


CS162
Operating Systems and
Systems Programming
Lecture 5

Abstractions 3: IPC, Pipes and Sockets
A quick programmer's viewpoint


September 14th, 2020
Prof. John Kubitowicz
<http://cs162.eecs.Berkeley.edu>

Goals for Today: IPC and Sockets

- **Key Idea:** Communication between processes and across the world looks like File I/O
- Introduce Pipes and Sockets
- Introduce TCP/IP Connection setup for Webserver



`write(wfd, wbuf, wlen);`



`n = read(rfd, rbuf, rmax);`

9/14/20

Kubitowicz CS162 © UCB Fall 2020

Lec 5.2

Recall: Creating Processes with `fork()`

- `pid_t fork()` – copy the current process
 - State of original process **duplicated** in Parent and Child!
 - Address Space (Memory), File Descriptors, etc...
- Return value from `fork()`: pid (like an integer)
 - When > 0:
 - » Running in (original) **Parent** process
 - » return value is **pid** of new child
 - When = 0:
 - » Running in new **Child** process
 - When < 0:
 - » Error! Must handle somehow
 - » Running in original process

```
int status;
pid_t tcpid;
...
cpid = fork();
if (cpid > 0) {
    mypid = getpid();
    printf("[%d] parent of [%d]\n", mypid, cpid);
    tcpid = wait(&status);
    printf("[%d] bye %d(%d)\n", mypid, tcpid, status);
} else if (cpid == 0) {
    mypid = getpid();
    printf("[%d] child\n", mypid);
    exit(42);
}
...
```

- **WHY FORK?**
 - (mostly true) without `fork()`, you cannot create new processes!
 - Fork was the original mechanism for creating concurrency in UNIX (long before Linux!)
 - See, however, `Linux clone()` which gives you more flexibility

9/14/20

Kubitowicz CS162 © UCB Fall 2020

Lec 5.3

Recall: Key Unix I/O Design Concepts

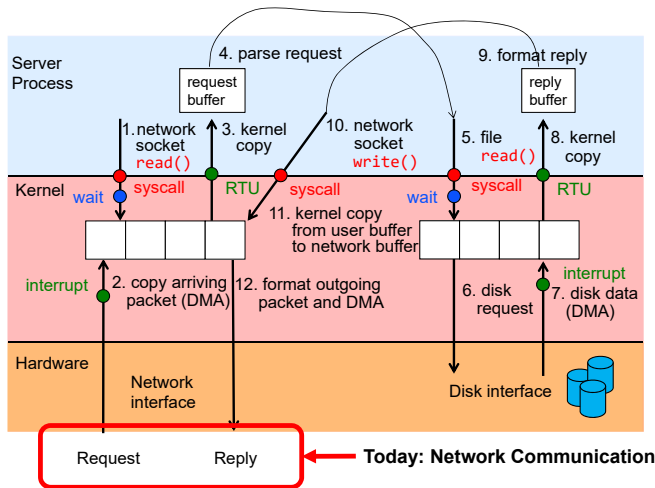
- **Uniformity – Everything Is a File!**
 - file operations, device I/O, and interprocess communication through `open`, `read/write`, `close`
 - Allows simple composition of programs
 - » `find | grep | wc ...`
- **Open before use**
 - Provides opportunity for access control and arbitration
 - Sets up the underlying machinery, i.e., data structures
- **Byte-oriented**
 - Even if blocks are transferred, addressing is in bytes
- **Kernel buffered reads**
 - Streaming and block devices looks the same, read blocks yielding processor to other task
- **Kernel buffered writes**
 - Completion of out-going transfer decoupled from the application, allowing it to continue
- **Explicit close**

