

## Network Programming

---

- Network Programming as Programming across Machine Boundaries
  - The Sockets API
  - Reliable Communication Channels: TCP
  - Dangerous at any Speed; connectionless communication and UDP
  - Server Design
  
  - Reading: R&R, Ch 18
- 

## Naming in a Networked Environment

---

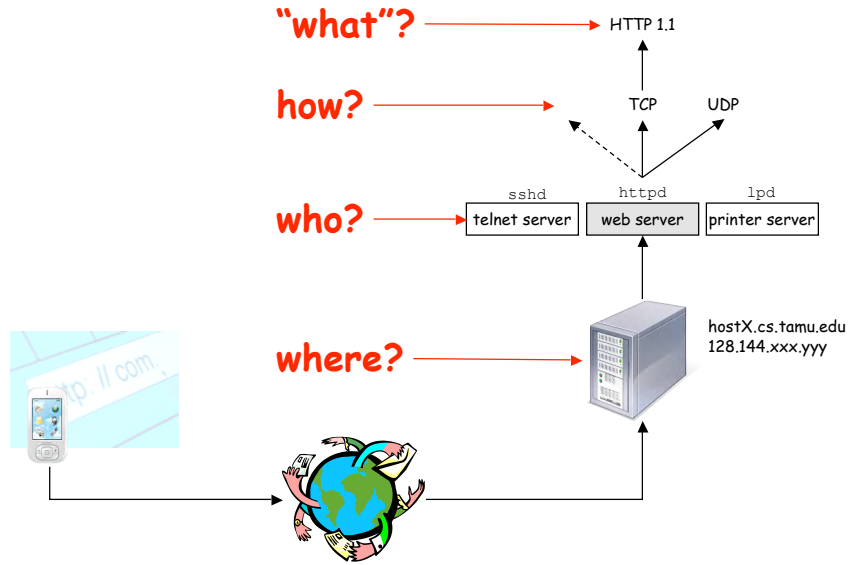
Reminder:

- Naming within a **single address space**?
  - addresses, duh!
- Naming **across address spaces**?
  - file descriptors
  - filenames (use file system as name space)
  - keys

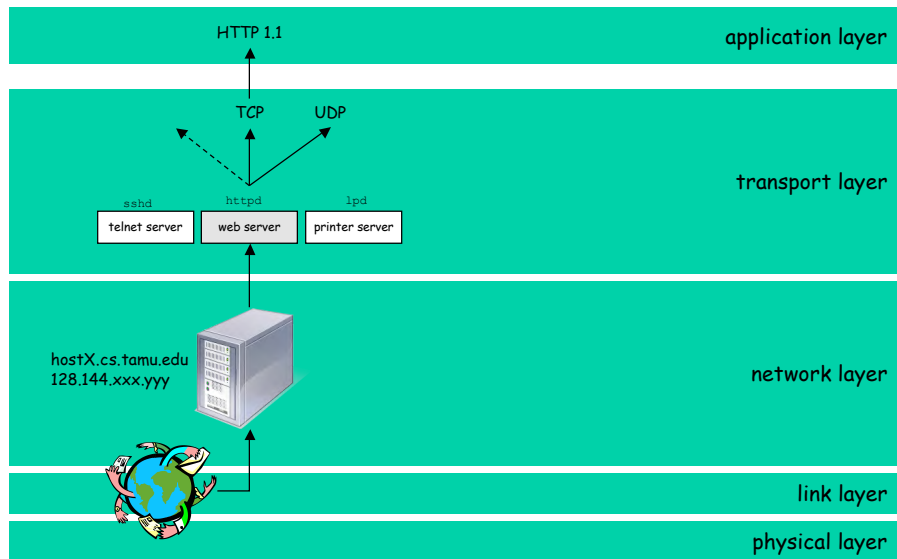
New:

- Naming **across machine boundaries**?
-

## Naming in a Networked Environment (II)



## Protocols for a Networked Environment

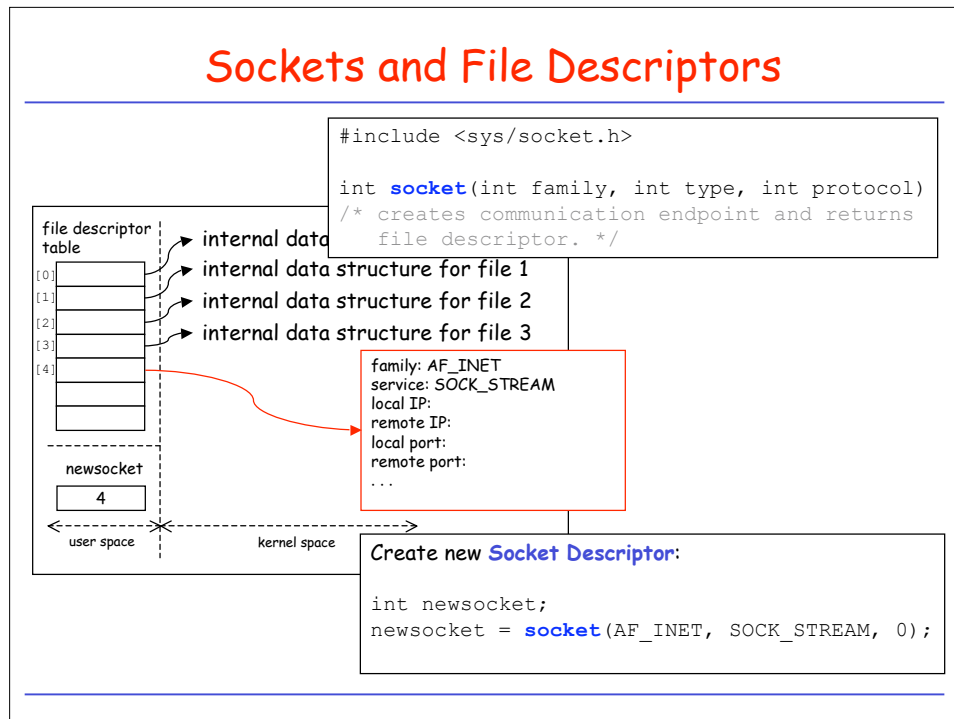


## The Socket API

- |   |   |
|---|---|
| <ul style="list-style-type: none"> <li>• What does an API need to support? [Comer&amp;Stevens]             <ul style="list-style-type: none"> <li>- allocate local resources (buffers)</li> <li>- specify local and remote communication endpoints</li> <li>- initiate a connection</li> <li>- wait for an incoming connection</li> <li>- send or receive data</li> <li>- determine when data arrives</li> <li>- generate urgent data</li> <li>- handle incoming urgent data</li> <li>- terminate a connection gracefully</li> <li>- handle connection termination from the remote site</li> <li>- abort communication</li> <li>- handle error conditions or a connection abort</li> <li>- release local resources when communication finishes</li> </ul> </li> </ul> | <ul style="list-style-type: none"> <li>• Existing TCP/IP APIs:             <ul style="list-style-type: none"> <li>- <b>Berkeley Socket API</b> <ul style="list-style-type: none"> <li>• aka socket API, socket interface, sockets</li> <li>• adapted by Linux and others</li> </ul> </li> <li>- <b>Windows Sockets</b></li> <li>- <b>System V Unix TLI</b> (Transport Layer Interface)</li> </ul> </li> </ul> |
|---|---|

## Specifying a Protocol Interface

- Reality Check: "All" networks today are based on TCP/IP.
- How to define a network API, then?!
  - **Approach 1:** Define functions that specifically support TCP/IP communication.
    - e.g. `makeTCPconnection(int32 host, int16 portno);`
  - **Approach 2:** Define functions that support network communication in general, and use parameters to handle TCP/IP as a special case.
- The socket API provides **generalized functions** that support network communication using **many possible protocols**.
- The programmer specifies the **type of service required** rather than the name of a **specific network protocol**.



