

CS 43: Computer Networks

Fall 2025, Week 2

Substitute Instructor: Kevin Webb



Slides adapted from Kurose & Ross, Kevin Webb, Vasanta Chaganti

Five-Layer Internet Model

Application: the application (e.g., the Web, Email)

Transport: end-to-end connections, reliability

Network: routing

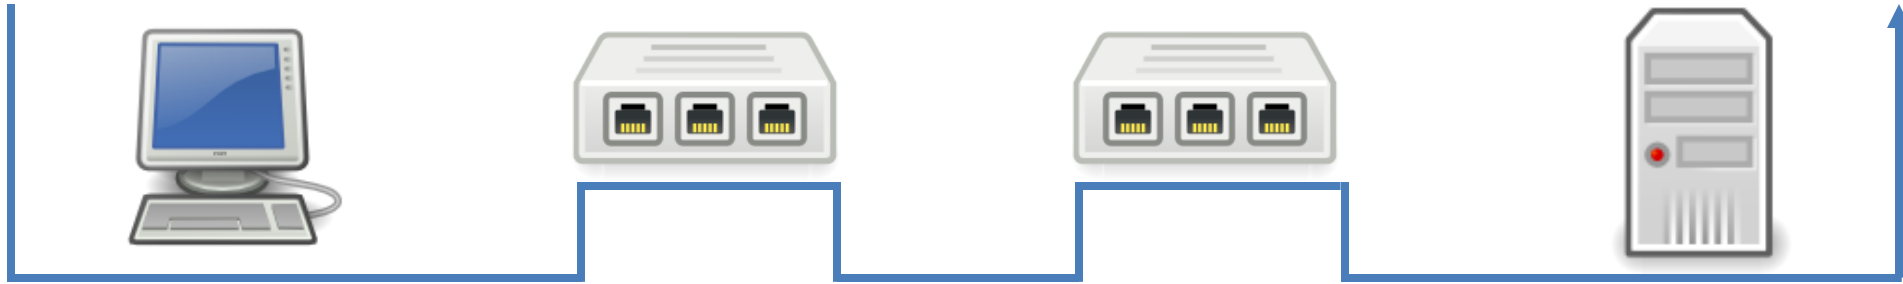
Link (data-link): framing, error detection

Physical: 1's and 0's/bits across a medium (copper, the air, fiber)

Networks have many concerns, such as reliability, error checking, naming and data ordering. Who/what should be responsible for addressing them? (Why? Which ones belong in which location?)

- A. The network should take care of these for us.
- B. The communicating hosts should handle these.
- C. Some other entity should solve these problems.

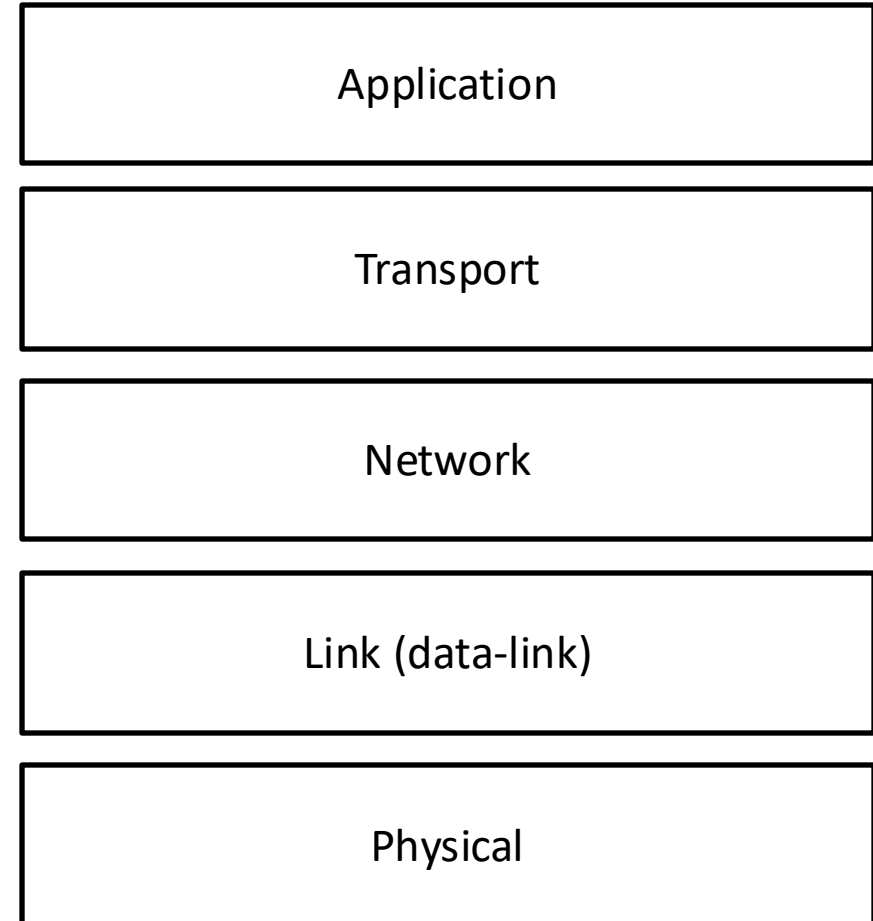
The “End-to-End” Argument



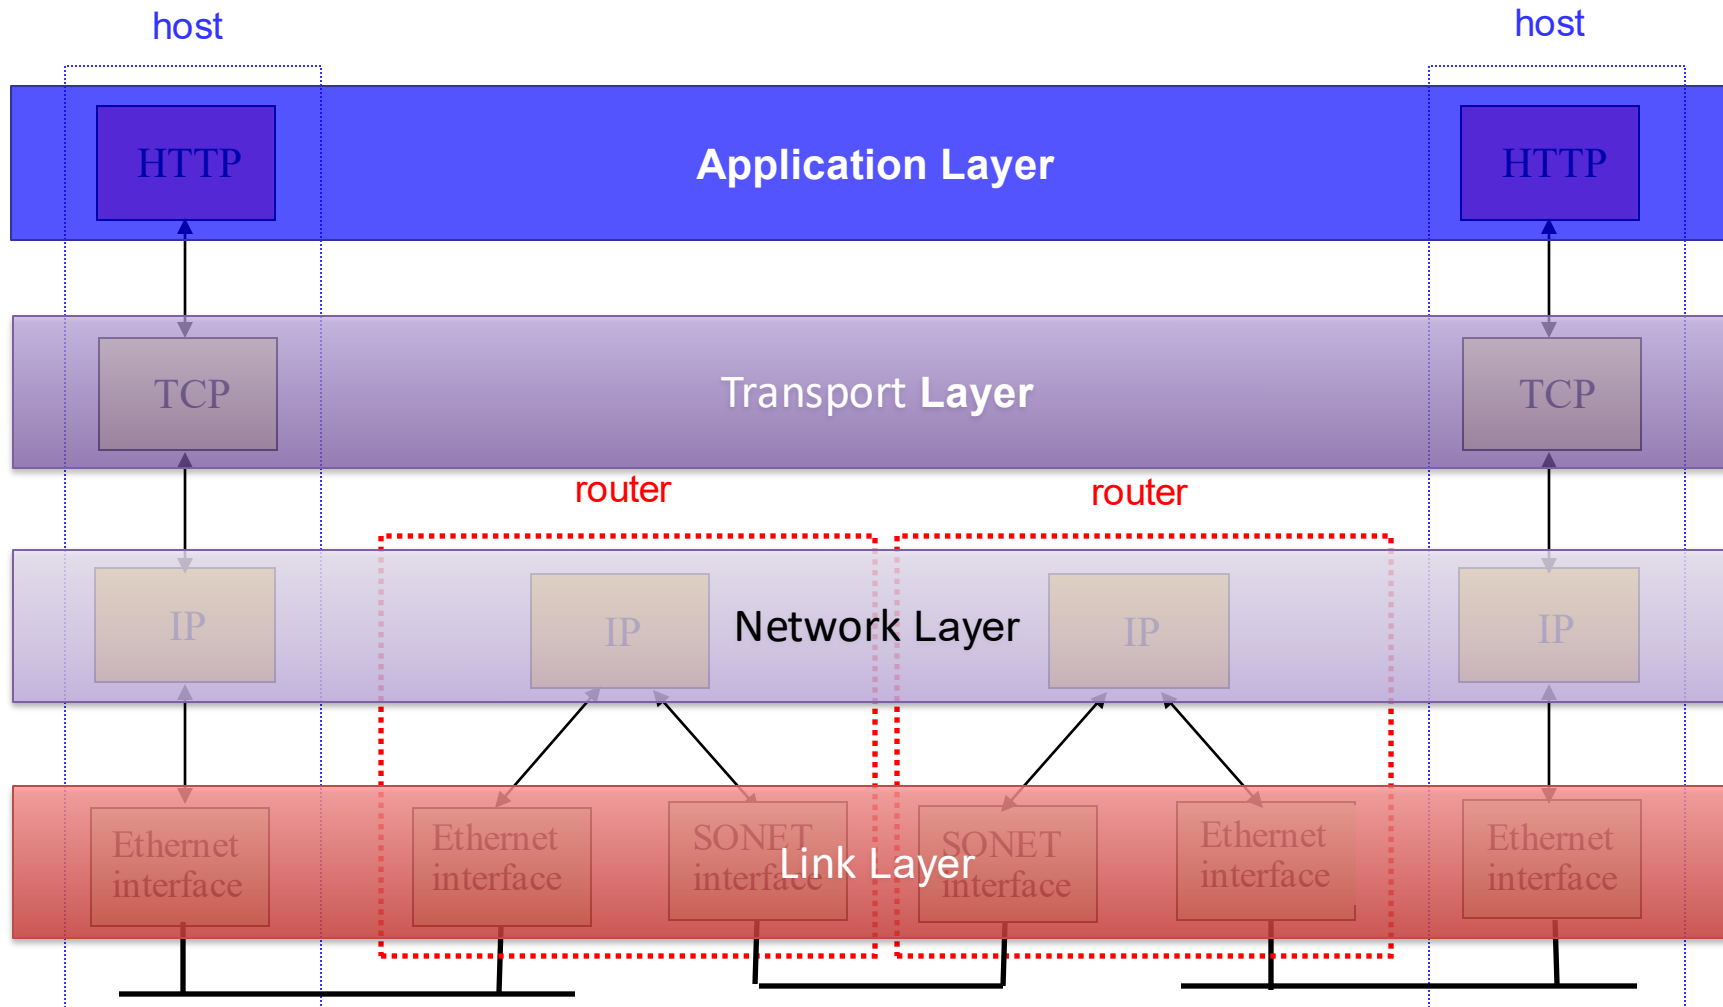
- Don't provide a function at lower level of abstraction (layer) if you have to do it at higher layer anyway - *unless there is a very good performance reason to do so.*
- Examples: error control, quality of service
- Reference: Saltzer, Reed, Clark, “End-To-End Arguments in System Design,” ACM Transactions on Computer Systems, Vol. 2 (4), 1984.

Which layers should routers participate in? (Getting data from host to host.) Why?

- A. All of Them
- B. Transport through Physical
- C. Network, Link and Physical
- D. Link and Physical

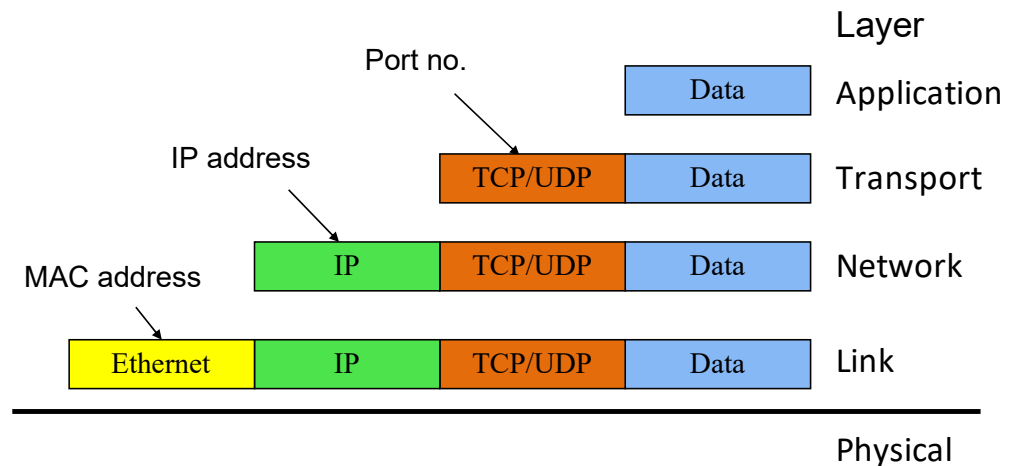


TCP/IP Protocol Stack



Application Layer (HTTP, FTP, SMTP, Zoom)

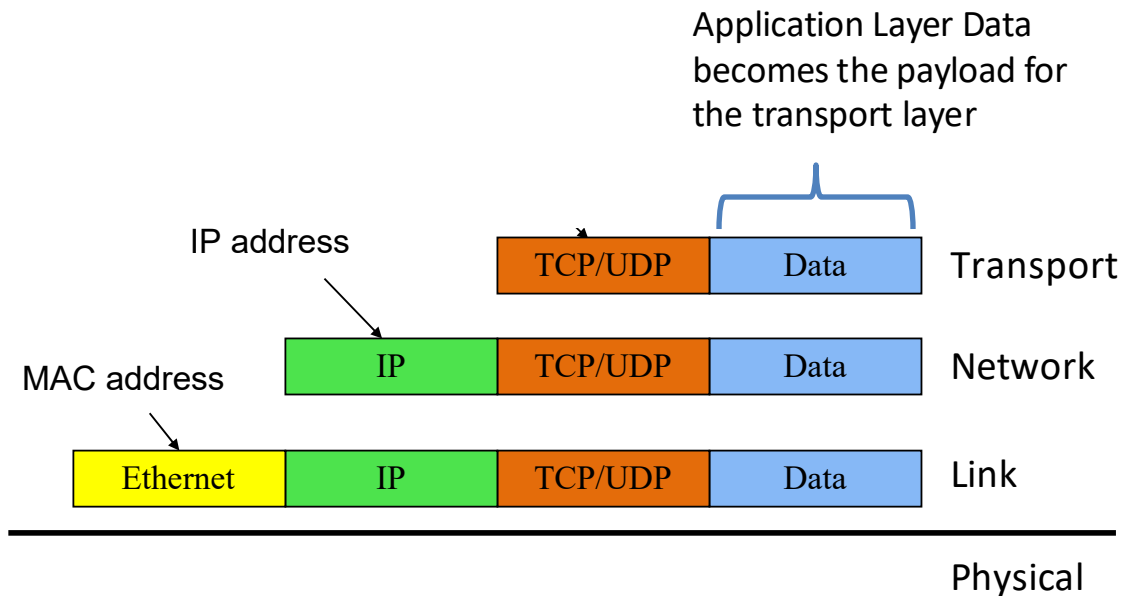
- Does whatever an application does!



Transport Layer (TCP, UDP)

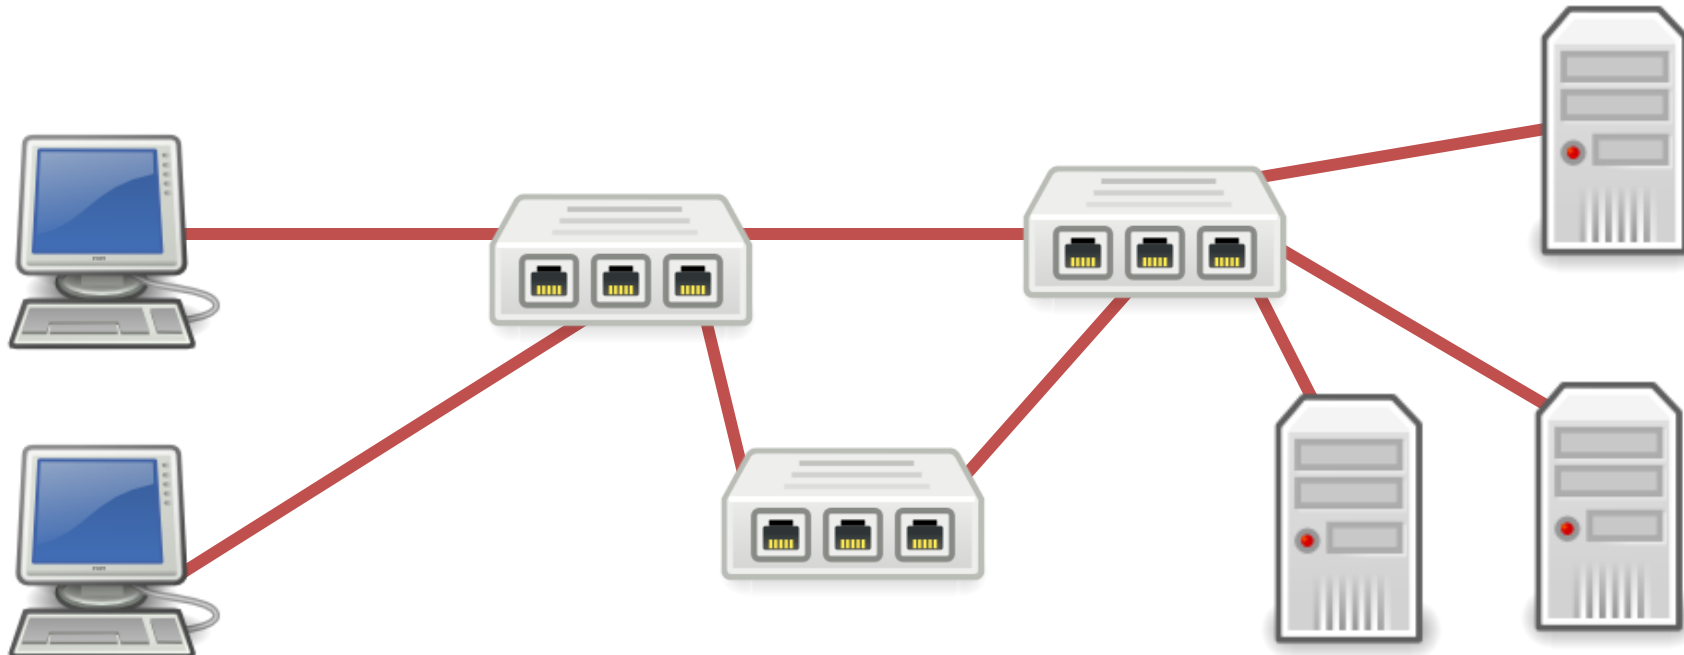
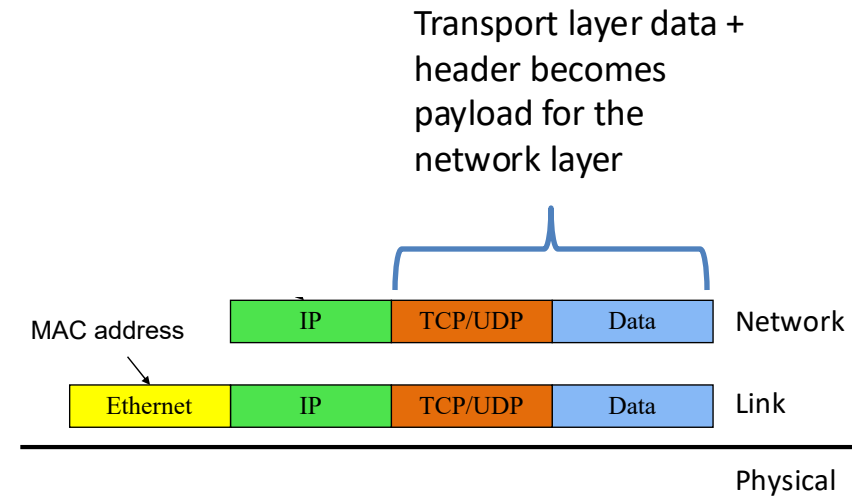
- Provides
 - Ordering
 - Error checking
 - Delivery guarantee
 - Congestion control
 - Flow control

- Or doesn't!



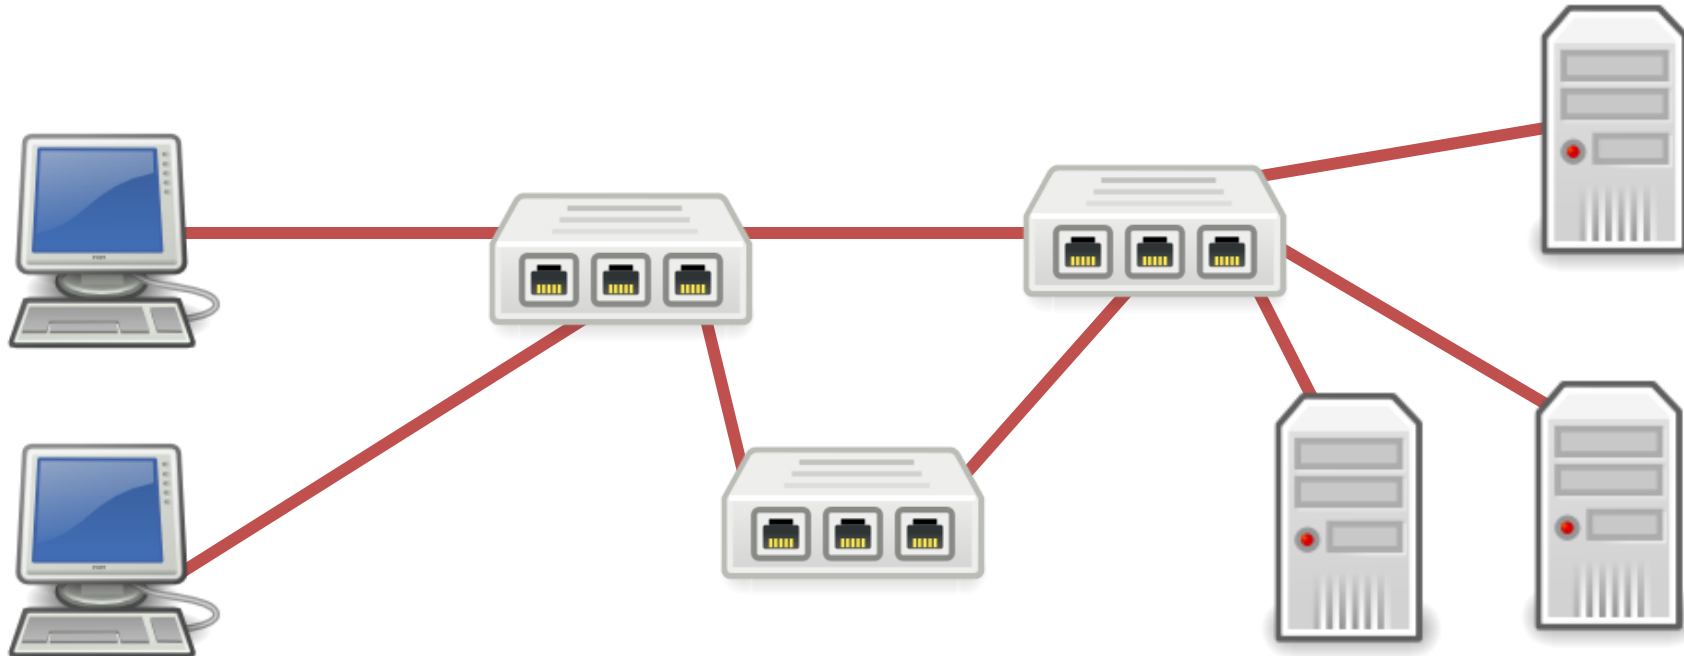
Network Layer (IP)

- **Routers:** choose paths through network



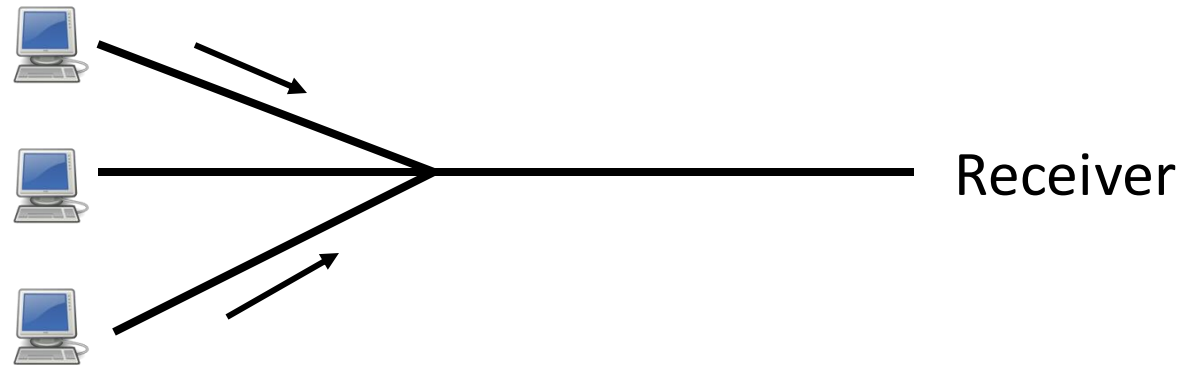
Network Layer (IP)

- **Routers:** chooses paths through network
 - *Circuit switching:* guaranteed channel for a session (Telephone system)
 - *Packet switching:* statistical multiplexing of independent pieces of data (Internet)

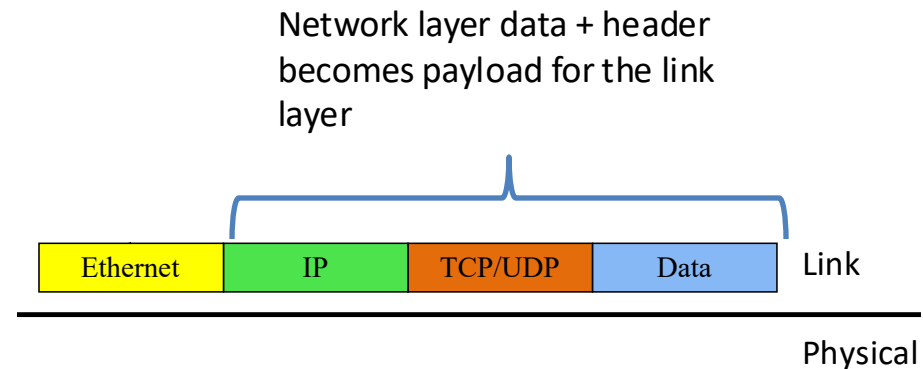


(Data) Link Layer (Ethernet, WiFi, DOCSIS)

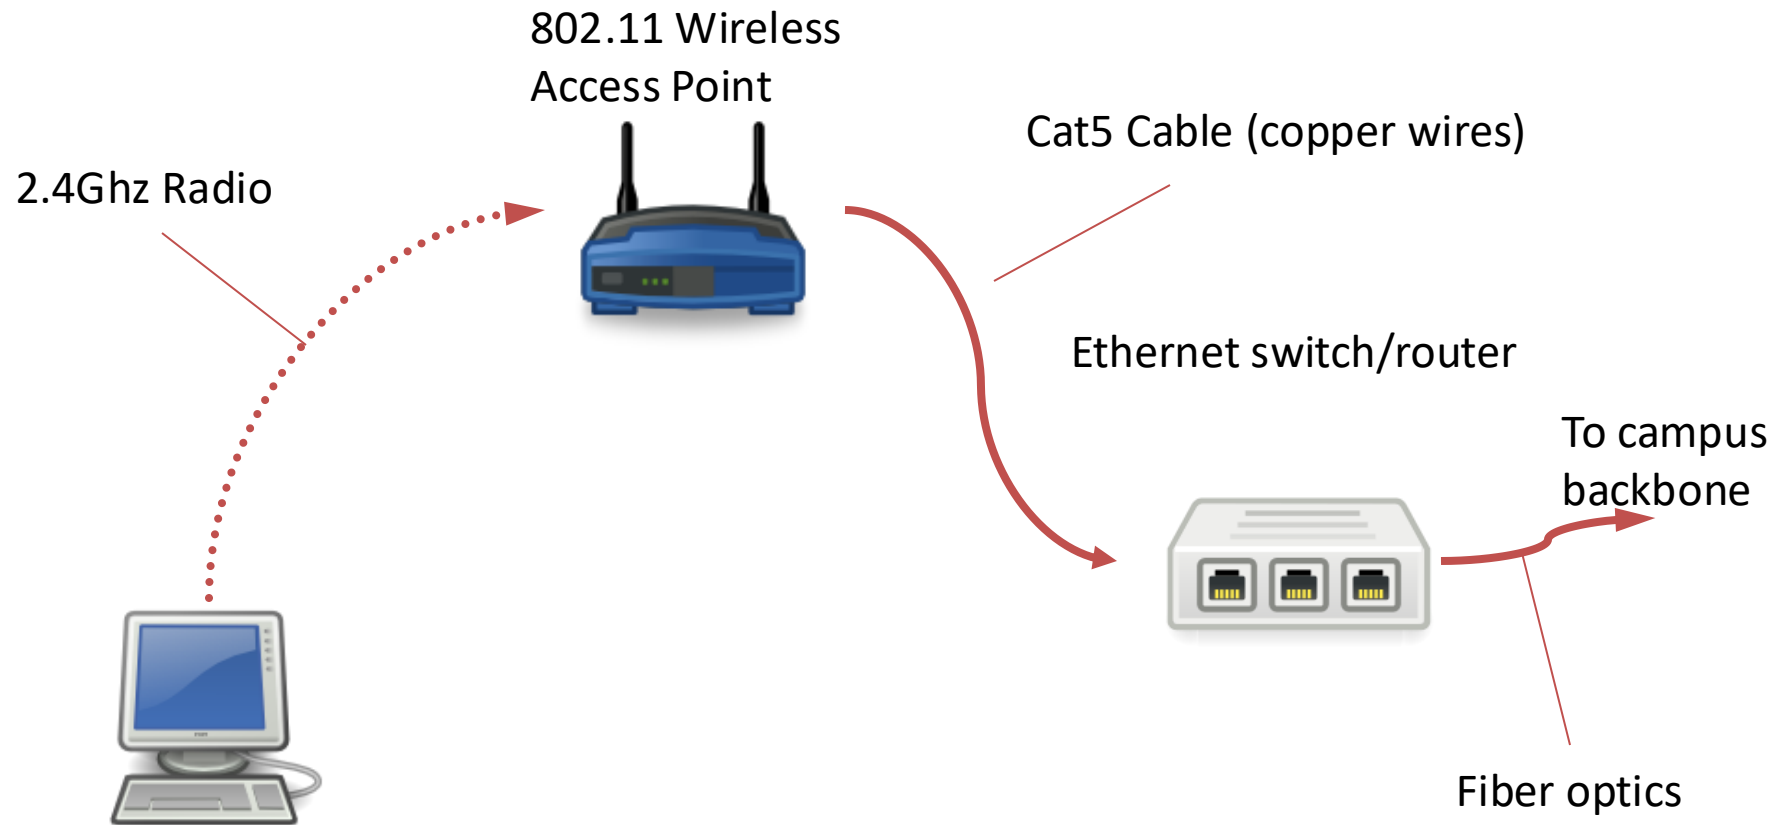
- Break message into chunks (frames) to send over physical medium
- Media access: can it send the frame now?



- Send frame, handle “collisions”



Physical layer (Copper, Coax, Air, Fiber Optics)



Because of our layering abstractions, we can use any technology we want, at any layer (as long as it doesn't interfere with the other layers).
(Why or why not?)

- A. Always
- B. Usually
- C. Sometimes
- D. Never

Internet Protocol Suite

HTTP

FTP

...