9/14/20

Kubitowicz CS162 © UCB Fall 2020

Lec 5.4

Putting it together: web server

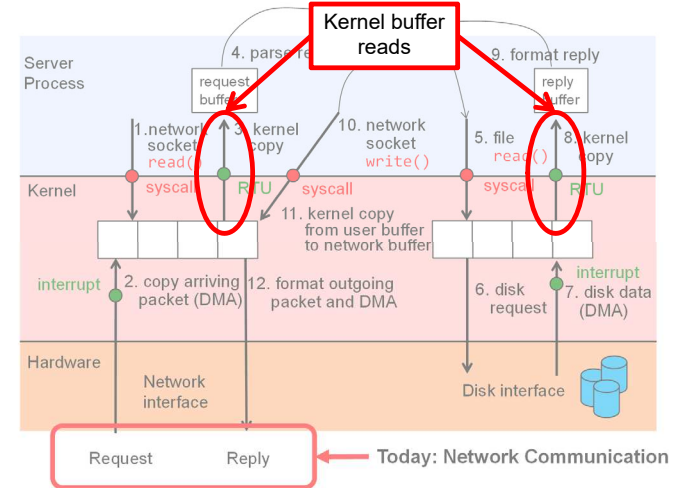


9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.5

Putting it together: web server

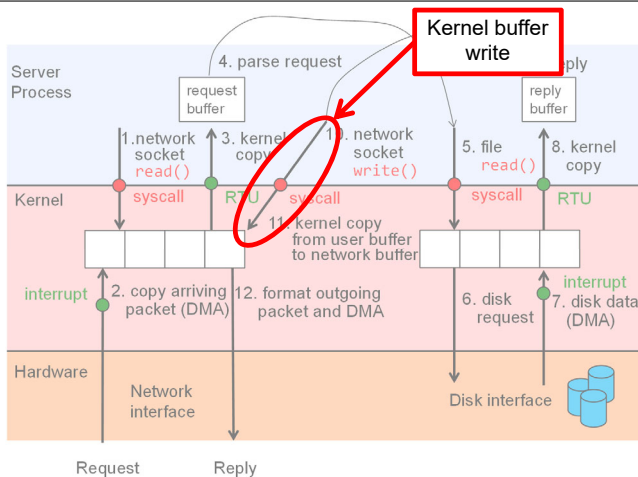


9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.6

Putting it together: web server



9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.7

Recall: C High-Level File API – Streams

- Operates on “streams” – unformatted sequences of bytes (with text or binary data), with a position:



```
#include <stdio.h>
FILE *fopen( const char *filename, const char *mode );
int fclose( FILE *fp );
```

Mode	Text	Binary	Descriptions
r		rb	Open existing file for reading
w		wb	Open for writing; created if does not exist
a		ab	Open for appending; created if does not exist
r+		rb+	Open existing file for reading & writing.
w+		wb+	Open for reading & writing; truncated to zero if exists, create otherwise
a+		ab+	Open for reading & writing. Created if does not exist. Read from beginning, write as append

- Open stream represented by pointer to a FILE data structure
 - Error reported by returning a NULL pointer
 - Pointer used in subsequent operations on the stream
 - Data buffered in user space