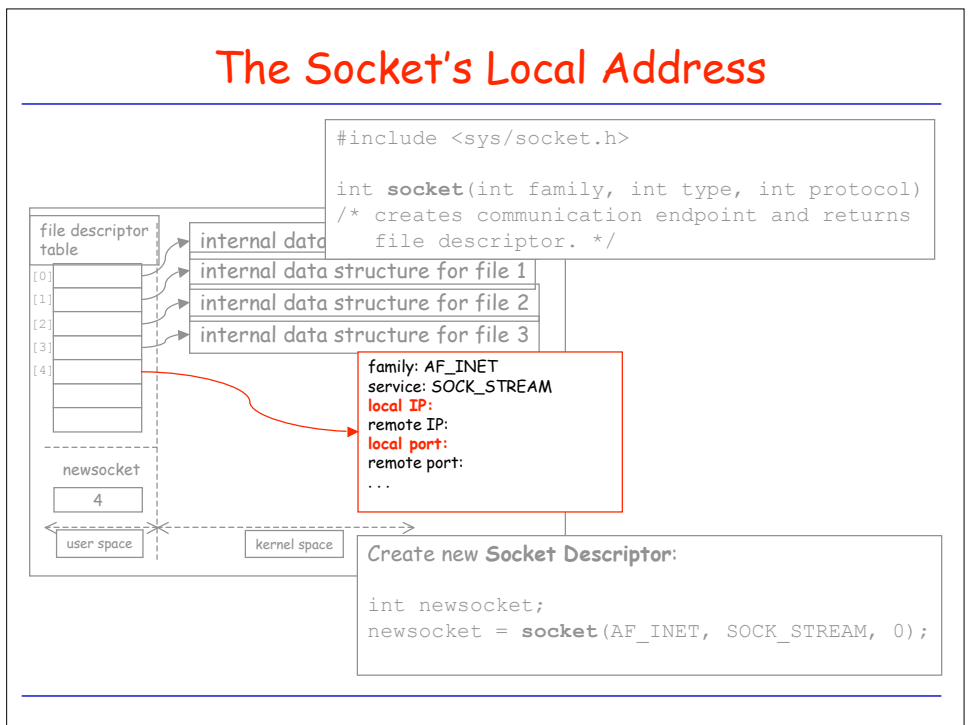
## TCP? UDP?!!

- **Connection-oriented Style** (`SOCK_STREAM`) of communication
- Implemented in TCP/IP as the **Transport Control Protocol (TCP)**.
- TCP provides full **reliability**:
  - verifies that data arrives, with **automatic retransmission**
  - computes **checksums** to detect corruption
  - uses sequence numbers to **guarantee ordering** of received packets
  - automatically **eliminates duplicated packets**
  - provides **flow control** (ensures that sender does not send more packets than the receiver can handle)
  - informs sender if network becomes inoperable for any reason
- TCP **protects network resources**:
  - provides **congestion control** (throttles transmission when it detects that network is congested)
- What is the **cost** of all this?! Connection establishment overhead.

## TCP? UDP?!!

- **Connectionless Style** (SOCK\_DGRAM) of communication
- Implemented in TCP/IP as the **User Datagram Protocol (UDP)**.
- UDP provides **no guarantee about reliable delivery**:
  - packets can be **lost, duplicated, delayed, or delivered out of order**
- UDP works well if the underlying network works well, e.g. local network
  
- **In practice**, programmers use UDP only when:
  1. The application requires that UDP must be used. (The application has been designed to handle reliability and delivery errors.)
  2. The application relies on hardware broadcast or multicast for delivery.
  3. Application runs in reliable, local environment, and overhead for reliability is unnecessary.

## The Socket's Local Address



## Defining the Socket's Local Address (server)

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

int bind(      int          socket,
              const struct sockaddr * address,
              socketlen_t   address_len);

/* associates socket file descriptor with
   communication endpoint address. */
```

Generalized socket address:

*(address family, endpoint address in that family)*

Examples:

- Address family **AF\_UNIX** has named pipes. Endpoint address of named pipes?

- Socket Addresses in **AF\_INET**?

```
struct sockaddr_in {          /* struct to hold an address */
    u_char    sin_len;       /* total length */
    u_short   sin_family;    /* type of address */
    u_short   sin_port;      /* protocol port number */
    struct in_addr sin_addr; /* IP address (declared to be */
                                /* u_long in some systems) */
    char      sin_zero[8];   /* unused (set to zero) */
};
```

## Host vs. Network Byte Order

- Big-endian vs. Little-endian.
- Network representation requires big-endian.
- Portability?!

```
#include <arpa/inet.h>

uint32_t htonl(uint32_t hostlong);
uint32_t ntohl(uint32_t netlong);
uint16_t htons(uint16_t hostshort);
uint16_t ntohs(uint16_t netshort);
```

Associate port 8652 with a socket:

```
struct sockaddr_in server;

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

server.sin_family      = AF_INET;
server.sin_addr.s_addr = htonl(INADDR_ANY);
server.sin_port        = htons((short)8652);

bind(sock, &server, sizeof(server));
```

## Prepare to Accept Incoming Connections (server)

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

int listen(int socket, int backlog);
/* specify willingness to accept incoming
   connections at given socket, with given
   queue limit.
   Connections are then accepted with
   accept(). */
```

When request for connection from client comes in, both client and server execute **hand-shake procedure** to set up the connection.