Zoom

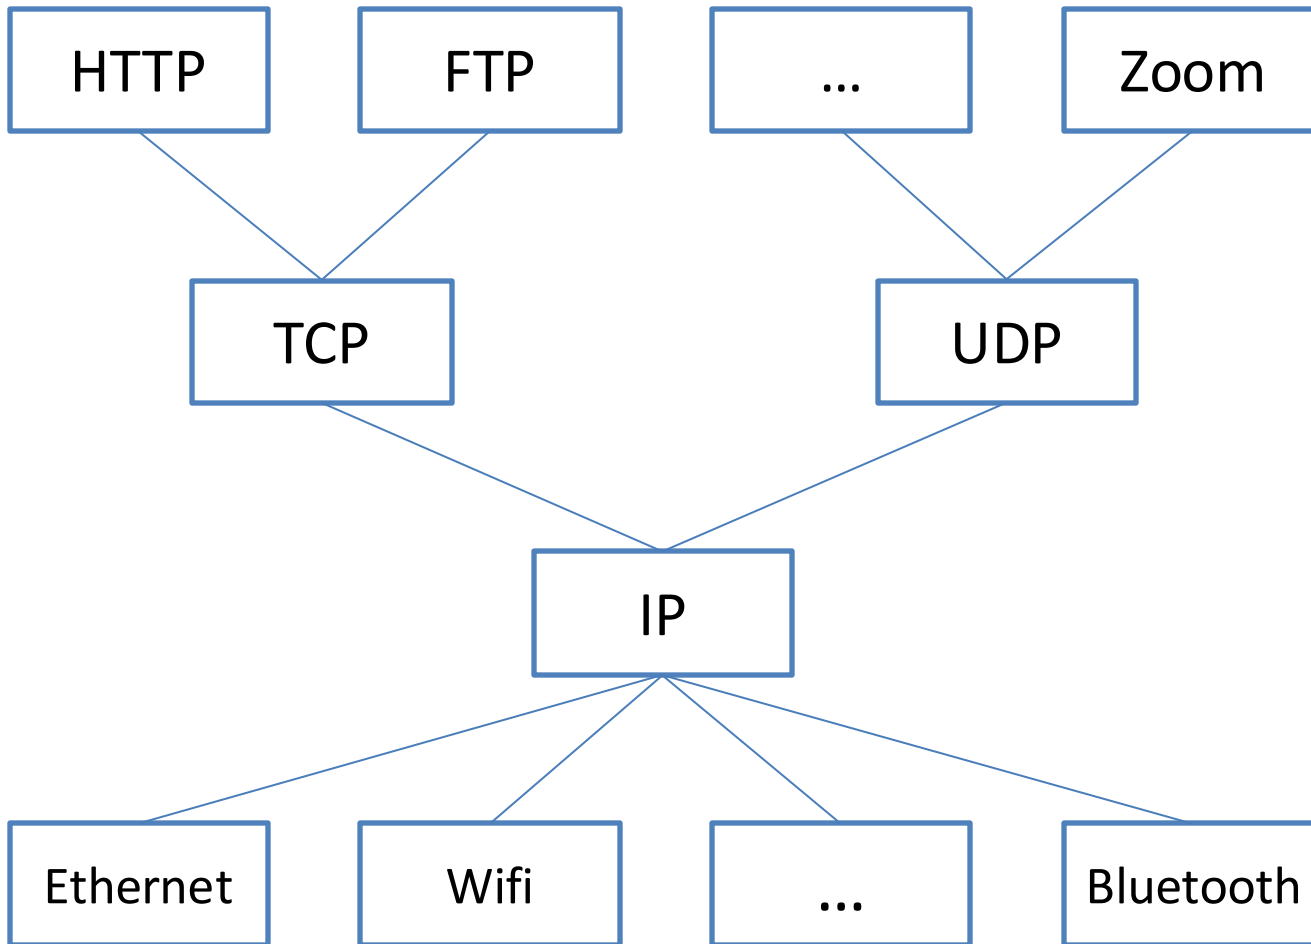
Ethernet

Wifi

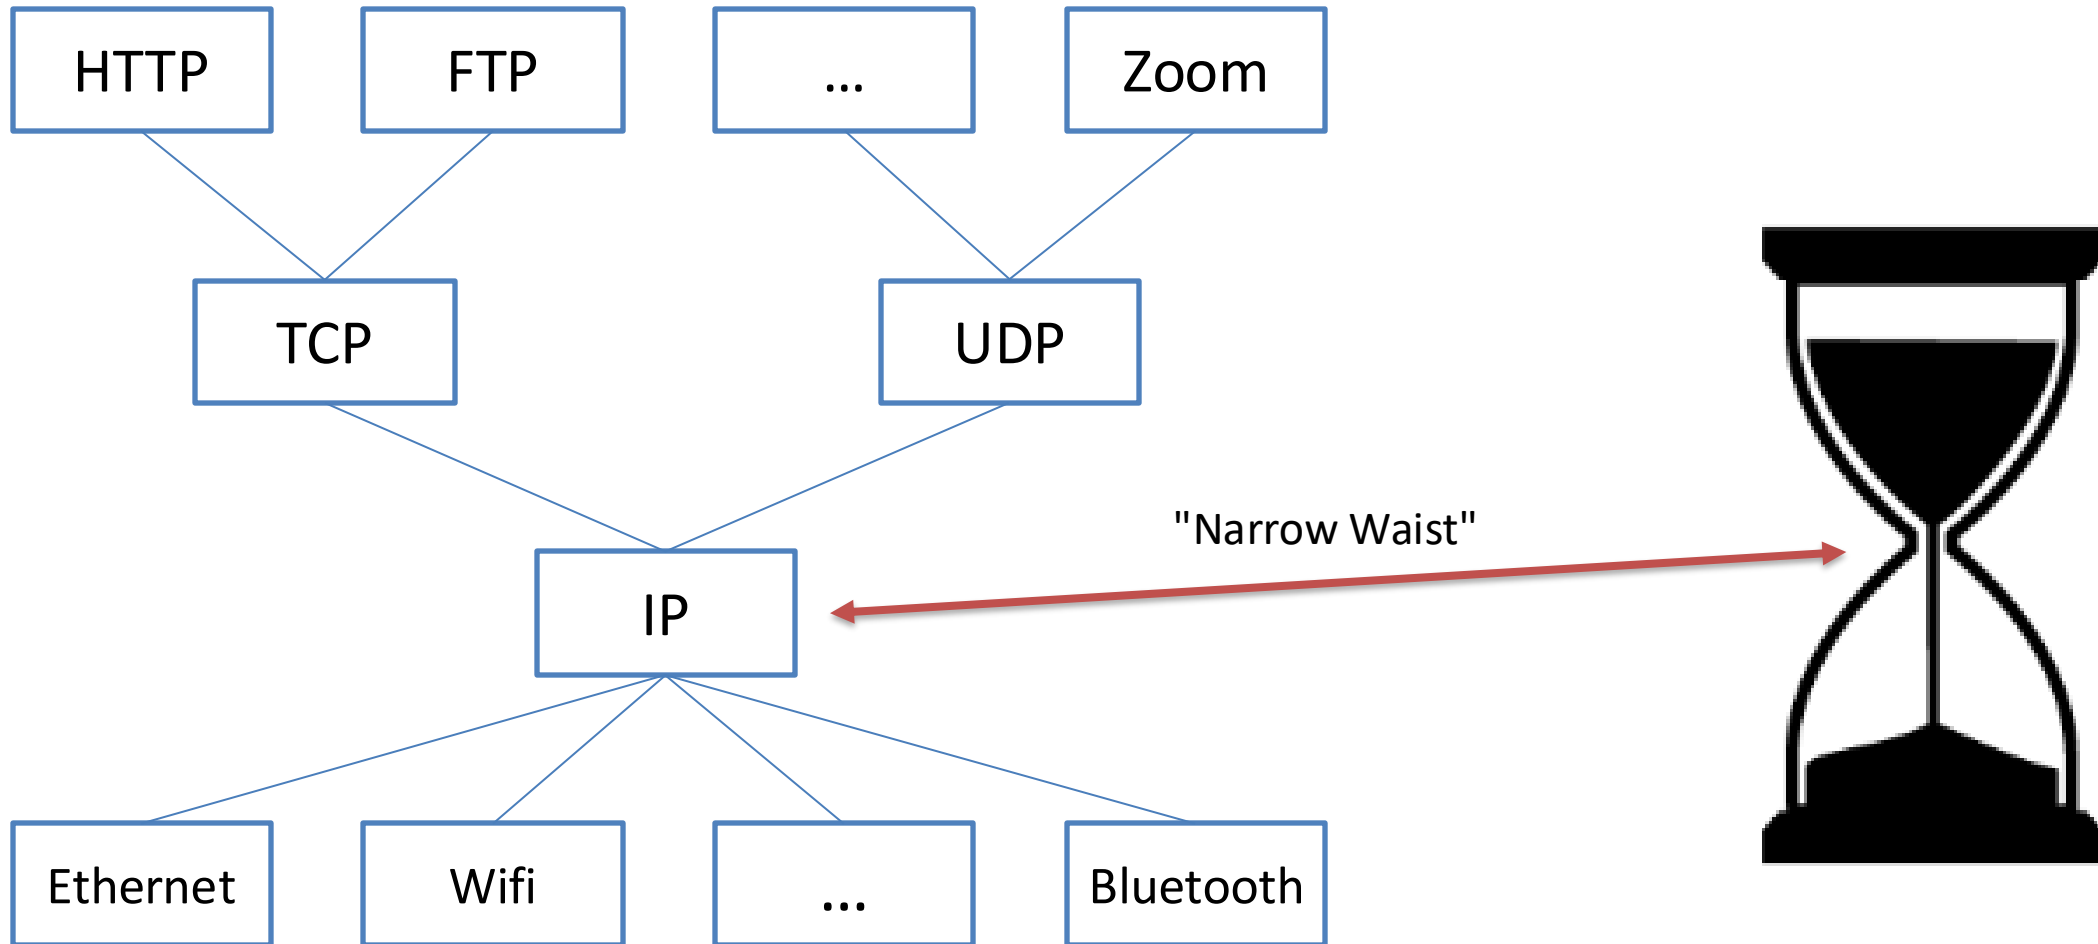
...

Bluetooth

Internet Protocol Suite



Internet Protocol Suite ("Hourglass model")



Putting this all together

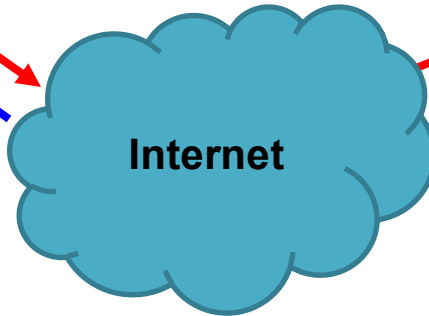
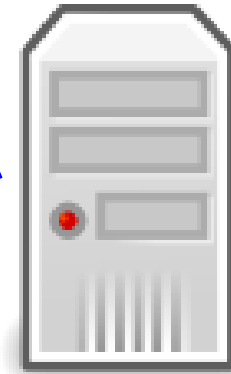
- **ROUGHLY**, what happens when I click on a Web page from Swarthmore?

My computer



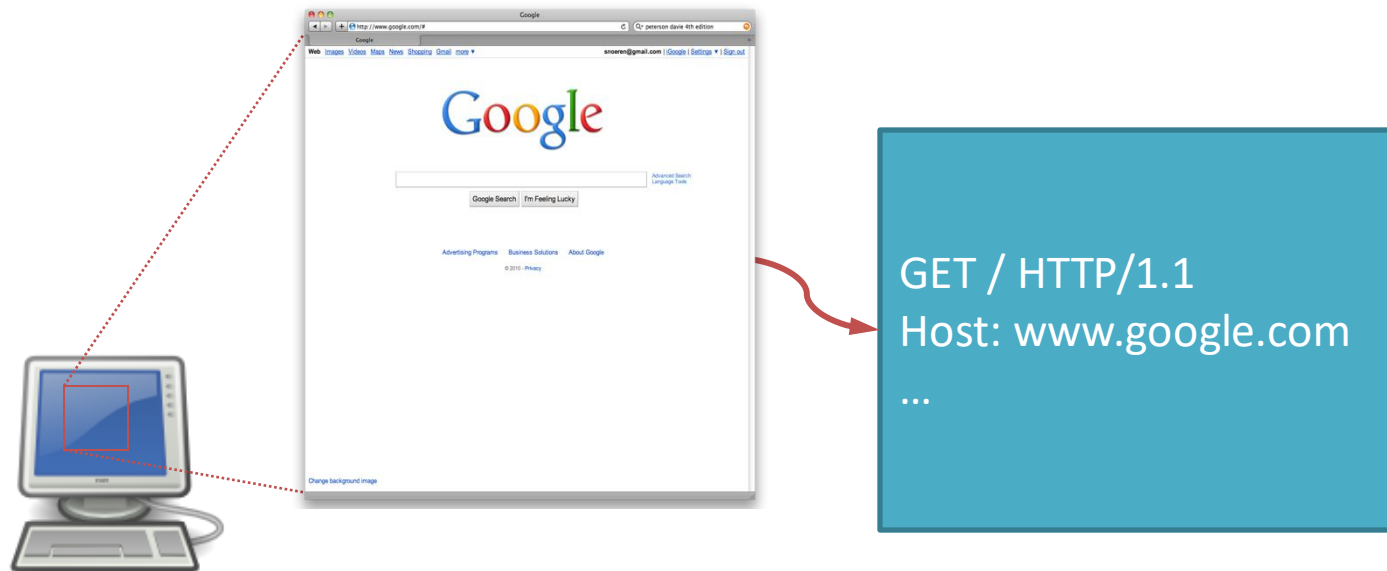
?

www.google.com



Web request (HTTP)

- Turn click into HTTP request



Name resolution (DNS)

- Where is `www.google.com`?

My computer
(130.58.68.164)



What's the address for `www.google.com`



Local DNS server
(130.58.68.10)

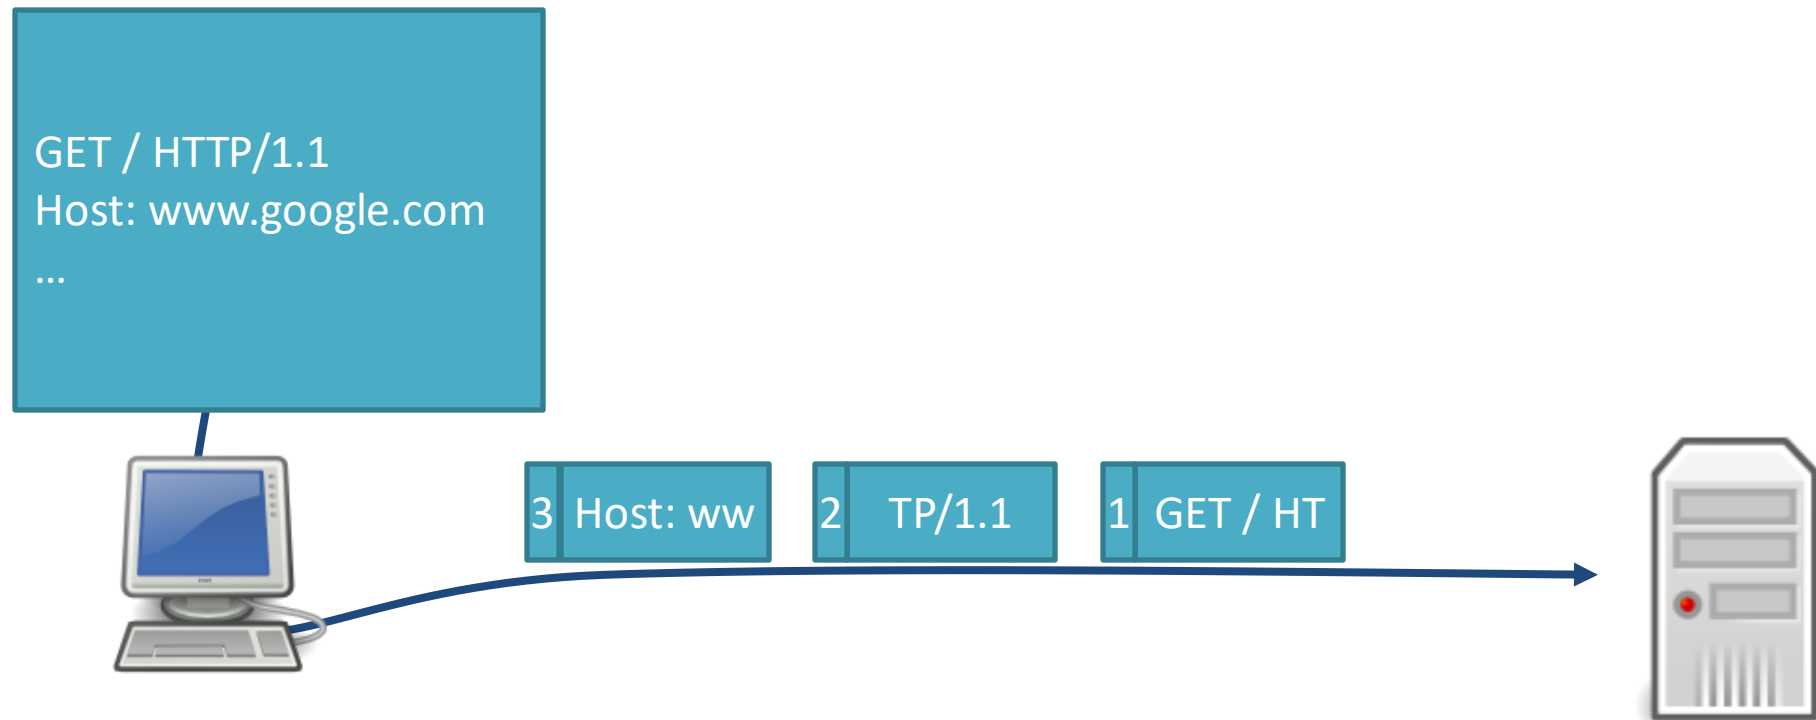


Oh, you can find it at 142.250.65.238



Transport (TCP)

- Break message into chunks (TCP segments)
- Should be delivered reliably & in-order



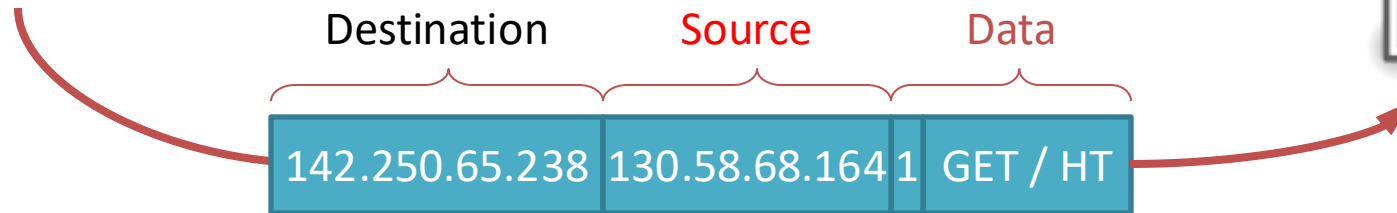
Global Network Addressing

- Add IP header, address each IP packet so it can traverse network and arrive at destination.

My computer
(130.58.68.164)

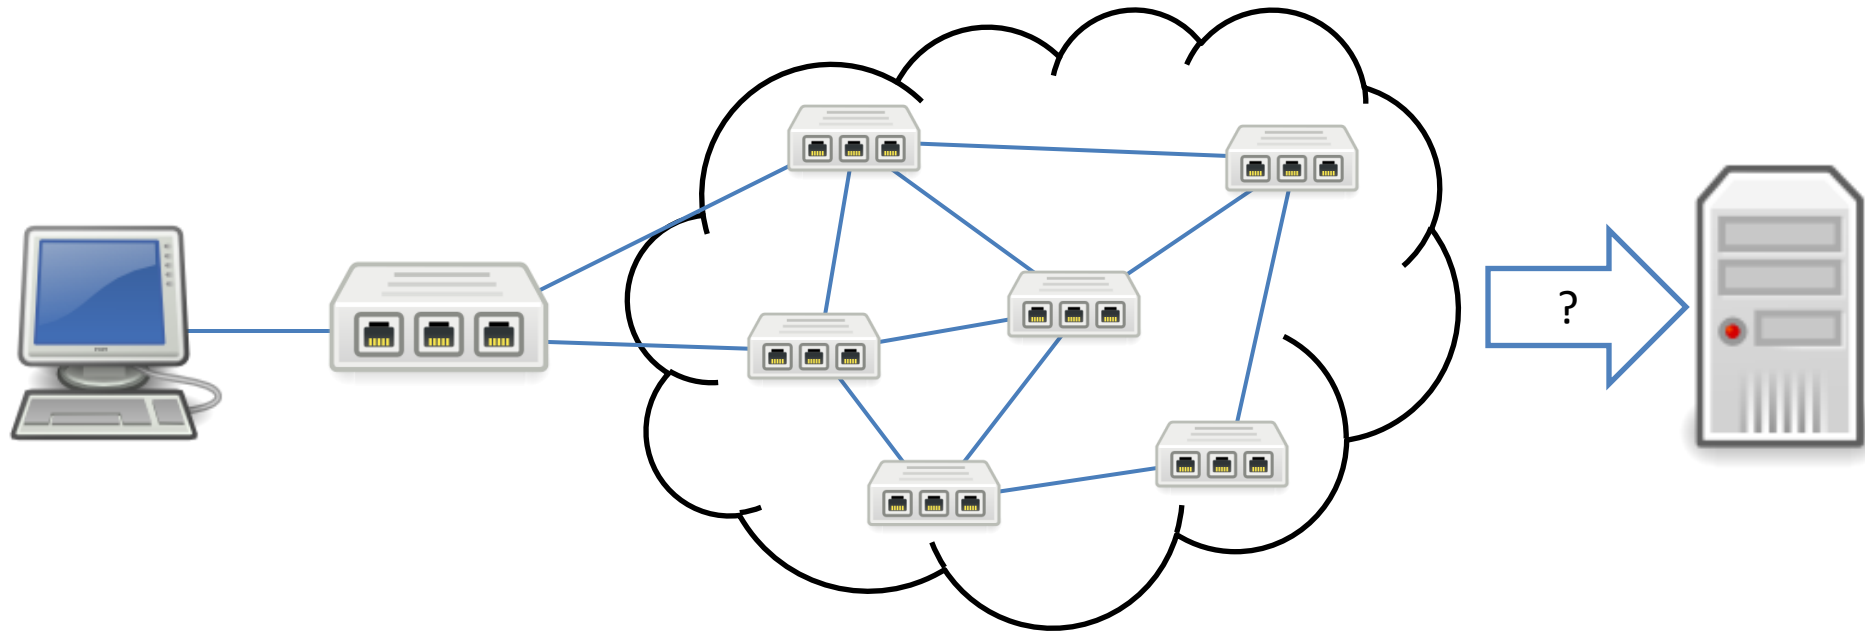


www.google.com
(142.250.65.238)



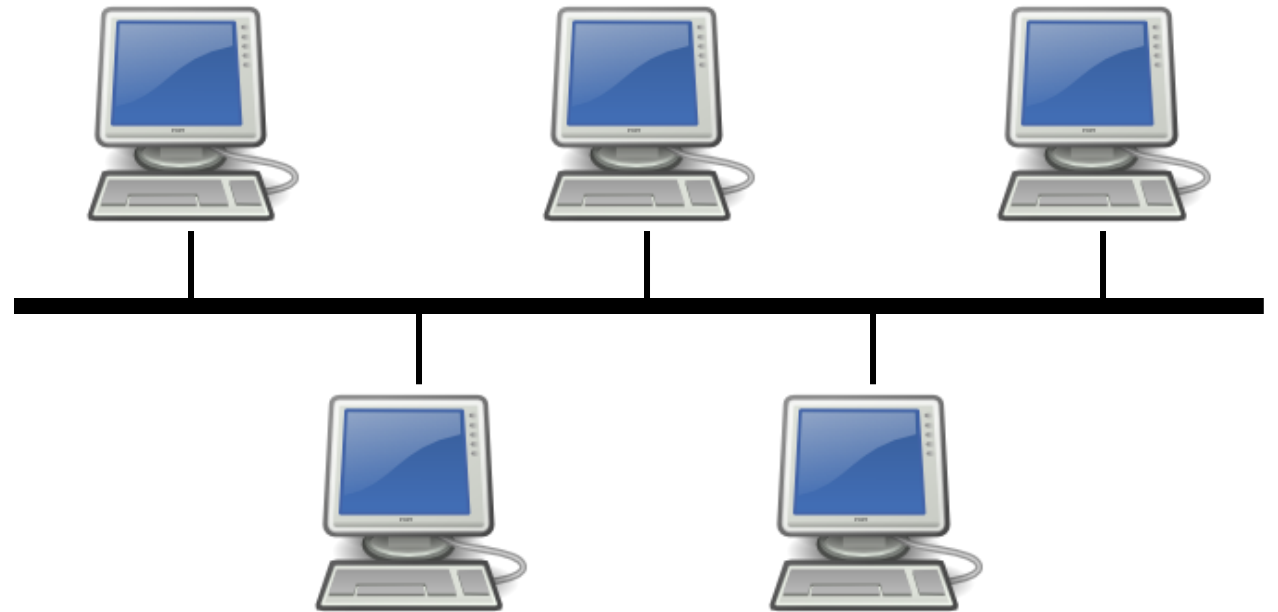
(IP) At Each Router

- Where do I send this to get it closer to Google?
- Which is the best route to take?



Link & Physical Layers

- Forward to the next node!
- Share the physical medium.
- Detect errors.



Summary

- Layers of abstraction divide up responsibility for network functionality
- End-to-end principle: do work at end hosts when possible
- Protocol governs message format and transfer procedure
- Messages encapsulated by protocol headers at each layer

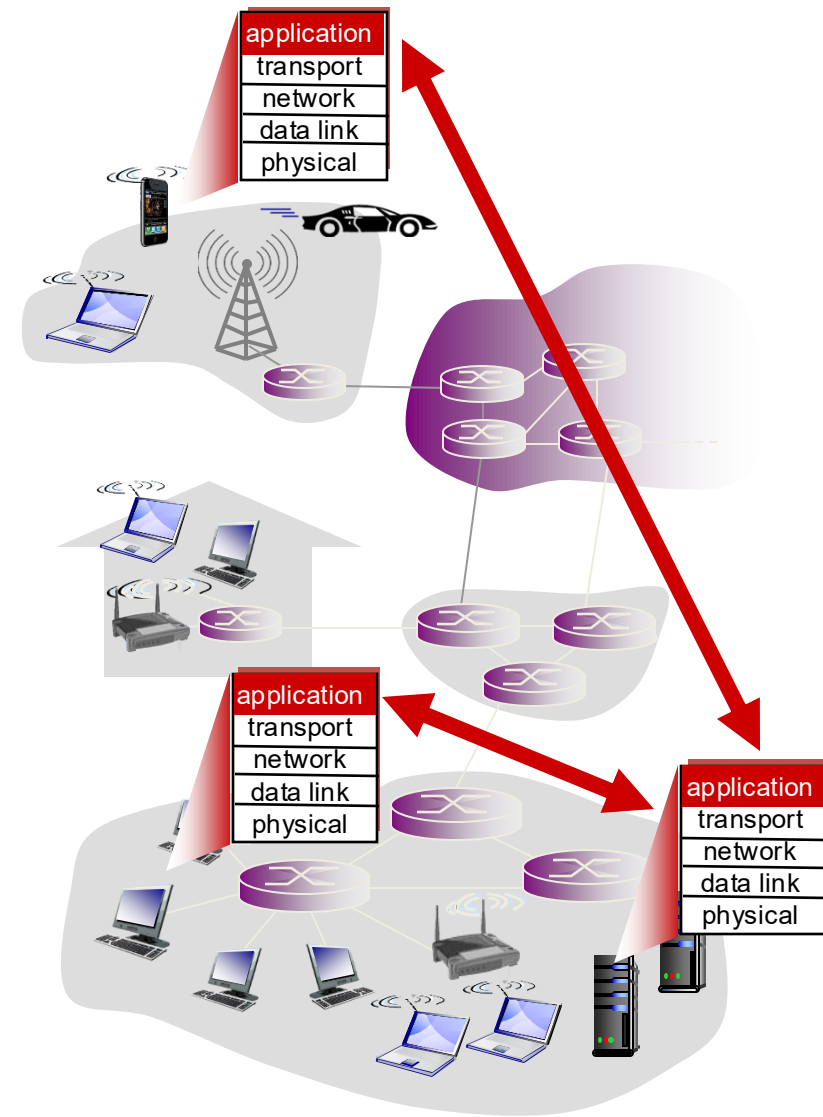
Up Next: Prep For Lab 1

- You need to know a bit about HTTP
- You need to know a bit about sockets
- After we get these lab prerequisites out of the way, we'll go into more depth about HTTP.

Creating a network app

write programs that:

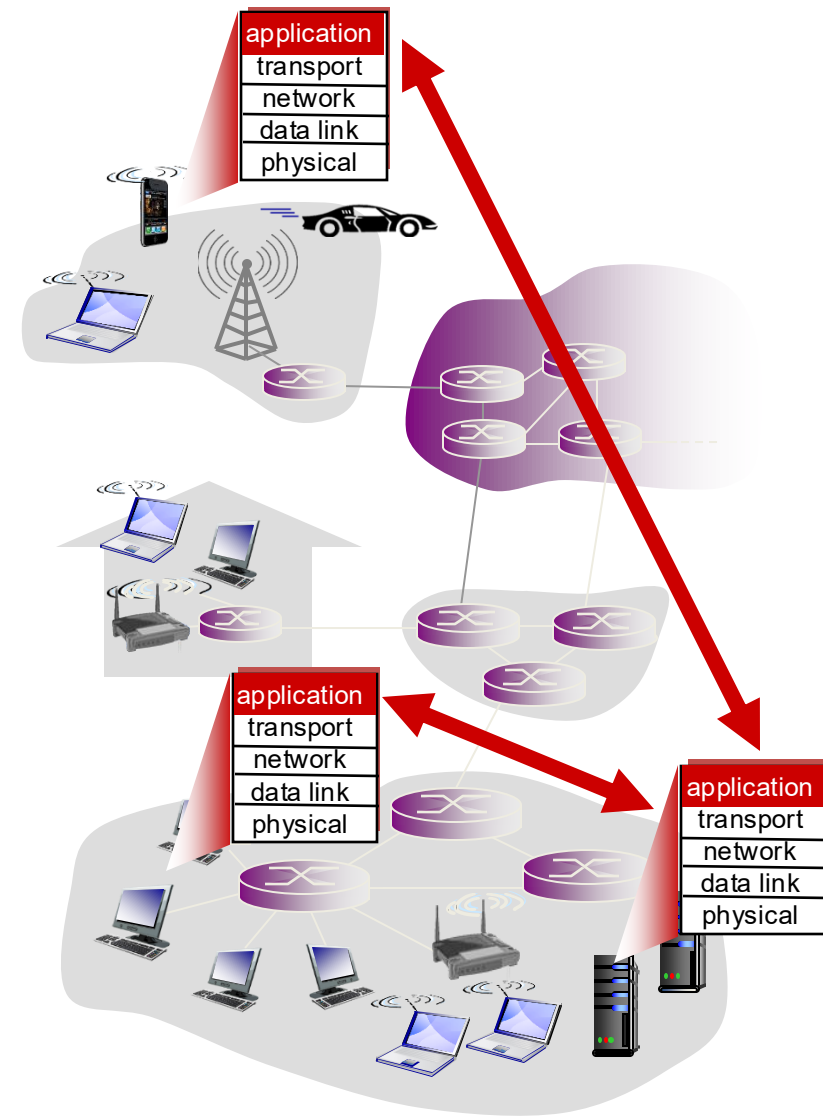
- run on (different) *end systems*
- communicate over network



Creating a network app

no need to write software for network-core devices!

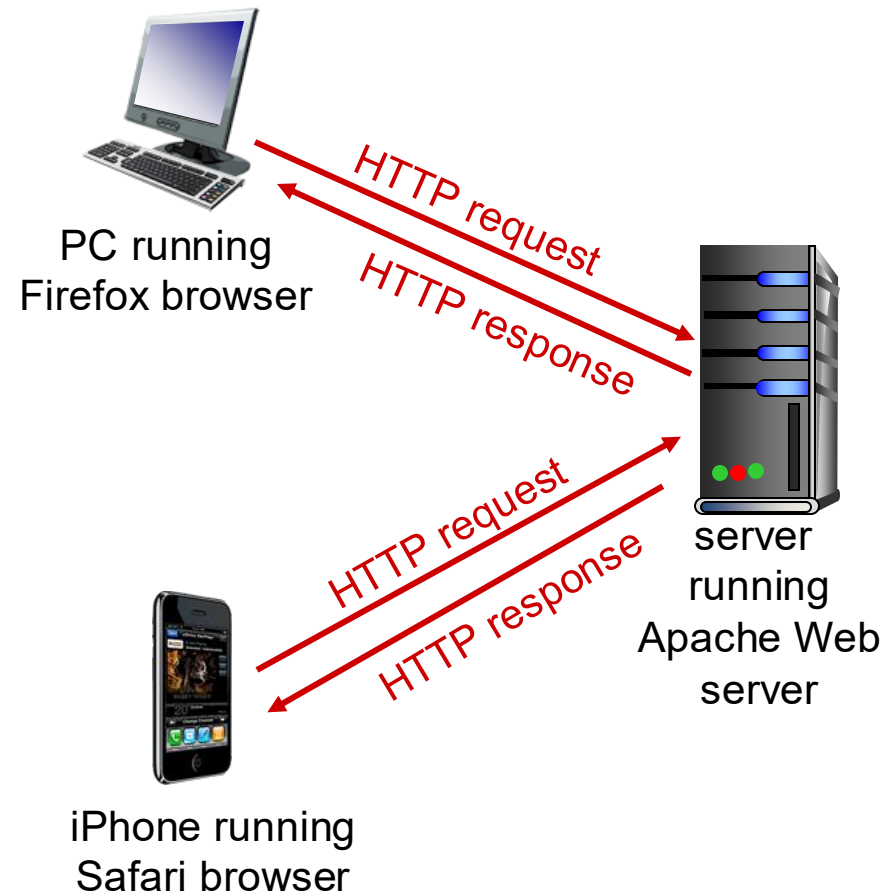
- network-core devices do not run user applications
- applications on end systems
 - rapid app development, propagation



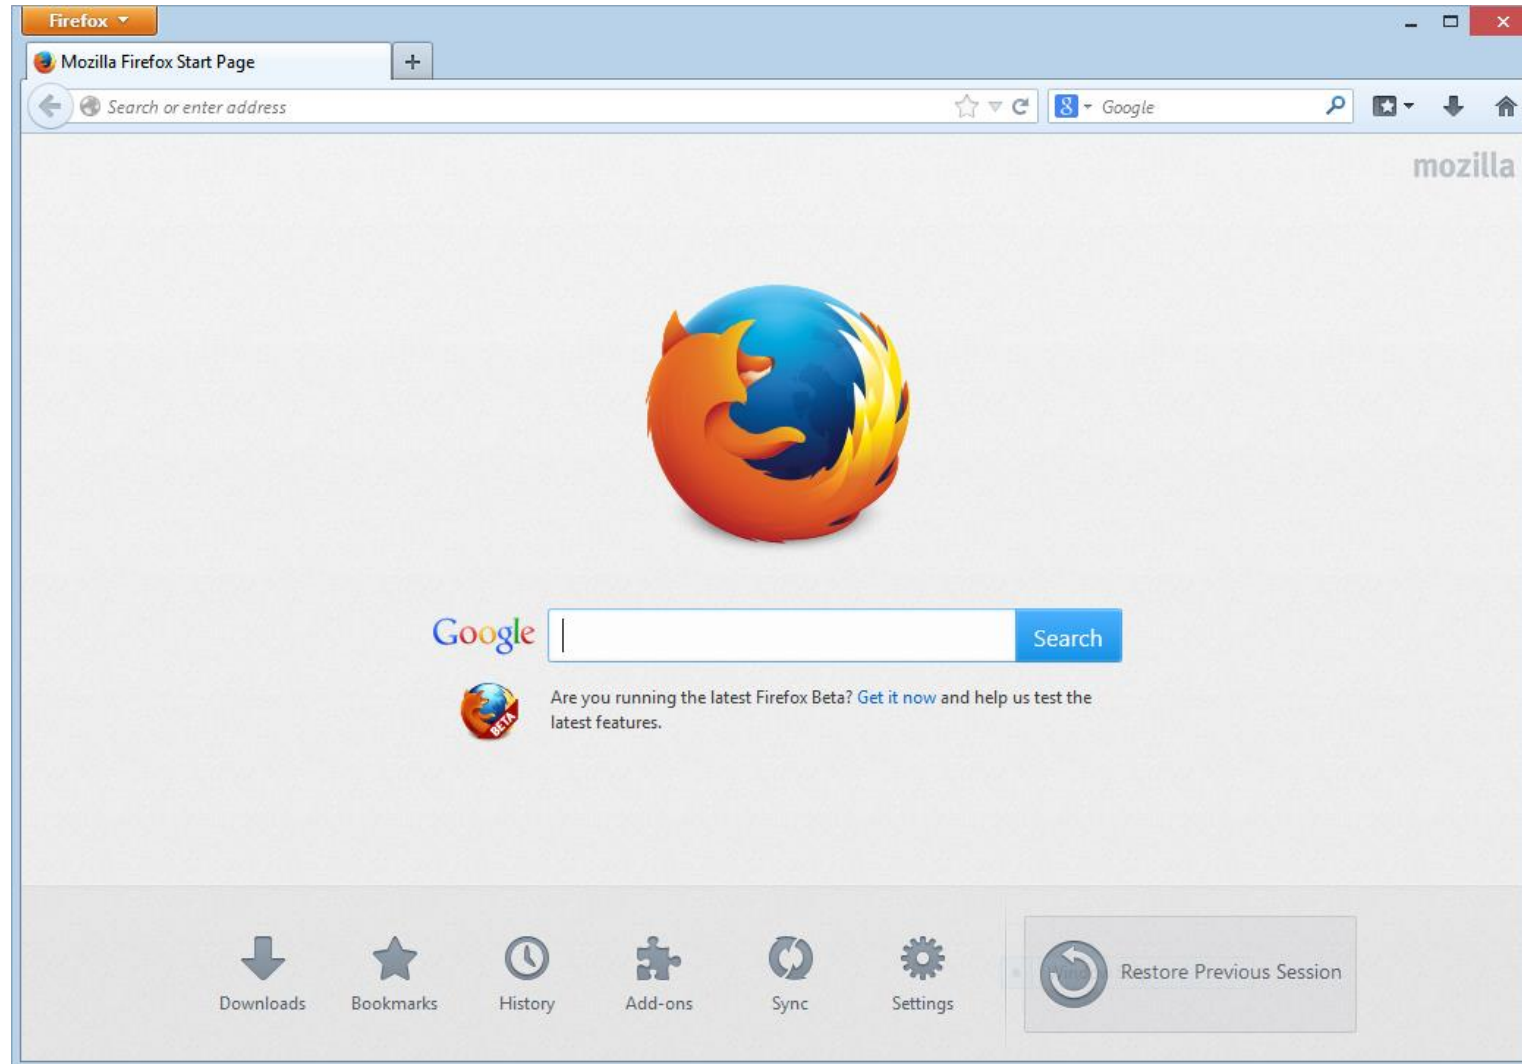
HTTP: HyperText Transfer Protocol

Client/Server model

- **client:** browser that uses HTTP to request, and receive Web objects.
- **server:** Web server that uses HTTP to respond with requested object.



