CS118 Discussion 1B, Week 1

Taqi Raza
BUNCHE 1209B, Fridays 12:00pm to 1:50pm

TA

- Taqi, PhD student in Computer Networking
- Discussion (1B): Bunche 1209, Fri 12:00 – 1:50 p.m.
- Office hours: Boelter Hall 2432 Fridays 9:45 – 11:45 a.m.
- TA website: http://web.cs.ucla.edu/~taqi/teaching/winter18/CS118.html
- Emails: Use [CS118] in the title or may be flagged as spam

Logistics

- Submit your signed Academic Integrity Agreement
- Grade decomposition:
  - Homework: 20% (due 6 p.m. next Wednesday)
  - Project 1: 8% (due Friday, Feb. 2nd)
  - Project 2: 12% (due Friday, Mar. 18th)
  - Midterm: 30% (Thursday, Feb. 15th, in-class)
  - Final: 30% (Monday, 3–6 p.m. Mar. 19th)

Logistics: Homework

- Online submission to Gradescope only. DEMO
- Submission guidelines:
  - 1. Hard deadline on submission, so submit early! You can submit multiple times before the deadline, but the system will refuse to accept any submissions after that.
  - 2. Each homework problem will have a dedicated answering box immediately below. Do not write your answers outside the box. Any answer outside the dedicated area may not get graded.
  - 3. You are encouraged to work out the problem on the PDF file directly; or by typing and compiling your answers with the LaTeX template we provide (if available) into PDF.
  - 4. If you prefer handwriting or have to draw diagrams, you may scan the paper copy (e.g., using a smartphone app), convert it to a PDF file and then upload. It is YOUR responsibility to upload a clear copy in black and white. Inaccessible answers will get a low score.

Logistics: Project

- Two projects (C/C++ only):
  - simple web server — introduction to network programming;
  - reliable data transfer — a simple user-level TCP-like transport protocol
- Form a team of 2 persons ASAP.
- Test environment:
  - Vagrant-based Ubuntu virtual machine
  - Demo code included

Outline

- Crash course — getting prepared for Project 1
- Socket programming

Socket Programming

Network programming

- **What is the model for network programming?**
- Where are we programming?
- Which APIs can we use? How to use them?

Client-server model

- Asymmetric communication
  - Client — Requests data:
    - Initiates communication
    - Waits for server's response
  - Server (Daemon) — Responds data requests:
    - Well-known by clients (e.g. IP address + port)
    - Waits for clients connection
    - Processes requests, send replies

Client-server model

- Client and server are not disjoint
  - A client can be a server of another client
  - A server can be a client of another server
  - Example?
- Server's service model
  - Concurrent: server processes multiple clients' requests simultaneously
  - Sequential: server processes clients' requests one by one
  - Hybrid: server maintains multiple connections, but responses sequentially

Network programming

- What is the model for network programming?
- **Where are we programming?**
- Which APIs can we use? How to use them?

Which layer are we at?

- “Clients” and “servers” are programs at application layer
- Transport layer is responsible for providing communication services for application layer
- Basic transport layer protocols:
  - TCP
  - UDP

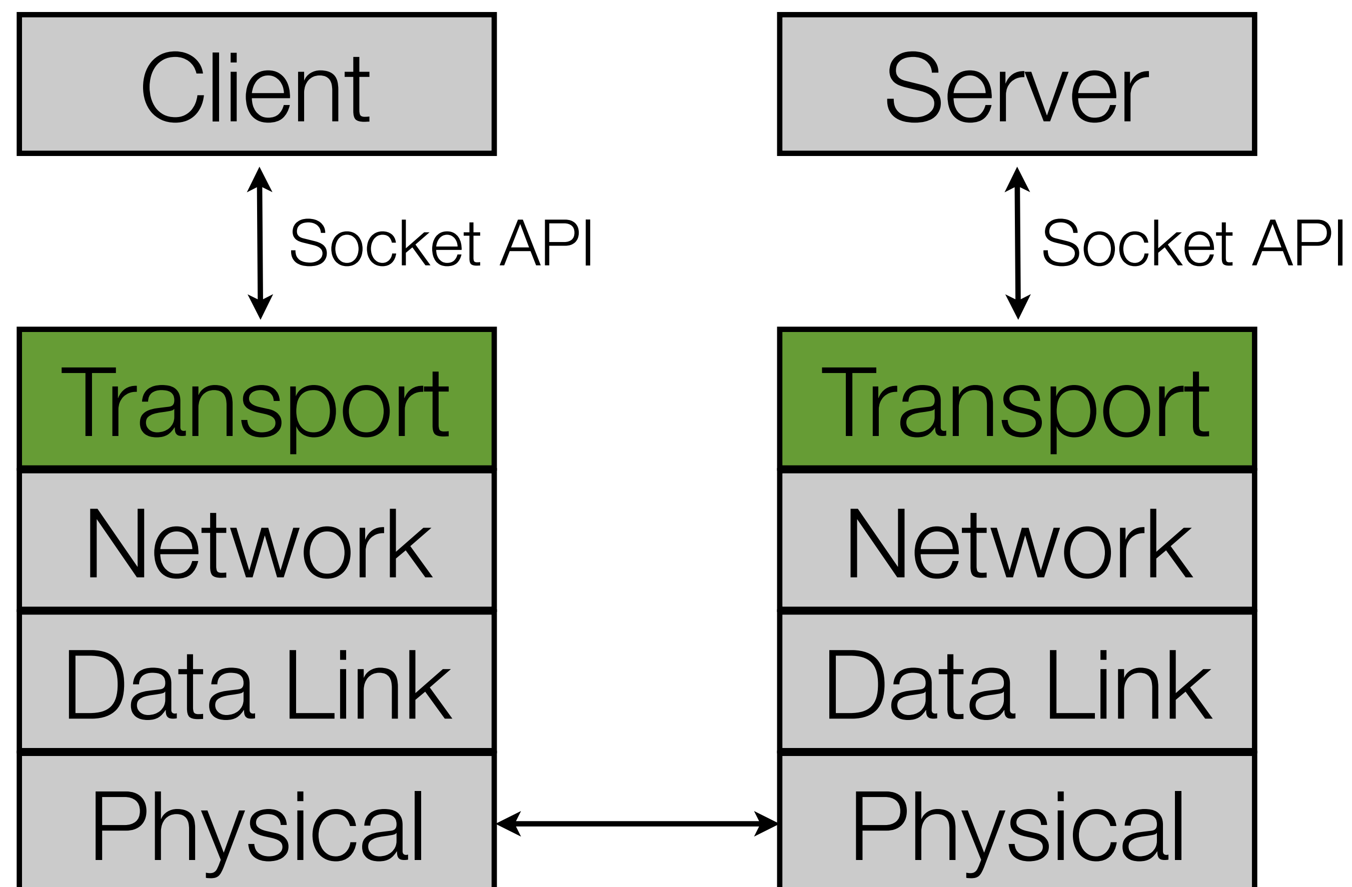

TCP: Transmission Control Protocol

- A connection is set up between client and server
- Reliable data transfer
  - Guarantee deliveries of all data
  - No duplicate data would be delivered to application
- Ordered data transfer
  - If A sends data D1 followed by D2 to B, B will also receive D1 before D2
- Data transmission: full-duplex byte stream
- Regulated data flow: flow control and congestion control

UDP: User Data Protocol

- Basic data transmission service
  - Unit of data transfer: datagram (in variable length)
- No reliability guarantee
- No ordered delivery guarantee
- No flow control / congestion control

Network programming

- What is the model for network programming?
- Where are we programming?
- **Which APIs can we use? How to use them?**

Our secret weapon: socket programming APIs

- From Wikipedia: “A network socket is an endpoint of an inter-process communication flow across a computer network”
- A socket is a tuple of <ip:port>
- Socket programming APIs help build the communication tunnel between applications and transport/network service
- We use TCP socket in this project

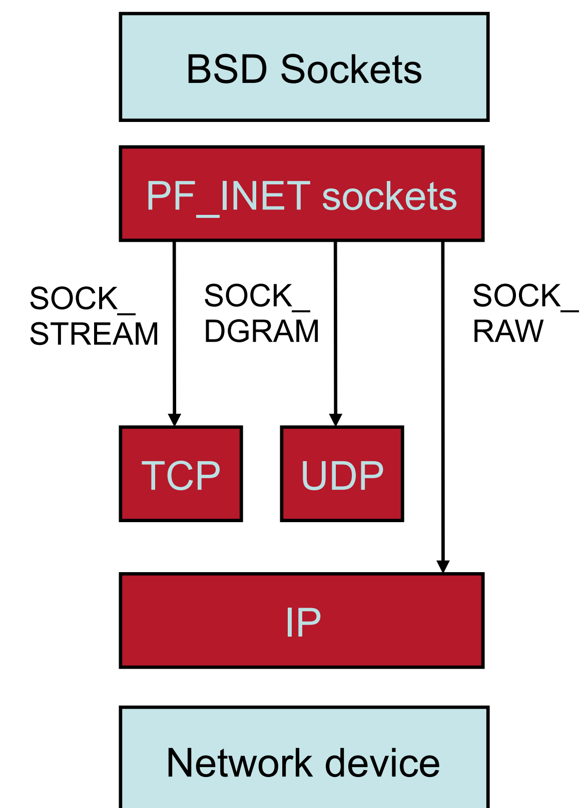

Socket: port number

- Port numbers are allocated and assigned by the IANA (Internet Assigned Numbers Authority)
- See RFC 1700 or https://www.ietf.org/rfc/rfc1700.txt

