


Socket Programming

Tran Giang Son, tran-giang.son@usth.edu.vn

ICT Department, USTH

Contents


- What's a socket?
- Why socket?
- What's in a socket?
- How to use sockets?



What & Why?



Layers



Why layering?

- Similar to application layering
- Application programmer
 - Doesn't care about routing
 - Doesn't care about Ethernet frame
 - Doesn't care about WiFi WPA2 encryption
 - Doesn't care about reliability implementations

Why layering?


- Application programmer
 - Passes the data down
 - Focus on the application



Why layering?

Lower layers:

- What does they need to know?



Why layering?

Lower layers:

- What does they need to know?
- Destination
 - Where to?
 - Hostname («resolved» with `gethostbyname()`)
 - IP address
 - Which service?
 - Indicated by port number

Socket: What?

- Endpoint of a two-way communication link between two networked programs
- Represented by a file descriptor after creation
 - Unix philosophy: Everything is a file
- *De facto* standard for TCP/UDP, replaced
 - NetBIOS / NetBEUI
 - IPX / SPX

Socket: Applications

Most network applications use sockets

- Send messages
- Share data: image, music, video, “cast”
- Interprocess Communication (IPC)

Socket: Which types?


- Stream sockets: connection-oriented with TCP
 - Make a connection
 - Transfer data
 - Close connection
 - Ensure sequence, error checking, etc. . .
 - Example: youtube video

Socket: Which types?

- Stream sockets: connection-oriented with TCP
 - Make a connection
 - Transfer data
 - Close connection
 - Ensure sequence, error checking, etc...
 - Example: youtube video
- Datagram sockets: connectionless with UDP
 - Transfer data without explicitly making a connection
 - Example: DNS

Socket: Why?

- The **standard** API for connecting processes
 - Local
 - Networked



Socket: Why?


- The **standard** API for connecting processes
 - Local
 - Networked
- Compatibility: widely supported
 - Linux / UNIX
 - Windows
 - macOS

Socket: Why?

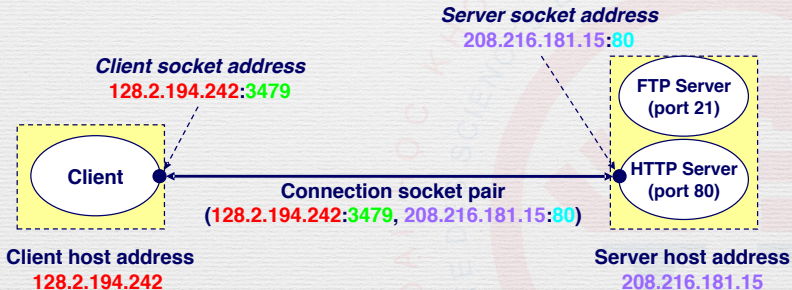
- The **standard** API for connecting processes
 - Local
 - Networked
- Compatibility: widely supported
 - Linux / UNIX
 - Windows
 - macOS
- Low level
 - Minimize amount of data transfer
 - Fast, very little overhead
 - Customizable, flexible, self-defined protocol

Why NOT socket over higher level libraries?

- Low level: more efforts
 - Define protocol
 - Message boundaries
 - Data representation
 - Security
- Session control
 - Authentication, etc.



What's in a socket?



Overview & Setup

Overview

Server

- Passively waits
- Passive socket

Client


- Initiates the connection
- Active socket

Overview


Steps

- Setup
 - Where is the remote host?
 - What service?
- Transfer Data
 - Send/Receive ~ write() / read()
- Close


Socket Overview



Server: Overview



Server socket



Landline phone

Socket: Important struct

```
struct sockaddr_in {
    short          sin_family;    // e.g. AF_INET
    unsigned short sin_port;     // e.g. htons(3490)
    struct in_addr sin_addr;     // see struct in_addr, below
    char          sin_zero[8];  // zero this if you want to
};

struct in_addr {
    unsigned long s_addr; // load with inet_aton()
};
```

Server: Setup socket()

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

- domain: AF_INET (IPv4) or AF_INET6 (IPv6)
- type: SOCK_STREAM (TCP) or SOCK_DGRAM (UDP)
- protocol: 0

For example:

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


Server: Binding `bind()`

```
int bind(int sockfd,  
        const struct sockaddr *bind_addr,  
        socklen_t addrlen);
```

- `sockfd`: file descriptor that `socket()` returned
- `bind_addr`: a «`struct sockaddr_in`» for IPv4
- `addrlen`: size of the struct pointed by `bind_addr`

Server: Example for `socket()` and `bind()`

```

struct sockaddr_in saddr;
int sockfd;
unsigned short port = 80;
if ((sockfd = socket(AF_INET, SOCK_STREAM, 0)) < 0) {
    printf("Error creating socket\n");
    ...
}

memset(&saddr, 0, sizeof(saddr));
saddr.sin_family = AF_INET;
saddr.sin_addr.s_addr = htonl(INADDR_ANY);
saddr.sin_port = htons(port);
if (bind(sockfd, (struct sockaddr *) &saddr, sizeof(saddr)) < 0) {
    printf("Error binding\n"); ...
}

```

Server: Listen to incoming connections `listen()`

```
int listen(int sockfd, int backlog);
```

- `sockfd`: file descriptor that `socket()` returned
- `backlog`: number of pending connections to queue

For example:

```
listen(sockfd, 10);
```

Server: Accept an incoming connection `accept()`

- Server must explicitly accept incoming connections

```
int accept(int sockfd,
           struct sockaddr *addr,
           socklen_t *addrlen)
```

- `sockfd`: file descriptor that `socket()` returned
- `addr`: pointer to store client address
- `addrlen`: size of `addr`
- Returns a file descriptor for the connected socket

For example:

```
int client = accept(sockfd,
                   (struct sockaddr_in *) &caddr, &crlen);
```

Server: Example for `listen()` and `accept()`

```
if (listen(sockfd, 5) < 0) {  
    printf("Error listening\n");  
    ...  
}  
  
crlen=sizeof(caddr);  
if ((clientfd=accept(sockfd,  
    (struct sockaddr *) &caddr, &crlen)) < 0) {  
    printf("Error accepting connection\n");  
    ...  
}
```

Server: Complete Example

```
int sockfd, clen, clientfd;
struct sockaddr_in saddr, caddr;
unsigned short port = 80;
if ((sockfd=socket(AF_INET, SOCK_STREAM, 0)) < 0) {
    printf("Error creating socket\n");
    ...
}

memset(&saddr, 0, sizeof(saddr));
saddr.sin_family = AF_INET;
saddr.sin_addr.s_addr = htonl(INADDR_ANY);
saddr.sin_port = htons(port);

if ((bind(sockfd, (struct sockaddr *) &saddr, sizeof(saddr)) < 0) {
    printf("Error binding\n");
    ...
}

if (listen(sockfd, 5) < 0) {
    printf("Error listening\n");
    ...
}

clen=sizeof(caddr);
if ((clientfd=accept(sockfd, (struct sockaddr *) &caddr, &clen)) < 0) {
    printf("Error accepting connection\n");
    ...
}
```

Practical Work 3: Server setup

- Write a new program in C
 - Name it « 03.practical.work.server.setup.c »
 - Write a server that:
 - listens to **TCP** port **8784** [USTH in a T9 dial pad!]
 - binds to all possible interfaces
 - prints a message when a client connects to it
- Deploy to your shiny VPS
- Test the connection
 - Use «telnet» or «nc»
- Push your C program to corresponding forked Github repository