When server is busy, two things can happen

- Connection request is queued (as long as **backlog** not exceeded)
- or
- Connection request is refused (client receives **ECONNREFUSED** error)

## Handle Incoming Connections (server)

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

int accept(int socket,
           struct sockaddr * address,
           socklen_t * address_len);
/* Accept a connection on a socket.
   Create a new socket with same properties of "socket" and
   return new file descriptor. */
```

- Extract first connection request on queue of pending connections.
- Create new socket with same properties of given socket.
- Returns new file descriptor for the socket.
- Blocks caller if no pending connections are present in queue.
- New socket may not be used to accept more connections.
- Original socket **socket** remains open.
- Argument **address** (result parameter) is filled in with the address of connecting entity.
- Argument **address\_len** (value-result parameter) initially contains length of space pointed to by **address**. On return it contains length of actual length of space.

## In the meantime, at the Client's End (client)

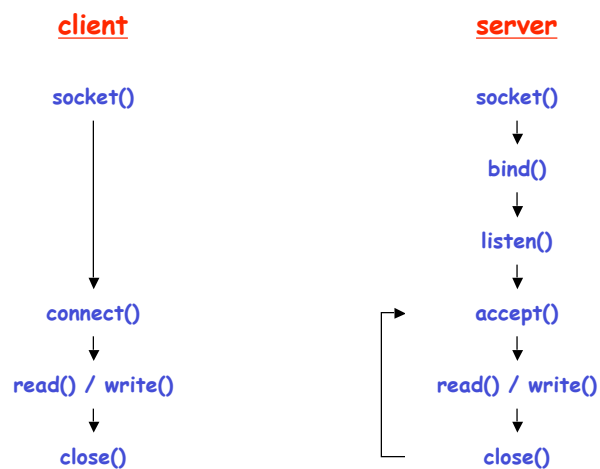
```
#include <sys/socket.h>

int connect(int socket,
            struct sockaddr * address,
            socklen_t address_len);
/* Initiate a connection on a socket.
   Structure address specifies the
   other endpoint of the connection. */
```

Note difference between `SOCK_DGRAM` and `SOCK_STREAM` sockets:

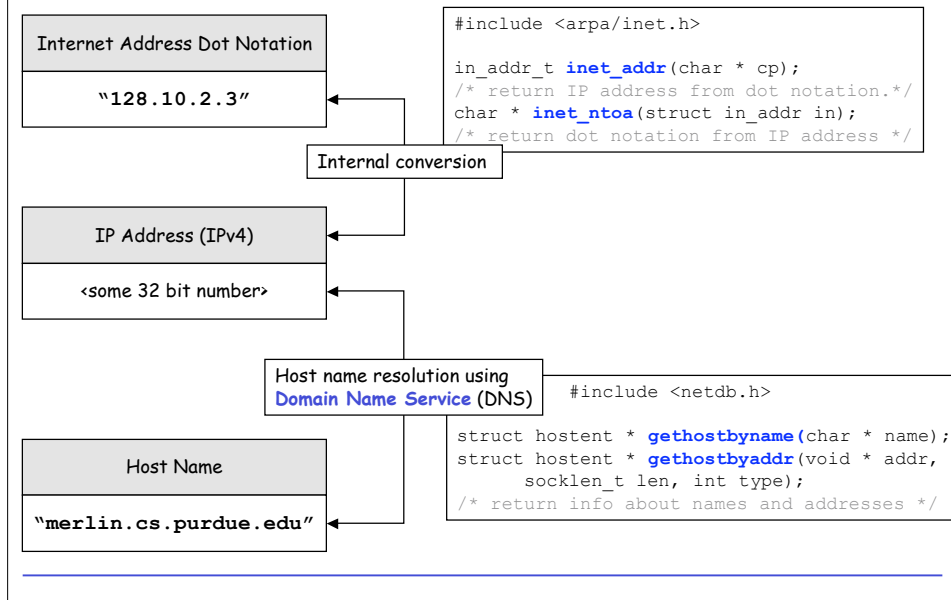
- For `SOCK_STREAM` sockets, `connect` attempts to establish connection with other socket.
  - Generally, stream sockets may connect only once.
- For `SOCK_DGRAM` sockets, `connect` specifies the peer with which to associate socket.
  - Datagram sockets may dissolve association by connecting to invalid address, such as null address.
- For `address`, fill in `family`, `address`, and `port number`.

## Connection Establishment: Summary





## Addressing: IP Address vs. Dotted vs. Name



## Host Names and IP Addresses

### Host Information:

```
#include <netdb.h>

struct hostent {
    char * h_name; /* canonical name of host */
    char ** h_aliases; /* alias list */
    int h_addrtype; /* host address type */
    int h_length; /* length of address */
    char ** h_addr_list; /* list of addresses */
}
```

Example: Translate a host name into IP address for use in `connect()` call:

```
struct hostent * hp;
struct sockaddr_in the_server;

if ((hp = gethostbyname("www.cs.tamu.edu")) == NULL)
    fprintf(stderr, "Failed to resolve host name\n");
else
    memcpy((char*)&the_server.sin_addr.s_addr,
        hp->h_addr_list[0], hp->h_length);
```

## Example TCP Client: DAYTIME client [Comer]

```
#define LINELEN 128

/* forward */
int connectTCP(const char * host, const char * service);

/* main program */
int main(argc, char * argv) {

    char * host    = "localhost"; /* use local host if none supplied */
    char * service = "daytime";  /* default service port */

    if (argc > 1) host    = argv[1];
    if (argc > 2) service = argv[2];

    int s = connectTCP(host, service);

    while ( (int n = read(s, buf, LINELEN)) > 0) {
        buf[n] = '\0'; /* ensure null terminated */
        (void) fputs(buf, stdout);
    }
}
```

## Example TCP Client: (cont...)

```
#define LINELEN 128

int connectTCP(const char * host , const char * service) {
    struct sockaddr_in sin; /* Internet endpoint address */
    memset(&sin, 0, sizeof(sin));
    sin.sin_family = AF_INET;

    /* Map service name to port number */
    if (struct servent * pse = getservbyname(service, "tcp") )
        sin.sin_port = pse->s_port;
    else if ((sin.sin_port = htons((unsigned short)atoi(service))) == 0)
        fprintf(stderr, "can't get <%s> service entry\n", service);

    /* Map host name to IP address, allowing for dotted decimal */
    if (struct hostent * phe = gethostbyname(host) )
        memcpy(&sin.sin_addr, phe->h_addr, phe->h_length);
    else if ((sin.sin_addr.s_addr = inet_addr(host)) == INADDR_NONE )
        fprintf(stderr, "can't get <%s> host entry\n", host);

    /* Allocate socket */
    int s = socket(AF_INET, SOCK_STREAM, 0);
    if (s < 0) fprintf(stderr, "can't create socket: %s\n", strerror(errno));

    /* Connect the socket */
    if (connect(s, (struct sockaddr *)&sin, sizeof(sin)) < 0)
        fprintf(stderr, "can't connect to %s.%s: %s\n", host, service, strerror(errno));

    return s;
}

/* forward */
int connectTCP(const char * host , const char * service) {
    struct sockaddr_in sin; /* Internet endpoint address */
    memset(&sin, 0, sizeof(sin));
    sin.sin_family = AF_INET;

    /* Map service name to port number */
    if (struct servent * pse = getservbyname(service, "tcp") )
        sin.sin_port = pse->s_port;
    else if ((sin.sin_port = htons((unsigned short)atoi(service))) == 0)
        fprintf(stderr, "can't get <%s> service entry\n", service);

    /* Map host name to IP address, allowing for dotted decimal */
    if (struct hostent * phe = gethostbyname(host) )
        memcpy(&sin.sin_addr, phe->h_addr, phe->h_length);
    else if ((sin.sin_addr.s_addr = inet_addr(host)) == INADDR_NONE )
        fprintf(stderr, "can't get <%s> host entry\n", host);

    /* Allocate socket */
    int s = socket(AF_INET, SOCK_STREAM, 0);
    if (s < 0) fprintf(stderr, "can't create socket: %s\n", strerror(errno));