9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.8

Recall: Low-Level File I/O: The RAW system-call interface

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

int open (const char *filename, int flags [, mode_t mode])
int creat (const char *filename, mode_t mode)
int close (int filedes)
```

Bit vector of:

- Access modes (Rd, Wr, ...)
- Open Flags (Create, ...)
- Operating modes (Appends, ...)

Bit vector of Permission Bits:

- User|Group|Other X R|W|X

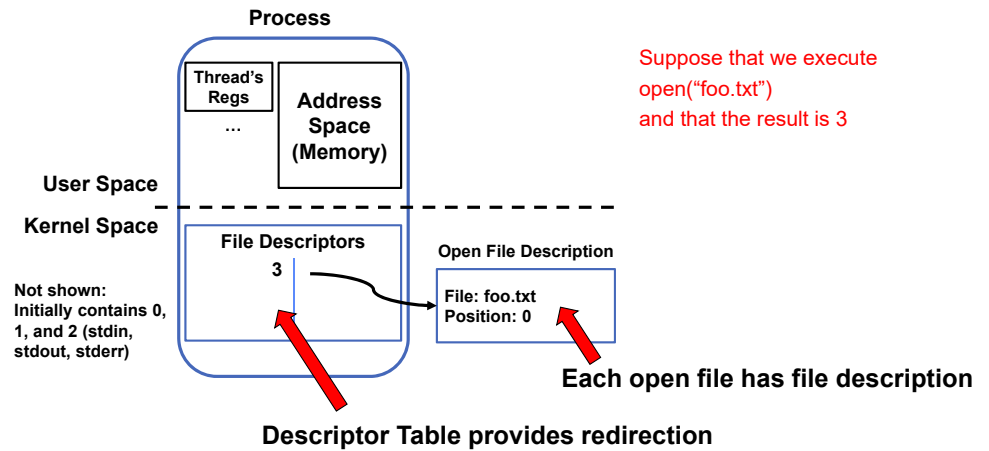
- Integer return from `open()` is a **file descriptor**
 - Error indicated by return < 0 : the global `errno` variable set with error
 - File Descriptor used in subsequent operations on the file
- Streams (opened with `fopen()`) have a file descriptor *inside of them!*
 - Retrievable with `fileno(FILE *stream)` \Rightarrow internal file descriptor

9/14/20

Kubiawicz CS162 © UCB Fall 2020

Lec 5.9

Recall: Representation of a Process (inside kernel!)

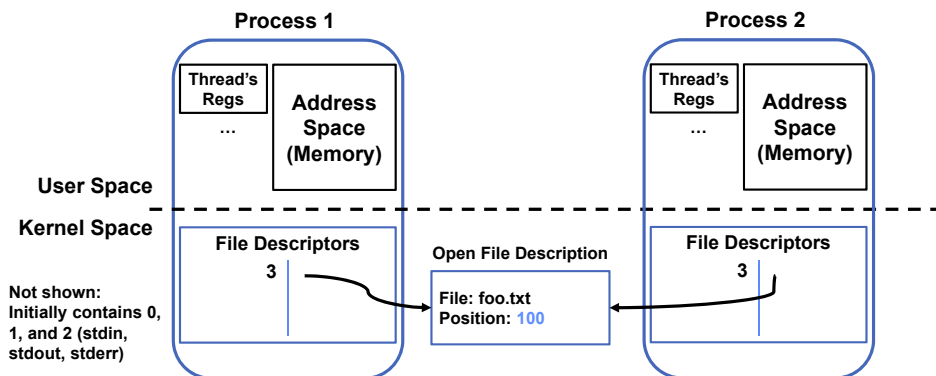


9/14/20

Kubiawicz CS162 © UCB Fall 2020

Lec 5.10

Recall: What Happens on fork()?



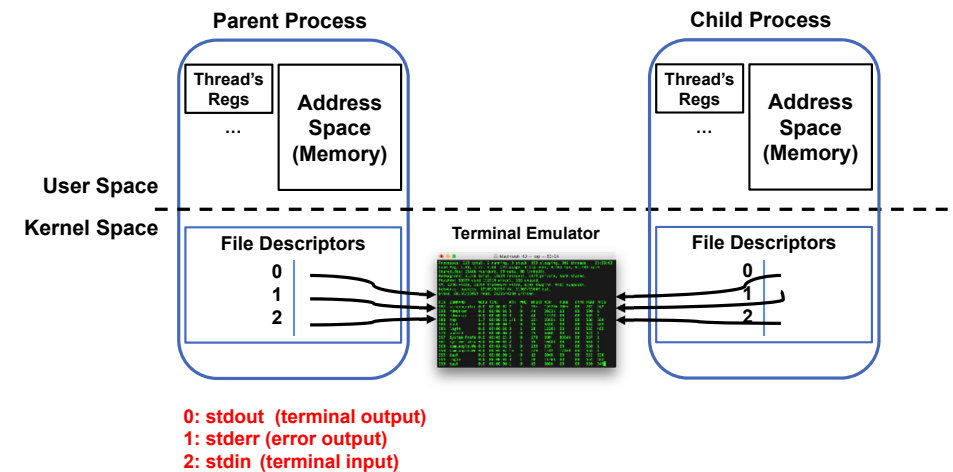
- After `fork()`:
 - File Descriptors *copied*: child has same descriptor table as parent!
 - File Descriptions *shared*: child and parent can both manipulate/change open files

9/14/20

Kubiawicz CS162 © UCB Fall 2020

Lec 5.11

Recall standard file descriptors: 0, 1, 2



9/14/20

Kubiawicz CS162 © UCB Fall 2020

Lec 5.12

Administrivia

- Homework 1 due Wednesday
- Project 1 in full swing!
 - We expect that your design document will give intuitions behind your designs, not just a dump of pseudo-code
 - Think of this you are in a company and your TA is your manager
- Should be attending your permanent discussion section!
 - Remember to turn on your camera in Zoom
 - Discussion section attendance is mandatory
- Midterm 1: October 1st, 5-7PM (Three weeks from tomorrow!)