What IS A Web Browser?



HTTP and the Web

- **web page** consists of **objects**
- object can be: an HTML file (index.html)

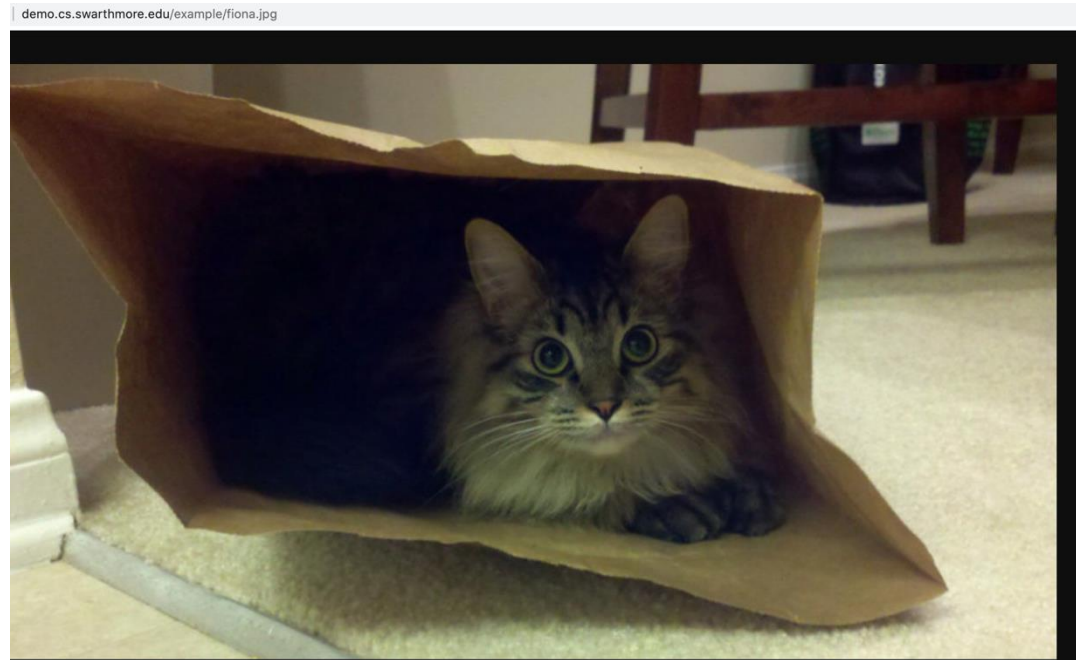
demo.cs.swarthmore.edu/index.html

This is the root page of the demo server. The interesting examples live in the [/example](#) directory. They are:

- [/example/directory/](#): An example of a directory.
- [/example/fiona.jpg](#): An example image (one of Kevin's cats).
- [/example/hello.txt](#): A simple text file.
- [/example/index.html](#): An HTML file serving as the default page for the /example directory.
- [/example/pic.html](#): An HTML file that links to the cat picture.
- [/example/pride_and_prejudice.pdf](#): A large PDF (binary) file containing Jane Austen's "Pride and Prejudice".
- [/example/pride_and_prejudice.txt](#): A large text file containing Jane Austen's "Pride and Prejudice".

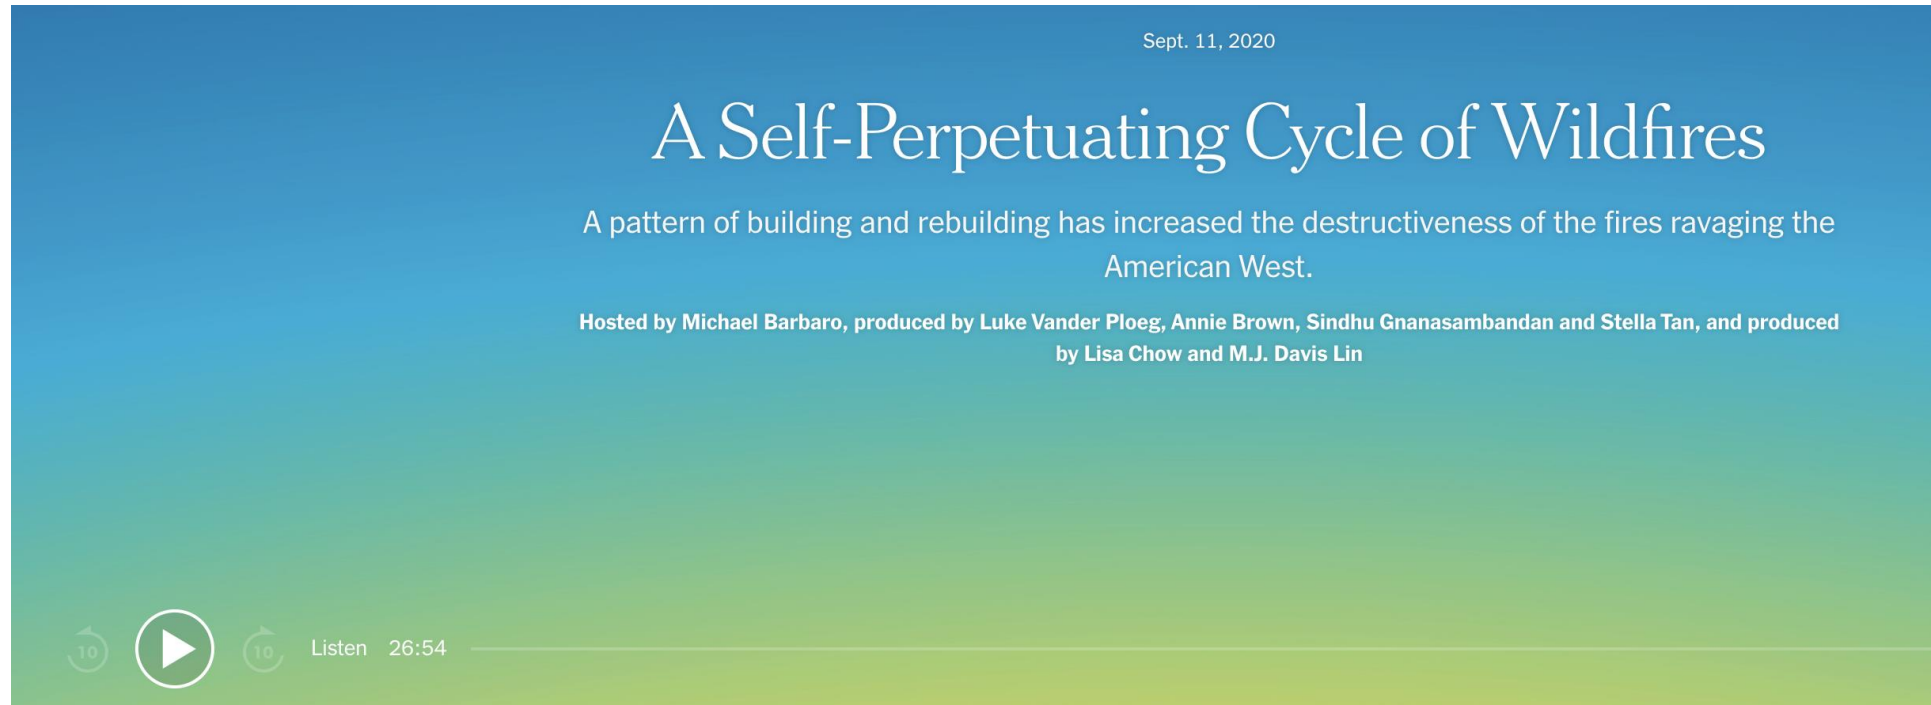
Web objects

- **web page** consists of **objects**
- object can be: JPEG image



Web objects

- web page consists of objects
- object can be: audio file



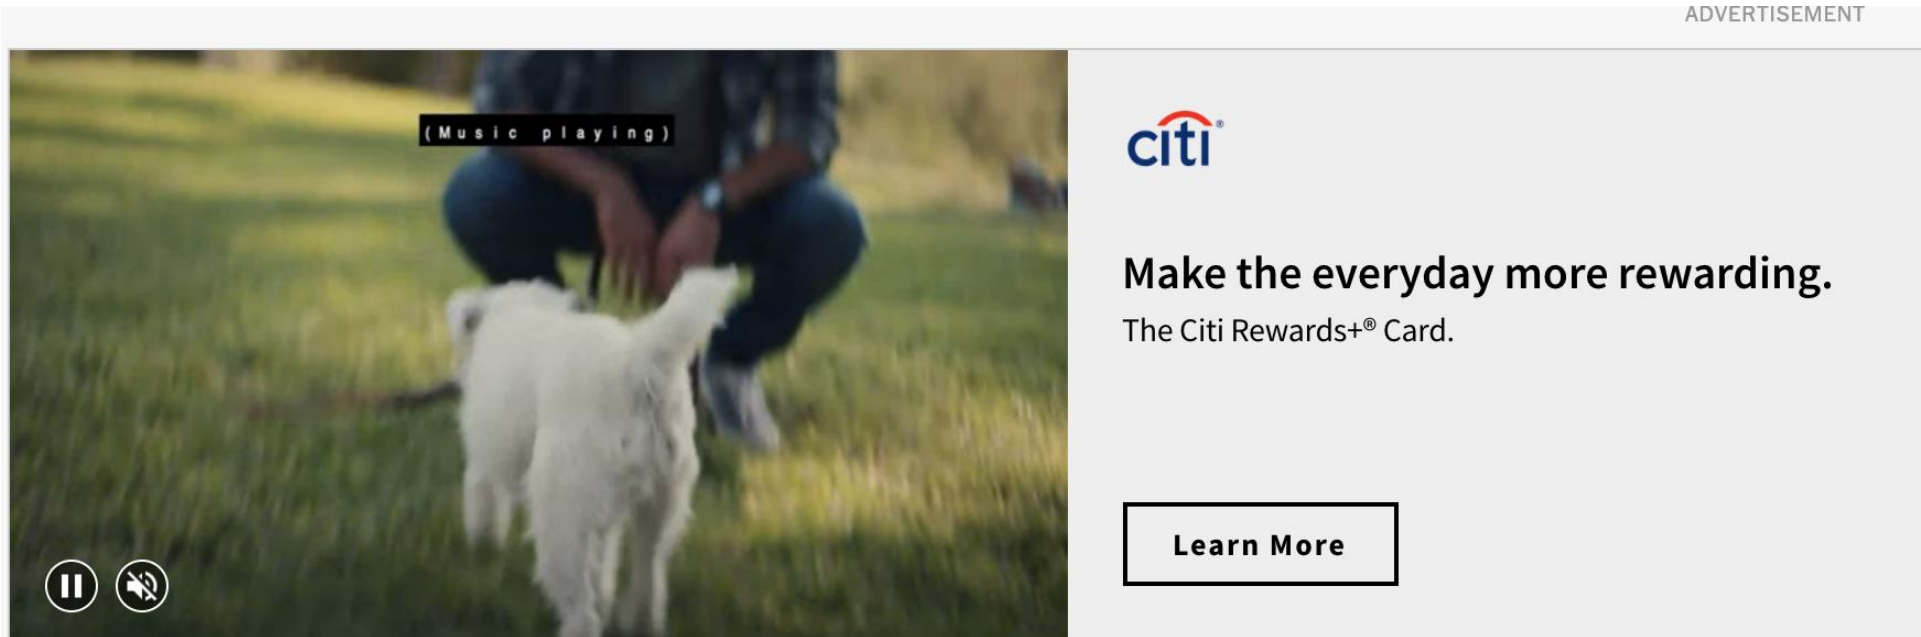
Courtesy: New York Times

Slide 32

Web objects

- **web page** consists of **objects**
- object can be: video, java applets, etc.

ADVERTISEMENT



The advertisement is a video player. The video shows a person crouching in a grassy field, petting a white dog. A black box with white text "(Music playing)" is overlaid on the video. In the bottom left corner of the video player, there are two circular icons: a pause icon and a mute icon. To the right of the video, on a light gray background, is the Citi logo. Below the logo, the text "Make the everyday more rewarding." is displayed in a bold, black font, followed by "The Citi Rewards+® Card." in a smaller, regular black font. At the bottom of this section is a black rectangular button with the text "Learn More" in white.

(Music playing)

citi

Make the everyday more rewarding.
The Citi Rewards+® Card.

Learn More

HTTP and the Web

- a web page consists of **base HTML-file** which includes **several referenced objects**
- each object is addressable by a **URL**, e.g.,

This is the root page of the demo server. The interesting examples live in the [/example](#) directory. They are:

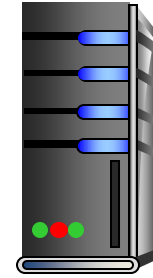
- [/example/directory/](#): An example of a directory.
- [/example/fiona.jpg](#): An example image (one of Kevin's cats).
- [/example/hello.txt](#): A simple text file.
- [/example/index.html](#): An HTML file serving as the default page for the /example directory.
- [/example/pic.html](#): An HTML file that links to the cat picture.
- [/example/pride_and_prejudice.pdf](#): A large PDF (binary) file containing Jane Austen's "Pride and Prejudice".
- [/example/pride_and_prejudice.txt](#): A large text file containing Jane Austen's "Pride and Prejudice".

`demo.cs.swarthmore.edu/example/pic.html`

host name

path name

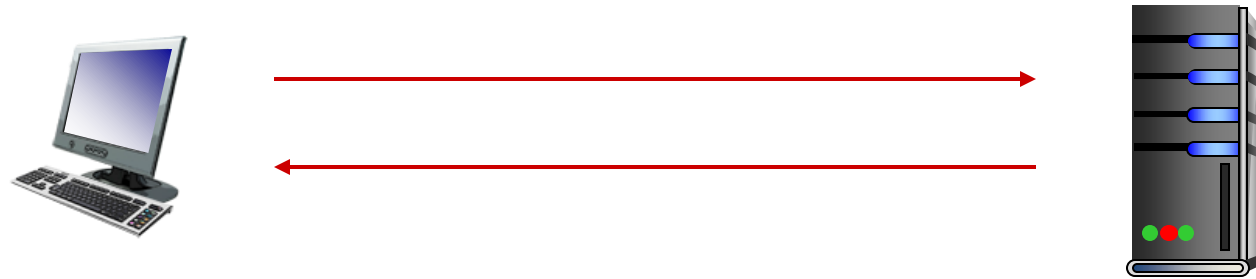
HTTP Overview



1. User types in a URL.

`http://some.host.name.tld/directory/name/file.ext`

HTTP Overview

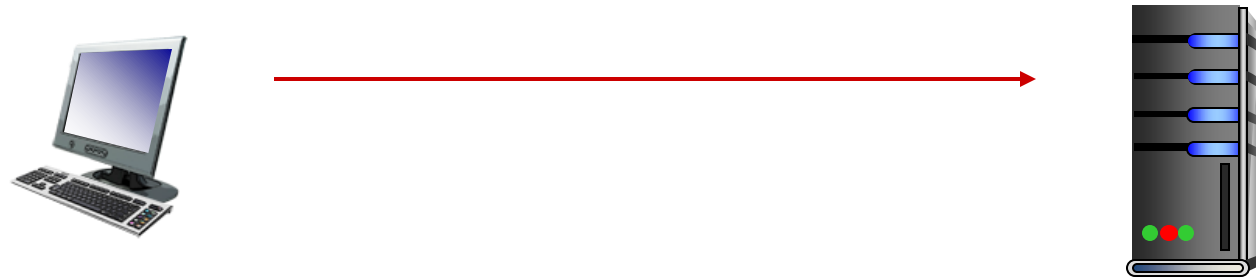


2. Browser establishes connection with server.

Looks up “some.host.name.tld”

Calls connect()

HTTP Overview



3. Browser requests the corresponding data.

GET /directory/name/file.ext HTTP/1.0

Host: some.host.name.tld

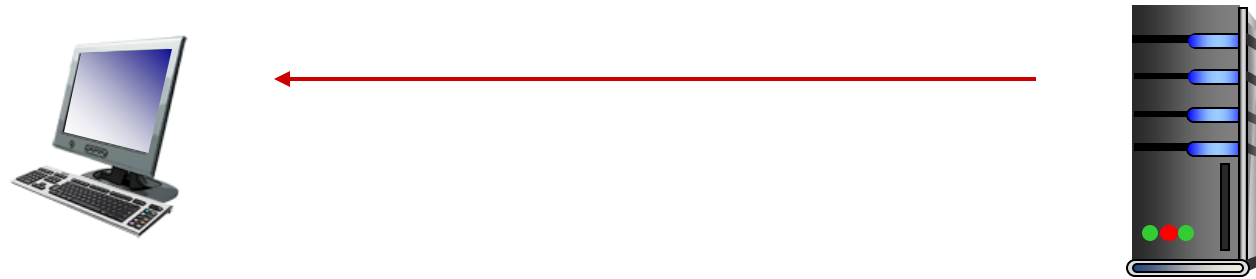
[other optional fields, for example:]

User-agent: Mozilla/5.0 (Windows NT 6.1; WOW64)

Accept-language: en

[Blank line]

HTTP Overview



4. Server responds with the requested data.

HTTP/1.0 200 OK

Content-Type: text/html

Content-Length: 1299

Date: Sun, 01 Sep 2013 21:26:38 GMT

[Blank line]

(Data data data data...)

HTTP Overview



5. Browser renders the response, fetches any additional objects, and waits for server to close the connection.

HTTP Overview



5. Browser renders the response, including any additional objects, and waits for the connection to close.

```
<html>
  <head>
    <title>Page title!</title>
  </head>

  <body>
    <p>a paragraph of text</p>

    
    
    ...
  </body>
</html>
```


HTTP Overview

1. User types in a URL.
2. Browser establishes connection with server.
3. Browser requests the corresponding data.
4. Server responds with the requested data.
5. Browser renders the response, fetches any additional objects, and waits for server to close the connection.

It's a document retrieval system, where documents point to (link to) each other, forming a "web".

HTTP Overview (Lab 1)

1. User types in a URL.
2. Browser establishes connection with server.
3. Browser requests the corresponding data.
4. Server responds with the requested data.
5. Browser ~~renders the response, fetches any additional objects, and~~ waits for server to close the connection.

It's a document retrieval system, where documents point to (link to) each other, forming a "web".

Trying out HTTP (client side) for yourself

1. Telnet to your favorite Web server:

```
telnet demo.cs.swarthmore.edu 80
```

Opens TCP connection to port 80 (default HTTP server port) at example server. Anything typed is sent to server on port 80 at demo.cs.swarthmore.edu

2. Type in a GET HTTP request:

```
GET / HTTP/1.0  
Host: demo.cs.swarthmore.edu  
(blank line)
```

By typing this in (hit enter twice), you send this minimal (but complete) GET request to the HTTP server.

3. Look at response message sent by HTTP server!

Example (live demo)

Example

```
kwebb@sesame ~ $ telnet demo.cs.swarthmore.edu 80
Trying 130.58.68.26...
Connected to demo.cs.swarthmore.edu.
Escape character is '^]'.
GET /example/hello.txt HTTP/1.0
Host: demo.cs.swarthmore.edu

HTTP/1.0 200 OK
Content-Type: text/plain; charset=utf-8
ETag: "914817348"
Last-Modified: Mon, 24 Feb 2020 06:06:27 GMT
Content-Length: 40
Connection: close
Date: Thu, 20 Jan 2022 18:03:33 GMT
Server: lighttpd/1.4.59

Hello, you found the example text file!
```

Request

Response
headers

Response body

(This is what you should be
saving to file in lab 1.)

Note!

```
kwebb@sesame ~ $ telnet demo.cs.swarthmore.edu 80
```

```
Trying 130.58.68.26...
```

```
Connected to demo.cs.swarthmore.edu.
```

```
Escape character is '^['.
```

```
GET /example/hello.txt
```

```
Host: demo.cs.swarthmore.edu
```

This server is intentionally NOT using encryption, to make it easier to work with for lab 1!

```
HTTP/1.0 200 OK
```

```
Content-Type: text/plain; charset=utf-8
```

```
ETag: "914817348"
```

```
Last-Modified: Mon, 24 Feb 2020 06:06:27 GMT
```

```
Content-Length: 40
```

```
Connection: close
```

```
Date: Thu, 20 Jan 2022 18:03:33 GMT
```

```
Server: lighttpd/1.4.59
```

```
Hello, you found the example text file!
```

HTTPS (live demo)

- Telnet transfers unencrypted data ("clear text")
 - Great for learning
 - Not so great for real world security / privacy
- For a similar (interactive) command line experience with encryption:
 - `openssl s_client -crLf -connect server.name:443`

HTTP request message

- two types of HTTP messages: *request, response*
- **HTTP request message:**
 - ASCII (human-readable format)

The diagram illustrates the structure of an HTTP request message. It consists of a request line followed by header lines, each terminated by a carriage return (CR) and a line-feed (LF) character. The request line is annotated with a blue arrow pointing to the text 'request line (GET, POST, HEAD, etc. commands)'. The header lines are grouped by a blue bracket and labeled 'header lines'. The final carriage return and line feed are annotated with a blue arrow pointing to the text 'carriage return, line feed'. The message is as follows:

```
GET /~kwebb/index.html HTTP/1.1\r\nHost: web.cs.swarthmore.edu\r\nUser-Agent: Firefox/3.6.10\r\nAccept: text/html,application/xhtml+xml\r\nAccept-Language: en-us,en;q=0.5\r\nAccept-Encoding: gzip,deflate\r\nAccept-Charset: ISO-8859-1,utf-8;q=0.7\r\nKeep-Alive: 115\r\nConnection: keep-alive\r\n\r\n
```

Annotations:

- request line (GET, POST, HEAD, etc. commands)
- header lines
- carriage return, line feed
- carriage return character (CR)
- line-feed character (LF)

Why do we have these `\r\n` (CRLF) things all over the place?

```
GET /~kwebb/index.html HTTP/1.1\r\n
Host: web.cs.swarthmore.edu\r\n
User-Agent: Firefox/3.6.10\r\n
Accept: text/html,application/xhtml+xml\r\n
Accept-Language: en-us,en;q=0.5\r\n
Accept-Encoding: gzip,deflate\r\n
Accept-Charset: ISO-8859-1,utf-8;q=0.7\r\n
Keep-Alive: 115\r\n
Connection: keep-alive\r\n
\r\n
```

- A. They're generated when the user hits 'enter'.
- B. They signal the end of a field or section.
- C. They're important for some other reason.
- D. They're an unnecessary protocol artifact.

How else might we delineate messages?

(What are the good/bad properties of each of these ideas?)

- A. There's not much else we can do.
- B. Force all messages to be the same size.
- C. Send the message size prior to the message.
- D. Some other way (discuss).

HTTP is all text...

- Makes the protocol simple
 - Easy to delineate message (\r\n)
 - (Relatively) human-readable
 - No worries about encoding or formatting data
 - Variable length data
- Not the most efficient
 - Many protocols use binary fields
 - Sending “12345678” as a string is 8 bytes
 - As an integer, 12345678 needs only 4 bytes
 - The headers may come in any order
 - Requires string parsing / processing

HTTP is all text...

- The HTTP **PROTOCOL** is all text
 - That is, the messages that are required (request and response)
 - All headers are text
- The BODY of a message might **NOT** be text
- This distinction is critically important for lab 1!
 - Fine to use string functions on HTTP messages
 - You better not use string functions on body data

Visualizing HTTP: telnet

```
kwebb@sesame ~ $ telnet demo.cs.swarthmore.edu 80
Trying 130.58.68.26...
Connected to demo.cs.swarthmore.edu.
Escape character is '^]'.
GET /example/hello.txt HTTP/1.0
Host: demo.cs.swarthmore.edu

HTTP/1.0 200 OK
Content-Type: text/plain; charset=utf-8
ETag: "914817348"
Last-Modified: Mon, 24 Feb 2020 06:06:27 GMT
Content-Length: 40
Connection: close
Date: Thu, 20 Jan 2022 18:03:33 GMT
Server: lighttpd/1.4.59

Hello, you found the example text file!
```

Visualizing HTTP: wireshark

File Edit View Go Capture Analyze Statistics Telephony Tools Internals Help

Filter: **http** Expression... Clear Apply

No.	Time	Source	Destination	Protocol	Length	Info
14	14.211168	130.58.68.164	130.58.68.137	HTTP	68	GET /~kwebb/ HTTP/1.1
16	14.268895	130.58.68.137	130.58.68.164	HTTP	5447	HTTP/1.1 200 OK (text/html)

▶ Frame 16: 5447 bytes on wire (43576 bits), 5447 bytes captured (43576 bits)

▶ Ethernet II, Src: SuperMic_74:d5:b8 (00:25:90:74:d5:b8), Dst: Hewlett-_94:5f:1e (d8:d3:85:94:5f:1e)

▶ Internet Protocol Version 4, Src: 130.58.68.137 (130.58.68.137), Dst: 130.58.68.164 (130.58.68.164)

▶ Transmission Control Protocol, Src Port: http (80), Dst Port: 35736 (35736), Seq: 1, Ack: 55, Len: 5381

▼ Hypertext Transfer Protocol

▼ HTTP/1.1 200 OK\r\n

▶ [Expert Info (Chat/Sequence): HTTP/1.1 200 OK\r\n]

Request Version: HTTP/1.1

Status Code: 200

Response Phrase: OK

Date: Mon, 02 Sep 2013 20:10:28 GMT\r\n

Server: Apache/2.2.22 (Ubuntu)\r\n

Last-Modified: Sat, 31 Aug 2013 20:44:44 GMT\r\n

ETag: "c3f7c-1401-4e54468c06210"\r\n

Accept-Ranges: bytes\r\n

▼ Content-Length: 5121\r\n

[Content length: 5121]

Vary: Accept-Encoding\r\n

Content-Type: text/html\r\n

\r\n

▼ Line-based text data: text/html

<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Strict//EN" "http://www.w3.org/TR/xhtml1/DTD/xhtml1-strict.dtd" >\r\n

<html xmlns="http://www.w3.org/1999/xhtml" xml:lang="en" lang="en" >\r\n

<head>\r\n

0040 fe fc 48 54 54 50 2f 31 2e 31 20 32 30 30 20 4f ..HTTP/1 .1 200 0

0050 4b 0d 0a 44 61 74 65 3a 20 4d 6f 6e 2c 20 30 32 K..Date: Mon, 02

0060 20 53 65 70 20 32 30 31 33 20 32 30 3a 31 30 3a Sep 201 3 20:10:

0070 32 38 20 47 4d 54 0d 0a 53 65 72 76 65 72 3a 20 28 GMT.. Server:

HTTP Response Status Code (... : Packets: 20 Displayed: 2 Marked: 0 Load time: 0:00.000 Profile: Default

- There's more to say about HTTP, but for lab 1, let's talk a bit about sockets too...

What is a socket?

An inter-process communication (IPC) abstraction through which an application may send and receive data.

Behaves similarly to way an open file handle allows an application to read and write data to storage.

Recall Inter-process Communication (IPC)

- Processes must communicate to cooperate
- Must have two mechanisms:
 - Data transfer
 - Synchronization

Inter-process Communication (IPC)

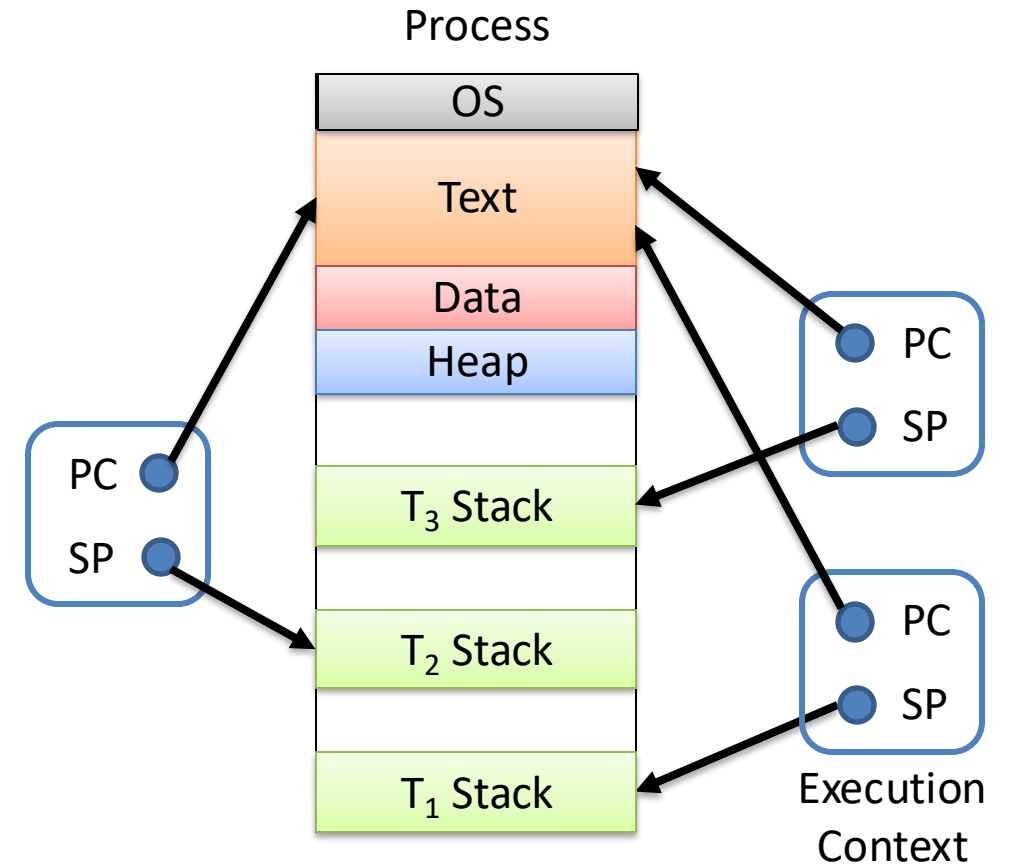
- Operating systems provide several IPC mechanisms (Take CS 45)
 - files
 - shared memory (in several ways)
 - pipes
 - ...
 - sockets
- Broadly, these fall into two categories:
 1. Shared memory
 2. Message passing

Only works on one computer (shared hardware).

Also, this is what you're most familiar with.

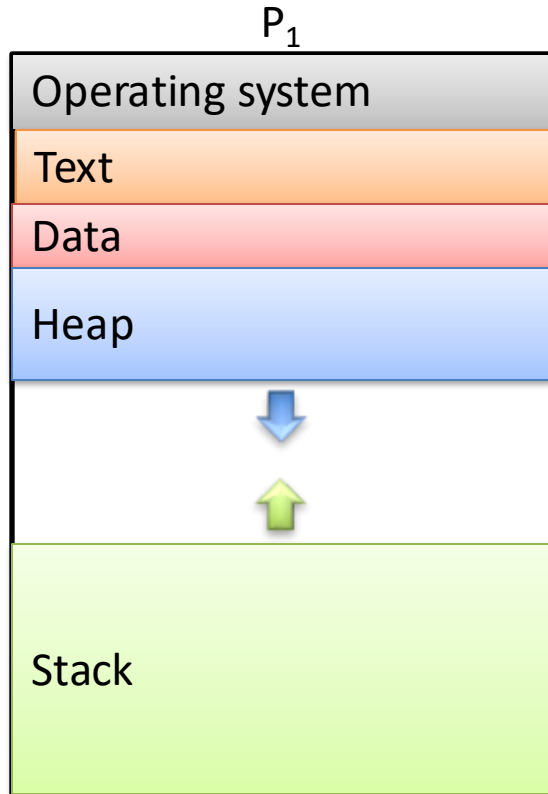
Thread Model (Shared Memory)

- Single process with multiple copies of execution resources.
- ONE shared virtual address space!
 - All process memory shared by every thread.
 - Threads coordinate by sharing variables (typically on heap)



Note: this is technically not IPC (there's only one process), but this is the most common form of shared memory today.