    /* Connect the socket */
    if (connect(s, (struct sockaddr *)&sin, sizeof(sin)) < 0)
        fprintf(stderr, "can't connect to %s.%s: %s\n", host, service, strerror(errno));

    return s;
}

/* main program */
int main(argc, char * argv) {

    char * host    = "localhost"; /* use local host if none supplied */
    char * service = "daytime";  /* default service port */

    if (argc > 1) host    = argv[1];
    if (argc > 2) service = argv[2];


    int s = connectTCP(host, service);

    while ( (int n = read(s, buf, LINELEN)) > 0) {
        buf[n] = '\0'; /* ensure null terminated */
        (void) fputs(buf, stdout);
    }
}
```

## Server Software Design: Issues [Comer]

- **Concurrent vs. Iterative Servers:**  
The term **concurrent server** refers to whether the server permits multiple requests to proceed concurrently, **not** to whether the underlying implementation uses multiple, concurrent threads of execution.  
Iterative server implementations are easier to build and understand, but may result in poor performance because they make clients wait for service.
- **Connection-Oriented vs. Connectionless Access:**  
Connection-oriented (TCP, typically) servers are easier to implement, but have resources bound to connections.  
Reliable communication over UDP is not easy!
- **Stateful vs. Stateless Servers:**  
How much information should the server maintain about clients? (What if clients crash, and server does not know?)

## Example: Iterative, Connection-Oriented Server



```

int passiveTCPsock(const char * service, int backlog) {
    struct sockaddr_in sin;          /* Internet endpoint address */
    memset(&sin, 0, sizeof(sin));    /* Zero out address */
    sin.sin_family = AF_INET;
    sin.sin_addr.s_addr = INADDR_ANY;

    /* Map service name to port number */
    if (struct servent * pse = getservbyname(service, "tcp"))
        sin.sin_port = pse->s_port;
    else if ((sin.sin_port = htons((unsigned short)atoi(service))) == 0)
        perror("can't get <service> service entry\n", service);

    /* Allocate socket */
    int s = socket(AF_INET, SOCK_STREAM, 0);
    if (s < 0) perror("can't create socket: %s\n", strerror(errno));

    /* Bind the socket */
    if (bind(s, (struct sockaddr *)&sin, sizeof(sin)) < 0)
        perror("can't bind to ...");

    /* Listen on socket */
    if (listen(s, backlog) < 0)
        perror("can't listen on ...");

    return s;
}
  
```

### Example: Iterative, Connection-Oriented Server

The diagram illustrates the flow of an iterative server. On the left, a vertical flowchart labeled 'server' shows the sequence of operations: `socket()`, `bind()`, `listen()`, `accept()`, `read() / write()`, and `close()`. A red box highlights the `accept()`, `read() / write()`, and `close()` steps, with a feedback arrow indicating a loop. On the right, the C code implements this logic. The `main` function sets up a master socket (`m_sock`) and enters a loop where it repeatedly calls `accept()` to handle incoming connections. Each connection is processed by reading/writing data and then closing the socket (`s_sock`).

```

int main(int argc, char * argv[]) {
    char * service = "daytime"; /* service name or port number */
    int m_sock, s_sock; /* master and slave socket */
    service = argv[1];

    int m_sock = passiveTCPsock(service, 32);

    for (;;) {
        s_sock = accept(m_sock, (struct sockaddr*)&fsin, sizeof(fsin));
        if (s_sock < 0) errexit("accept failed: %s\n", strerror(errno));

        time_t now;
        time(&now);
        char * pts = ctime(&now);
        write(s_sock, pts, strlen(pts));

        close(s_sock);
    }
}
    
```

### Example: Concurrent, Connection-Oriented Server

The diagram illustrates the flow of a concurrent server. On the left, a vertical flowchart labeled 'server' shows the sequence of operations: `socket()`, `bind()`, `listen()`, `accept()`, `fork()`, and `close()`. A red box highlights the `socket()`, `bind()`, and `listen()` steps. A feedback arrow loops from the `close()` step back to `accept()`. On the right, the C code implements this logic. The `main` function sets up a master socket (`m_sock`) and enters a loop where it repeatedly calls `accept()`. For each connection, it uses `fork()` to create a child process that handles the request, while the parent process continues to accept new connections. The child process handles the request and then exits, and the parent process closes the socket (`s_sock`).