 - We understand that this partially conflicts with CS170, but those of you in CS170 can start that exam after 7PM (according to CS170 staff)
 - Video Proctored, No curve, Use of computer to answer questions
 - More details as we get closer to exam
- Start Planning on how your group will collaborate on projects!
 - Virtual Coffee Hours with your group (with camera)
 - Regular Brainstorming meetings?
 - Try to meet multiple times a week



9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.13

Today: Communication Between Processes

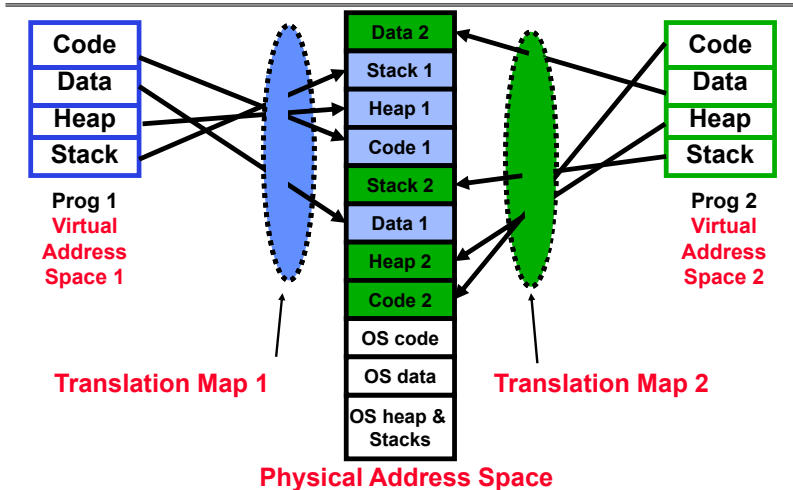
- What if processes wish to communicate with one another?
 - Why? Shared Task, Cooperative Venture with Security Implications
- Process Abstraction Designed to Discourage Inter-Process Communication!
 - Prevent one process from interfering with/stealing information from another
- So, must do something special (and agreed upon by both processes)
 - Must “Punch Hole” in security
- This is called “Interprocess Communication” (or IPC)

9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.14

Recall: Processes Protected from each other



9/14/20


Kubiatowicz CS162 © UCB Fall 2020

Lec 5.15

Communication Between Processes

- Producer (writer) and consumer (reader) may be distinct processes
 - Potentially separated in time
 - How to allow selective communication?
- Simple option: use a file!
 - We have already shown how parents and children share file descriptions:

`write(wfd, wbuf, wlen);`



`n = read(rfd, rbuf, rmax);`

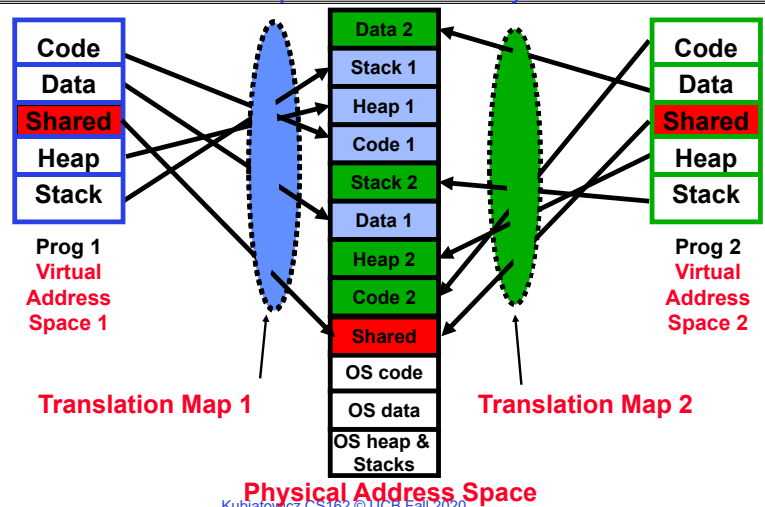
- Why might this be wasteful?
 - Very expensive if you only want transient communication (non-persistent)

9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.16

Shared Memory: Better Option? Topic for another day!



9/14/20


Kubiatowicz CS162 © UCB Fall 2020

Lec 5.17

Communication Between Processes (Another Option)

- Suppose we ask Kernel to help?
 - Consider an in-memory queue
 - Accessed via system calls (for security reasons):