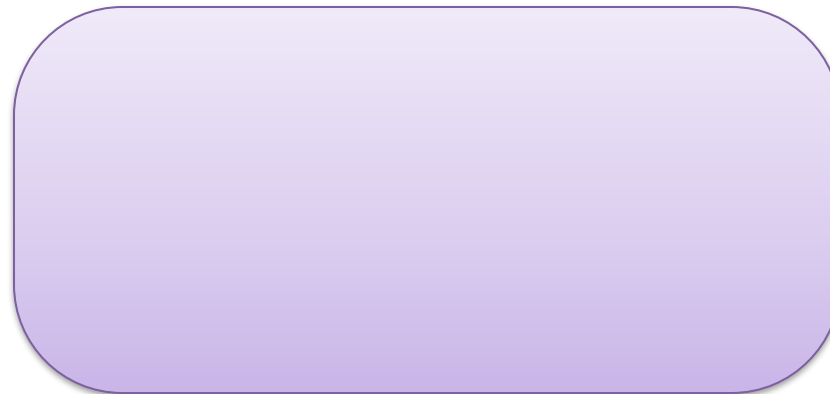
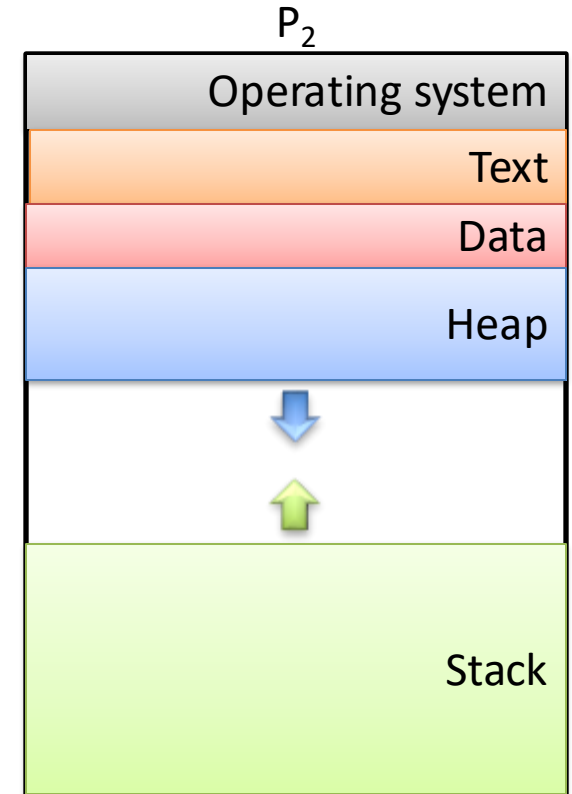
Message Passing IPC (Pipe)



Let's say process P_1 wants to send data to process P_2 .

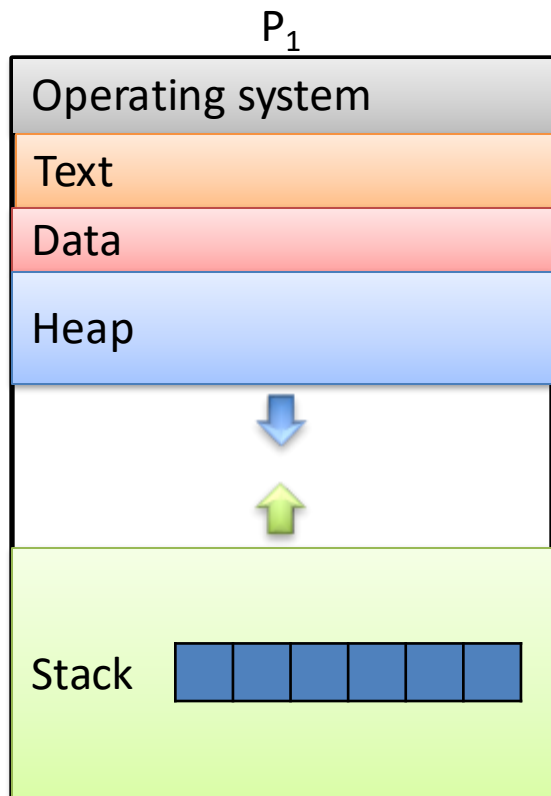
They execute on the same hardware and share an operating system.

They do NOT directly share any memory.



OS kernel

Message Passing IPC (Pipe)

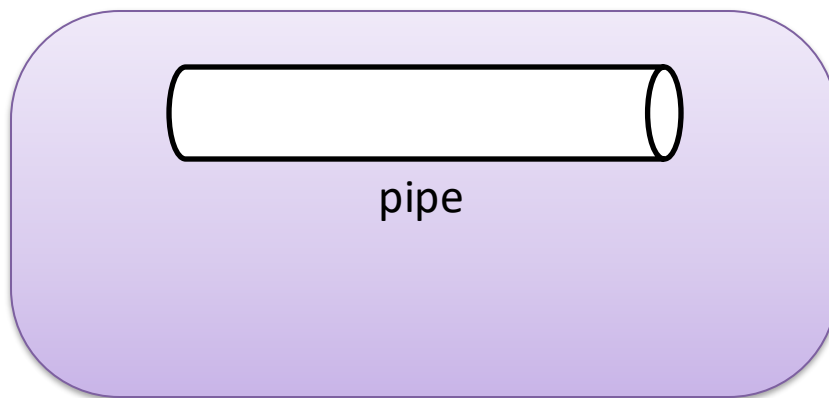


P_1 can send data into the pipe by calling:

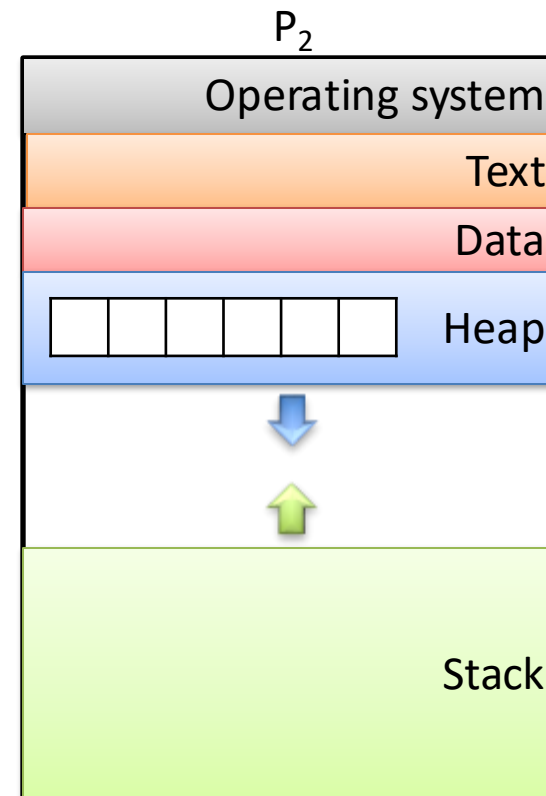
`write(..., data pointer, count)`

data pointer: the start of data to copy

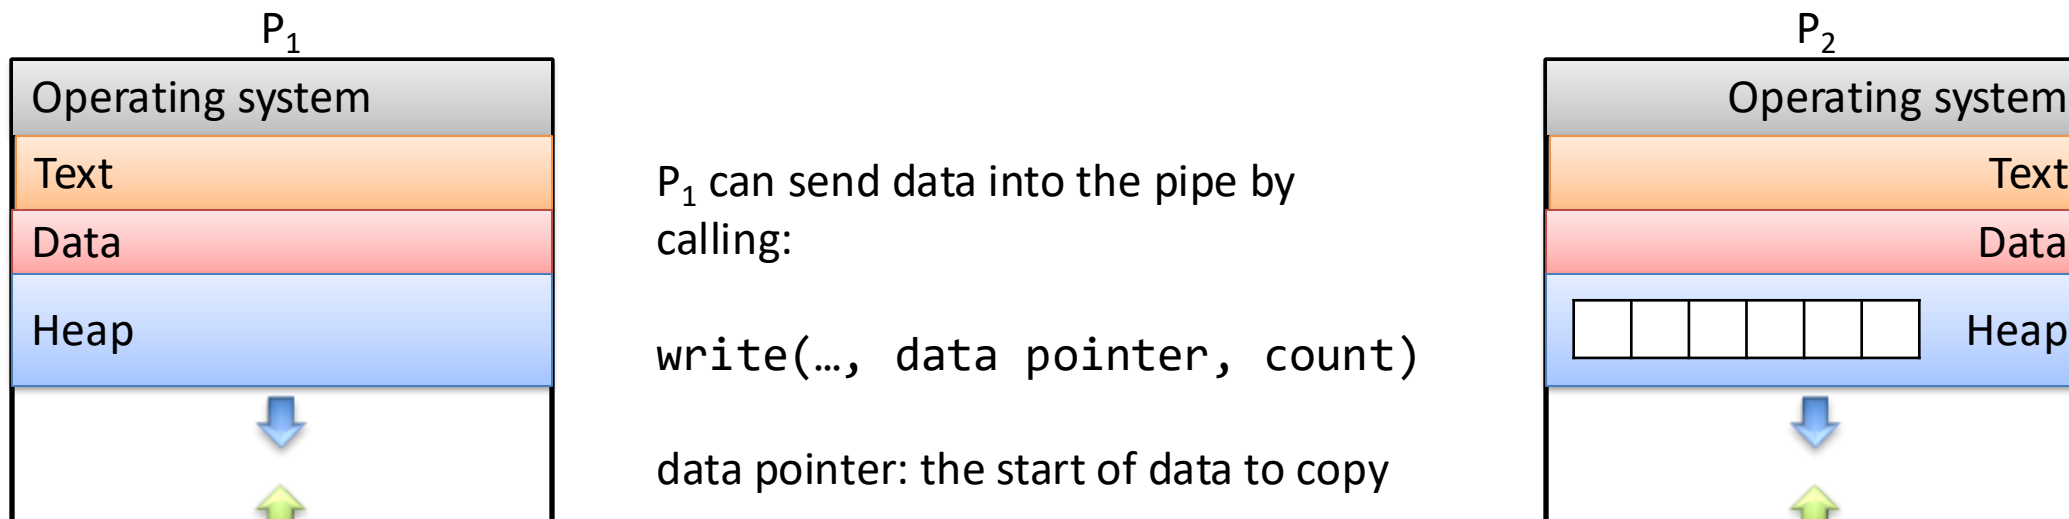
count: how many bytes to copy (at most)



OS kernel



Message Passing IPC (Pipe)



NAME

`write` - write to a file descriptor

SYNOPSIS

```
#include <unistd.h>
```

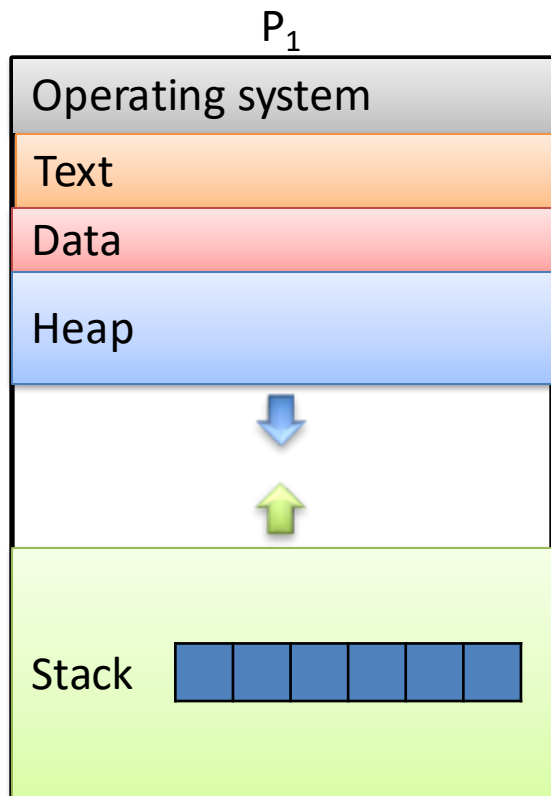
```
ssize_t write(int fd, const void *buf, size_t count);
```

DESCRIPTION

`write()` writes up to count bytes from the buffer starting at buf to the file referred to by the file descriptor fd.

OS kernel

Message Passing IPC (Pipe)

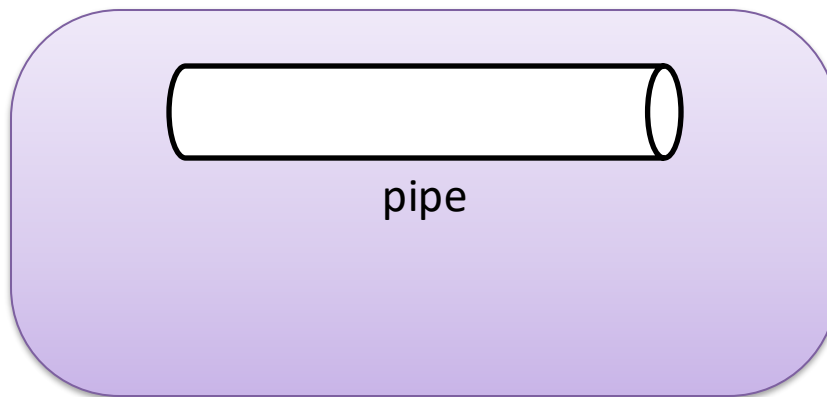
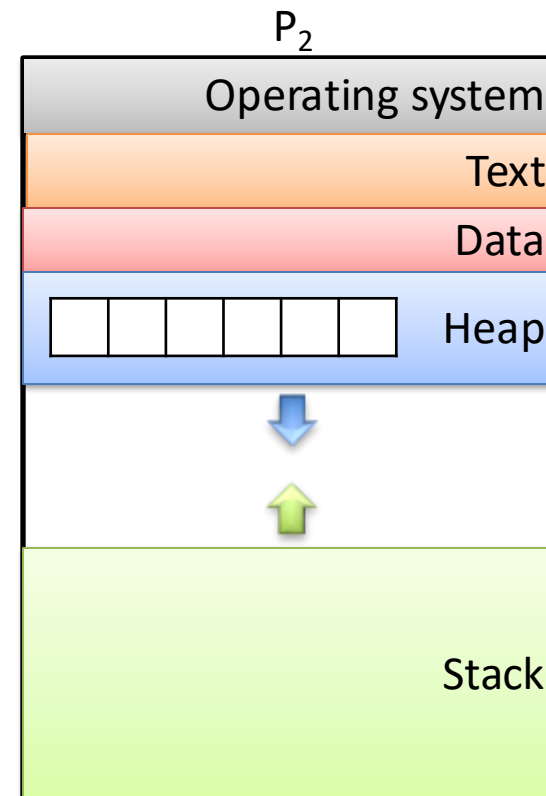


P_1 can send data into the pipe by calling:

`write(fd, data pointer, count)`

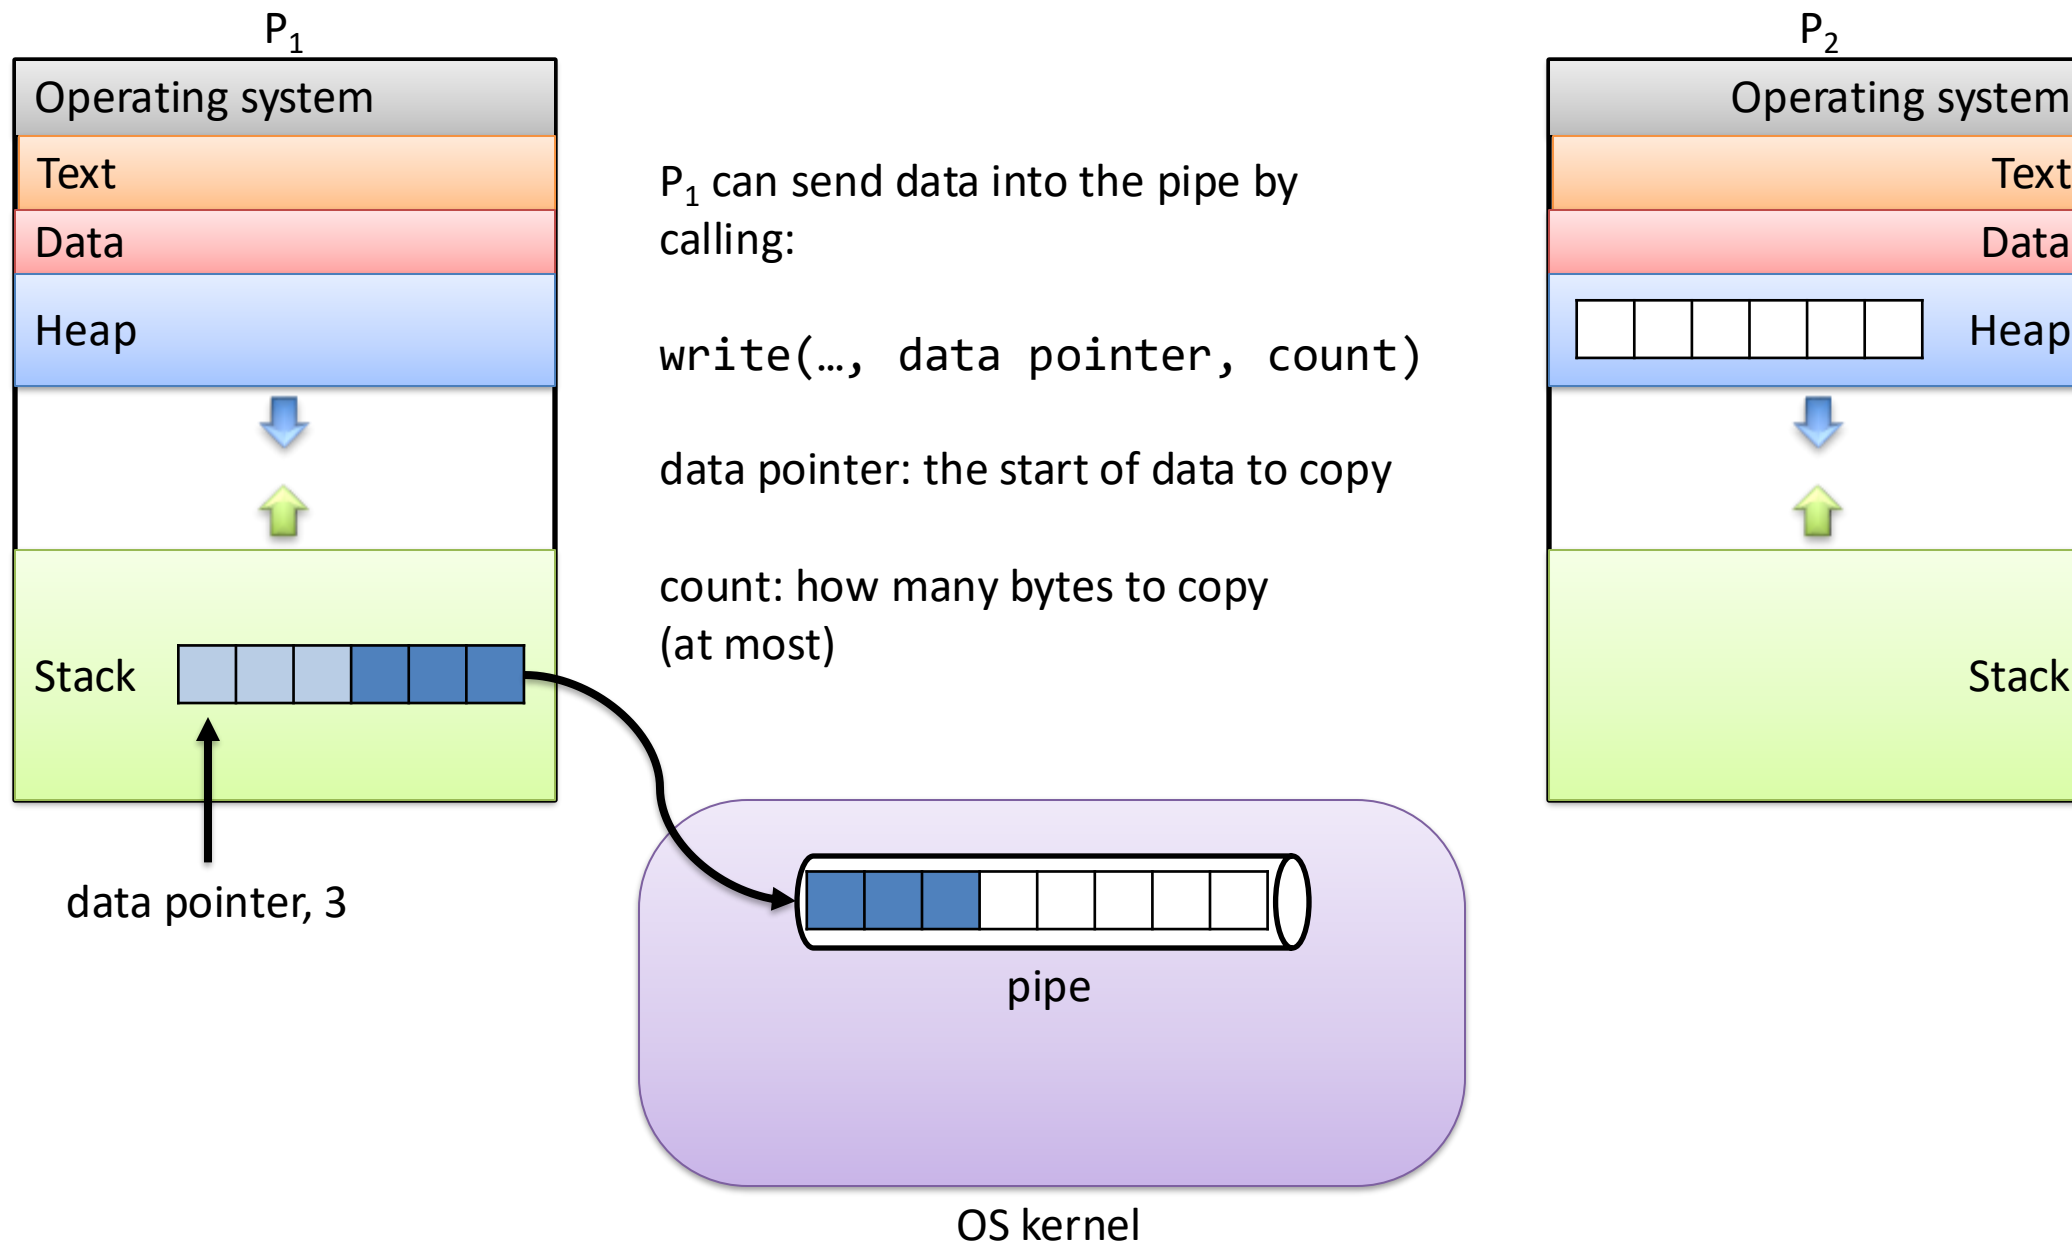
data pointer: the start of data to copy

count: how many bytes to copy

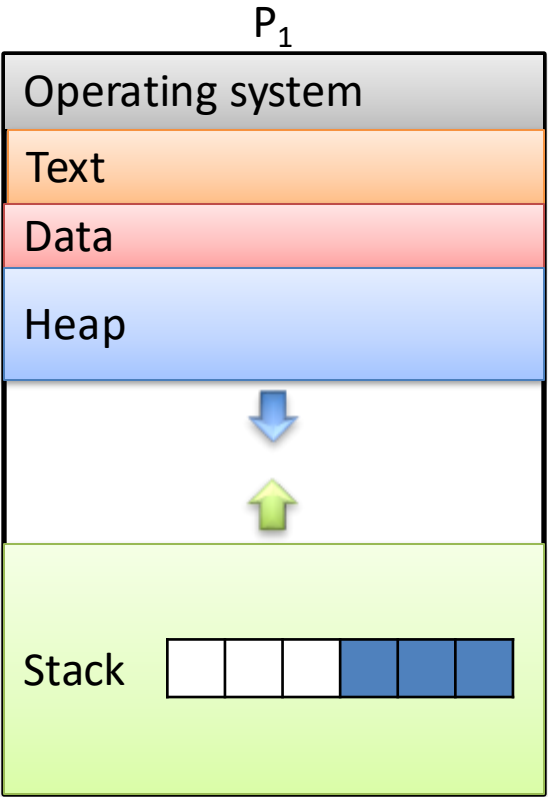


OS kernel

Message Passing IPC (Pipe)



Message Passing IPC (Pipe)

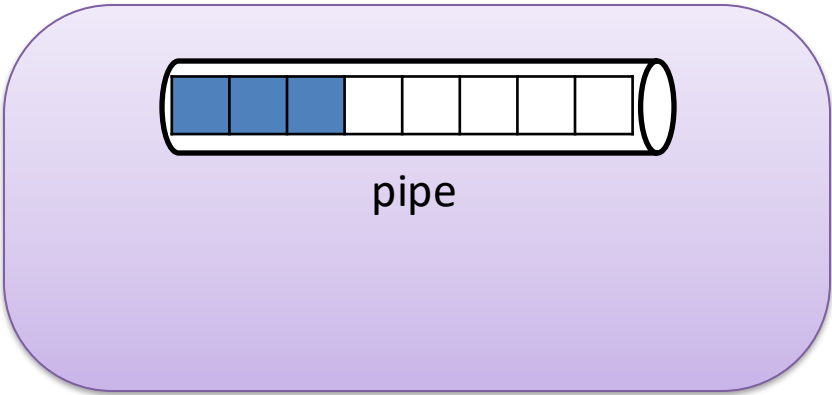
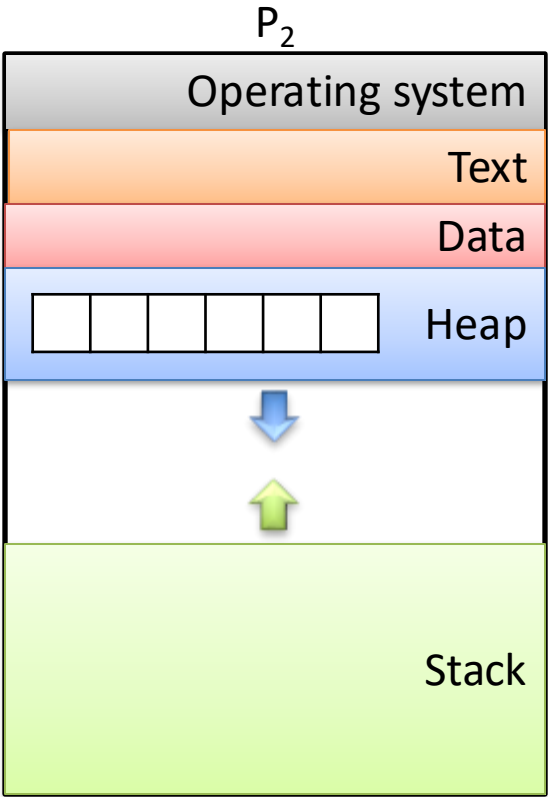


P_2 can receive data from the pipe by calling:

`read(..., data pointer, count)`

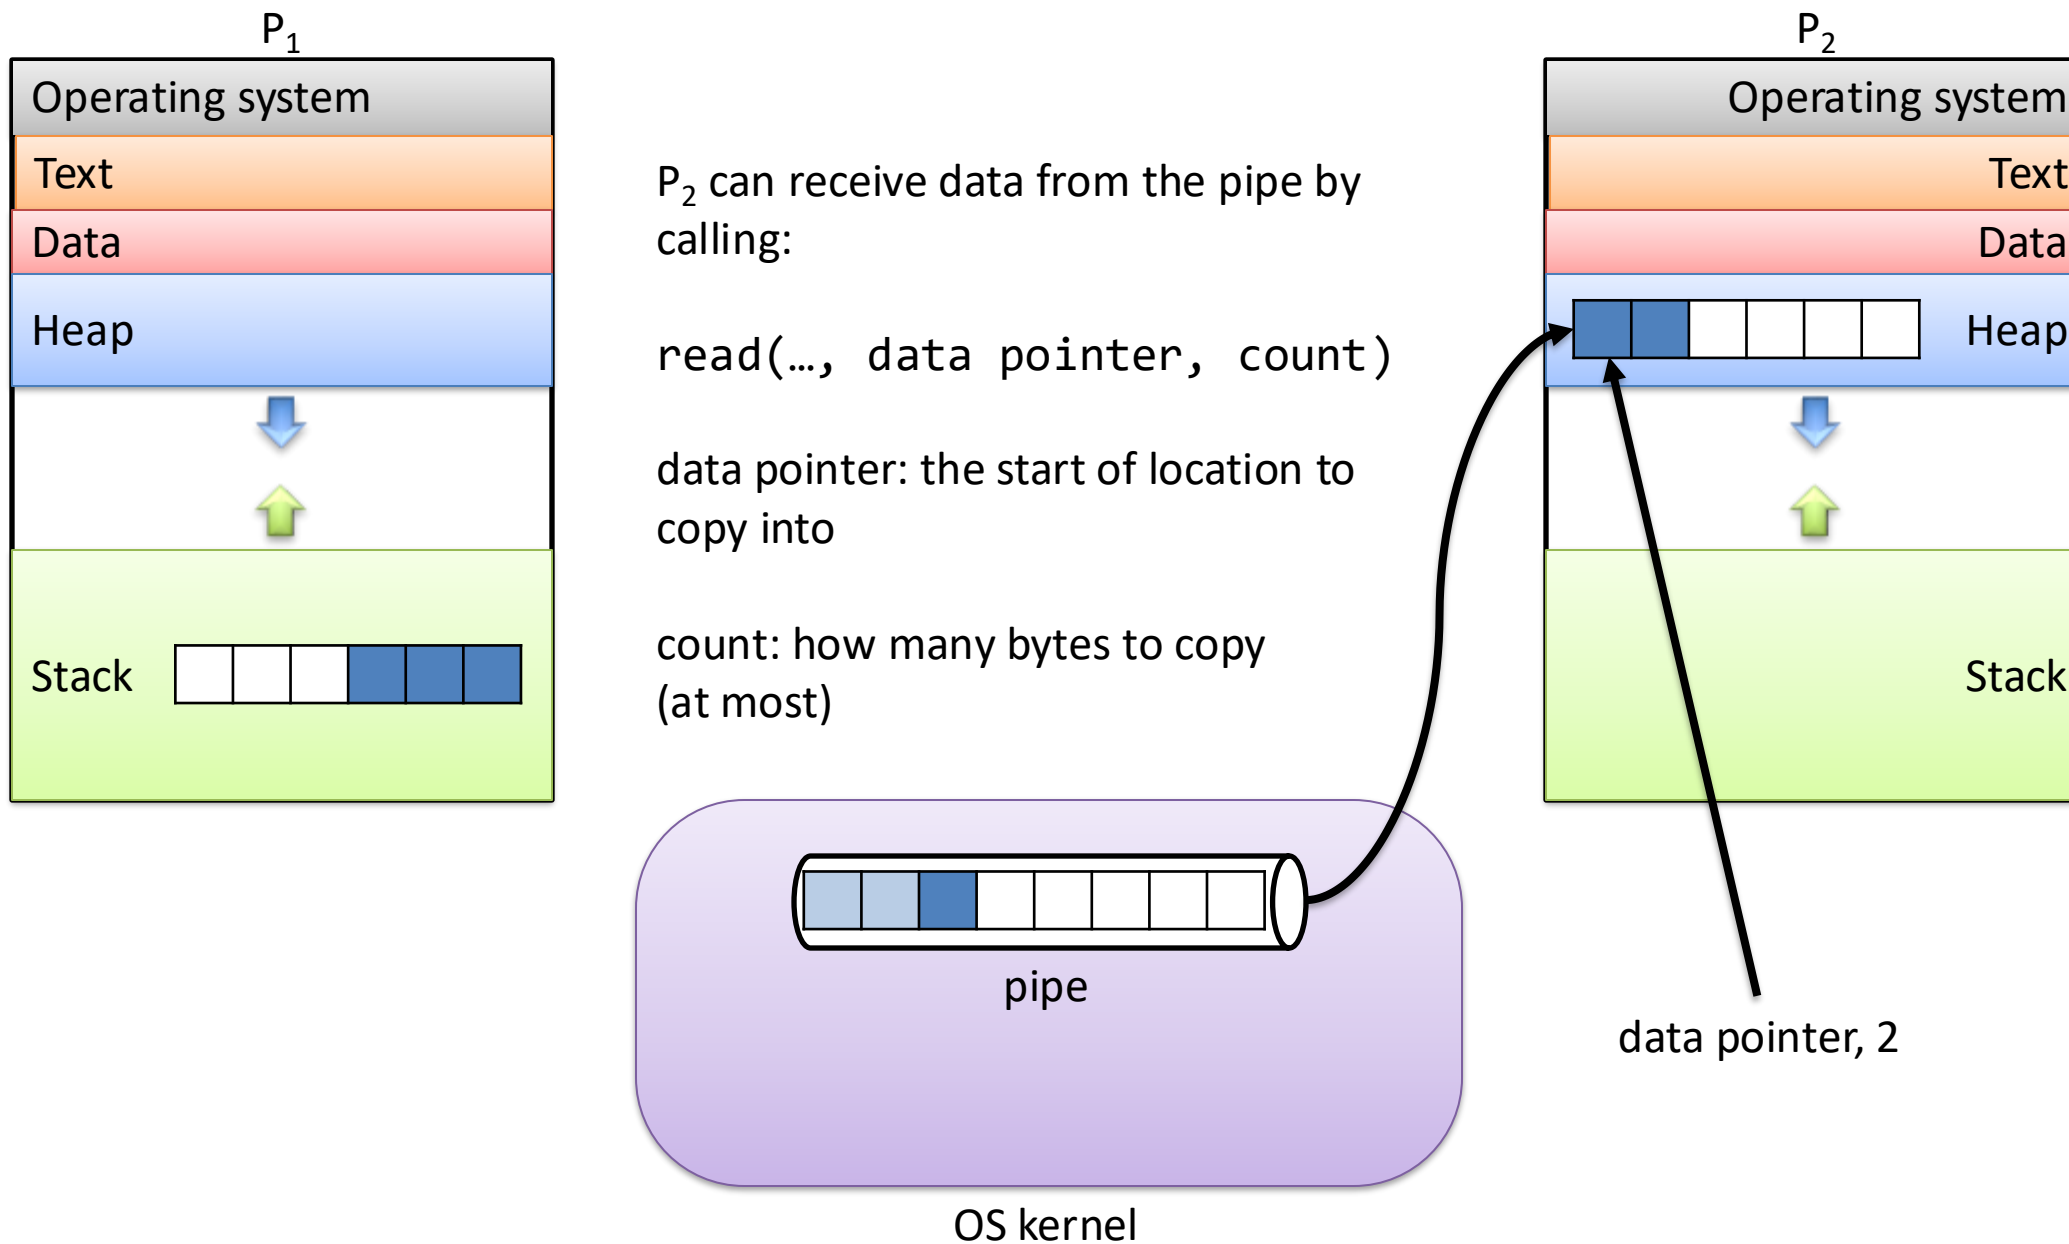
data pointer: the start of location to copy into

count: how many bytes to copy (at most)