1-512	• standard services (see **`/etc/services`**) • super-user only
513-1023	• registered and controlled, also used for identity verification • super-user only
1024-49151	• registered services/ephemeral ports
49152-65535	• private/ephemeral ports

TCP socket: basic steps

- Create service
- Establish a TCP connection
- Send and receive data
- Close a TCP connection

TCP socket: service setup

TCP Client

TCP Server

TCP socket: service setup

TCP Client	**TCP Server**
	socket( )

TCP socket: service setup

TCP Client

TCP Server

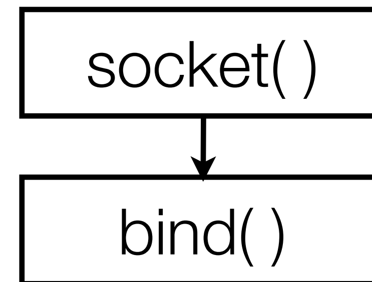

TCP socket: service setup

TCP Client

TCP Server

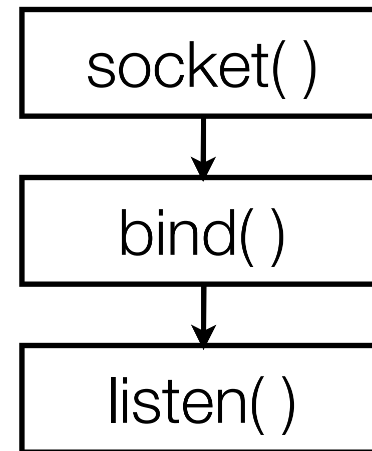

TCP socket: service setup

TCP Client

TCP Server

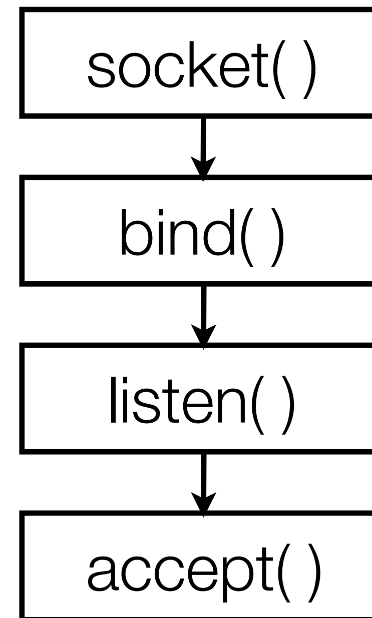

TCP socket: service setup

TCP Client

TCP Server

socket( ) → bind( ) → listen( ) → accept( )

blocked until connection from client

TCP socket: service setup

TCP Client

socket( )

TCP Server

socket( )

bind( )

listen( )

accept( )

blocked until connection from client

TCP socket: establish connection

TCP Client

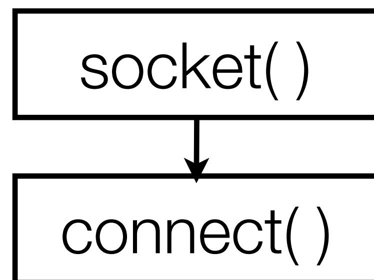

TCP Server

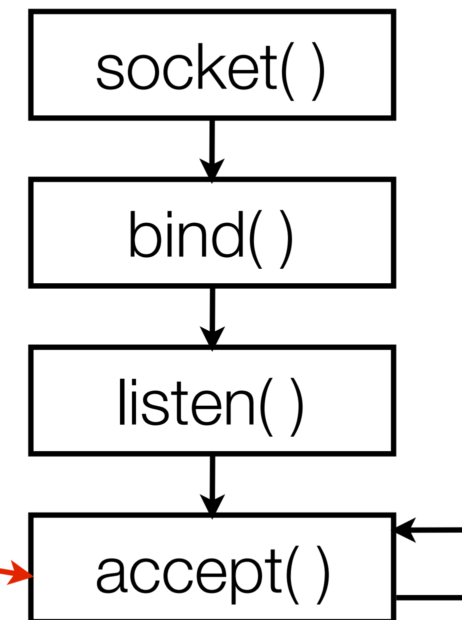

blocked until connection from client

TCP socket: send and receive data

TCP Client

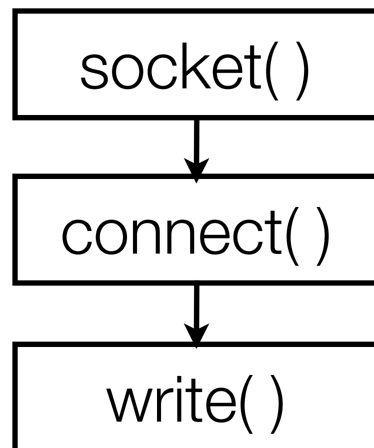

TCP Server

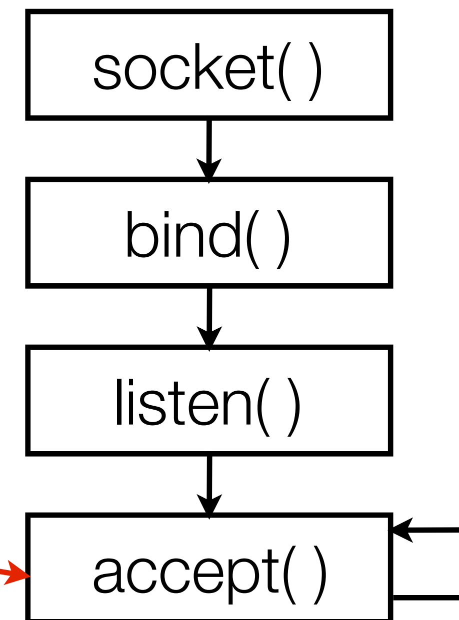

blocked until connection from client

TCP socket: send and receive data

TCP Client

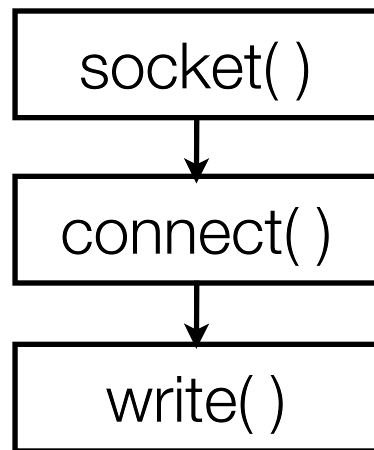

data (request)