```
write(wfd, wbuf, wlen);
```



```
n = read(rfd, rbuf, rmax);
```

- Data written by A is held in memory until B reads it
 - Same interface as we use for files!
 - Internally more efficient, since nothing goes to disk
- Some questions:
 - How to set up?
 - What if A generates data faster than B can consume it?
 - What if B consumes data faster than A can produce it?


9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.18

One example of this pattern: POSIX/Unix PIPE

```
write(wfd, wbuf, wlen);
```



```
n = read(rfd, rbuf, rmax);
```

- Memory Buffer is finite:
 - If producer (A) tries to write when buffer full, it *blocks* (Put sleep until space)
 - If consumer (B) tries to read when buffer empty, it *blocks* (Put to sleep until data)

```
int pipe(int fileds[2]);
```

- Allocates two new file descriptors in the process
- Writes to fileds[1] read from fileds[0]
- Implemented as a fixed-size queue

9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.19

Single-Process Pipe Example

```
#include <unistd.h>
int main(int argc, char *argv[])
{
    char *msg = "Message in a pipe.\n";
    char buf[BUFSIZE];
    int pipe_fd[2];
    if (pipe(pipe_fd) == -1) {
        fprintf(stderr, "Pipe failed.\n"); return EXIT_FAILURE;
    }
    ssize_t writelen = write(pipe_fd[1], msg, strlen(msg)+1);
    printf("Sent: %s [%ld, %ld]\n", msg, strlen(msg)+1, writelen);

    ssize_t readlen = read(pipe_fd[0], buf, BUFSIZE);
    printf("Rcvd: %s [%ld]\n", msg, readlen);