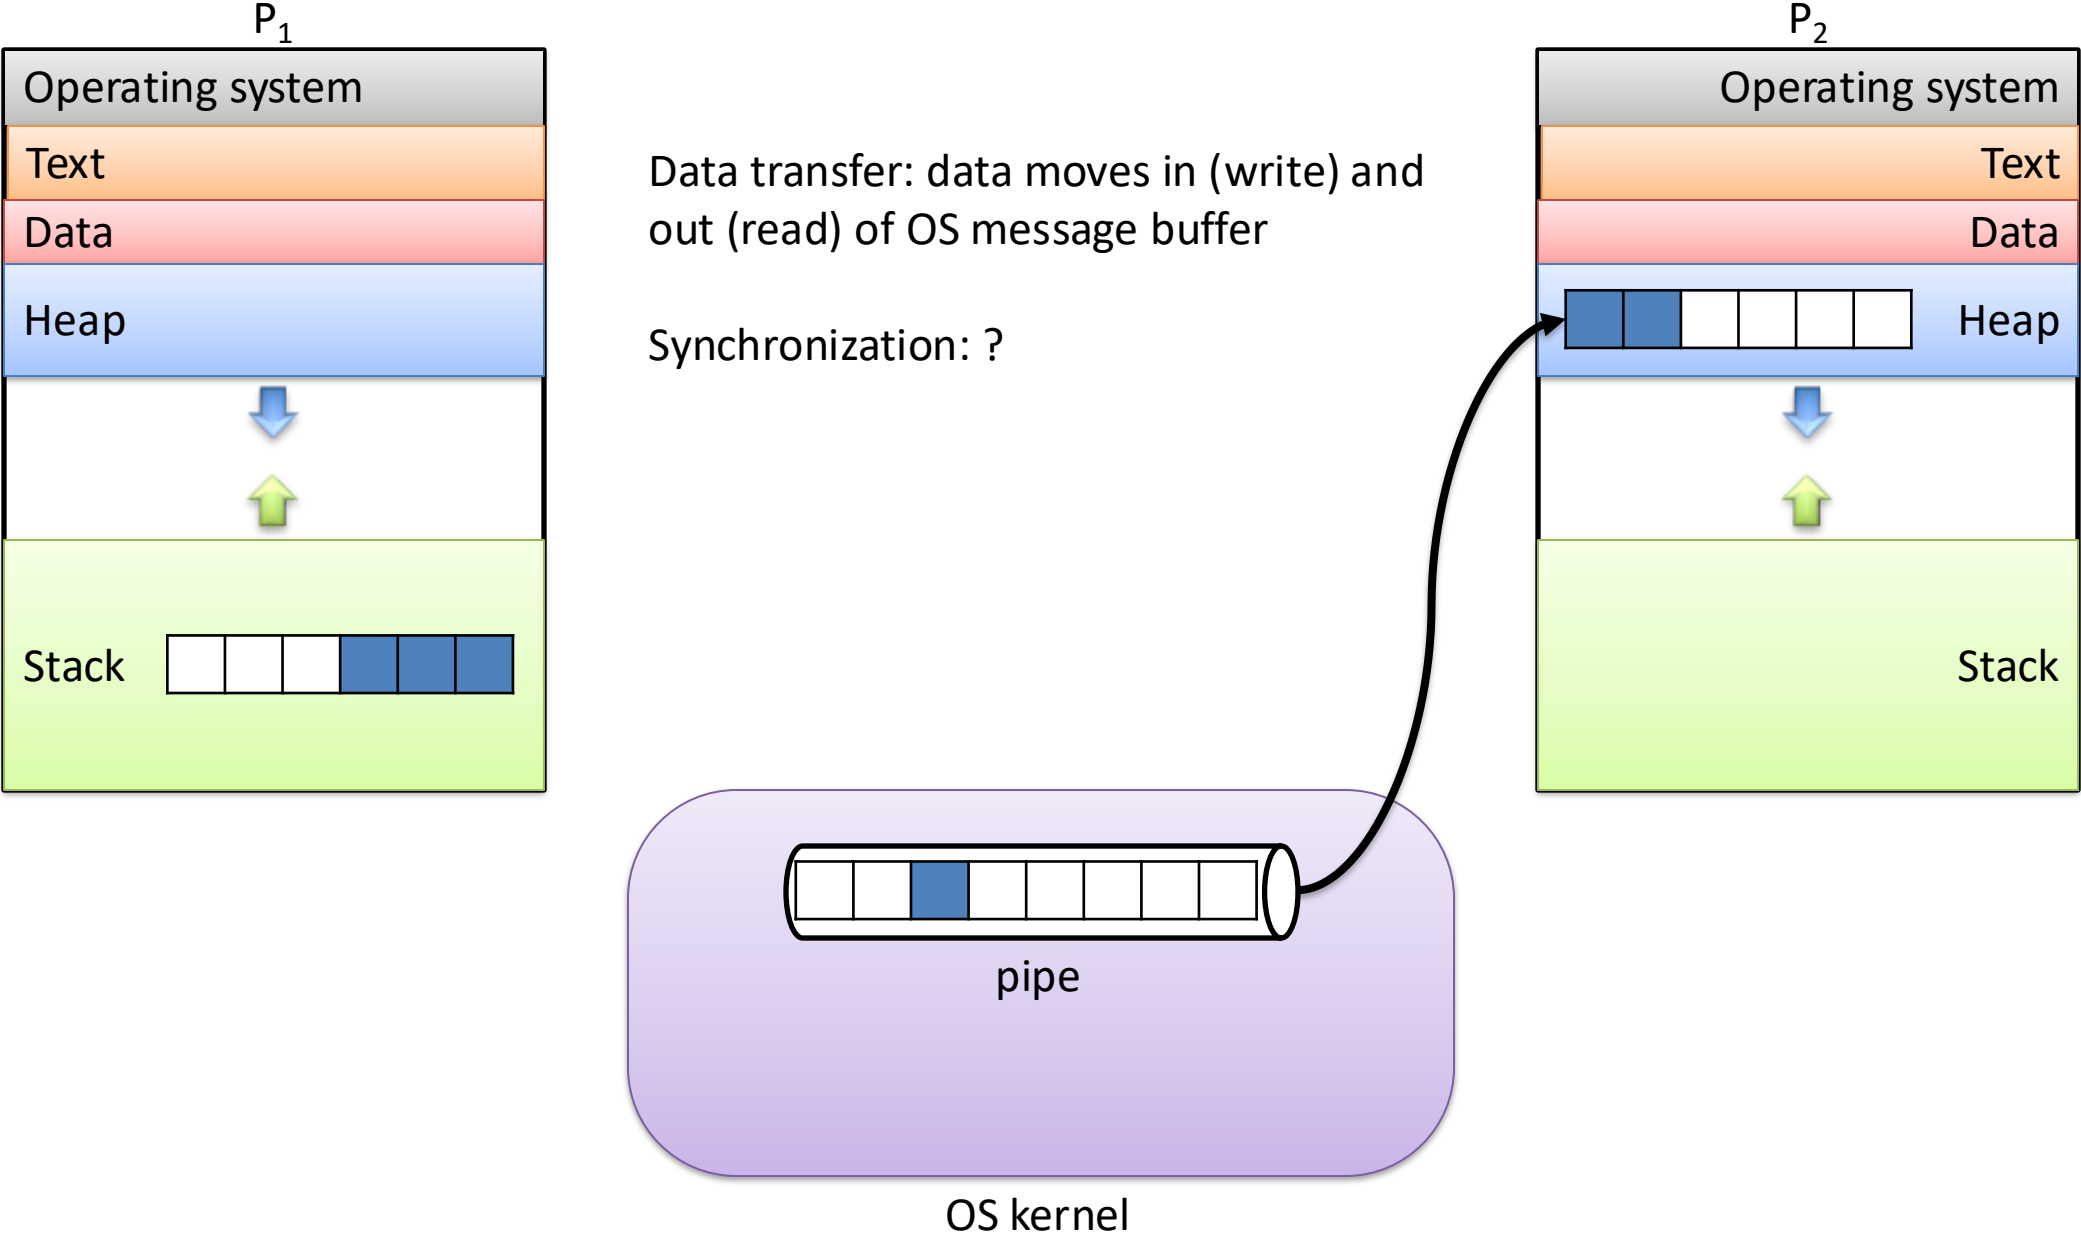


OS kernel

Message Passing IPC (Pipe)



Message Passing IPC (Pipe)

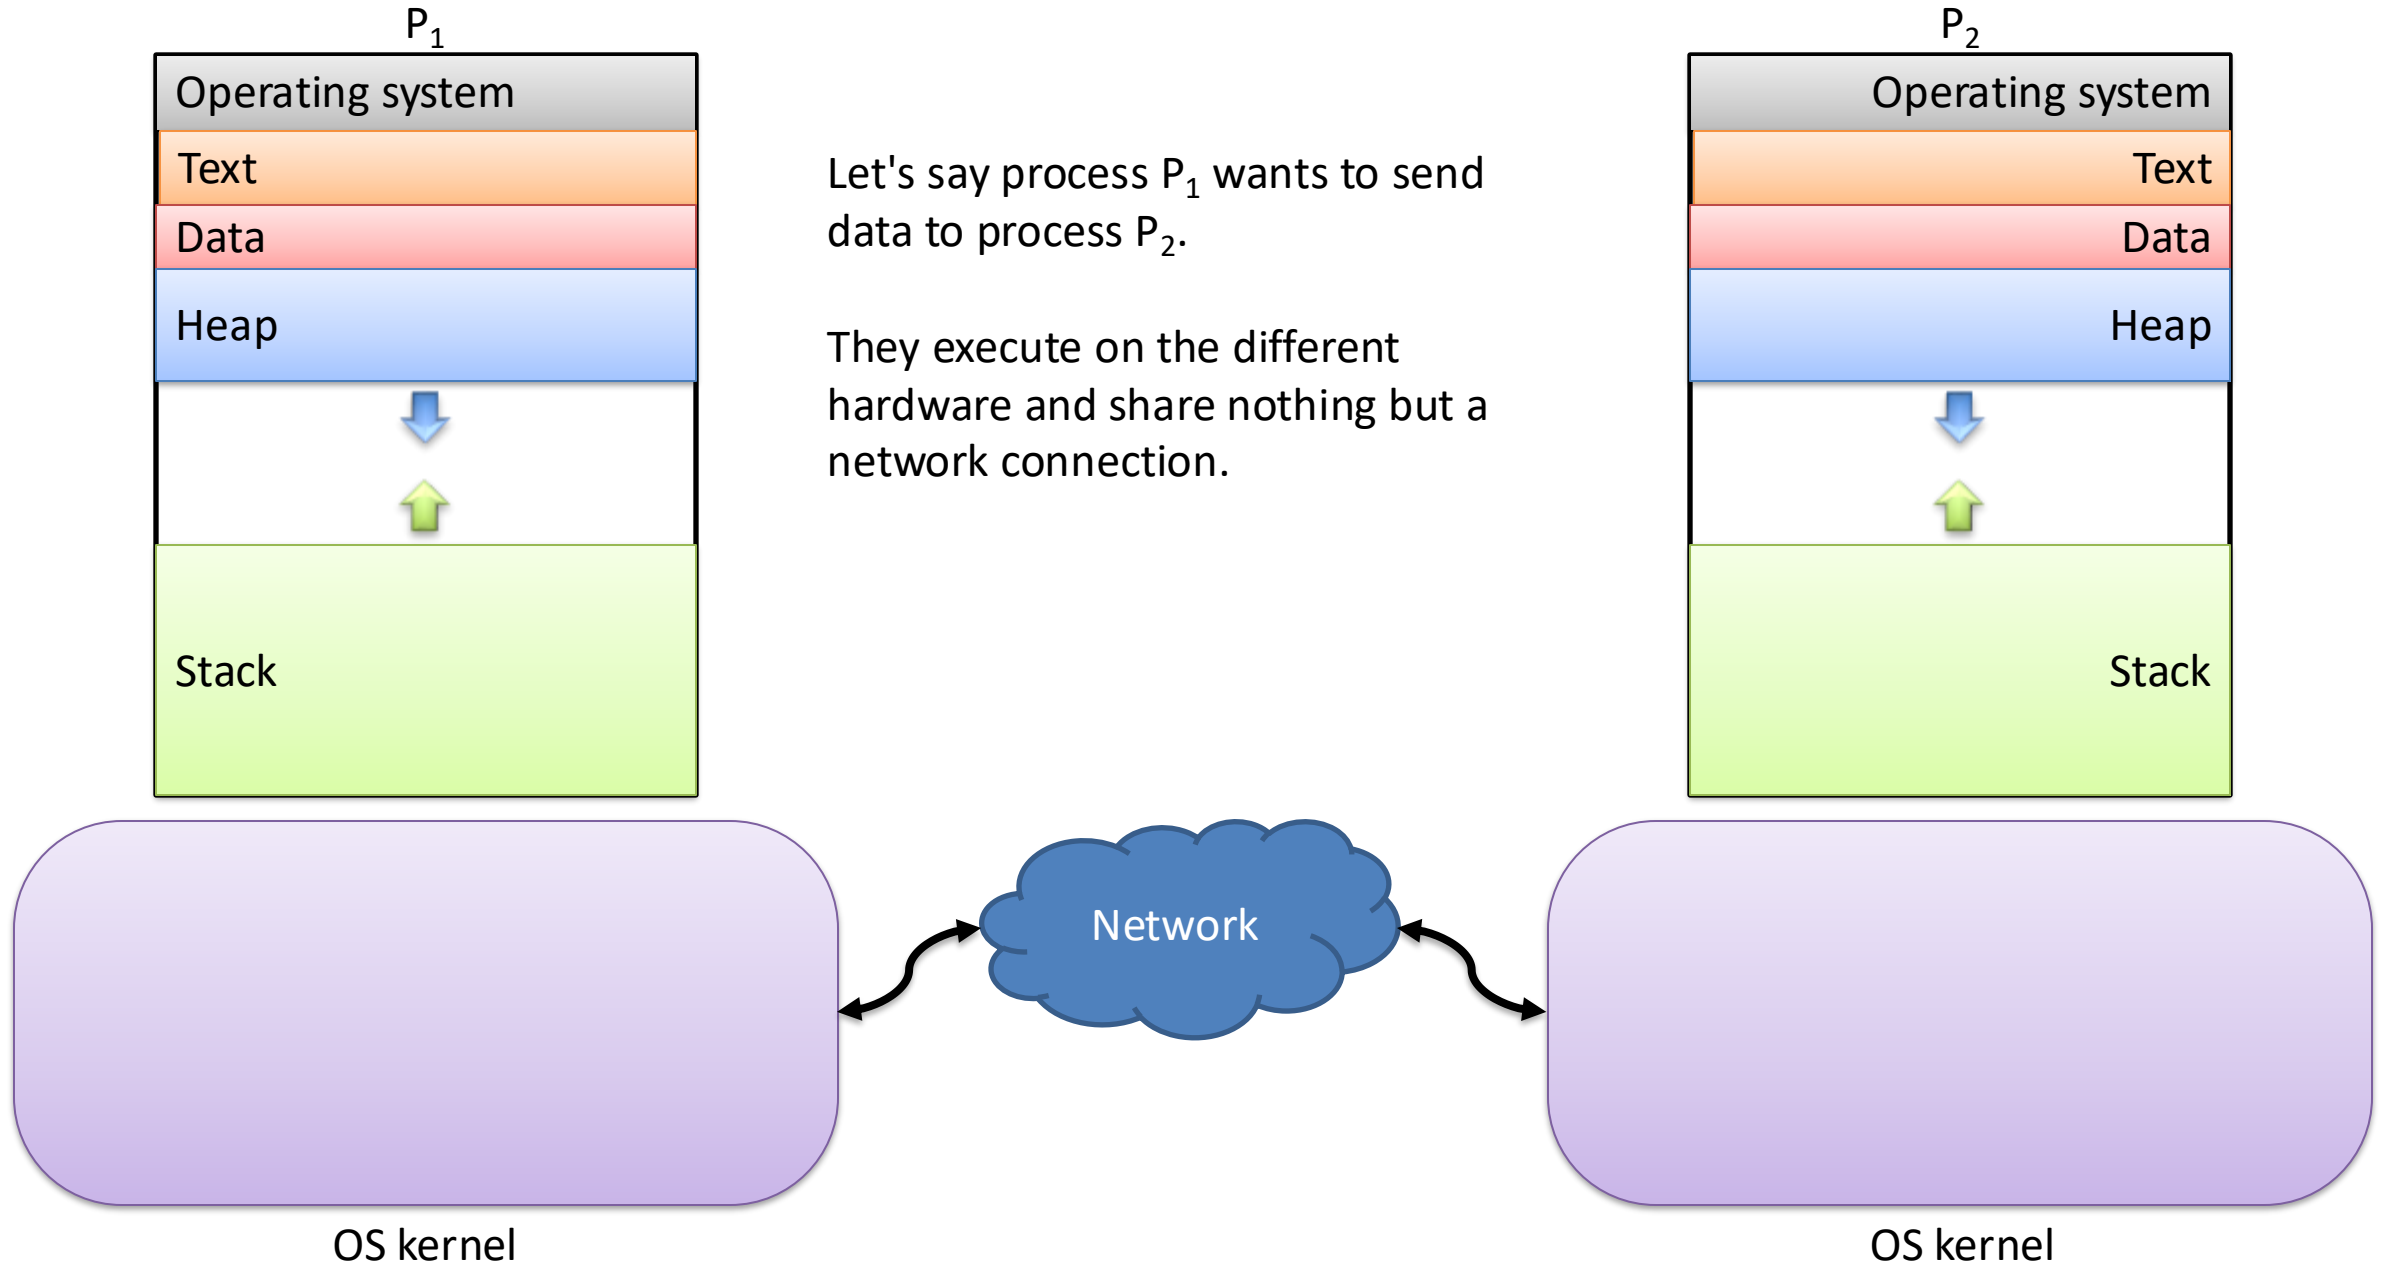


Where is the synchronization* in message passing IPC?

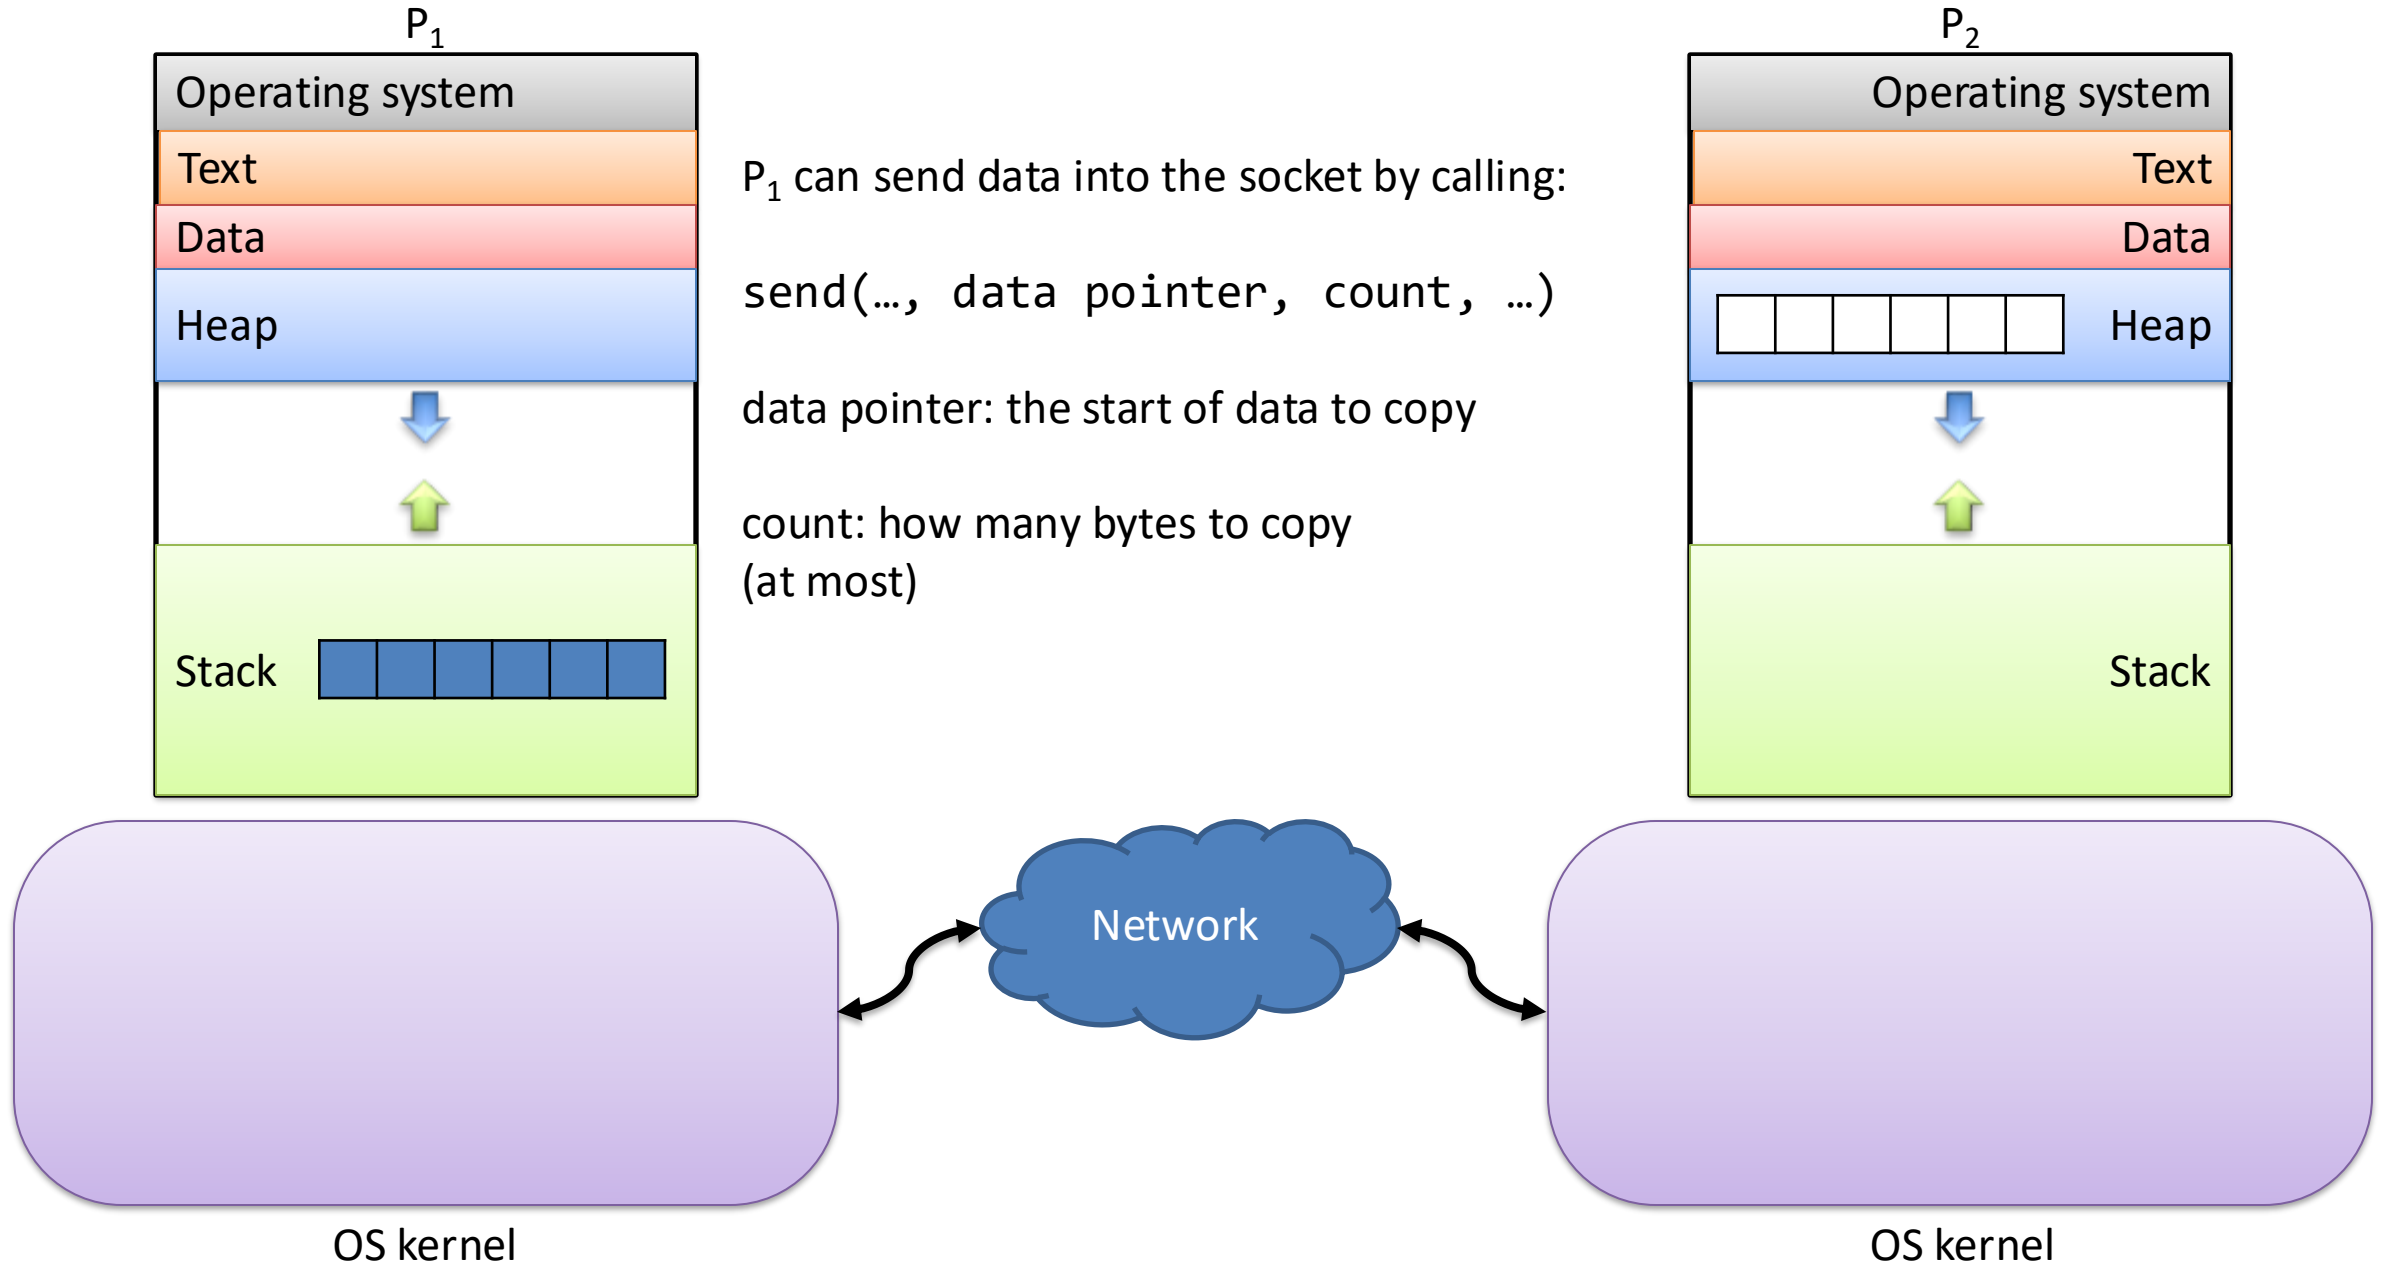
(*application synchronization)

- A. The OS adds synchronization.
- B. Synchronization is determined by the order of sends and receives.
- C. The communicating processes exchange synchronization messages (lock/unlock).
- D. There is no synchronization mechanism.

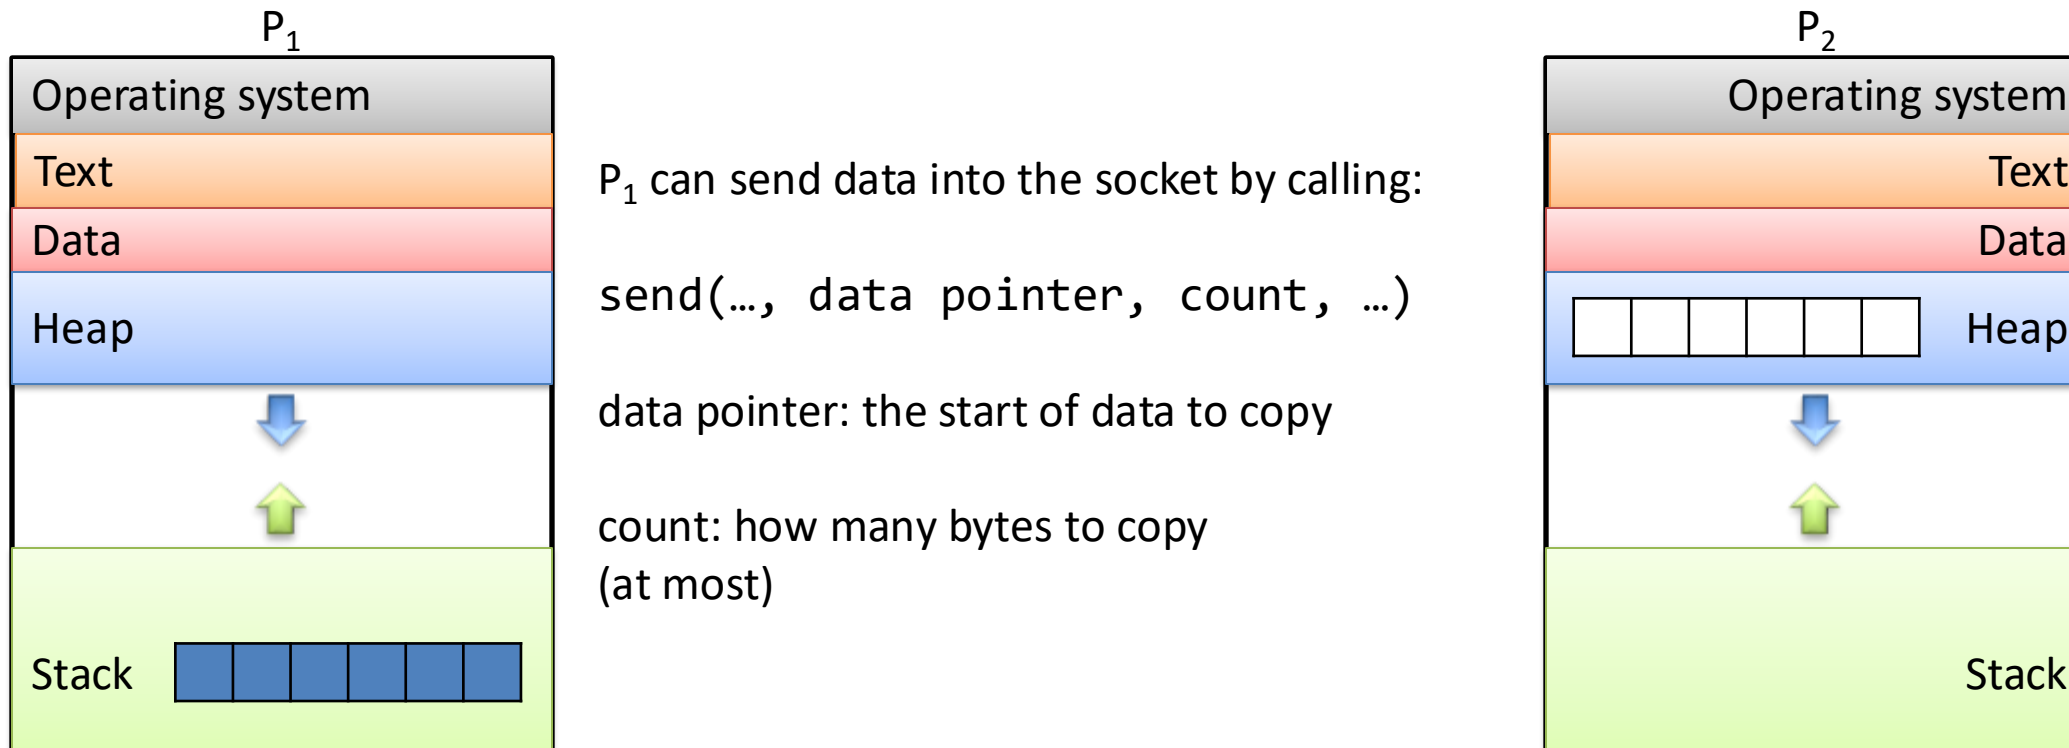
Message Passing IPC (Socket)



Message Passing IPC (Socket)



Message Passing IPC (Socket)



NAME

`send`, `sendto`, `sendmsg` – send a message on a socket

SYNOPSIS

```
#include <sys/types.h>
#include <sys/socket.h>
```

```
ssize_t send(int sockfd, const void *buf, size_t len, int flags);
```

Message Passing IPC (Socket)

NAME

write - write to a file descriptor

SYNOPSIS

```
#include <unistd.h>
```

```
ssize_t write(int fd, const void *buf, size_t count);
```

DESCRIPTION

write() writes up to count bytes from the buffer starting at buf to the file referred to by the file descriptor fd.