TCP Server

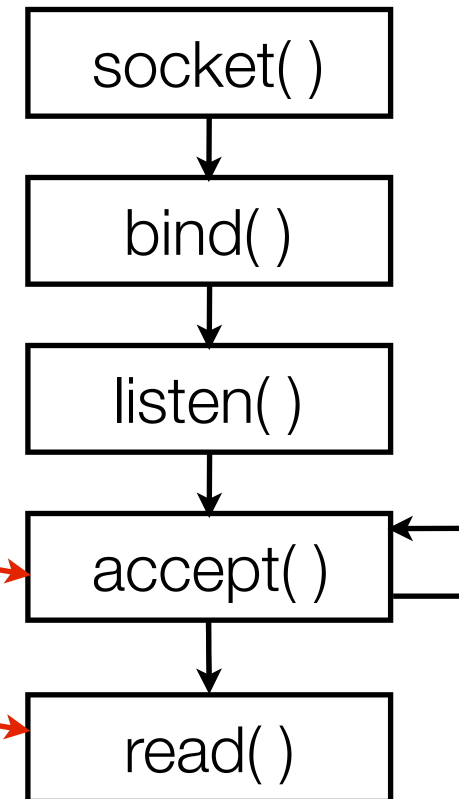

blocked until connection from client

TCP socket: send and receive data

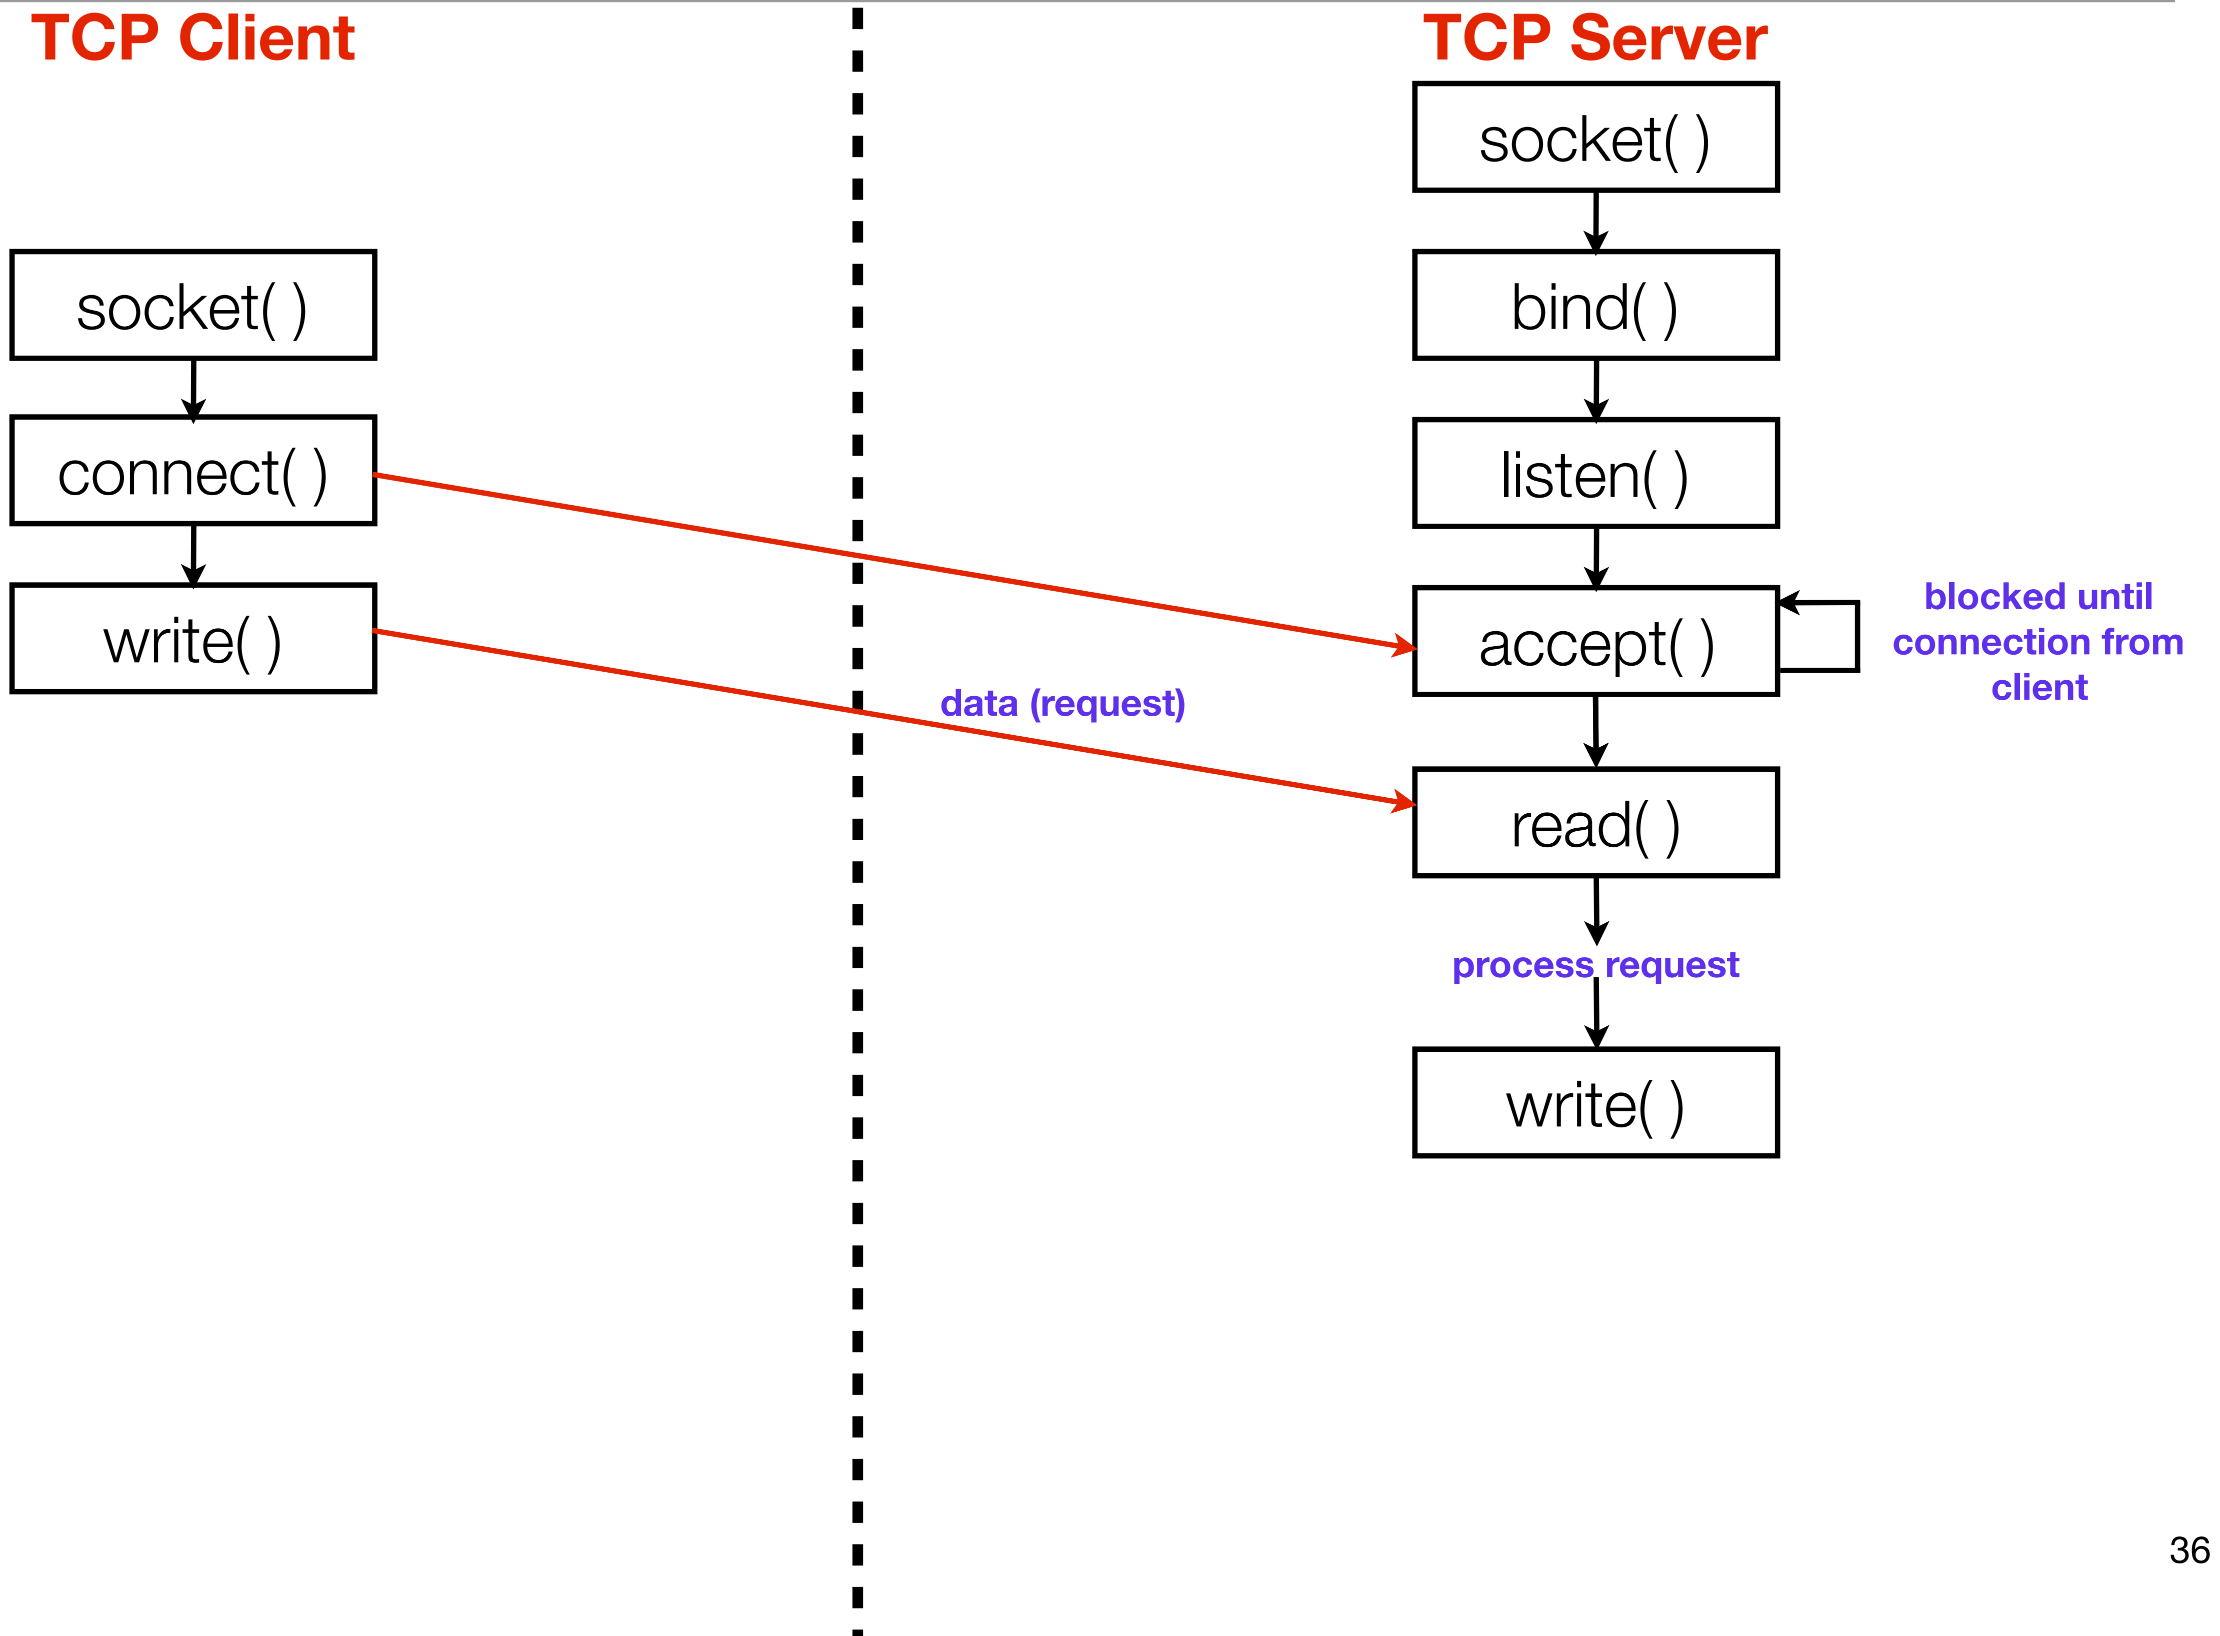

TCP socket: send and receive data

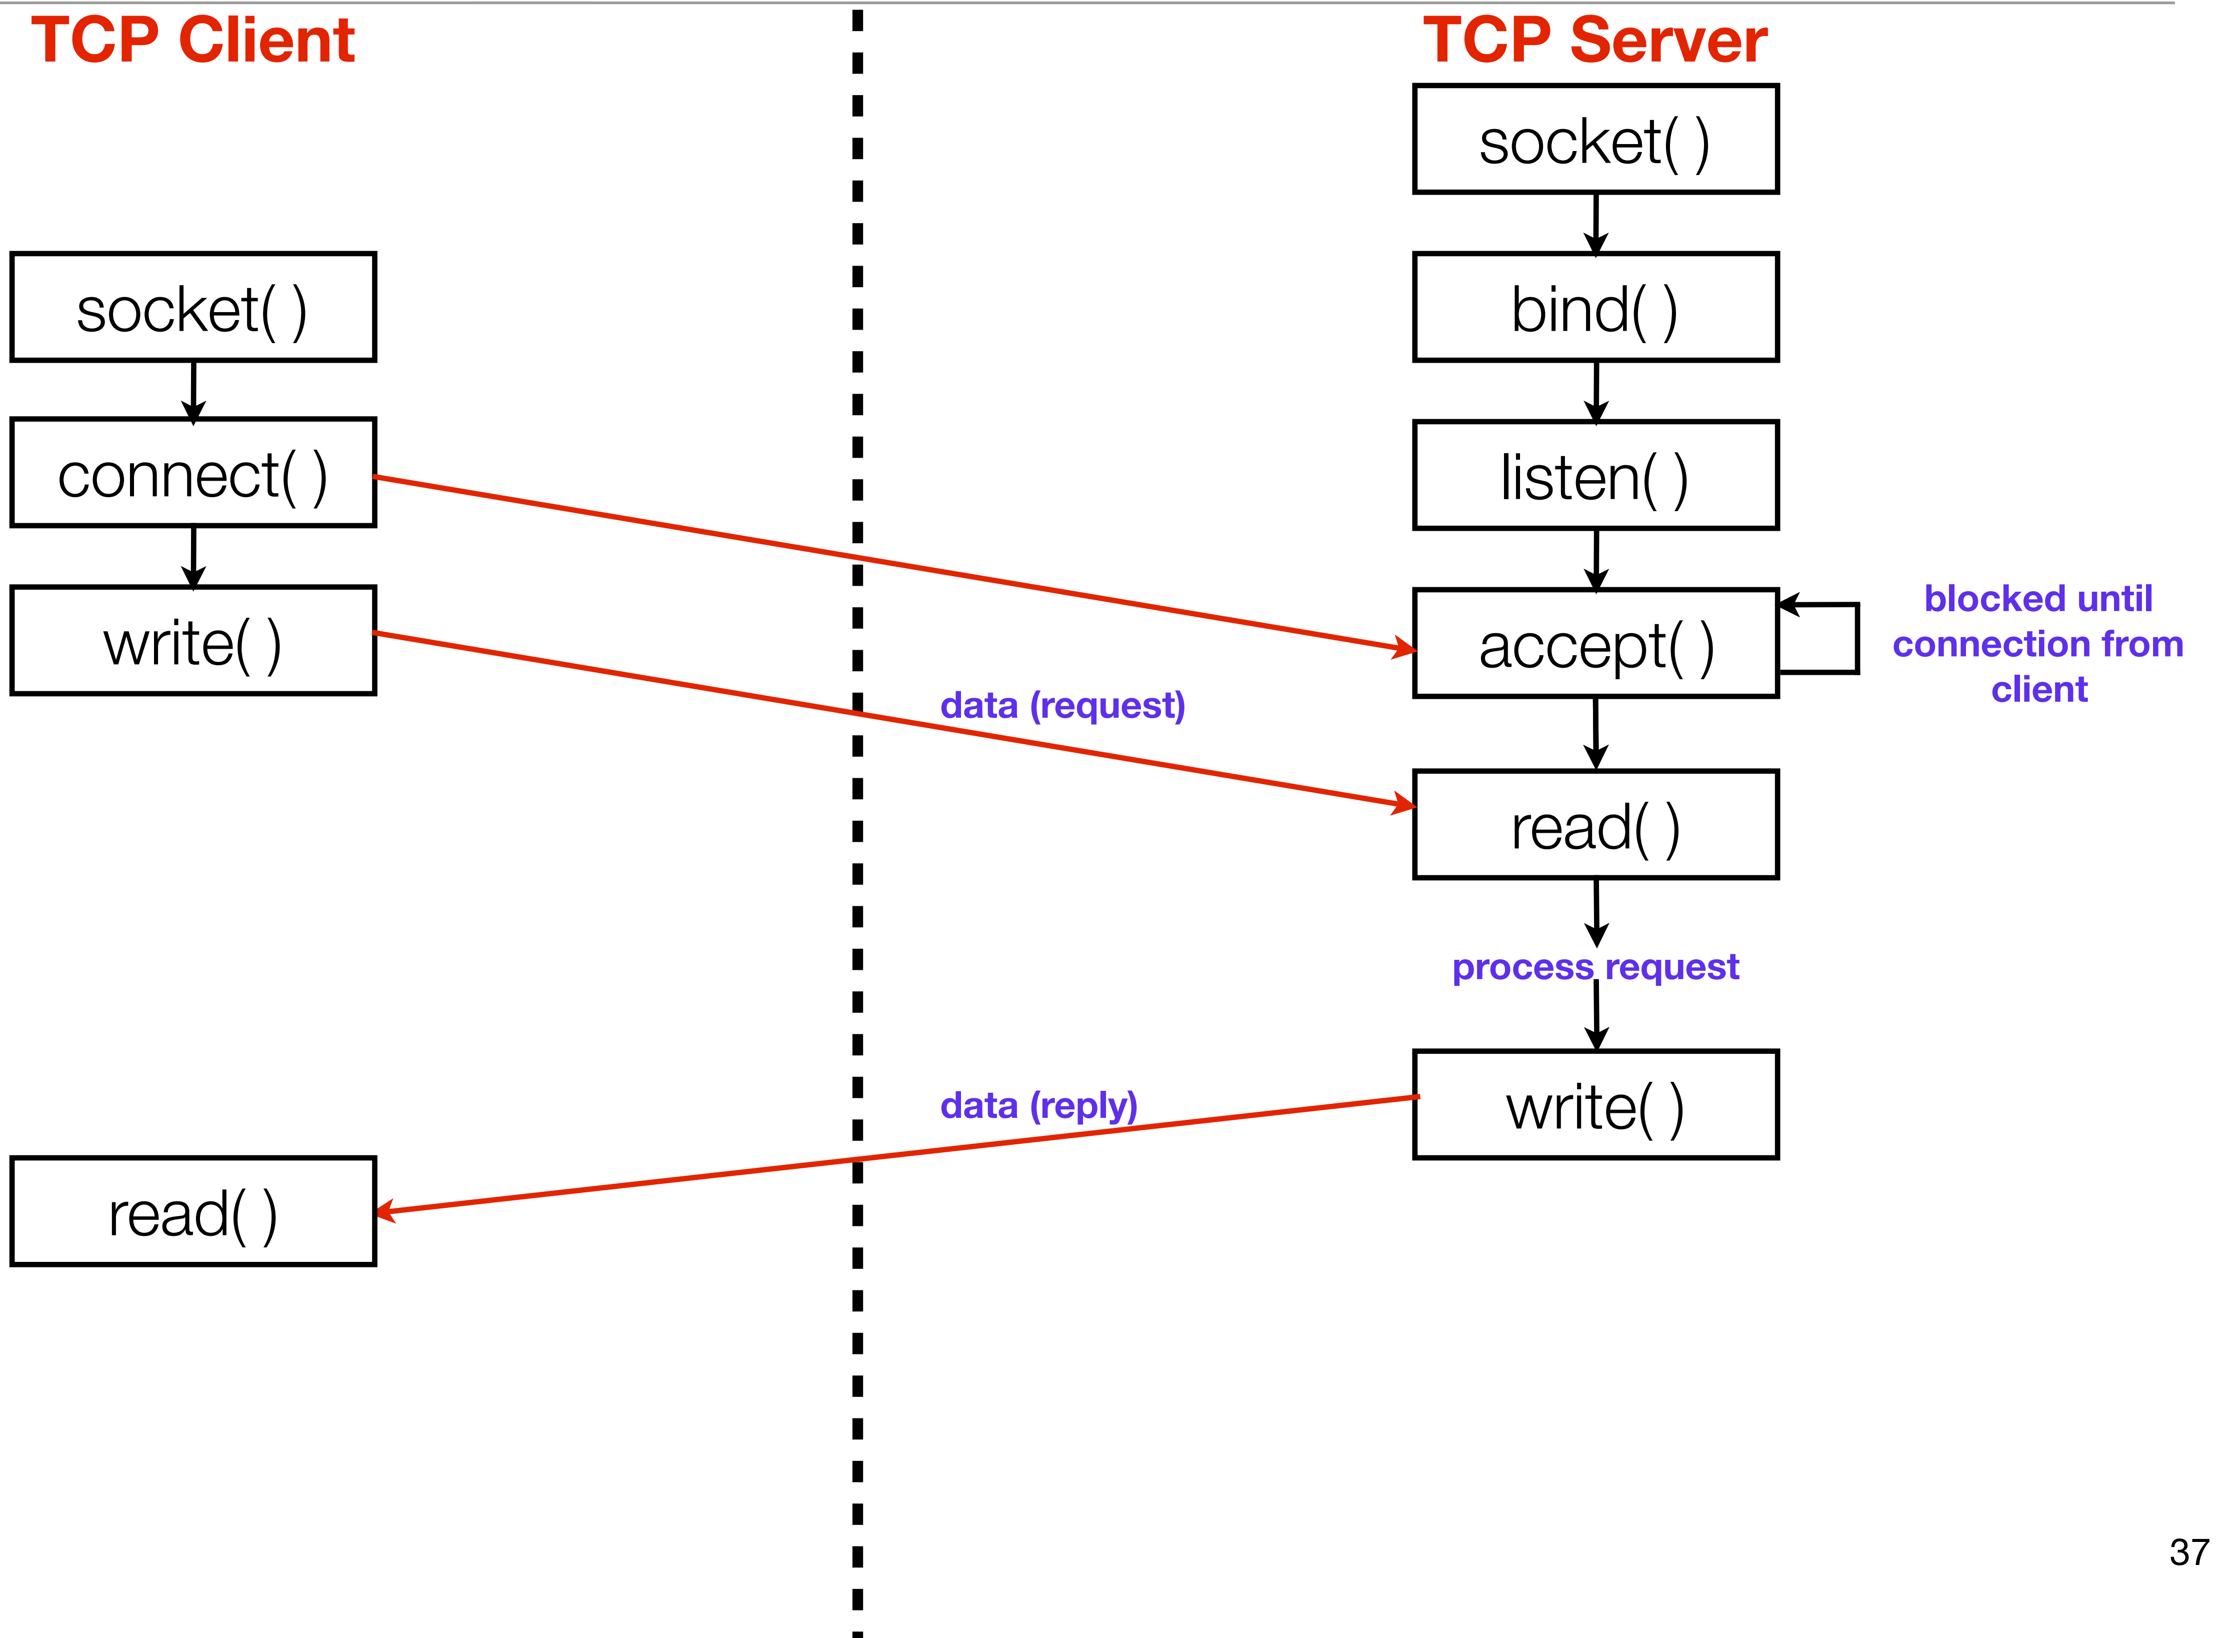

TCP socket: close connection

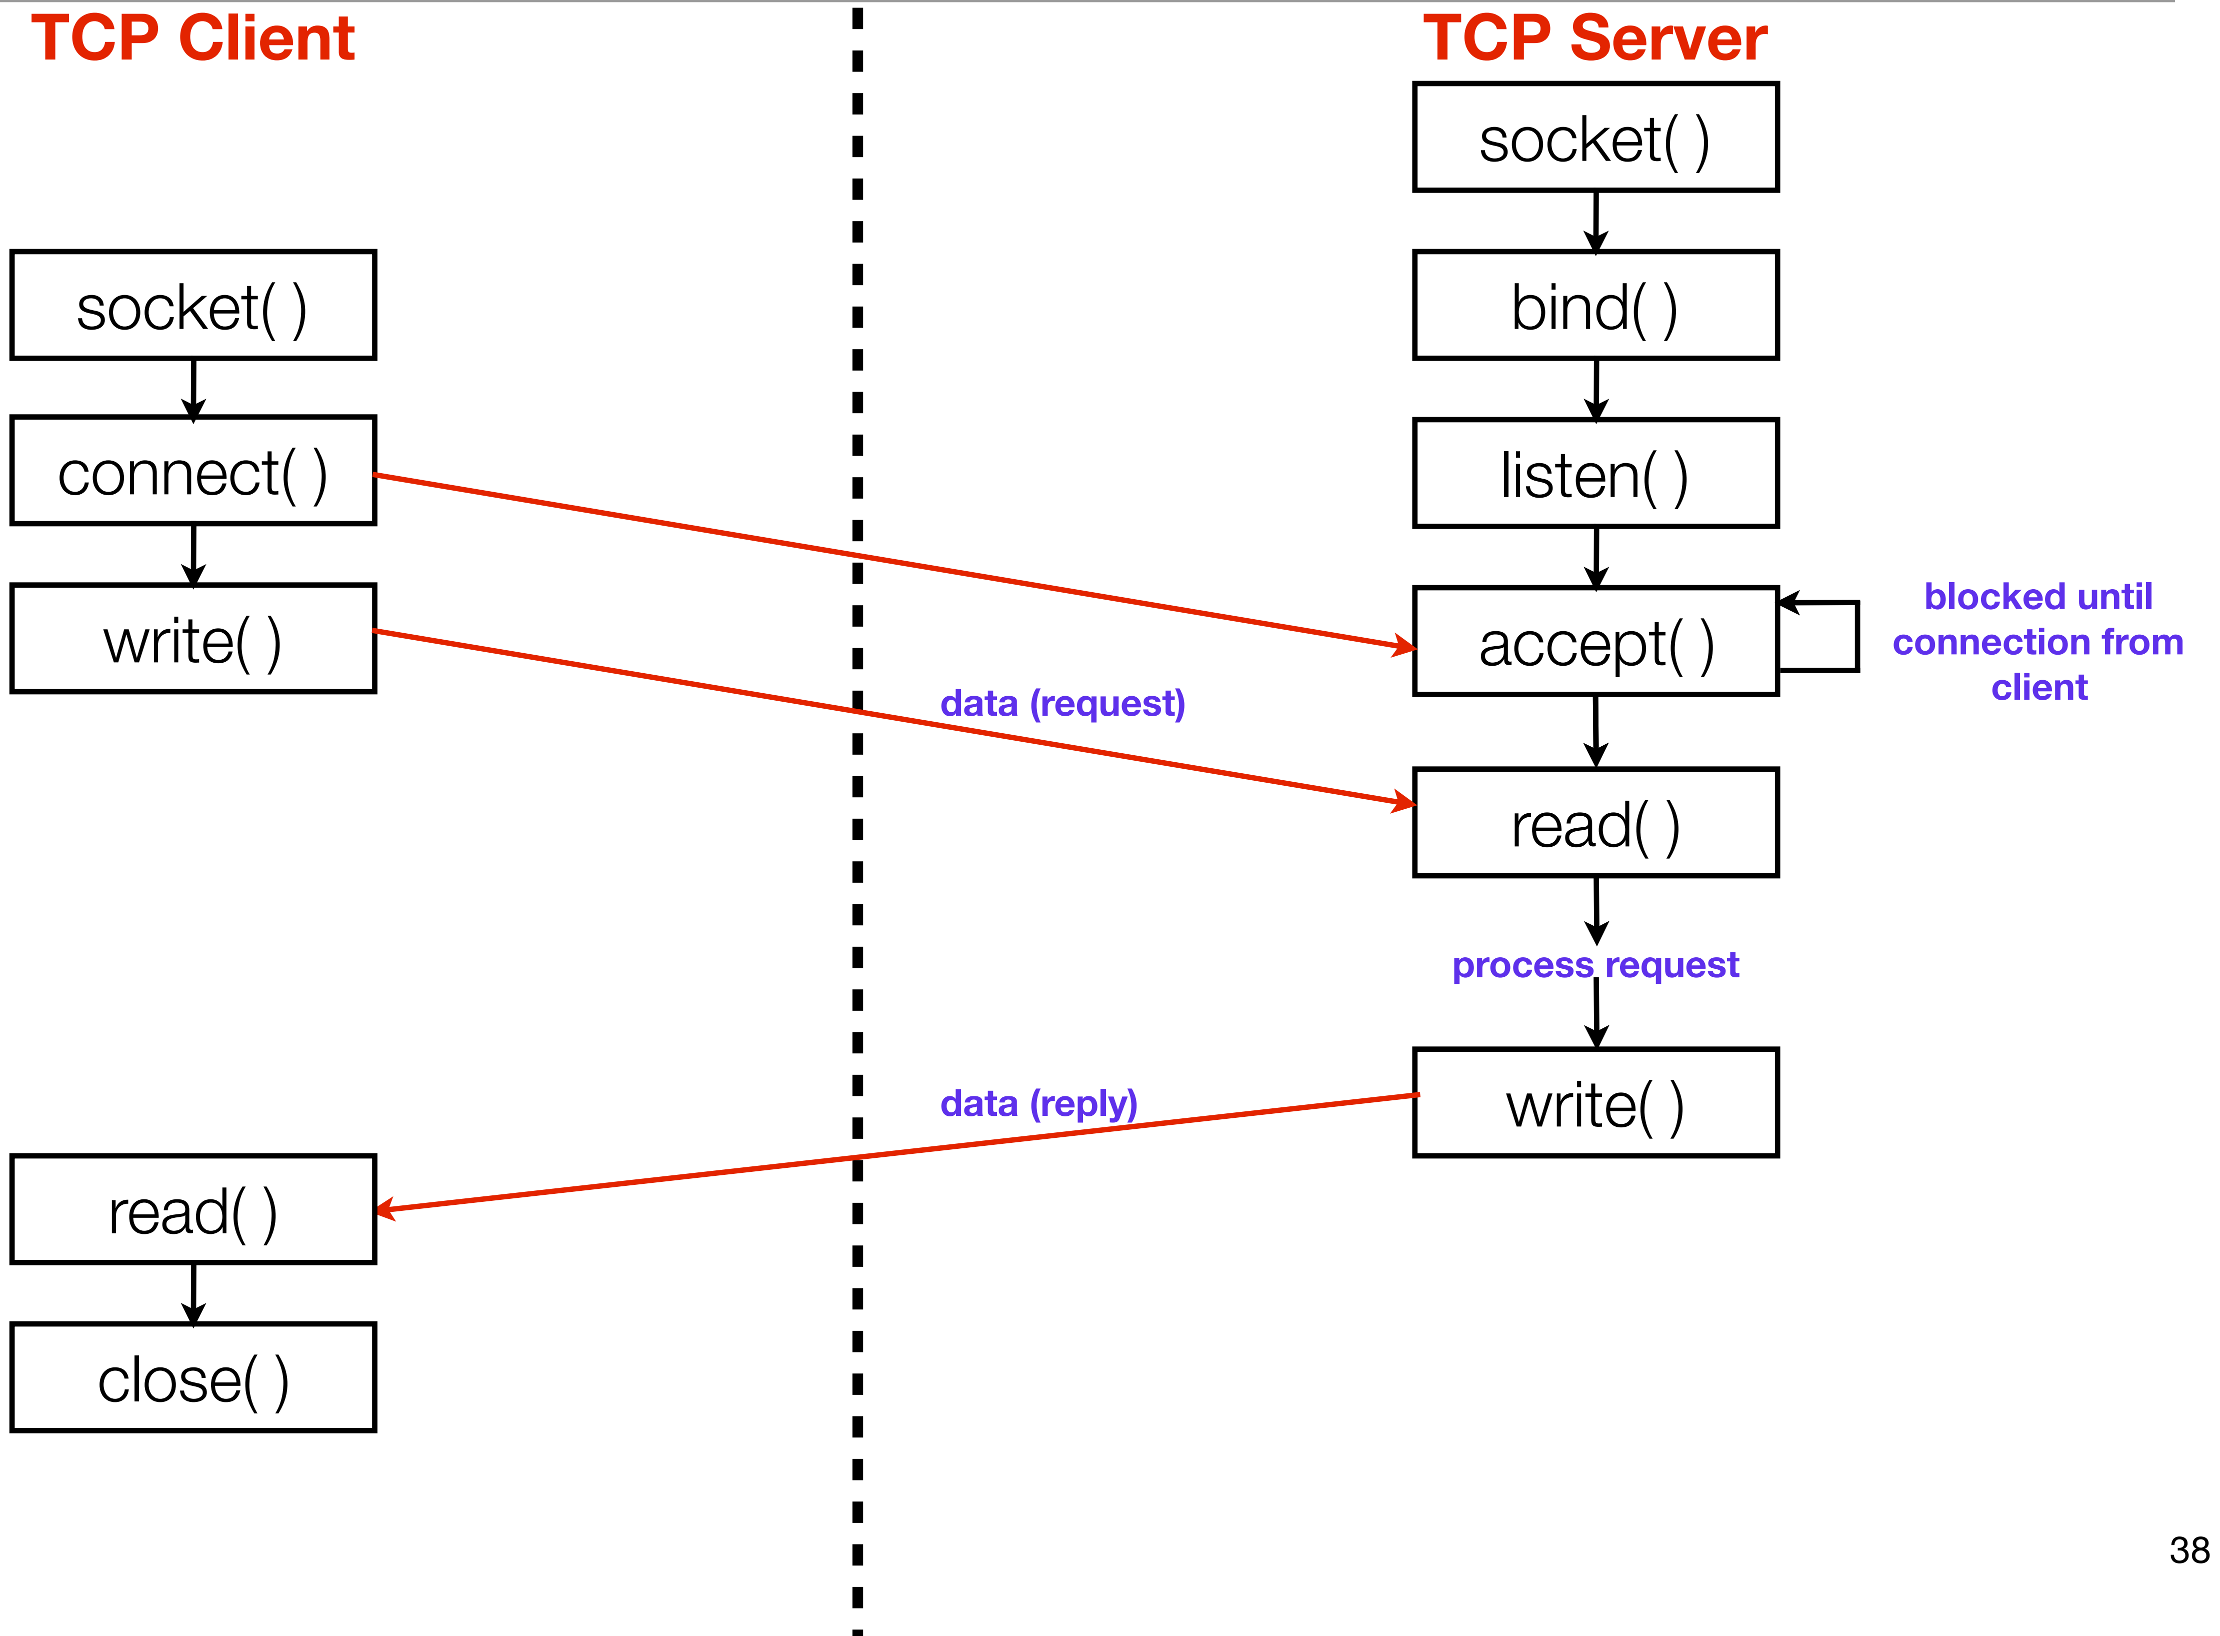

TCP socket: close connection

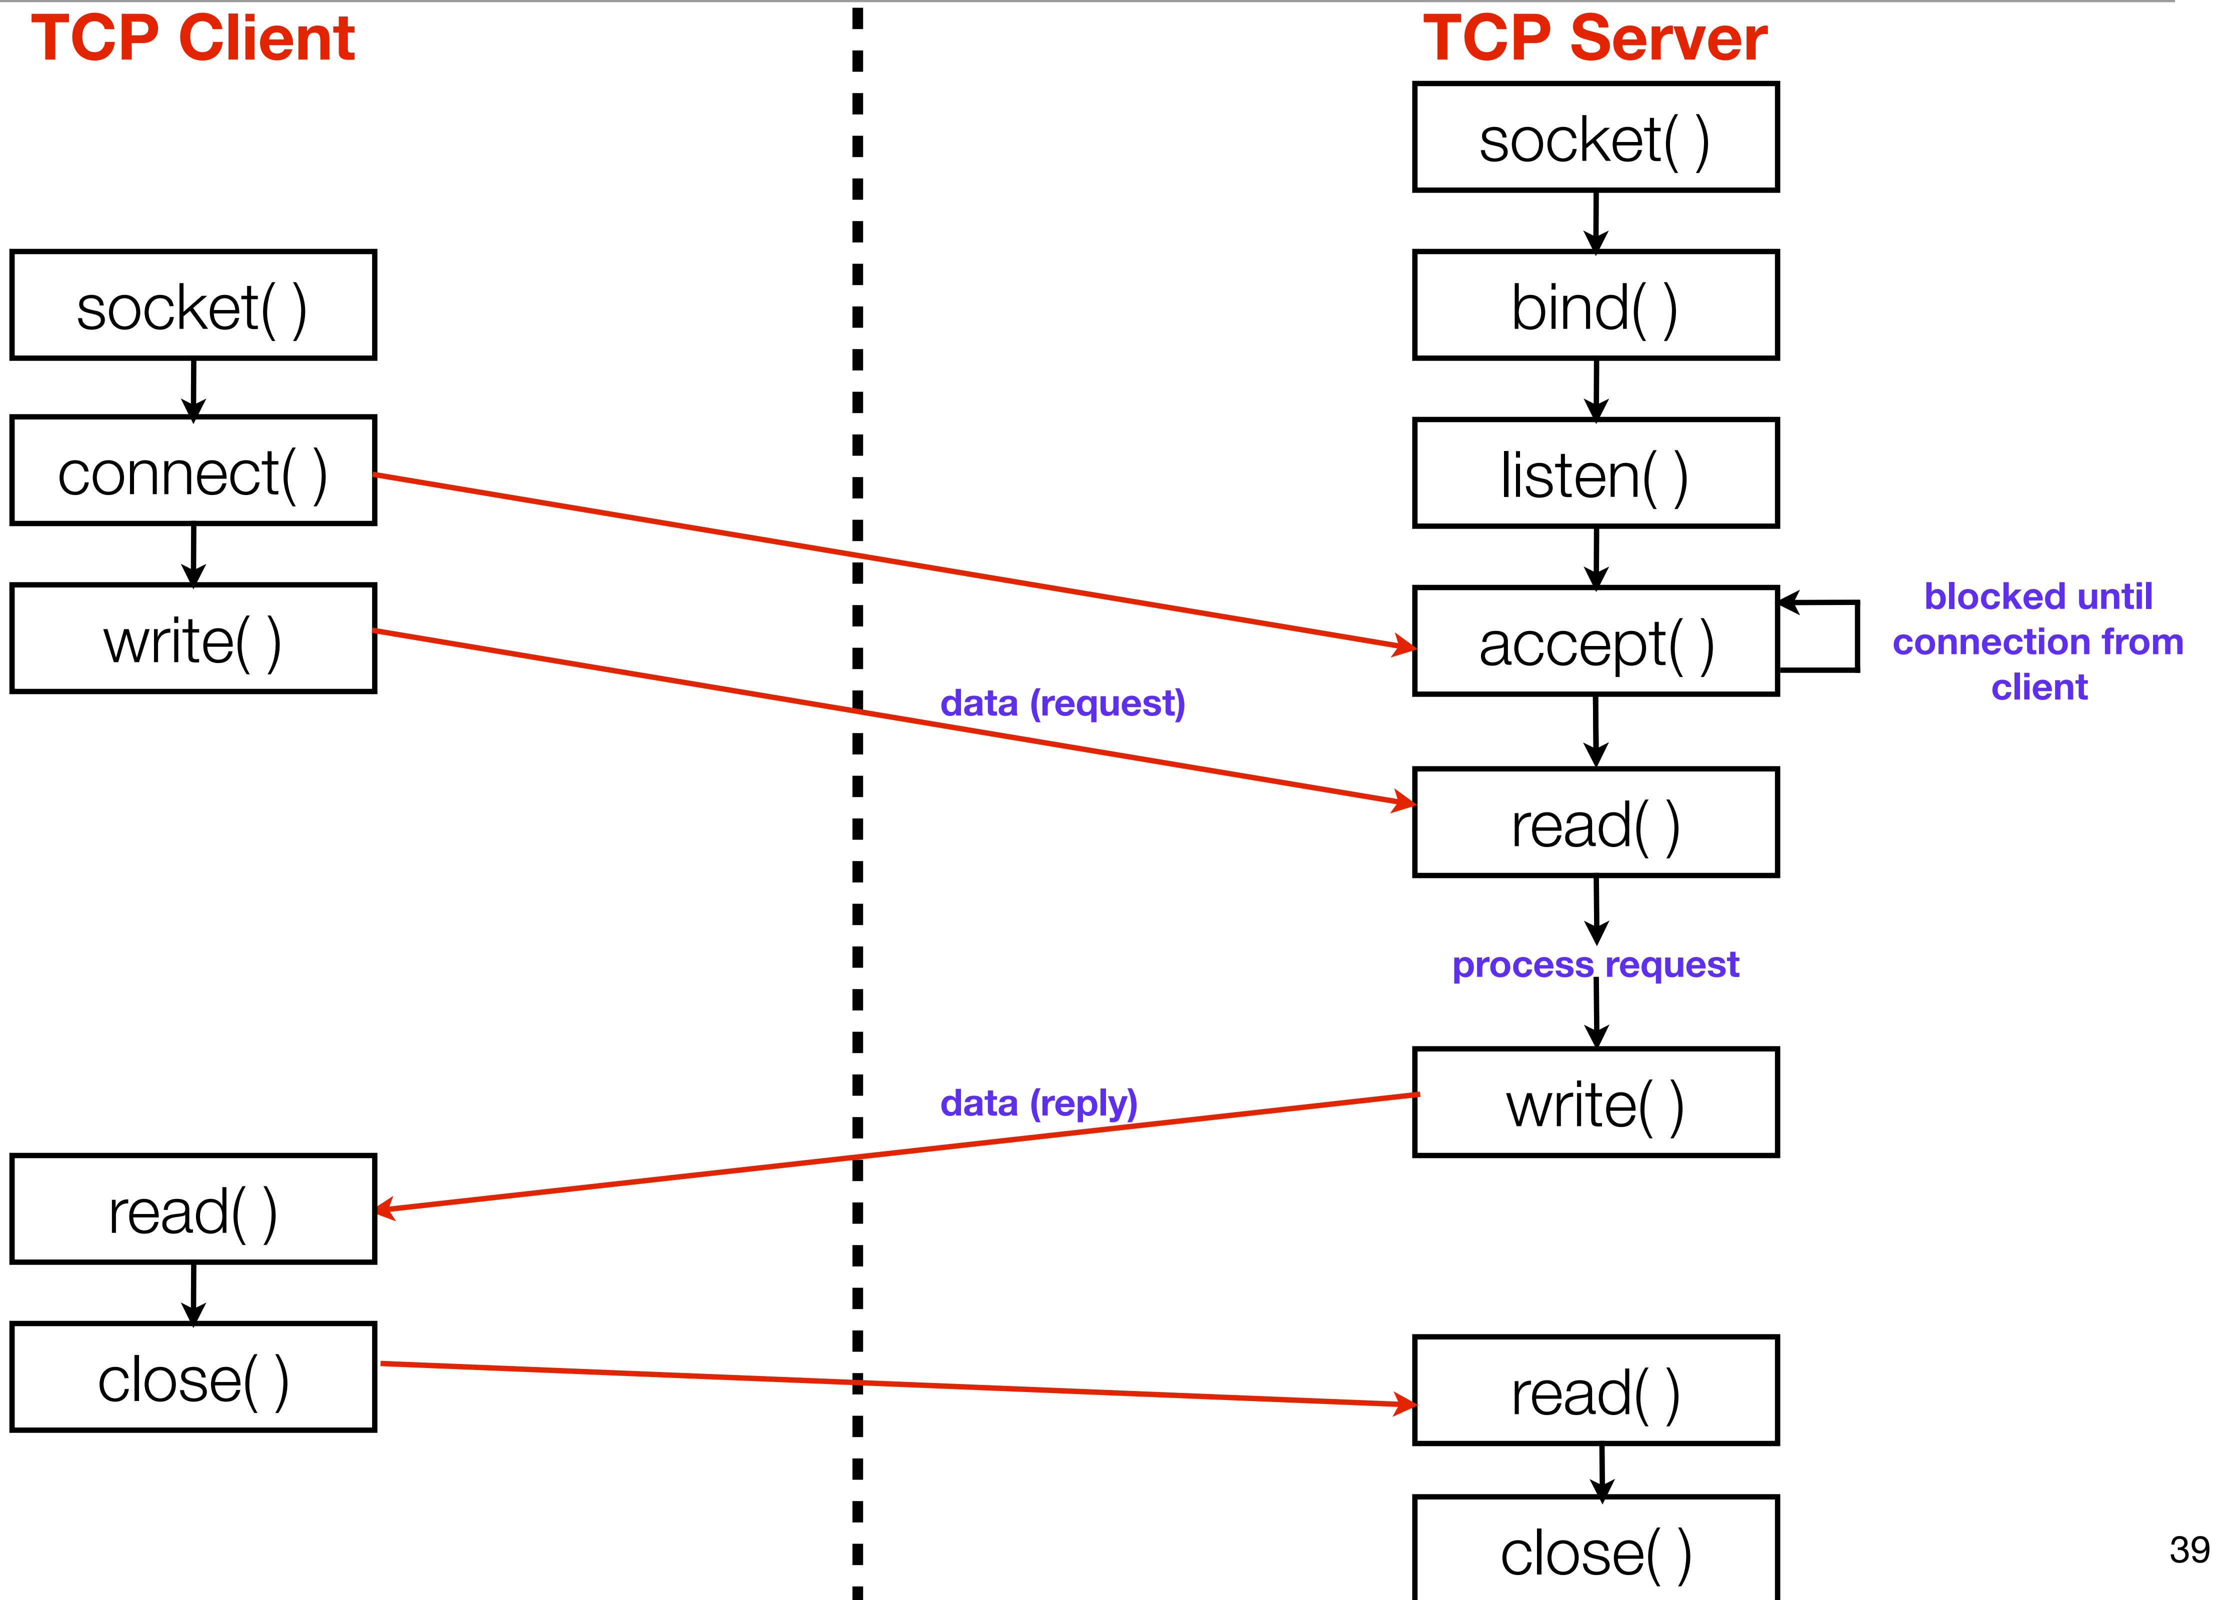

Socket programming API: syscalls

- **int socket(int domain, int type, int protocol);**
    - Create a socket
    - returns the socket descriptor or -1(failure). Also sets errno upon failure
    - **domain**: protocol family
        - **PF_INET** for IPv4, **PF_INET6** for IPv6, **PF_UNIX** or **PF_LOCAL** for Unix socket, **PF_ROUTE** for routing