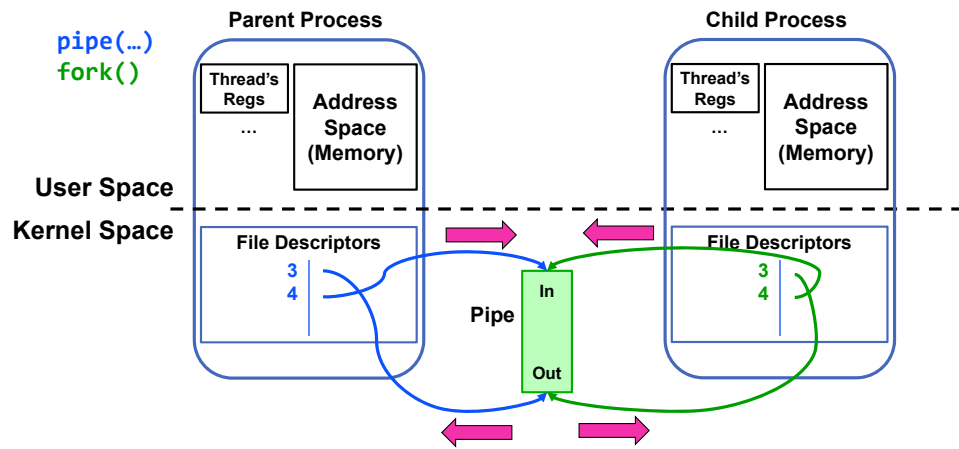
    close(pipe_fd[0]);
    close(pipe_fd[1]);
}
```

9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.20

Pipes Between Processes



9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.21

Inter-Process Communication (IPC): Parent ⇒ Child

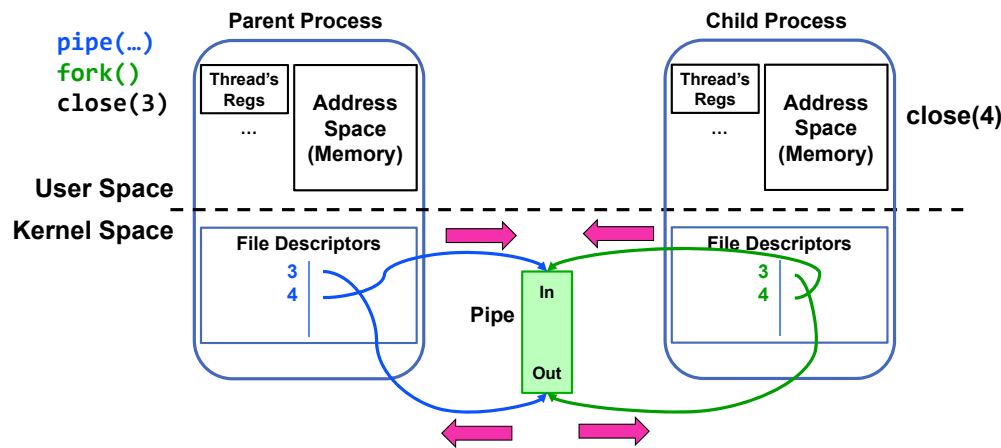
```
// continuing from earlier
pid_t pid = fork();
if (pid < 0) {
    fprintf(stderr, "Fork failed.\n");
    return EXIT_FAILURE;
}
if (pid != 0) {
    ssize_t writelen = write(pipe_fd[1], msg, msglen);
    printf("Parent: %s [%ld, %ld]\n", msg, msglen, writelen);
    close(pipe_fd[0]);
} else {
    ssize_t readlen = read(pipe_fd[0], buf, BUFSIZE);
    printf("Child Rcvd: %s [%ld]\n", msg, readlen);
    close(pipe_fd[1]);
}
```

9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.22

Channel from Parent ⇒ Child

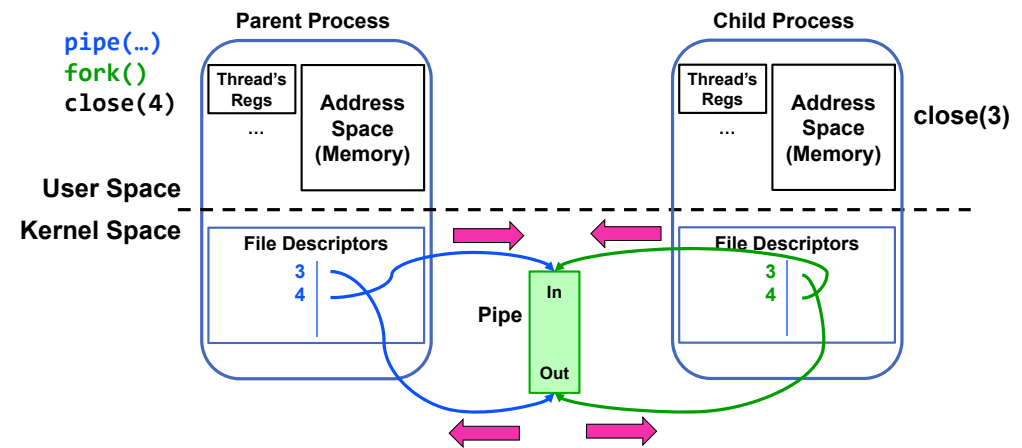


9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.23

Instead: Channel from Child ⇒ Parent



9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.24

When do we get EOF on a pipe?

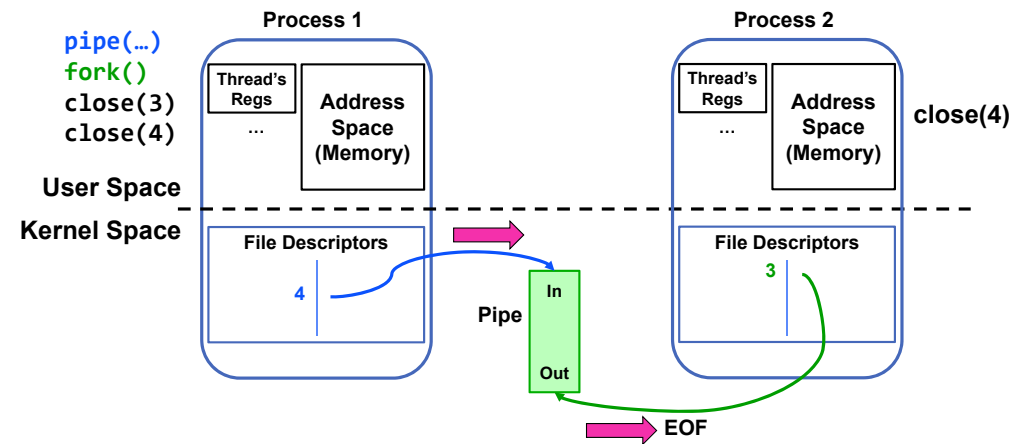
- After last “write” descriptor is closed, pipe is effectively closed:
 - Reads return only “EOF”
- After last “read” descriptor is closed, writes generate SIGPIPE signals:
 - If process ignores, then the write fails with an “EPIPE” error

9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.25

EOF on a Pipe



9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.26

Once we have communication, we need a *protocol*

- A protocol is an **agreement on how to communicate**
- Includes
 - **Syntax**: how a communication is specified & structured
 - » Format, order messages are sent and received
 - **Semantics**: what a communication means
 - » Actions taken when transmitting, receiving, or when a timer expires
- Described formally by a state machine
 - Often represented as a message transaction diagram
- In fact, across network may need a way to translate between different representations for numbers, strings, etc