NAME

send, sendto, sendmsg - send a message on a socket

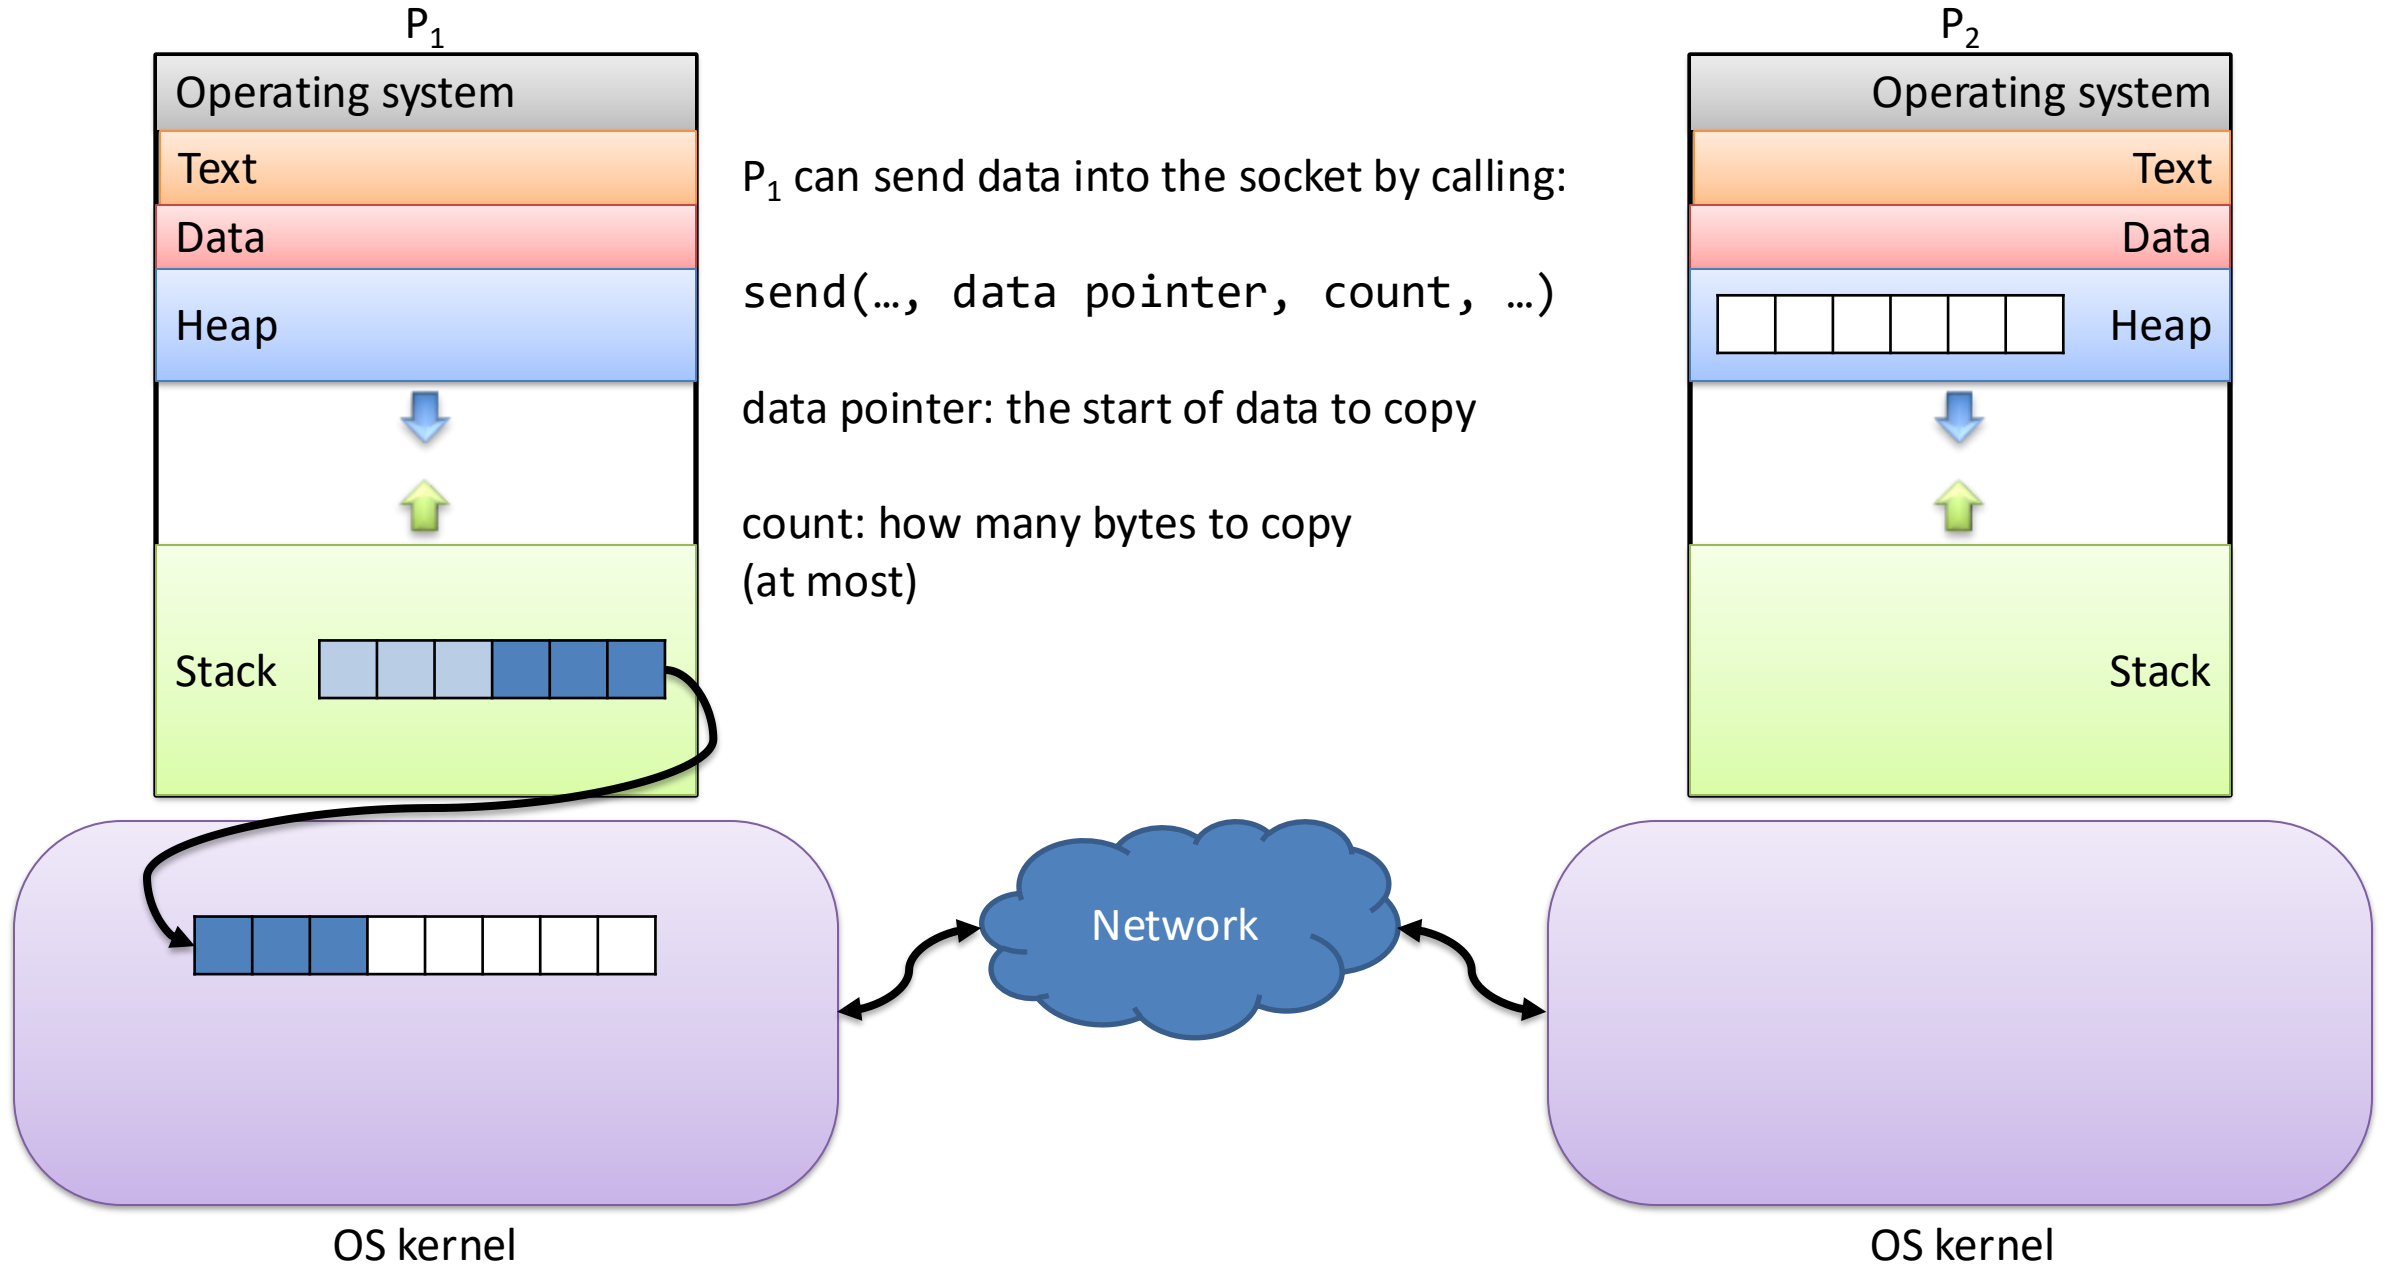
SYNOPSIS

```
#include <sys/types.h>
```

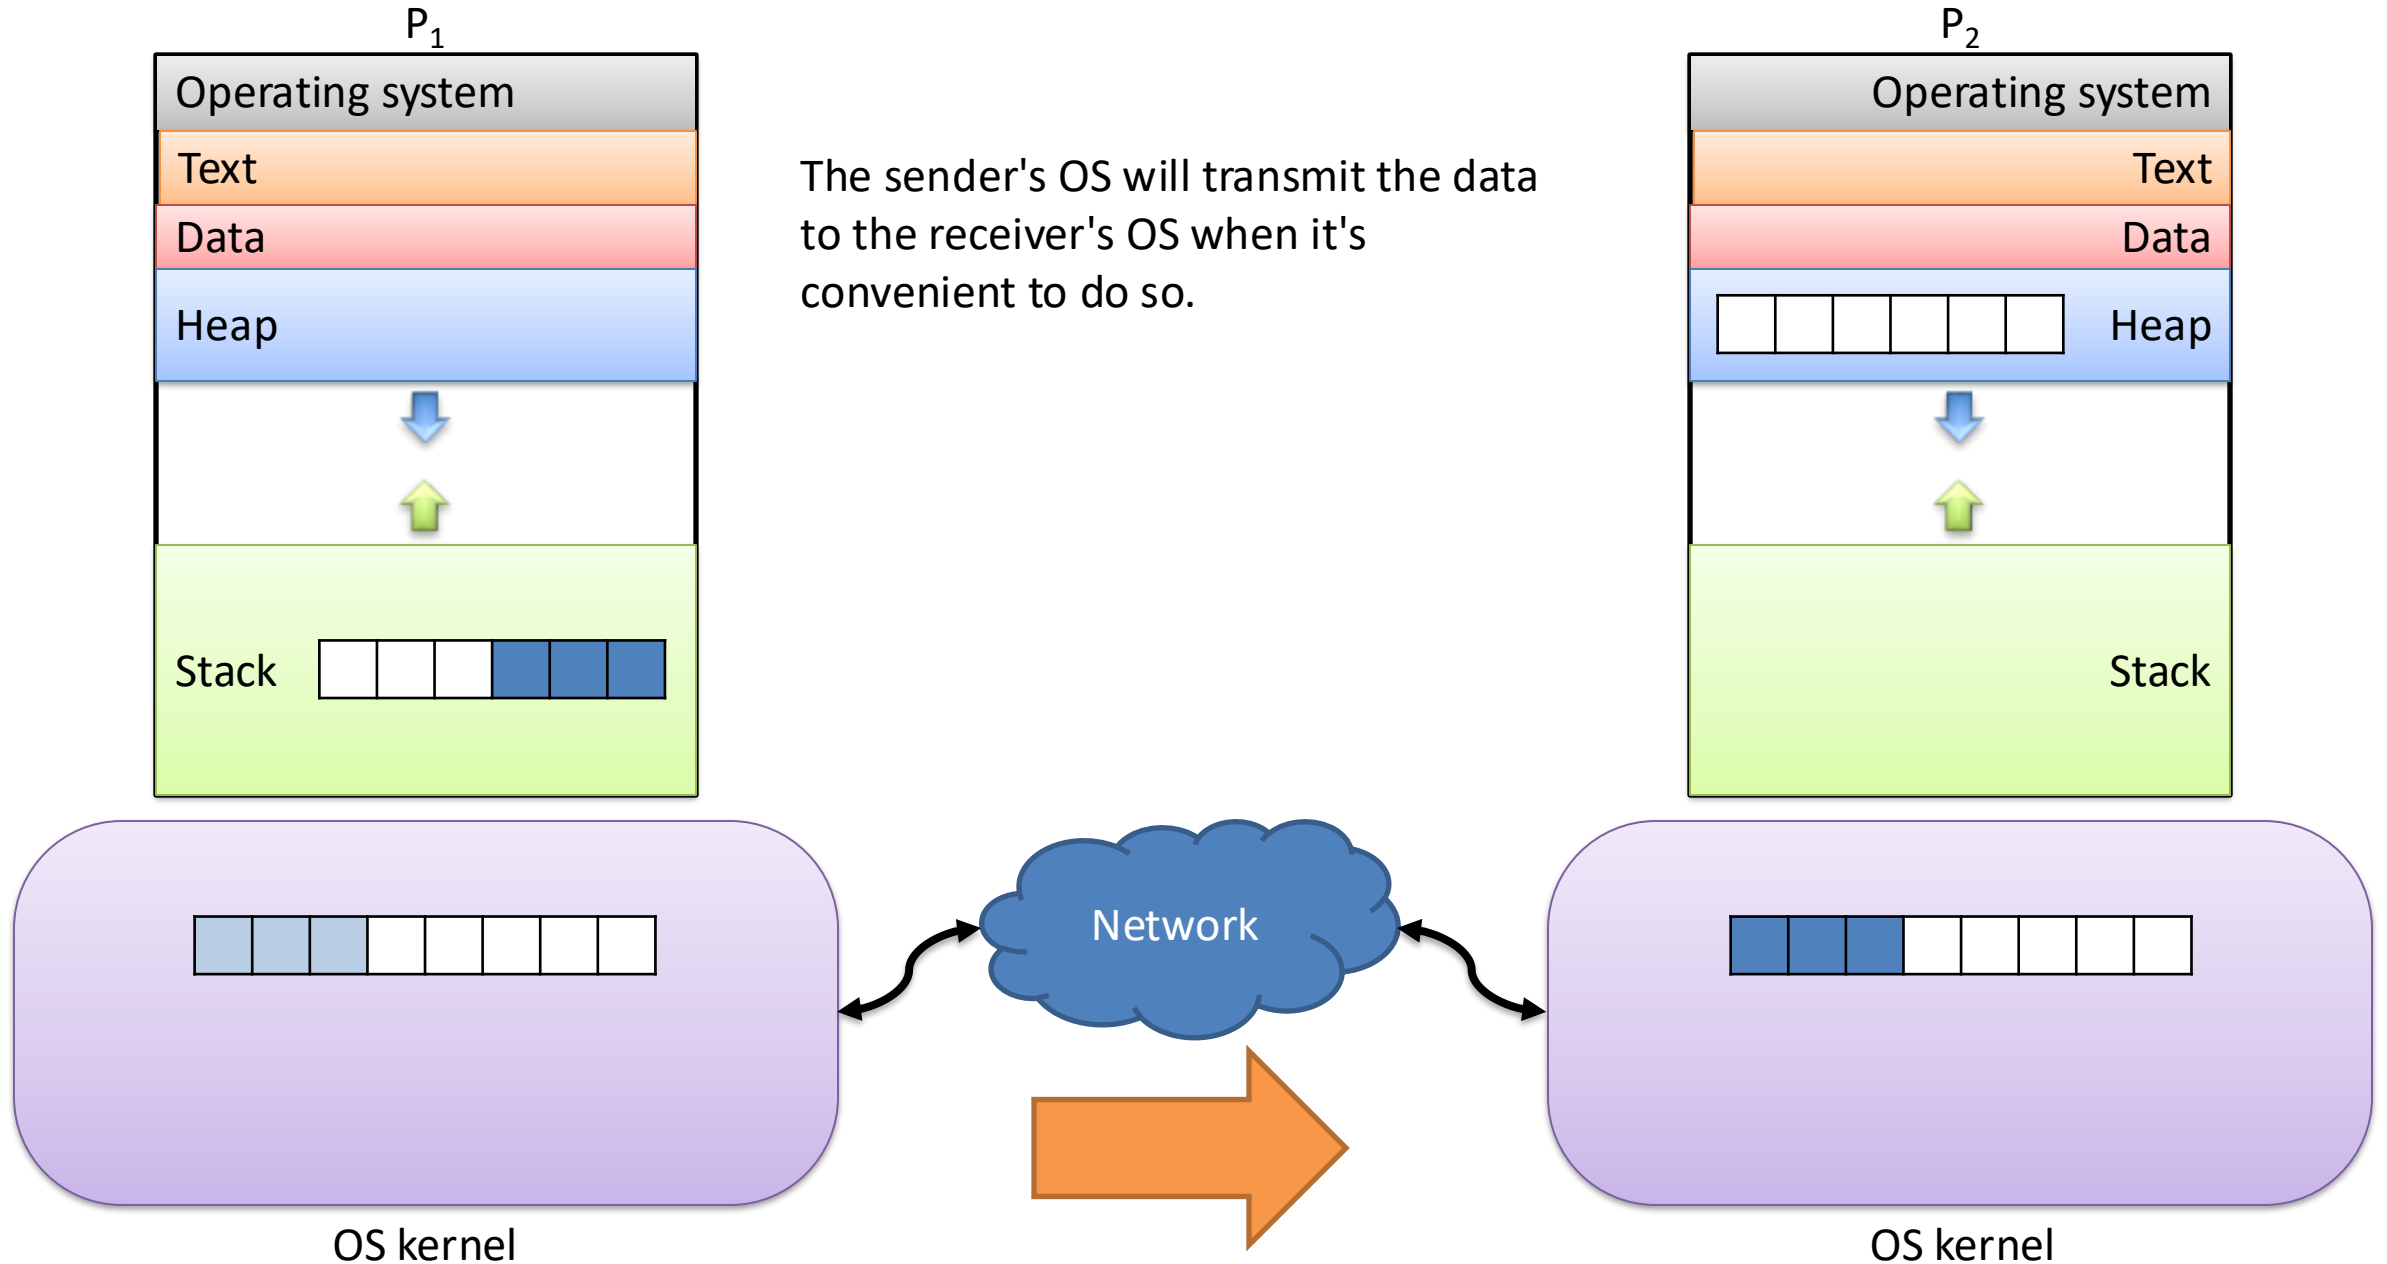
```
#include <sys/socket.h>
```

```
ssize_t send(int sockfd, const void *buf, size_t len, int flags);
```

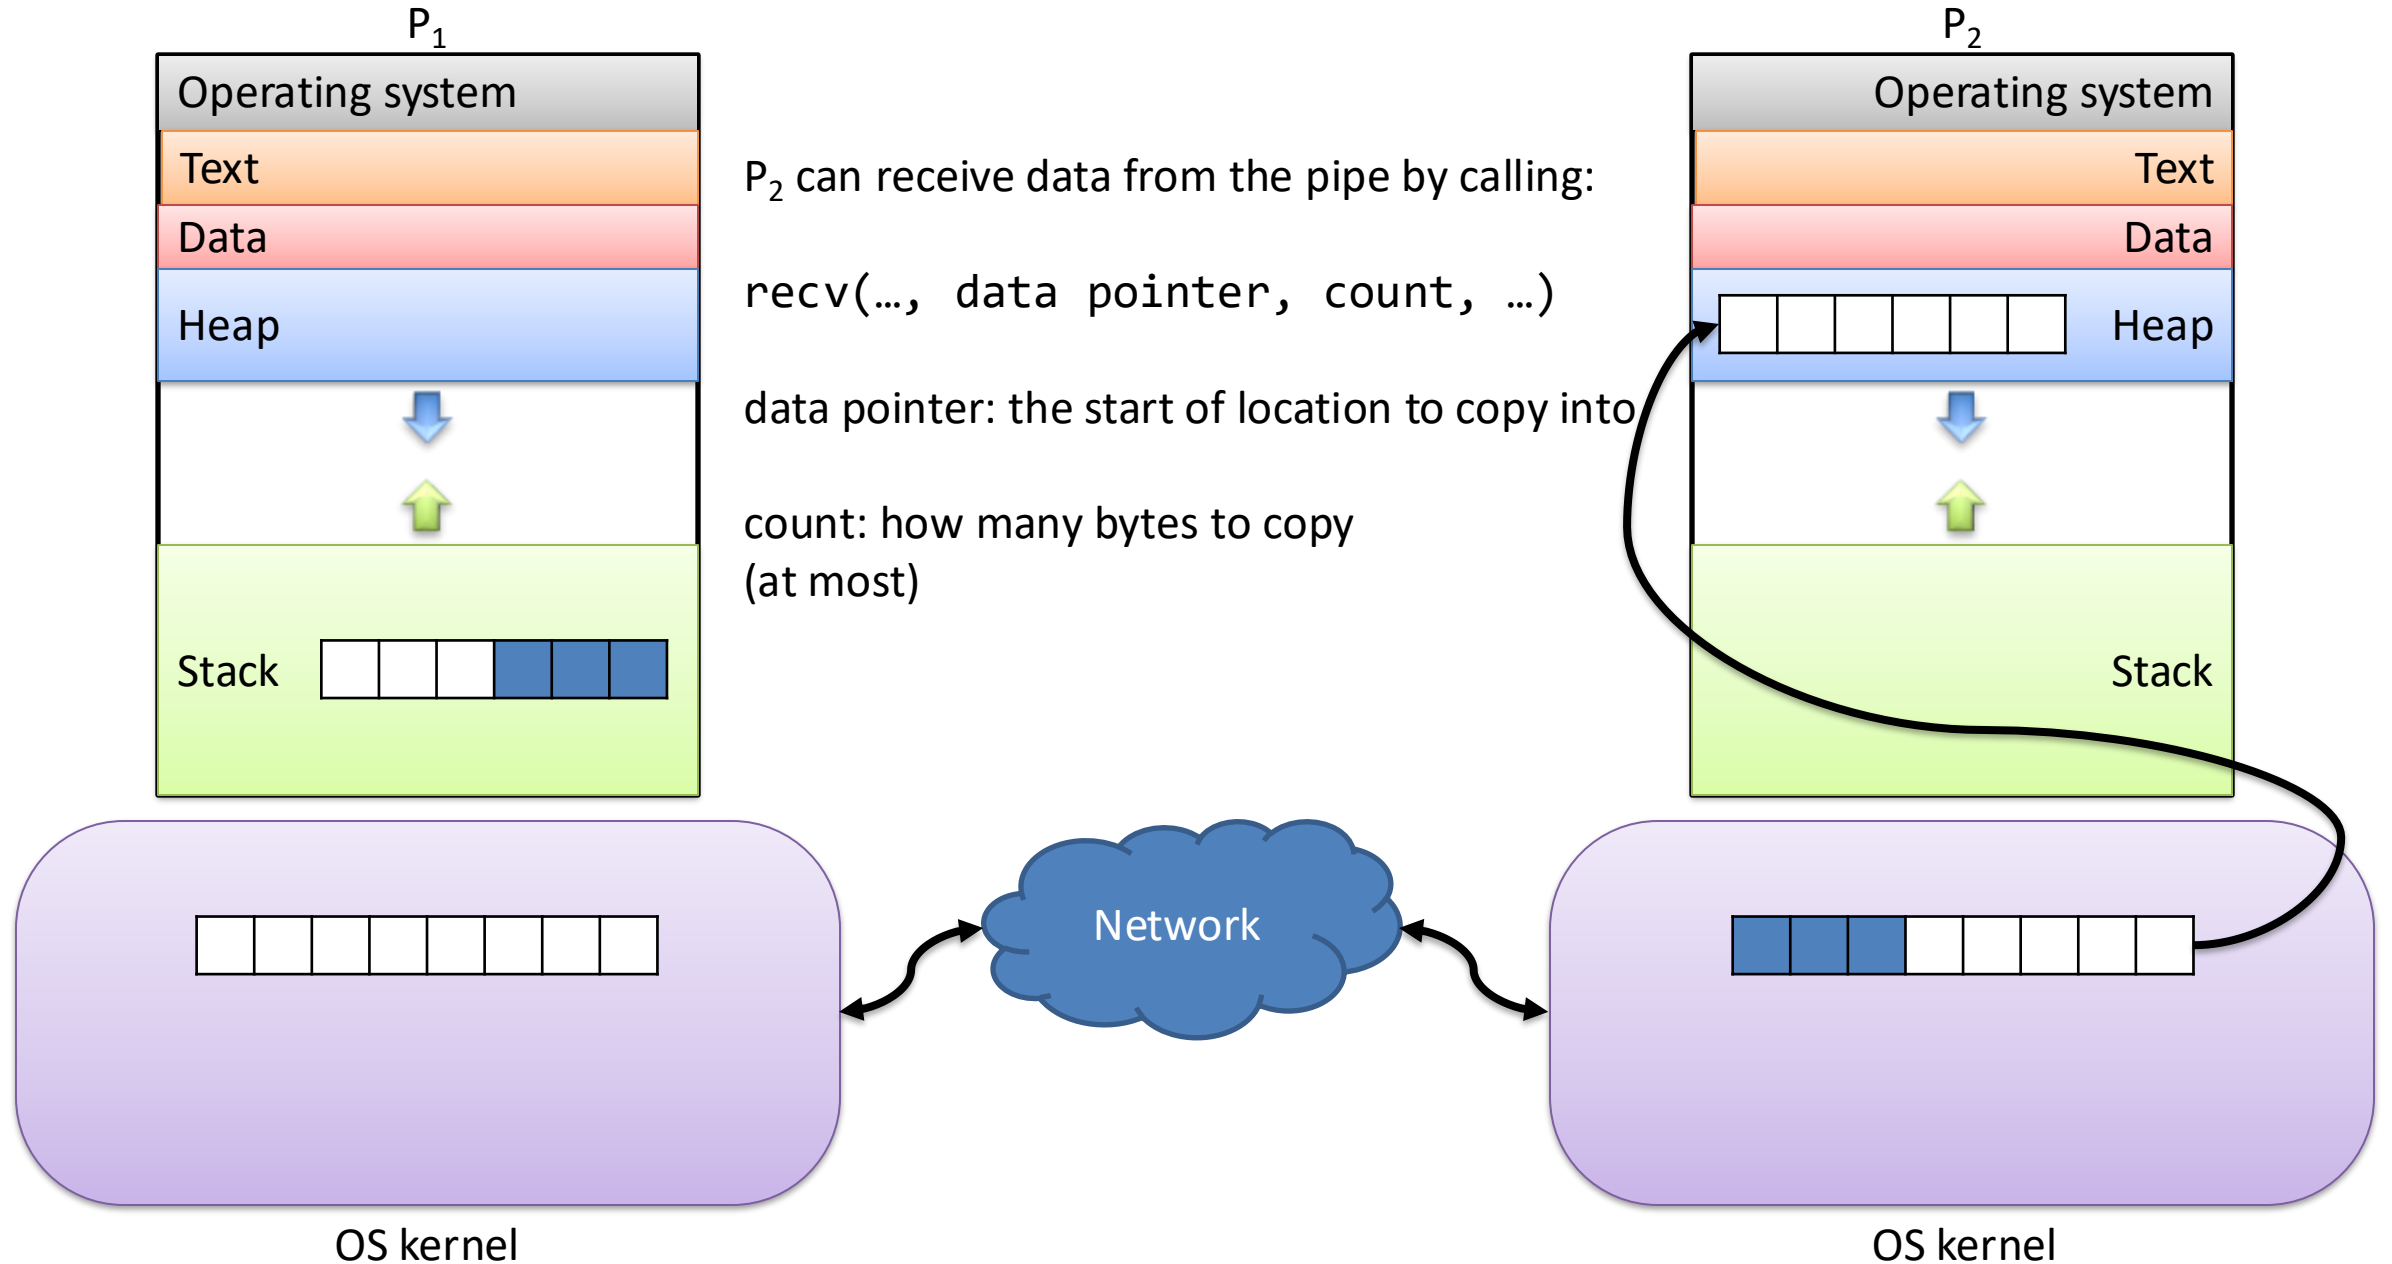

Message Passing IPC (Socket)



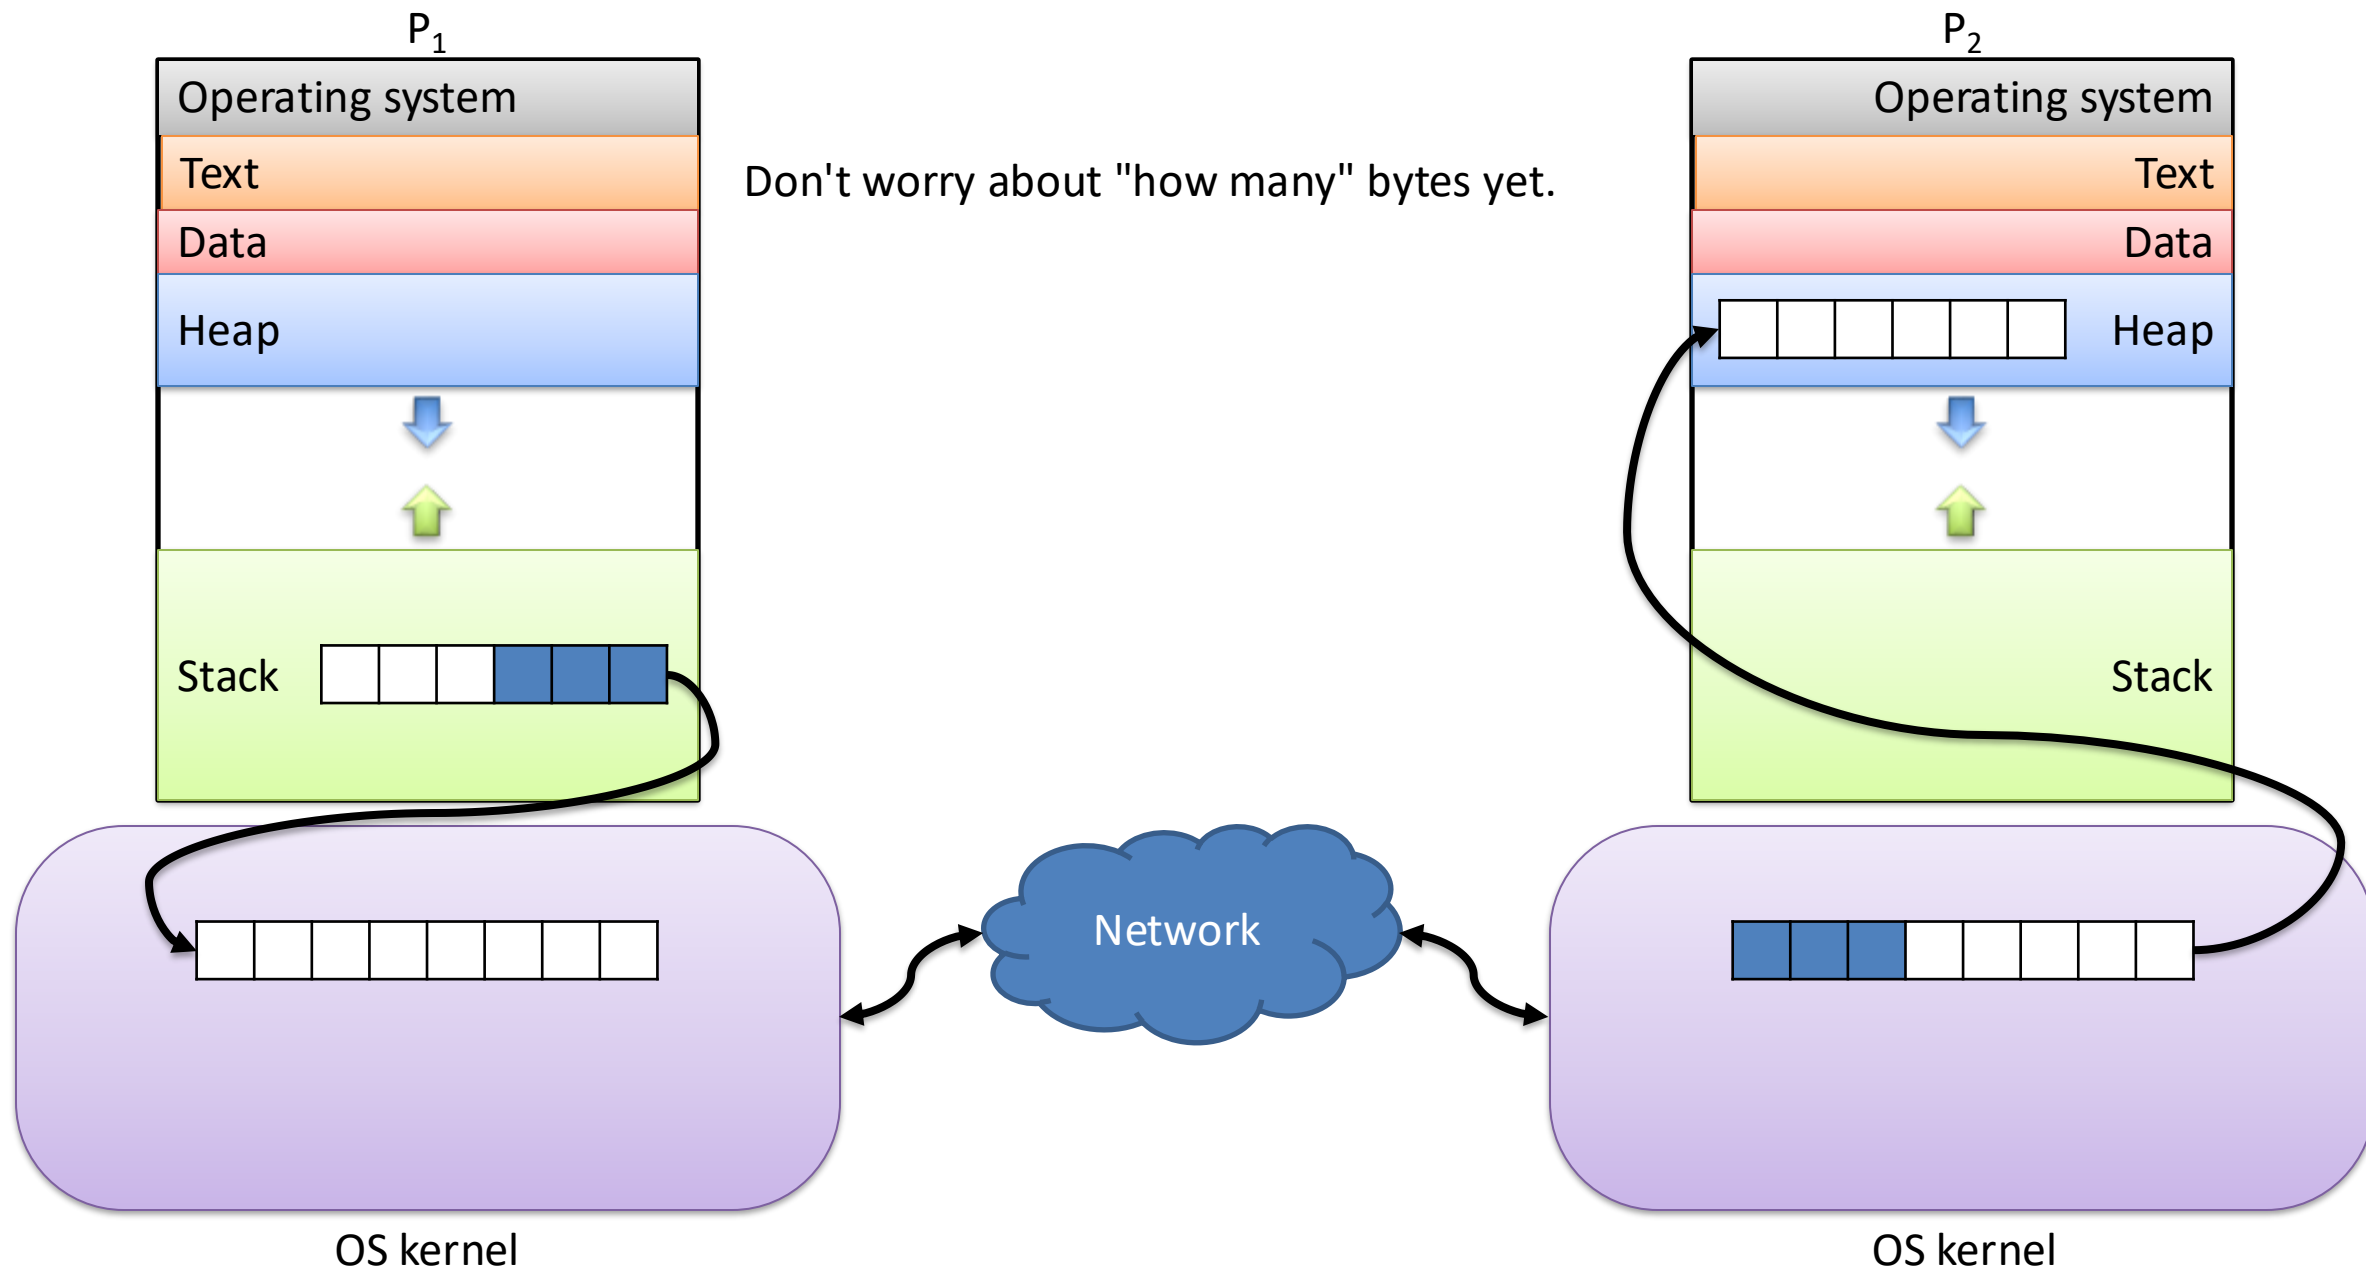
Message Passing IPC (Socket)