    - **type**: communication style
        - **SOCK_STREAM** for TCP (with **PF_INET**)
        - **SOCK_DGRAM** for UDP (with **PF_INET**)
    - **protocol**: protocol within family, which is typically set to 0

Socket programming API: syscalls

- **int bind(int sockfd, struct sockaddr* myaddr, int addrlen);**
  - Bind a socket to a local IP address and port number
  - returns 0 on success, -1 and sets errno on failure
  - **sockfd**: socket file descriptor returned by socket ( )
  - **myaddr**: includes IP address and port number
    - **NOTE**: sockaddr and sockaddr_in are of same size, use sockaddr_in and convert it to socketaddr
    - **sin_family**: protocol family, e.g. AF_INET
    - **sin_port**: port number assigned by caller
    - **sin_addr**: IP address
    - **sin_zero**: used for keeping same size as sockaddr
  - **addrlen**: sizeof(struct sockaddr_in)

```
struct sockaddr {
    short sa_family;
    char sa_data[14];
};

struct sockaddr_in {
    short sin_family;
    ushort sin_port;
    struct in_addr sin_addr;
    unsigned char  sin_zero[8];
};
```

Socket programming API: syscalls

- **int listen(int sockfd, int backlog);**
  - Put socket into passive state (wait for connections rather than initiating a connection)
  - returns 0 on success, -1 and sets errno on failure
  - **sockfd**: socket file descriptor returned by socket( )
  - **backlog**: the maximum number of connections this program can serve simultaneously