 - Such translation typically part of a **Remote Procedure Call (RPC)** facility
 - Don't worry about this now, but it is clearly part of the *protocol*

9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.27

Examples of Protocols in Human Interaction


1. Telephone
2. (Pick up / open up the phone)
3. Listen for a dial tone / see that you have service
4. Dial
5. Should hear ringing ...
6. Callee: "Hello?"
7. Caller: "Hi, it's John..."
Or: "Hi, it's me" (what's that about?)
8. Caller: "Hey, do you think ... blah blah blah ..." pause
9. Callee: "Yeah, blah blah blah ..."
10. Caller: Bye
11. Callee: Bye
12. Hang up

9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.28

Web Server




9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.29

Client-Server Protocols: Cross-Network IPC



- Many clients accessing a common server
- File servers, www, FTP, databases

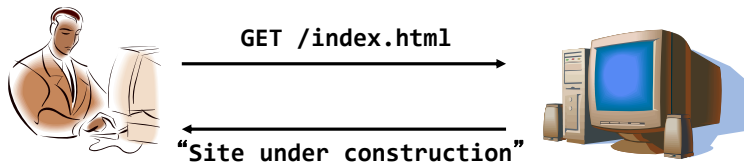
9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.30

Client-Server Communication

- Client is "sometimes on"
 - Sends the server requests for services when interested
 - E.g., Web browser on laptop/phone
 - Doesn't communicate directly with other clients
 - Needs to know server's address
- Server is "always on"
 - Services requests from many clients
 - E.g., Web server for `www.cnn.com`
 - Doesn't initiate contact with clients
 - Needs a fixed, well-known address



9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.31

What is a Network Connection?

- Bidirectional *stream* of bytes between two processes on possibly different machines
 - For now, we are discussing "TCP Connections"
- Abstractly, a connection between two endpoints A and B consists of:
 - A queue (bounded buffer) for data sent from A to B
 - A queue (bounded buffer) for data sent from B to A

9/14/20


Kubiatowicz CS162 © UCB Fall 2020

Lec 5.32

The Socket Abstraction: Endpoint for Communication

- **Key Idea:** Communication across the world looks like File I/O

