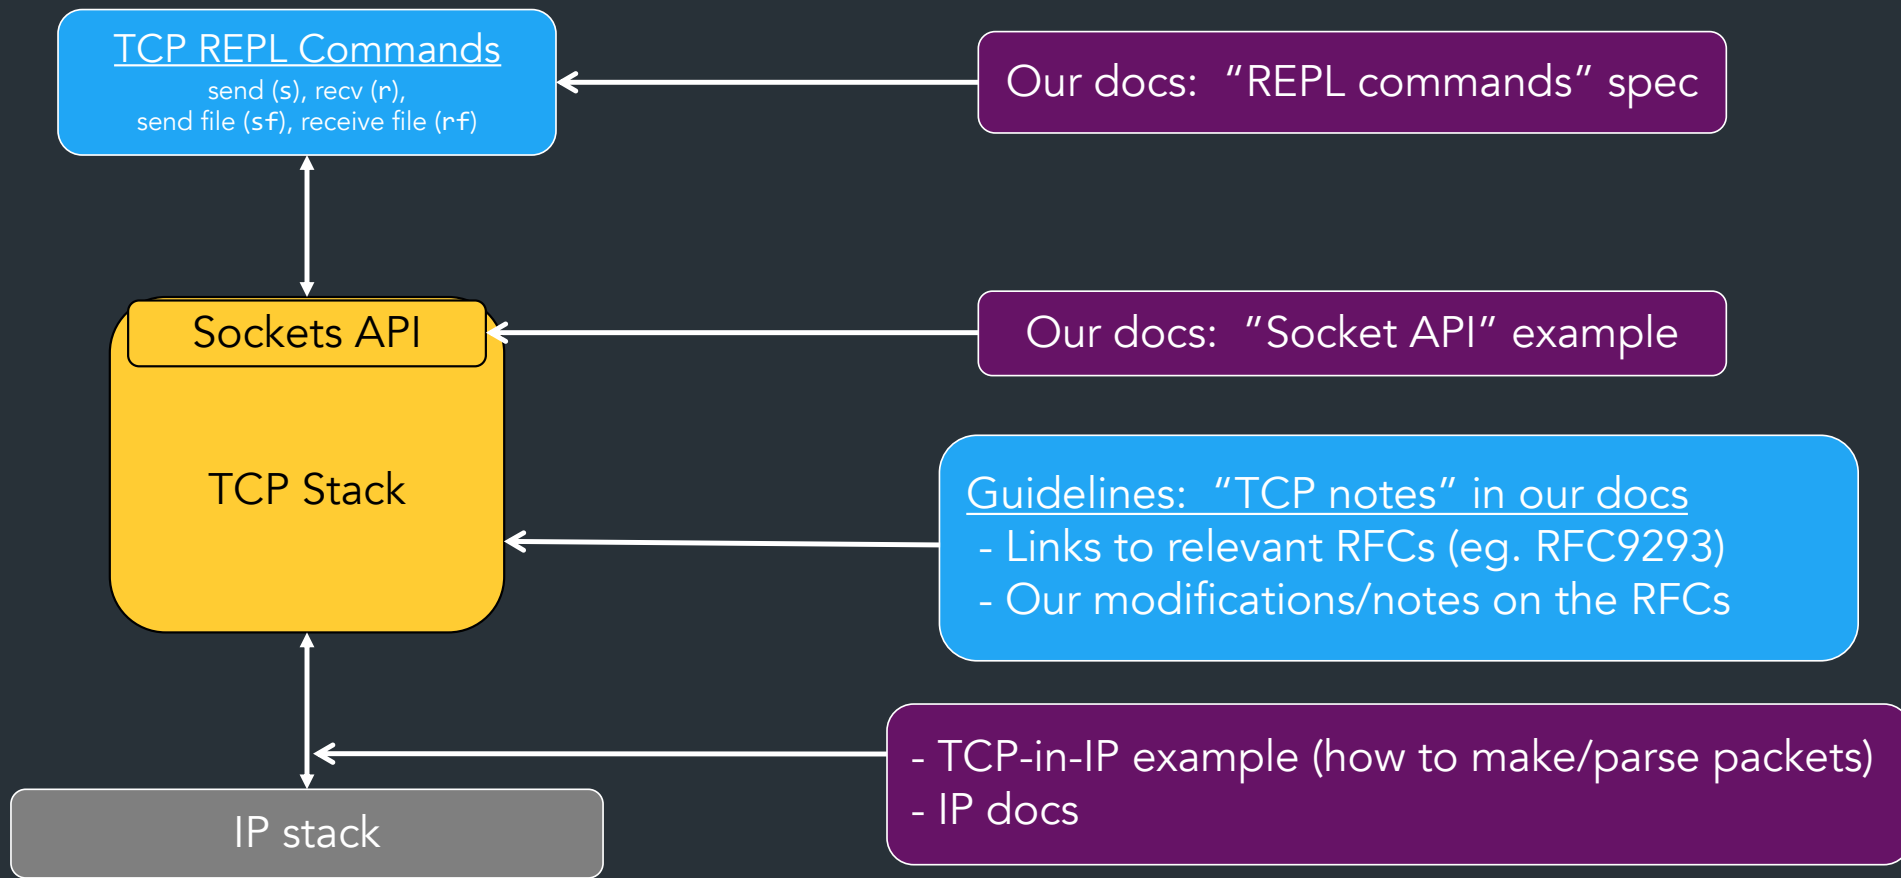


# Roadmap

## Final deadline

- Retransmissions (+ computing RTO from RTT)
- Zero-window probing
- Connection teardown
- Sending and receiving files (*sf*, *rf*)

# Where to get more info



# Closing thoughts

---

- Use your milestone time wisely!
- **Wireshark is the best way to test—use it!**
- Stuck? Don't know what's required? Just ask!  
(And see Ed FAQ)

We are here to help!