Socket programming API: syscalls

- **int accept(int sockfd, struct sockaddr* client_addr, int* addrlen);**
  - Accept a new connection
  - Return client's socket file descriptor or -1. Also sets errno on failure
  - **sockfd**: socket file descriptor for server, returned by socket( )
  - **client_addr**: IP address and port number of a client (returned from call)
  - **addrlen**: length of address structure = pointer to **int** set to **sizeof(struct sockaddr_in)**
  - **NOTE: client_addr and addrlen are result arguments**
    - i.e. The program passes empty client_addr and addrlen into the function, and the kernel will fill in these arguments with client's information **(why do we need them?)**

Socket programming API: syscalls

- **int connect (int sockfd, struct sockaddr* server_addr, int addrlen);**
  - Connecter to another socket (server)
  - Return 0 on success, -1 and sets errno on failure
  - **sockfd**: socket file descriptor (returned from socket)
  - **server_addr**: IP address and port number of the server
    - server's IP address and port number should be known in advance
  - **addrlen**: sizeof(struct sockaddr_in)

Socket programming API: syscalls

- **int write(int sockfd, char* buf, size_t nbytes);**
    - Write data to a TCP stream
    - Return the number of sent bytes or -1 on failures
    - **sockfd**: socket file descriptor from socket ( )
    - **buf**: data buffer
    - **nbytes**: the number of bytes that caller wants to send