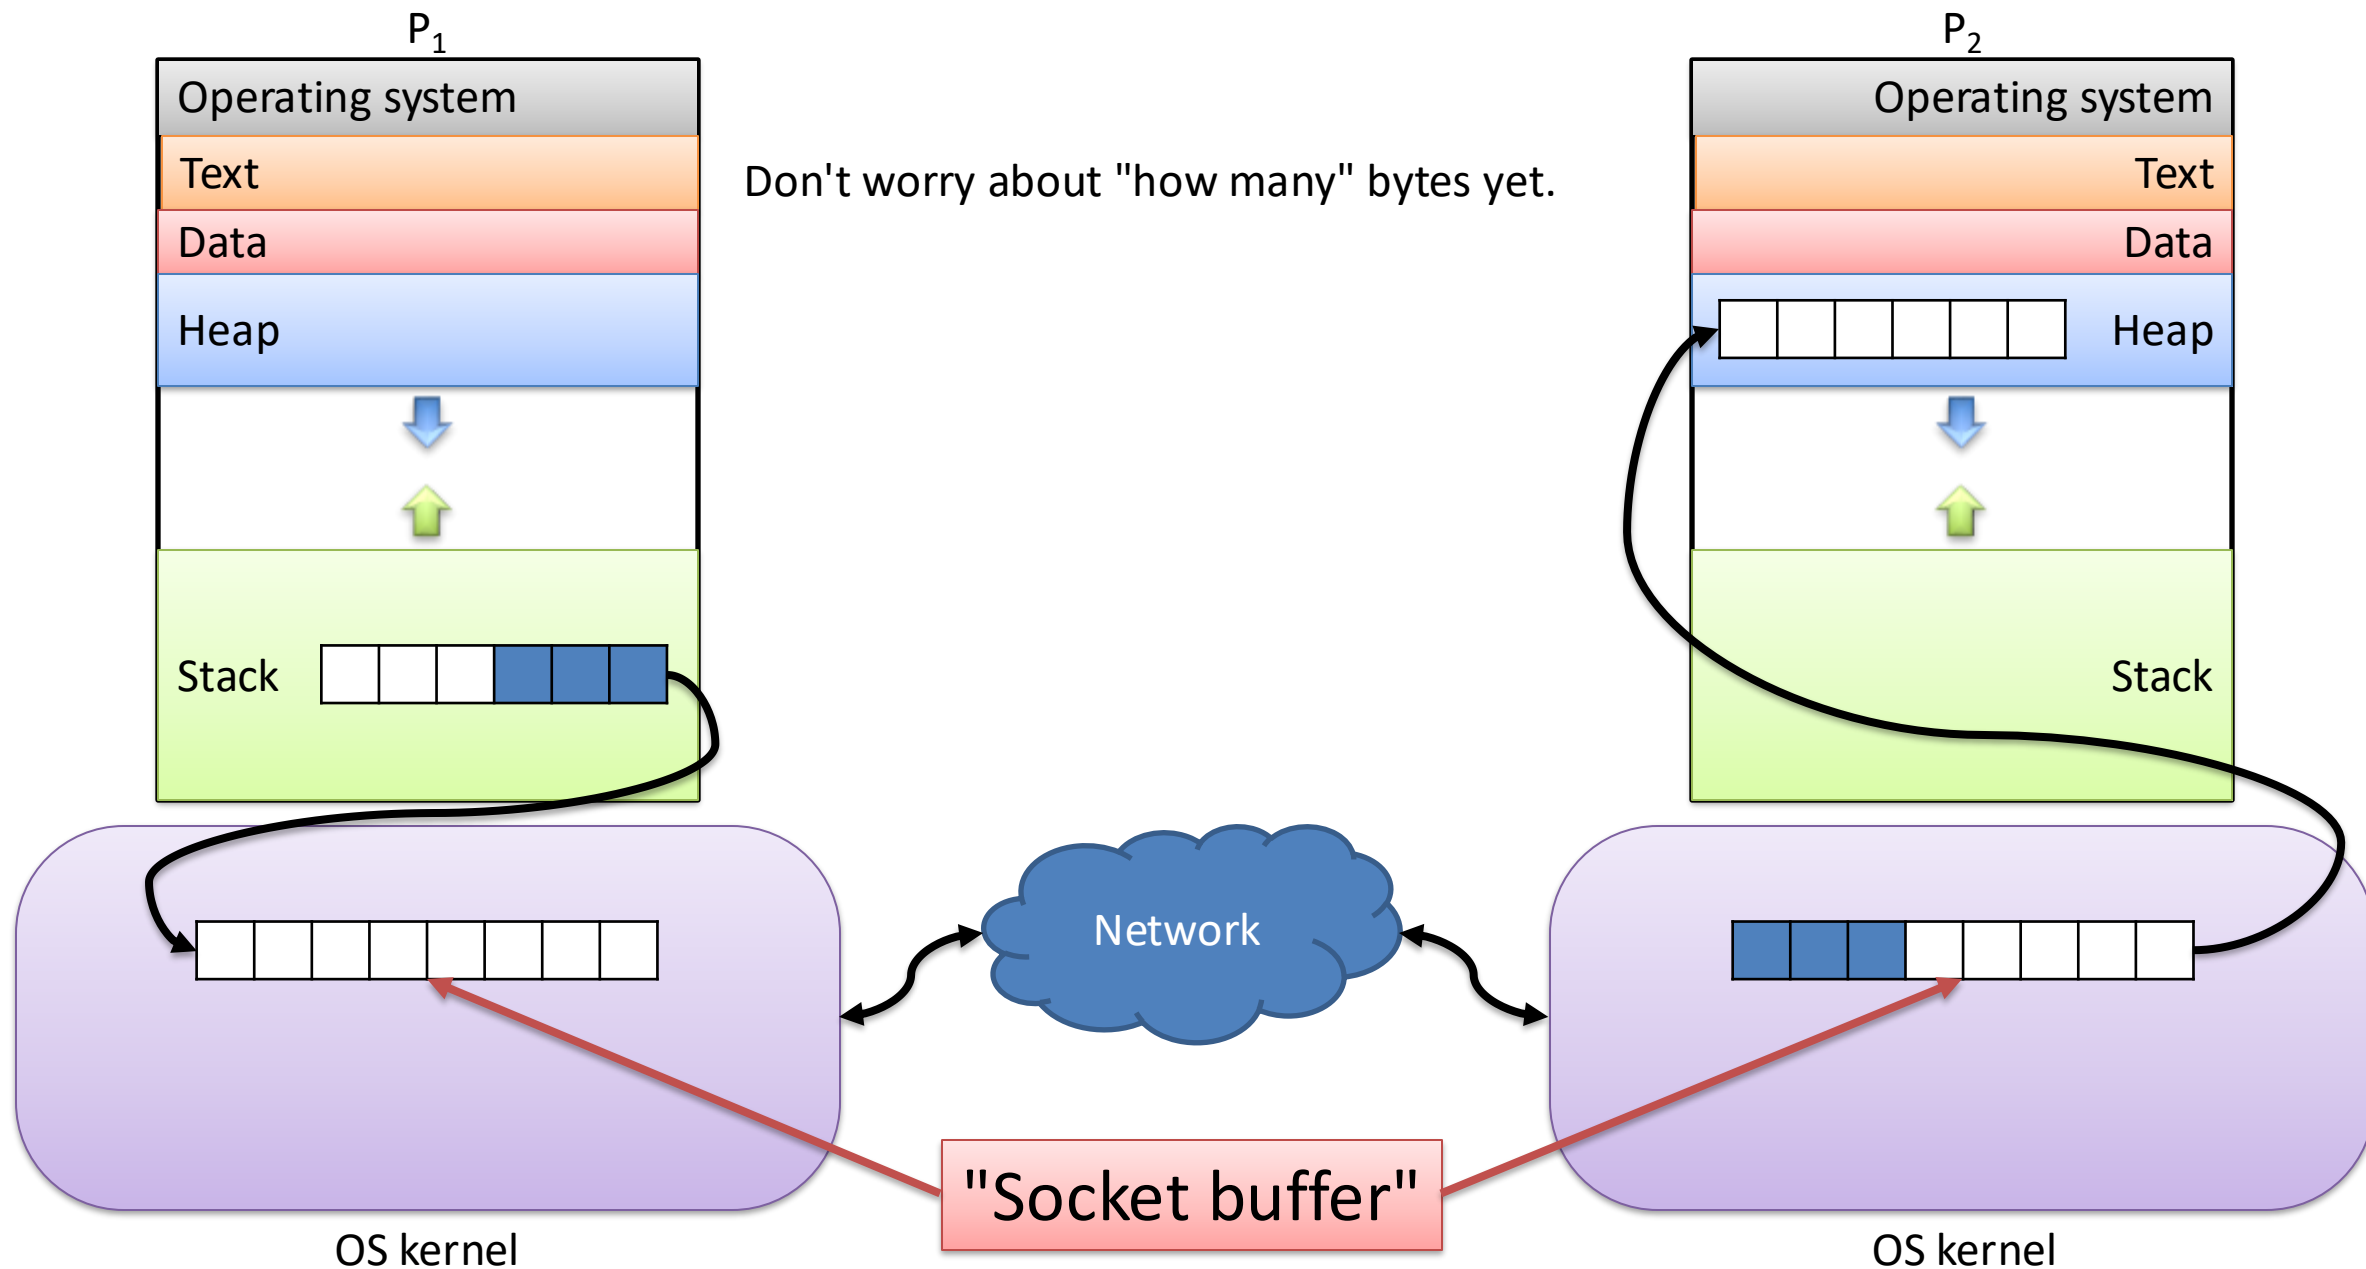
Message Passing IPC (Socket)



Questions about this model?

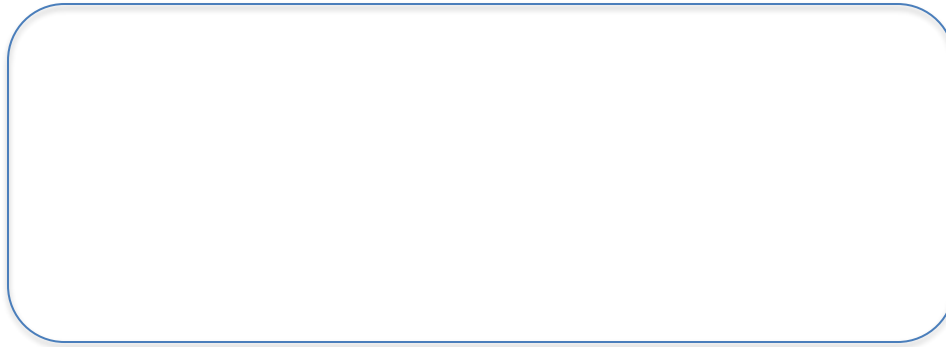


Questions about this model?

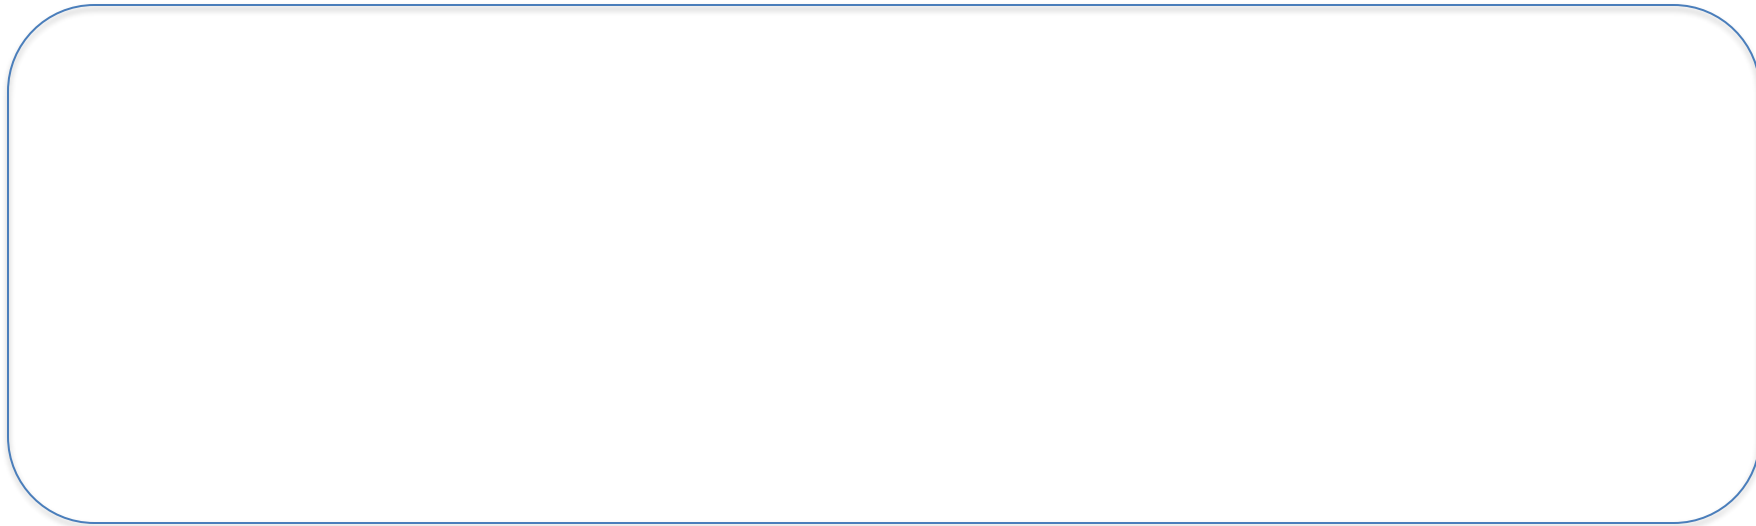


Descriptor Table

Process



- OS stores a table, per process, of descriptors



Kernel

Descriptors

Where do descriptors come from?

What are they?

```
OPEN(2)                                Linux Programmer's Manual
OPEN(2)

NAME
    open, openat, creat - open and possibly create a file

SYNOPSIS
#include <sys/types.h>
#include <sys/stat.h>
#include <fcntl.h>

int open(const char *pathname, int flags);
int open(const char *pathname, int flags, mode_t mode);
```

```
SOCKET(2)                               Linux Programmer's Manual    SOCKET(2)

NAME
    socket - create an endpoint for communication

SYNOPSIS
#include <sys/types.h>           /* See NOTES */
#include <sys/socket.h>

int socket(int domain, int type, int protocol);

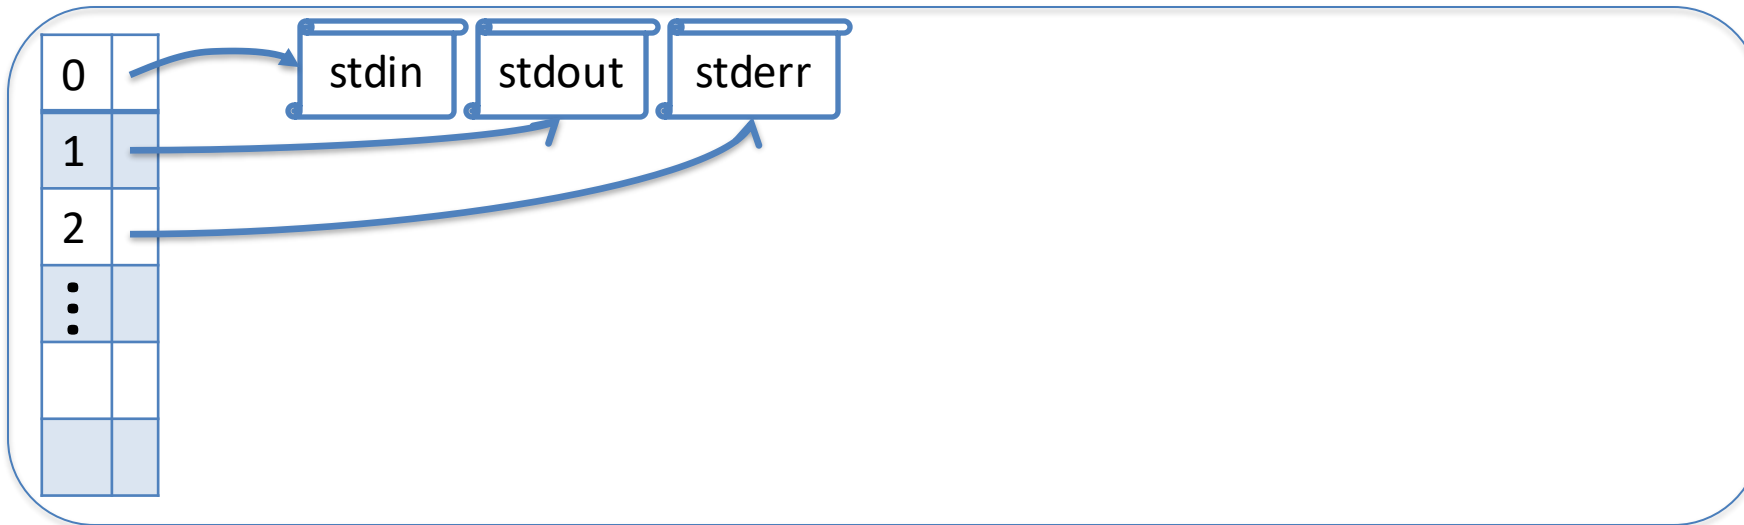
DESCRIPTION
    socket() creates an endpoint for communication and
    returns a descriptor.
```

Descriptor Table

Process



- OS stores a table, per process, of descriptors



Kernel

socket()

Process

```
int sock = socket(AF_INET, SOCK_STREAM, 0);
```

7

- `socket()` returns a socket descriptor
- Indexes into table

0	
1	
2	
⋮	
7	

stdin stdout stderr

Kernel

socket()

Process

```
int sock = socket(AF_INET, SOCK_STREAM, 0);
```

7

- OS stores details of the socket, connection, and pointers to buffers

0	
1	
2	
⋮	
7	

stdin

stdout

stderr

Family: AF_INET, Type: SOCK_STREAM
Local address: NULL, Local port: NULL
Send buffer , Receive buffer

Kernel

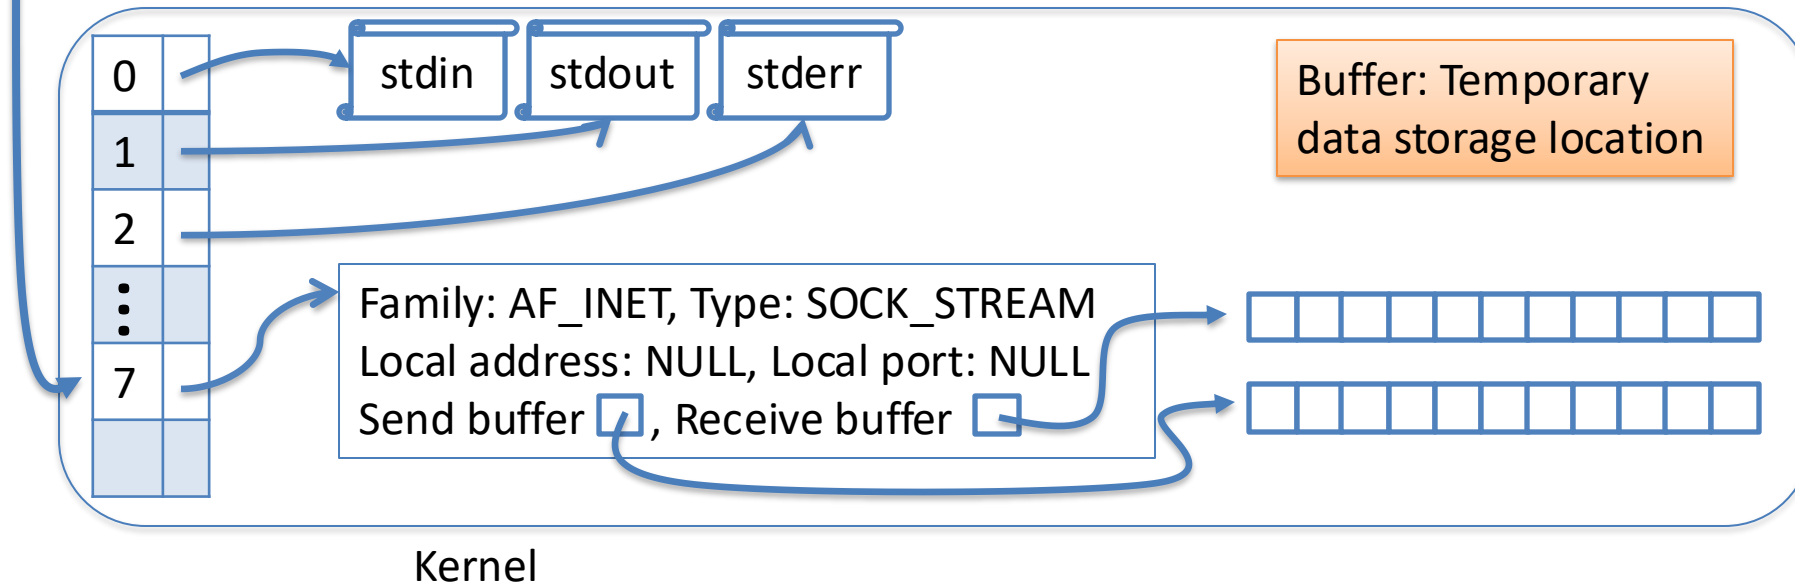
socket()

Process

```
int sock = socket(AF_INET, SOCK_STREAM, 0);
```

7

- OS stores details of the socket, connection, and pointers to buffers



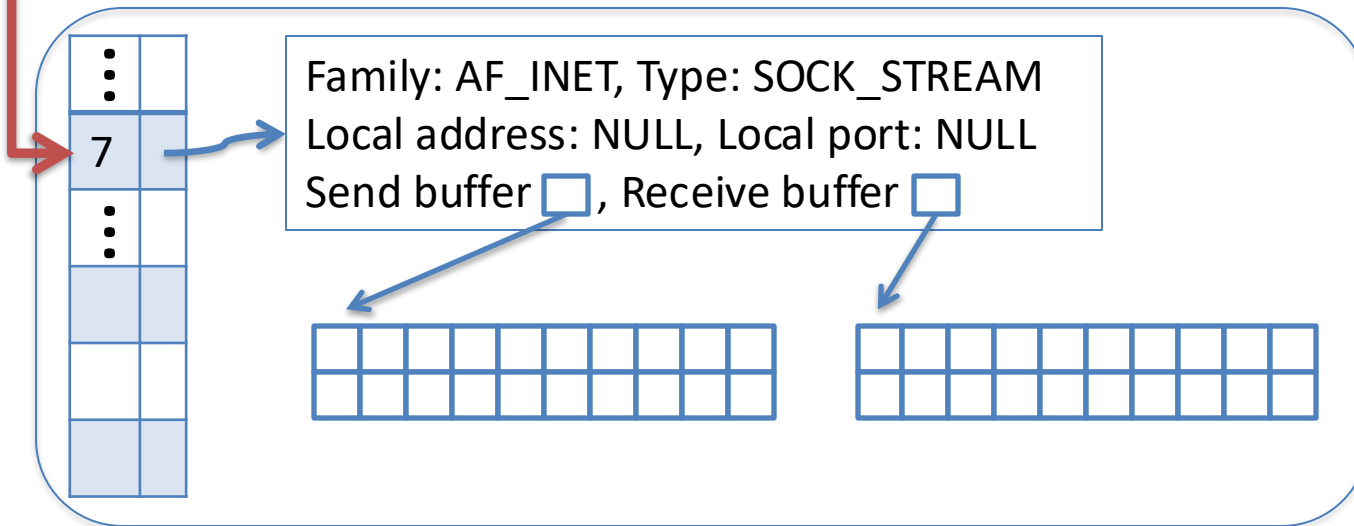
Socket Buffers

Process

```
int sock = socket(AF_INET, SOCK_STREAM, 0);
```

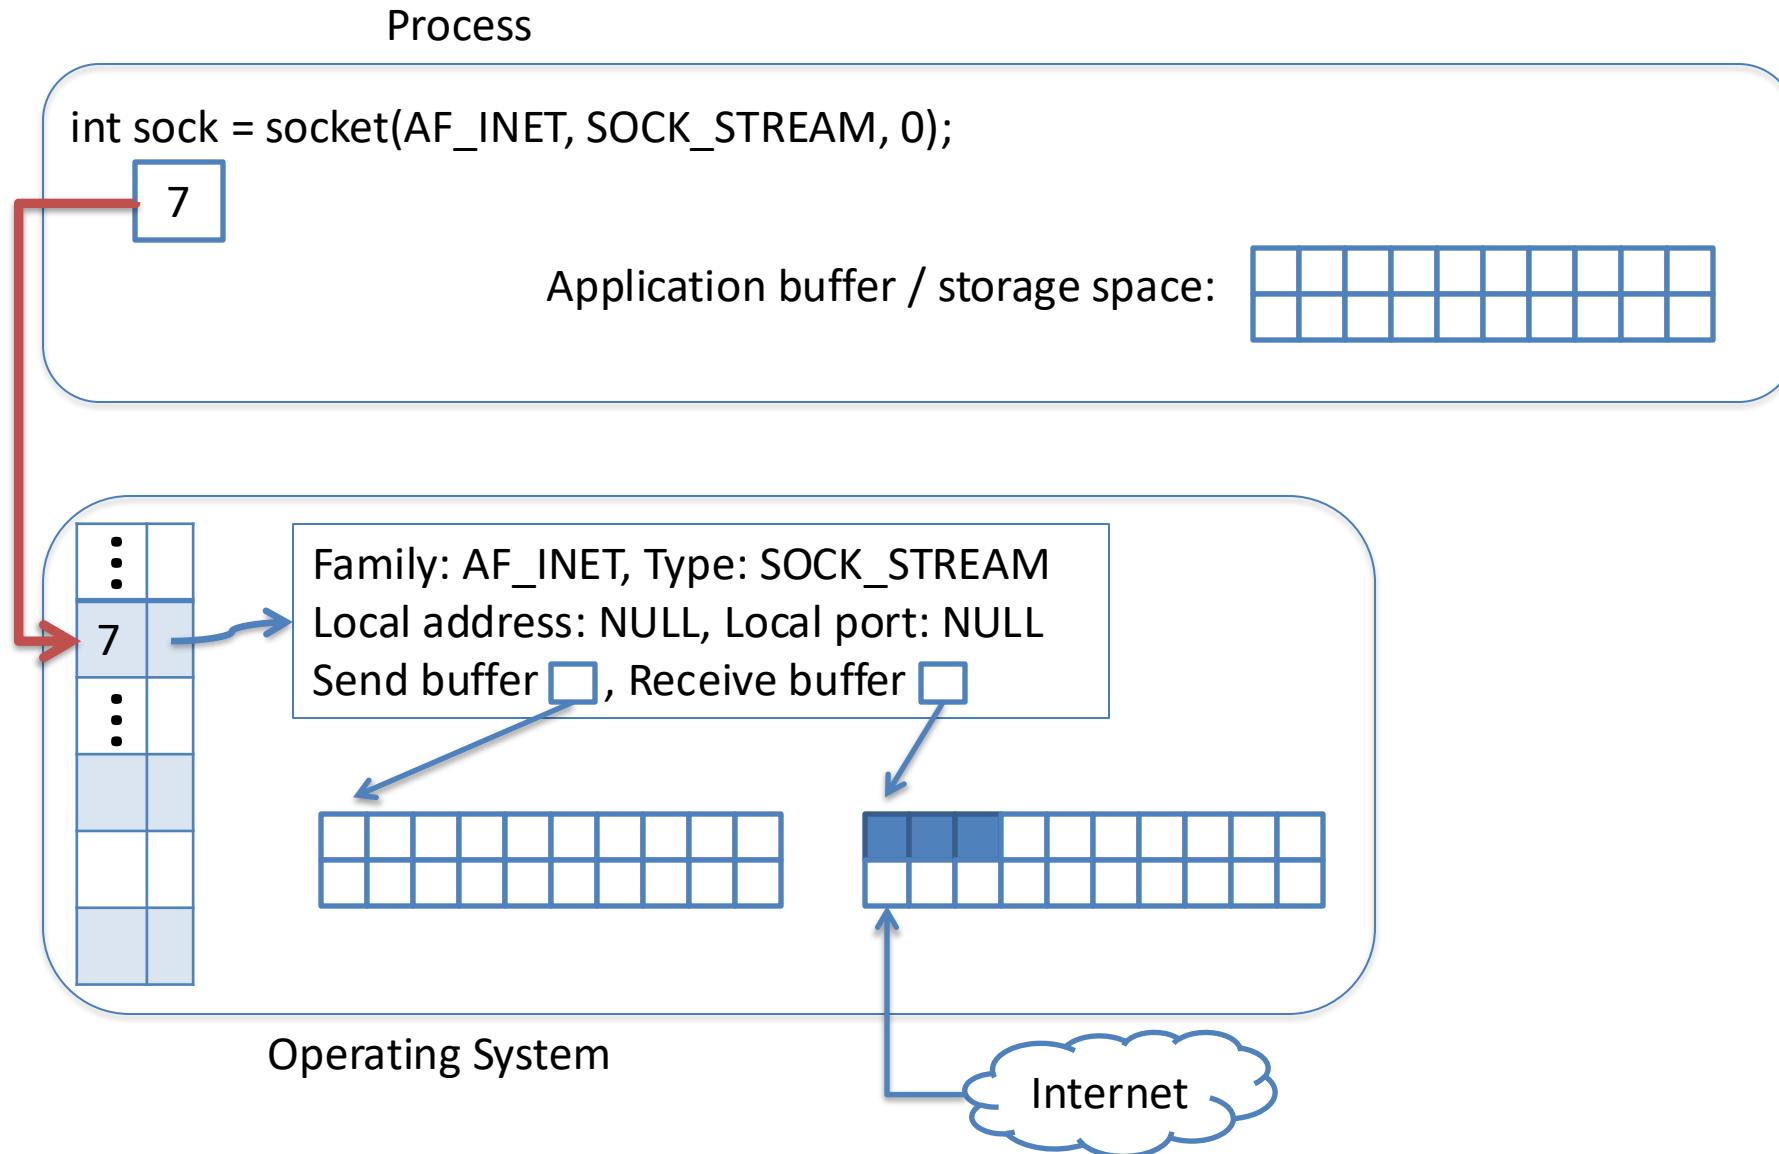
7

Application buffer / storage space:

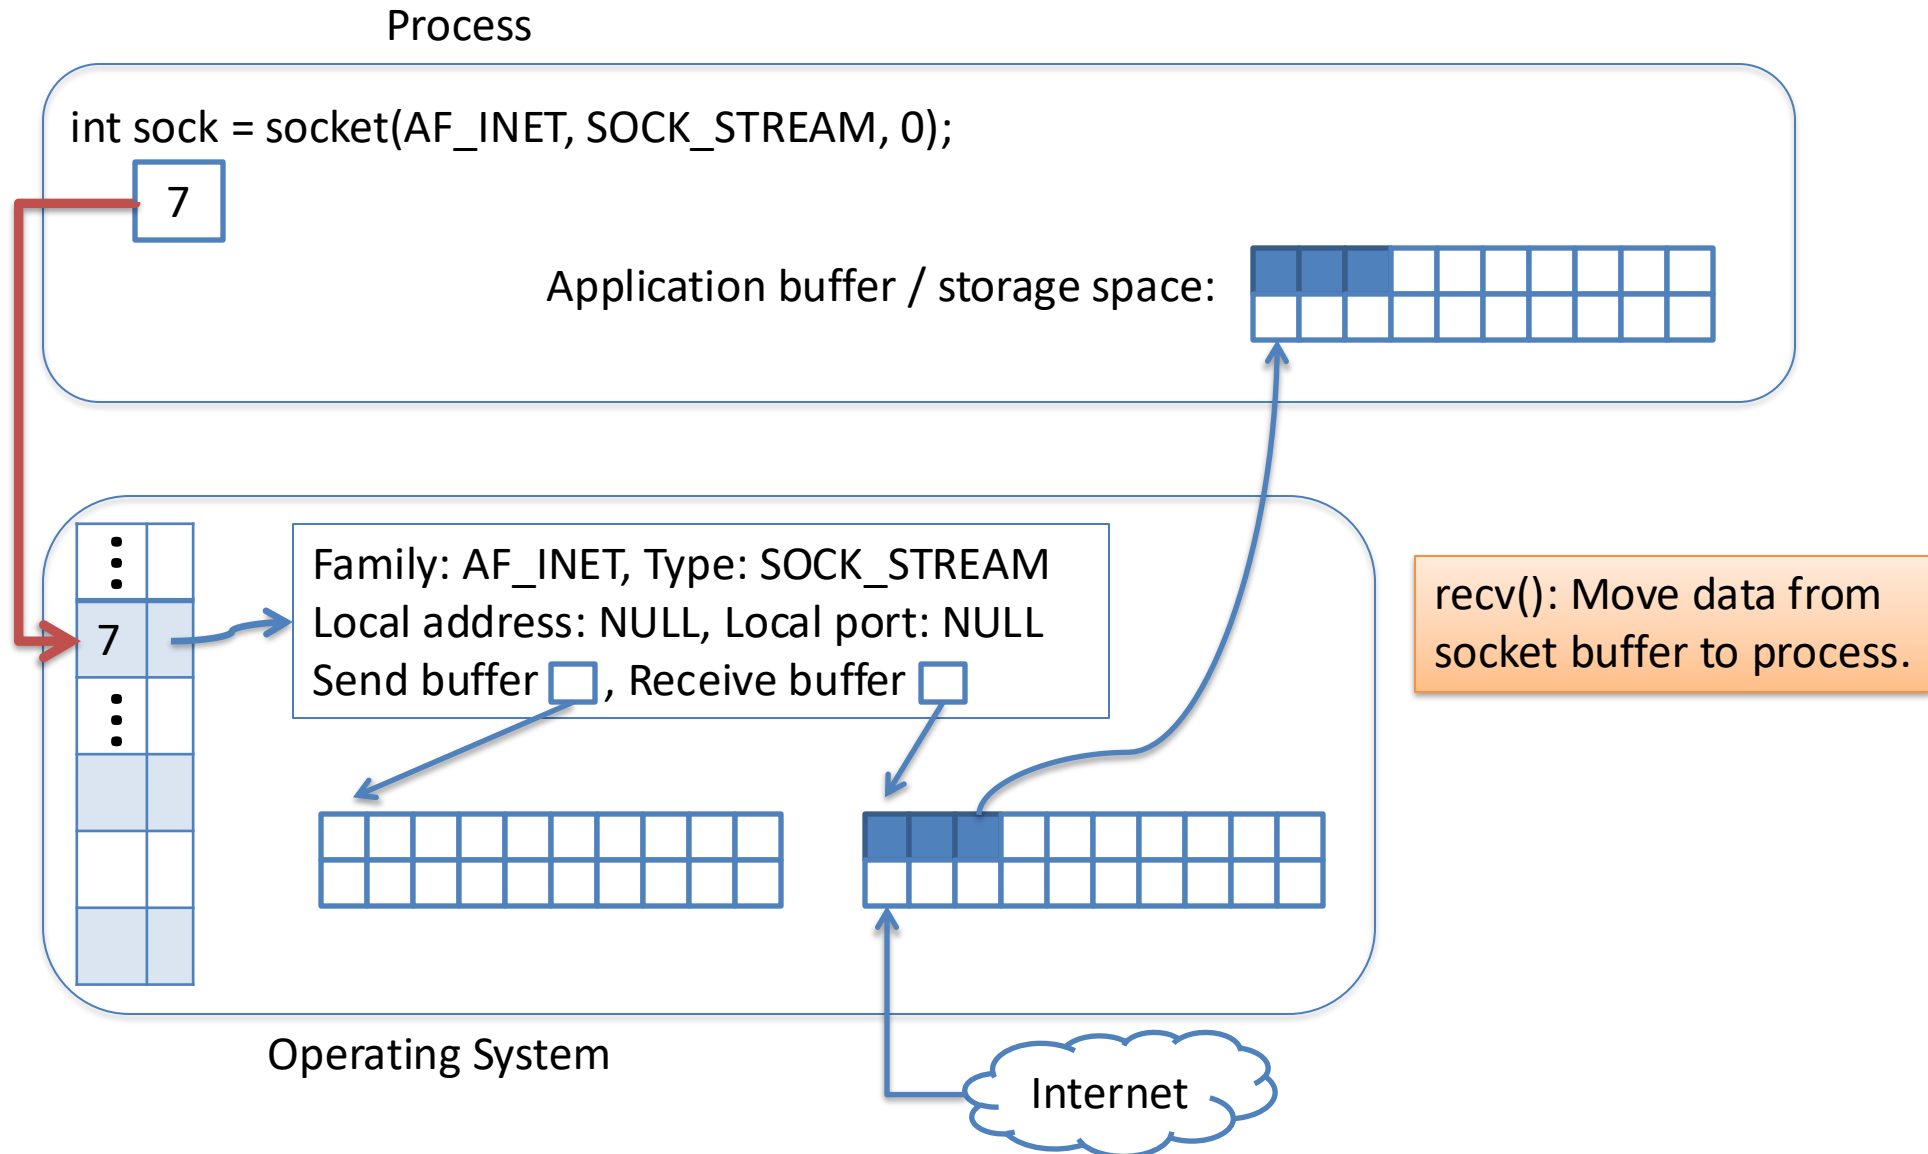


Operating System

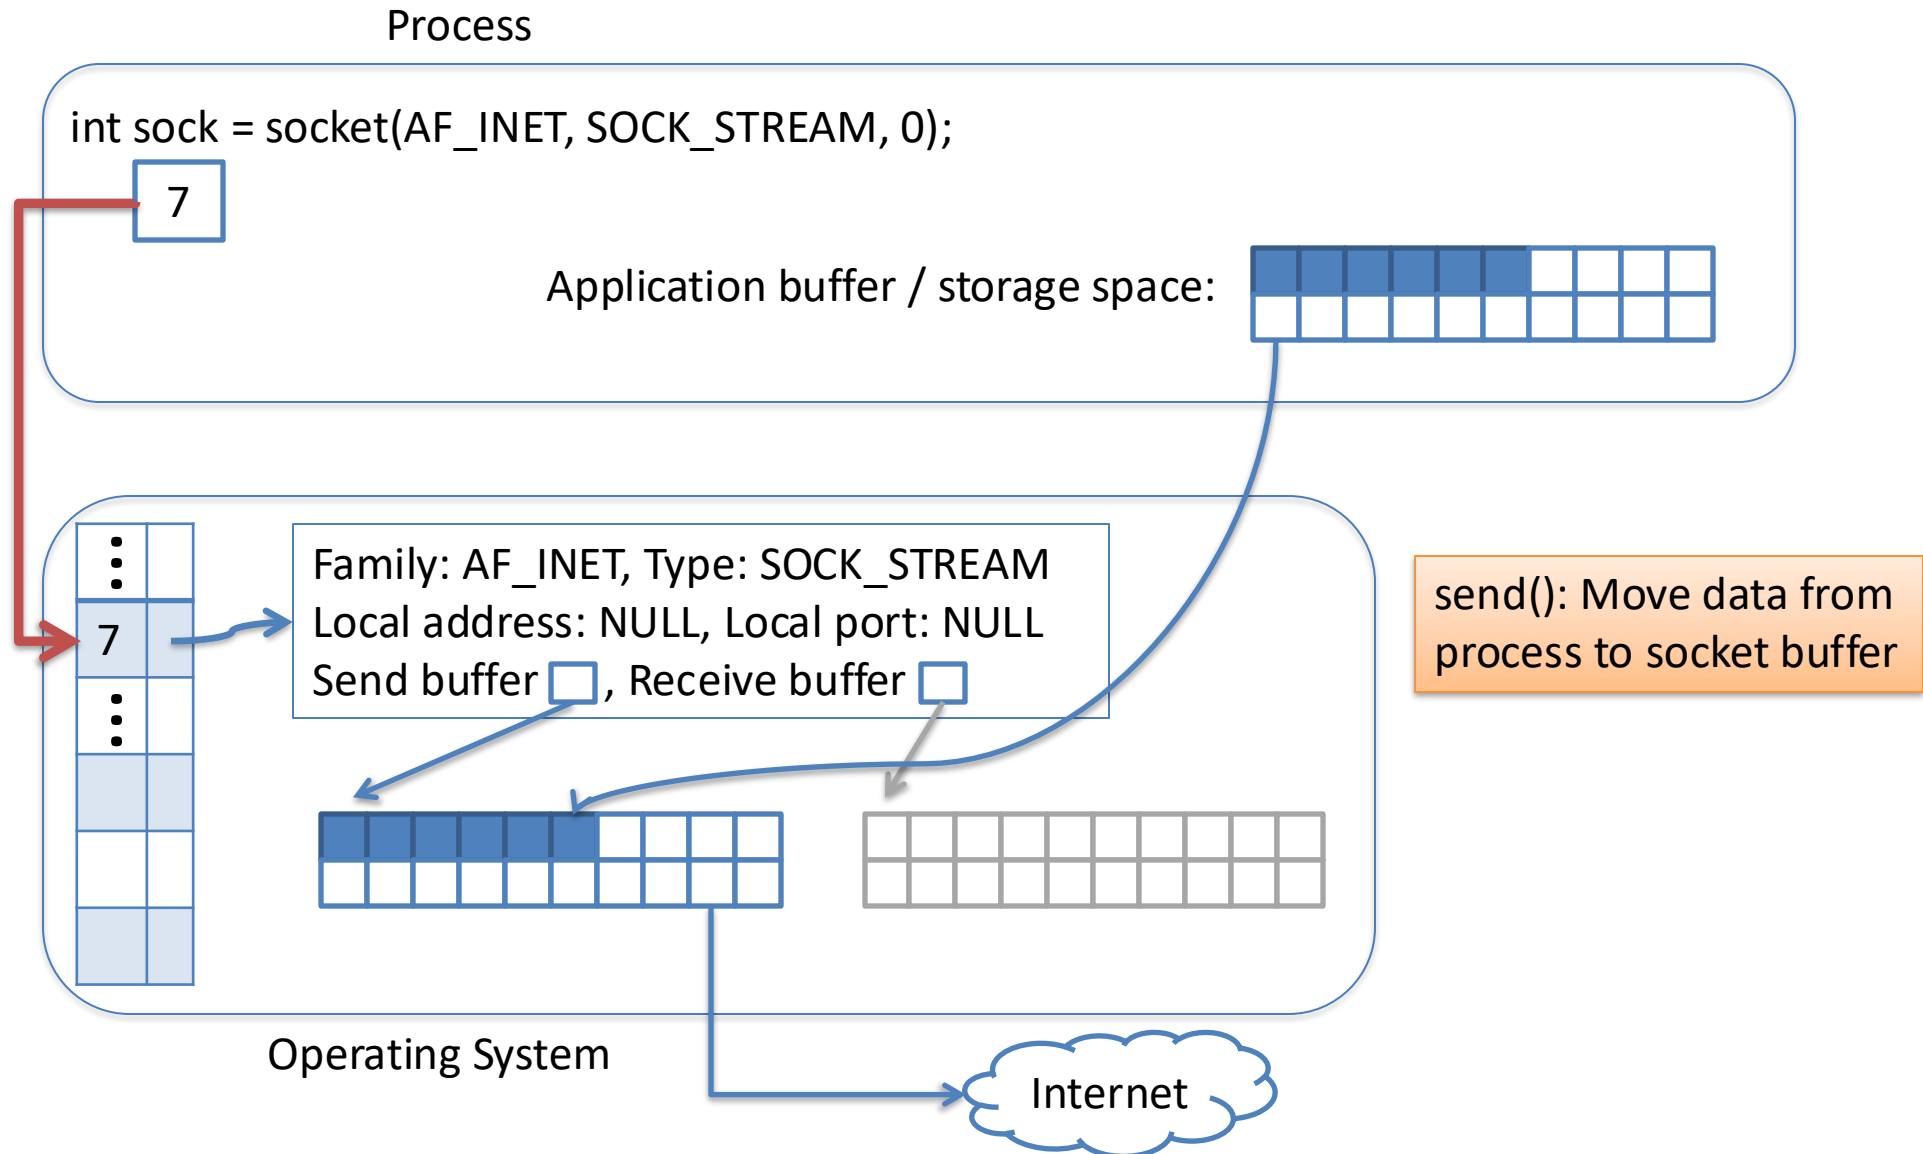
Socket Buffers



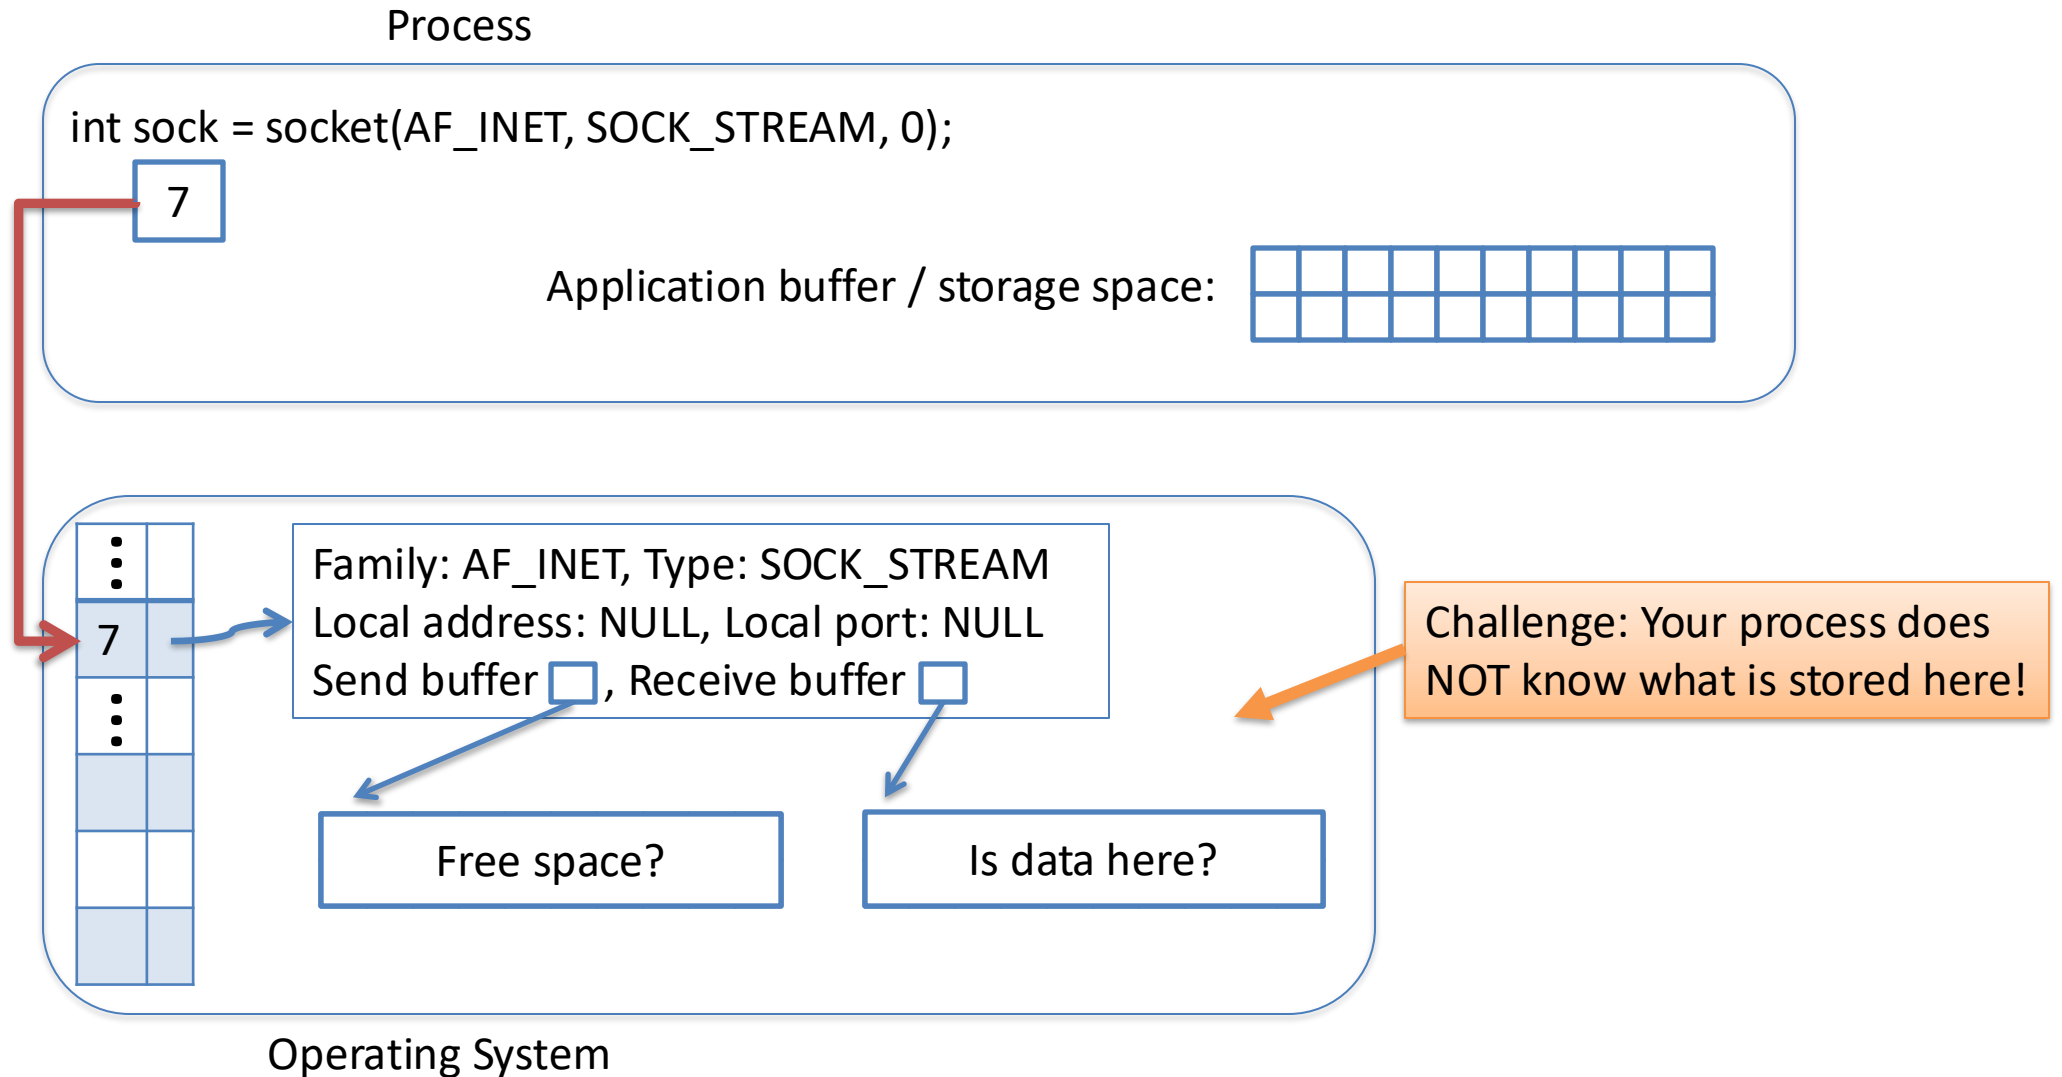
Socket Buffers



Socket Buffers



Socket Buffers



recv()

Process

```
int sock = socket(AF_INET, SOCK_STREAM, 0);  
    (assume we connect()ed here...)  
int recv_val = recv(sock, r_buf, 200, 0);
```

r_buf (size 200)



0	
1	
2	
⋮	
7	

Family: AF_INET, Type: SOCK_STREAM
Local address: ..., Local port: ...
Send buffer ☐, Receive buffer ☐



Is data here?

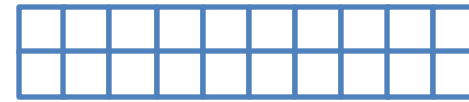
Kernel

What should we do if the receive socket buffer is empty? If it has 100 bytes?

Process

```
int sock = socket(AF_INET, SOCK_STREAM, 0);  
    (assume we connect()ed here...)  
int recv_val = recv(sock, r_buf, 200, 0);
```

r_buf (size 200)



Two Scenarios:

Socket buffer (receive)



Empty



100 bytes

Kernel

What should we do if the receive socket buffer is empty? If it has 100 bytes?

Process

```
int sock = socket(AF_INET, SOCK_STREAM, 0);  
    (assume we connect()ed here...)  
int rcv_val = recv(sock, r_buf, 200, 0);
```

r_buf (size 200)



Two Scenarios:

	Empty	100 Bytes
A	Block	Block
B	Block	Copy 100 bytes
C	Copy 0 bytes	Block
D	Copy 0 bytes	Copy 100 bytes
E	Something else	

Socket buffer (receive)



Empty



100 bytes

Kernel

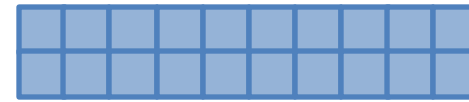
"Block" means pause the calling process.

What should we do if the send socket buffer is full? If it has 100 bytes?

Process

```
int sock = socket(AF_INET, SOCK_STREAM, 0);  
    (assume we connect()ed here...)  
int send_val = send(sock, s_buf, 200, 0);
```

s_buf (size 200)



Two Scenarios:

Socket buffer (send)



Full



100 bytes

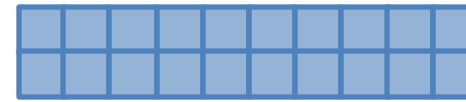
Kernel

What should we do if the send socket buffer is full? If it has 100 bytes?

Process

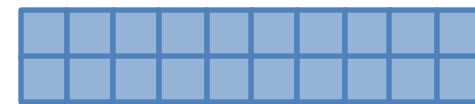
```
int sock = socket(AF_INET, SOCK_STREAM, 0);  
    (assume we connect()ed here...)  
int send_val = send(sock, s_buf, 200, 0);
```

s_buf (size 200)



Two Scenarios:

Socket buffer (send)



Full



100 bytes

	Full	100 Bytes
A	Return 0	Copy 100 bytes
B	Block	Copy 100 bytes
C	Return 0	Block
D	Block	Block
E	Something else	

Kernel

Blocking Implications

- DO NOT assume that you will `recv()` all of the bytes that you ask for.
 - DO NOT assume that you are done receiving.
 - ALWAYS receive in a loop!*
-
- DO NOT assume that you will `send()` all of the data you ask the kernel to copy.
 - Keep track of where you are in the data you want to send.
 - ALWAYS send in a loop!*

* Unless you're dealing with a single byte, which is rare.

ALWAYS check send() return value!

- When send() return value is less than the data size, **you are responsible for sending the rest.**

Data sent: 0
Data to send: 130

```
send(sock, data, 130, 0);
```



ALWAYS check send() return value!

- When send() return value is less than the data size, **you are responsible for sending the rest.**

Data sent: 0
Data to send: 130



60

```
send(sock, data, 130, 0);
```

Data sent: 60
Data to send: 130



ALWAYS check send() return value!

- When send() return value is less than the data size, **you are responsible for sending the rest.**

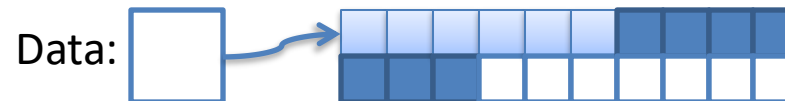
Data sent: 0
Data to send: 130



60

```
send(sock, data, 130, 0);
```

Data sent: 60
Data to send: 130



```
// Copy the 70 bytes starting from offset 60.  
send(sock, data + 60, 130 - 60, 0);
```

ALWAYS check send() return value!

- When send() return value is less than the data size, **you are responsible for sending the rest.**

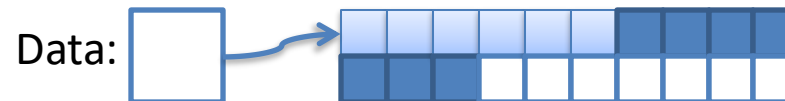
Data sent: 0
Data to send: 130



60

```
send(sock, data, 130, 0);
```

Data sent: 60
Data to send: 130



?

```
// Copy the 70 bytes starting from offset 60.  
send(sock, data + 60, 130 - 60, 0);
```

Repeat until all bytes are sent. (data_sent == data_to_send)...

Blocking Summary

send()

- Blocks when socket buffer for sending is full
- Returns less than requested size when buffer cannot hold full size

recv()

- Blocks when socket buffer for receiving is empty
- Returns less than requested size when buffer has less than full size

Always check the return value!

Create a TCP socket: `socket()`

```
int socket(int domain, int type, int protocol)
```

```
int sock = socket(AF_INET, SOCK_STREAM, 0);
```

- domain: communication domain of the socket: generic interface.
- type of socket: reliable vs. best-effort
- end-to-end protocol: TCP for a stream socket -
 - 0: default E2E for specified protocol family and type.

Create a TCP socket: socket()

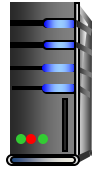
```
int socket(int domain, int type, int protocol)
int sock = socket(AF_INET, SOCK_STREAM, 0);
/* AF_INET: Communicate with IPv4 Address Family (AF),
   SOCK_STREAM: Stream-based protocol
   int sock: returns an integer-valued socket
   descriptor or handle
*/
if (sock < 0) { // If socket() fails, it returns -1
    perror("socket");
    exit(1);
}
```

Close a socket: close()

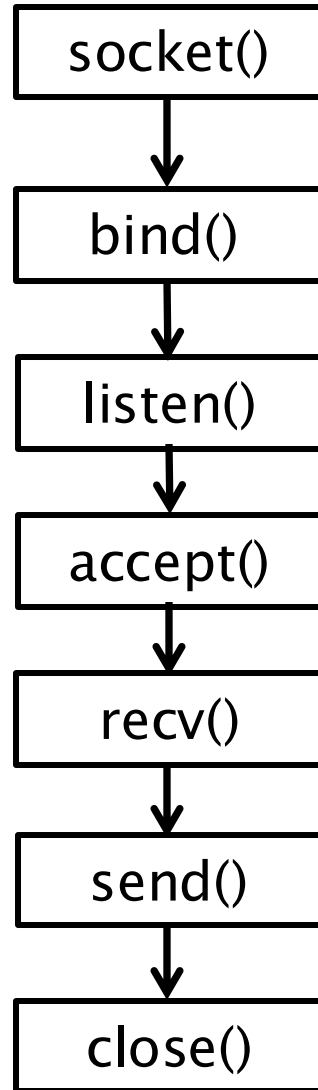
```
int close(int socket)
if (close(sock)) {
    perror("close");
    exit(1);
}
```

/* int socket: int socket descriptor is passed to close()*/

- Close operation similar to closing a file.
- initiate actions to shut down communication
- deallocate resources associated with the socket
- cannot send(), recv() after you close the socket.



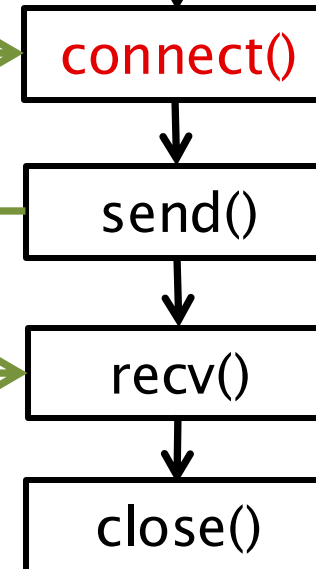
Server



connect()



Client



connect()

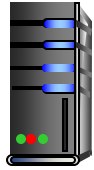
- Before you can communicate, a connection must be established.
- Client Initiates, Server waits.
- Once connect() returns, socket is connected and we can proceed with send(), recv()

```
int connect(int socket, const struct sockaddr  
*foreign_address, socklen_t address_length)
```

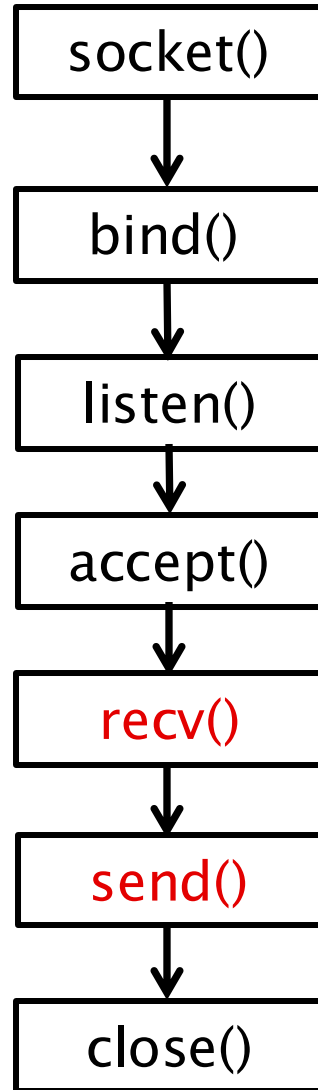

connect()

```
int connect(int socket, const struct sockaddr  
*foreign_address, socklen_t address_length)
```

```
struct sockaddr_in addr;  
int res = connect(sock, (struct sockaddr*)&addr, sizeof(addr));  
/* int socket: socket descriptor  
   foreignAddress: pointer to sockaddr_in containing Internet  
   address, port of server.  
   addressLength: length of address structure  
*/
```



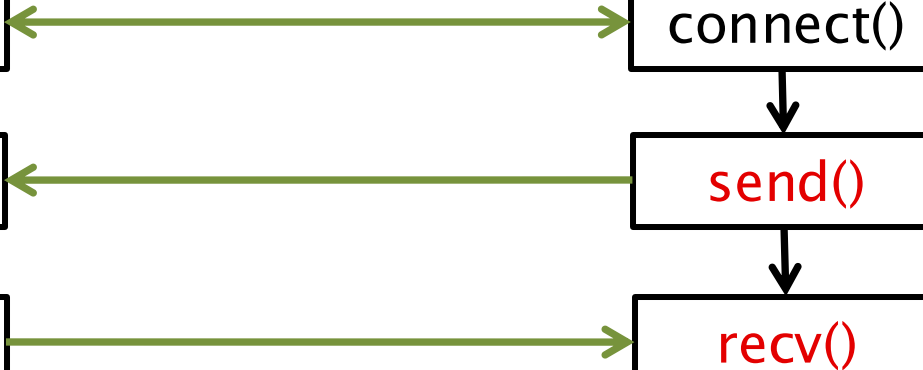
Server



send(), recv()



Client



send(), recv()

Socket is connected when:

- client calls connect()
- connected socket is returned by accept() on server

```
ssize_t send(int socket, const void *msg, msgLength, int flags)
```

```
ssize_t recv(int socket, void *rcvBuffer, size_t bufferLength, int flags)
```

```
/* int socket: socket descriptor
```

```
   return: # bytes sent/received or -1 for failure.
```

send()

send():

```
ssize_t send(int socket, const void *msg, msgLength, int flags)
```

```
/* int socket: socket descriptor
```

```
   send(): msg: sequence of bytes to be sent
```

```
   send(): msgLength: # bytes to send
```

send(), recv()

recv():

```
ssize_t recv (int socket, void *rcvBuffer, size_t bufferLength, int flags)  
int recv_count = recv(sock, buf, 255, 0);
```

/* int socket: socket descriptor

void *rcvBuffer: generally a char array

size_t bufferLength: length of buffer: max # bytes that can be received at once.

flags: setting flag to zero specifies default behavior.



Place all send() and recv() calls in a loop, until you are left with no more bytes to send or receive. One call to send()/recv(), irrespective of the buffer does not necessarily mean all your data will be received at once.