```

int passiveTCPsock(const char * service, int backlog);

int main(int argc, char * argv[]) {
    char * service = "daytime"; /* service name or port number */
    int m_sock, s_sock; /* master and slave socket */
    service = argv[1];

    int m_sock = passiveTCPsock(service, 32);

    for (;;) {
        s_sock = accept(m_sock, (struct sockaddr*)&fsin, sizeof(fsin));
        if (s_sock < 0) errexit("accept failed: %s\n", strerror(errno));

        if (fork() == 0) { /* child */
            close(m_sock);
            /* handle request here . . . */
            exit(error_code);
        }
        close(s_sock);
    }
}
    
```

### Example: Concurrent, Connection-Oriented Server

The diagram illustrates a concurrent server implementation using `fork()`. On the left, a vertical flowchart labeled "server" shows the sequence of operations: `socket()`, `bind()`, `listen()`, `accept()`, `fork()`, and `close()`. The `socket()`, `bind()`, and `listen()` steps are enclosed in a red box. Arrows indicate the flow from `accept()` to `fork()`, and from `fork()` to a separate `close()` box. A feedback arrow loops from the `close()` box back to `accept()`. On the right, the corresponding C code is shown. The `main` function calls `passiveTCPsock` to create a master socket, sets a signal handler for `SIGCHLD` to `cleanly_terminate_child`, and enters a loop where it repeatedly calls `accept`. For each accepted connection, it calls `fork` to create a child process to handle the request, then closes the master socket. The `cleanly_terminate_child` function is shown as a separate block that waits for the child process to terminate.

```

int passiveTCPsock(const char * service, int backlog);

int main(int argc, char * argv[] ) {
    char * service = "daytime"; /* service name or port number */
    int m_sock, s_sock; /* master and slave socket */
    service = argv[1];

    int m_sock = passiveTCPsock(service, 32);

    signal(SIGCHLD, cleanly_terminate_child);

    for (;;) {
        s_sock = accept(m_sock, (struct sockaddr*)&fsin, sizeof(fsin));
        if (s_sock < 0)
            if (errno == EINTR) continue;
            else errexit("accept failed: %s\n", strerror(errno));
        if (fork() == 0) { /* child */
            close(m_sock);
            /* handle request here . . . */
        }
        close(s_sock);
    }
}

void cleanly_terminate_child(int sig) {
    int status;
    while (wait3(&status, WNOHANG, NULL) >= 0)
    }
    
```

### Example: Concurrent, Connection-Oriented Server

The diagram illustrates a concurrent server implementation using `pthread_create()`. On the left, a vertical flowchart labeled "server" shows the sequence of operations: `socket()`, `bind()`, `listen()`, `accept()`, `pthread_create()`, and `close()`. The `socket()`, `bind()`, and `listen()` steps are enclosed in a red box. Arrows indicate the flow from `accept()` to `pthread_create()`, and from `pthread_create()` to a separate `close()` box. A feedback arrow loops from the `close()` box back to `accept()`. On the right, the corresponding C code is shown. The `main` function calls `passiveTCPsock` to create a master socket, then initializes `pthread` attributes and sets the state to `PTHREAD_CREATE_DETACHED`. It enters a loop where it repeatedly calls `accept`. For each accepted connection, it calls `pthread_create` to create a detached thread to handle the request, then closes the master socket. The `handle_request` function is shown as a separate block that simply closes the file descriptor.

```

int passiveTCPsock(const char * service, int backlog);

int main(int argc, char * argv[] ) {
    char * service = "daytime"; /* service name or port number */
    int m_sock, s_sock; /* master and slave socket */
    service = argv[1];

    int m_sock = passiveTCPsock(service, 32);

    pthread_t th; pthread_attr_t ta;
    pthread_attr_init(&ta);
    pthread_attr_setdetachstate(&ta, PTHREAD_CREATE_DETACHED);

    for (;;) {
        s_sock = accept(m_sock, (struct sockaddr*)&fsin, sizeof(fsin));
        if (s_sock < 0)
            if (errno == EINTR) continue;
            else errexit("accept failed: %s\n", strerror(errno));

        pthread_create(&th, &ta, handle_request, (void*)s_sock);
    }
}

int handle_request(int fd) {
    /* handle the request . . . */
    close(fd);
}
    
```

## Example: Concurrent, Connection-Oriented Server