Socket programming API: syscalls

- **int read(int sockfd, char* buf, size_t nbytes);**
  - Read data from TCP stream
  - Return the number of bytes read or -1 on failures
  - Return 0 if socket is closed
  - **sockfd**: socket file descriptor returned from socket ( )
  - **buf**: data buffer
  - **nbytes**: the number of bytes that caller can read (usually set as buffer size)

Socket programming API: syscalls

- **int close(int sockfd);**
  - close a socket
  - return 0 on success, or -1 on failure
  - After close, sockfd is no longer valid

Caveat: byte ordering matters

- Little Endian: least significant byte of word is stored in the lowest address
- Big Endian: most significant byte of word is stored in the lowest address
- Hosts may use different orderings, so we need byte ordering conversion
- **Network Byte Order = Big Endian**

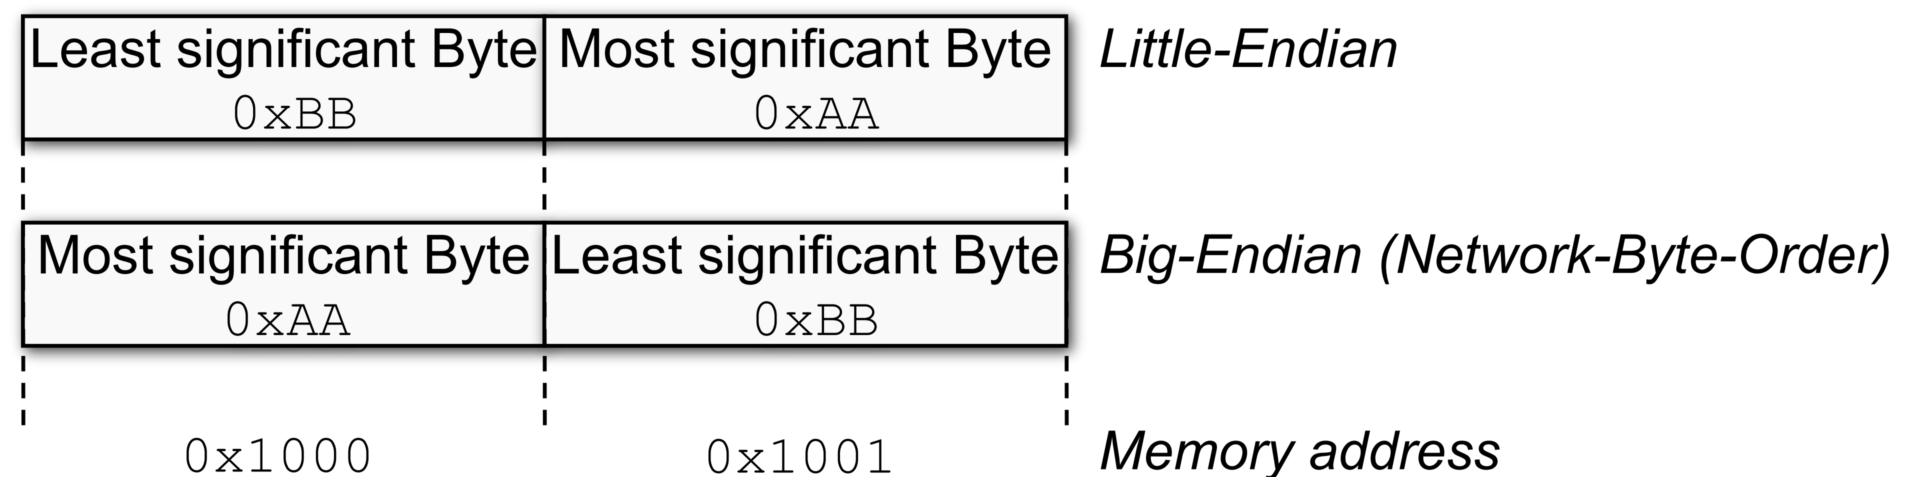

Caveat: byte ordering matters

- Byte ordering functions: used for converting byte ordering
- Example:

```
int m, n;
short int s,t;

m = ntohl (n)   net-to-host long (32-bit) translation
s = ntohs (t)   net-to-host short (16-bit) translation
n = htonl (m)   host-to-net long (32-bit) translation
t = htons (s)   host-to-net short (16-bit) translation
```

- Rule: for every int or short int
- Call htonl( ) or htons( ) before sending data
- Call ntohl( ) or ntohs( ) before reading received data

Address util functions

- All binary values are network byte ordered
- **struct hostent* gethostbyname (const char* hostname);**
  - Translate host name (e.g. “localhost”) to IP address (with DNS working)
- **struct hostent* gethostbyaddr (const char* addr, size_t len, int family);**
  - Translate IP address to host name
- **char* inet_ntoa (struct in_addr inaddr);**
  - Translate IP address to ASCII dotted-decimal notation (e.g. “192.168.0.1”)
- **int gethostname (char* name, size_namelen);**
  - Read local host’s name

Address util functions

- **in_addr_t inet_addr (const char* strptr);**
  - Translate dotted-decimal notation to IP address (network byte order)

```
struct sockaddr_in ina;
ina.sin_addr.s_addr = inet_addr("10.12.110.57");
```

- **int inet_aton (const char* strptr, struct in_addr *inaddr);**
  - Translate dotted-decimal notation to IP address

```
struct sockaddr_in my_addr;
my_addr.sin_family = AF_INET;         // host byte order
my_addr.sin_port = htons(MYPORT);     // short, network byte order
inet_aton("10.12.110.57",&(my_addr.sin_addr));
memset(&(my_addr.sin_zero), '\0', 8); // zero the rest of the struct
```

FYI: `struct hostent`

`char *h_name`	The real canonical host name.
`char **h_aliases`	A list of aliases that can be accessed with arrays—the last element is `NULL`
`int h_addrtype`	The result's address type, which really should be `AF_INET` for our purposes.
`int length`	The length of the addresses in bytes, which is 4 for IP (version 4) addresses.
`char **h_addr_list`	A list of IP addresses for this host. Although this is a `char**`, it's really an array of `struct in_addr*`s in disguise. The last element is `NULL`.
`h_addr`	A commonly defined alias for `h_addr_list[0]`. If you just want any old IP address for this host (they can have more than one) just use this field.

How to write a server: headers

```
/* PLEASE include these headers */
#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <errno.h>
#include <sys/types.h>
#include <sys/socket.h>
#include <sys/wait.h>
#include <netinet/in.h>
#define PORT 5000  /* Avoid reserved ports */
#define BACKLOG 10 /* pending connections queue size */
```

How to write a server: body (I)

```
int main()
{
    int sockfd, new_fd; /* listen on sock_fd, new connection on new_fd */
    struct sockaddr_in my_addr;  /* my address */
    struct sockaddr_in their_addr; /* connector addr */
    int sin_size;

    /* create a socket */
    if ((sockfd = socket(PF_INET, SOCK_STREAM, 0)) == -1) {
        perror("socket");
        exit(1);
    }
```

How to write a server: body (II)

```
// ...
/* bind the socket */
my_addr.sin_family = AF_INET;
my_addr.sin_port = htons(MYPORT); /* short, network byte order */
my_addr.sin_addr.s_addr = htonl(INADDR_ANY);
/* INADDR_ANY allows clients to connect to any one of the host’s IP address */

if (bind(sockfd, (struct sockaddr *) &my_addr,
         sizeof(struct sockaddr)) == -1) {
    perror("bind");
    exit(1);
}
```

How to write a server: body (III)

```
    // ...
        if (listen(sockfd, BACKLOG) == -1) {
        perror("listen");
        exit(1);
    }
    while(1) { /* main accept() loop */
        sin_size = sizeof(struct sockaddr_in);
        if ((new_fd = accept(sockfd, (struct sockaddr*)
                             &their_addr, &sin_size)) == -1) {
            perror("accept");
            continue;
        }
        printf("server: got connection from %s\n",
               inet_ntoa(their_addr.sin_addr));
        close(new_fd);
    }
}
```

How to write a client?

```c
/* include all the headers */
int main() {
    int sockfd, new_fd; /* listen on sock_fd, new connection on new_fd */
    struct sockaddr_in my_addr;  /* my address */
    struct sockaddr_in their_addr; /* connector addr */
    struct hostent* he;
    int sin_size;

    if ((sockfd = socket(PF_INET, SOCK_STREAM, 0)) == -1) {
        perror (“socket”);
        exit (1);
    }

    their_addr.sin_family = AF_INET; /* interp’d by host */
    their_addr.sin_port = htons (PORT);
    their_addr.sin_addr = *((struct in_addr*) he->h_addr);

    if (connect (sockfd, (struct sockaddr*) &their_addr,
                 sizeof (struct sockaddr)) == -1) {
        perror (“connect”);
        exit (1);
    }
    return 0;
}
```

Summary: what we have learned today

- What is the model for network programming?
  - **Client-Server model**
- Where are we programming?
  - **TCP and UDP in a nutshell**
- Which APIs can we use? How to use them?
  - **Socket programming**

Further Reading

- Stevens, W. Richard, Bill Fenner, and Andrew M. Rudoff. *UNIX Network Programming: The Sockets Networking API*. Vol. 1. Addison-Wesley Professional, 2004.

- Beej's Guide to Network Programming (http://beej.us/guide/bgnet)

- Socket Programming from Dartmouth, http://www.cs.dartmouth.edu/~campbell/cs60/socketprogramming.html

- C/C++ reference: http://en.cppreference.com

See you next time!

- TA: Taqi
- OH: BH2432 Fri 945–1145am
- Website:

http://web.cs.ucla.edu/~taqi/teaching/winter18/CS118.html