```
write(wfd, wbuf, wlen);
```



- Sockets: Endpoint for Communication
 - Queues to temporarily hold results
- Connection: Two Sockets Connected Over the network \Rightarrow IPC over network!
 - How to **open()**?
 - What is the namespace?
 - How are they connected in time?

Sockets: More Details

- **Socket:** An abstraction for one endpoint of a network connection
 - Another mechanism for **inter-process communication**
 - Most operating systems (Linux, Mac OS X, Windows) provide this, even if they don't copy rest of UNIX I/O
 - Standardized by POSIX
- First introduced in 4.2 BSD (Berkeley Standard Distribution) Unix
 - This release had some huge benefits (and excitement from potential users)
 - Runners waiting at release time to get release on tape and take to businesses
- Same abstraction for any kind of network
 - Local (within same machine)
 - The Internet (TCP/IP, UDP/IP)
 - Things "no one" uses anymore (OSI, Appletalk, IPX, ...)

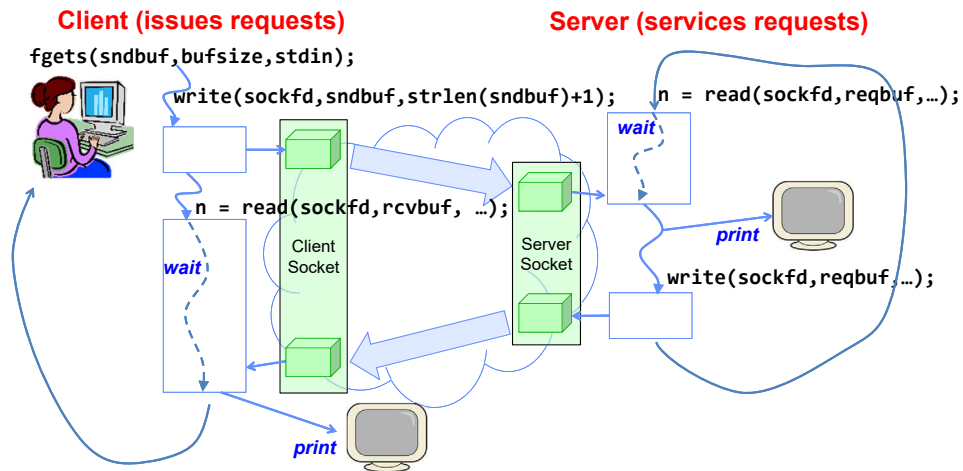
Sockets: More Details

- Looks just like a file with a **file descriptor**
 - Corresponds to a network connection (*two* queues)
 - **write** adds to output queue (queue of data destined for other side)
 - **read** removes from it input queue (queue of data destined for this side)
 - Some operations do not work, e.g. **lseek**
- How can we use sockets to support real applications?
 - A bidirectional byte stream isn't useful on its own...
 - May need messaging facility to partition stream into chunks
 - May need RPC facility to translate one environment to another and provide the abstraction of a function call over the network

Simple Example: Echo Server



Simple Example: Echo Server



9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.37

Echo client-server example

```

void client(int sockfd) {
    int n;
    char sndbuf[MAXIN]; char rcvbuf[MAXOUT];
    while (1) {
        fgets(sndbuf, MAXIN, stdin);          /* prompt */
        write(sockfd, sndbuf, strlen(sndbuf)+1); /* send (including null terminator) */
        memset(rcvbuf, 0, MAXOUT);           /* clear */
        n = read(sockfd, rcvbuf, MAXOUT);     /* receive */
        write(STDOUT_FILENO, rcvbuf, n);     /* echo */
    }
}
    
```

```

void server(int consockfd) {
    char reqbuf[MAXREQ];
    int n;
    while (1) {
        memset(reqbuf, 0, MAXREQ);
        len = read(consockfd, reqbuf, MAXREQ); /* Recv */
        if (n <= 0) return;
        write(STDOUT_FILENO, reqbuf, n);
        write(consockfd, reqbuf, n); /* echo */
    }
}
    
```

9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.38

What Assumptions are we Making?

- **Reliable**
 - Write to a file => Read it back. Nothing is lost.
 - Write to a (TCP) socket => Read from the other side, same.
 - Like pipes
- **In order (sequential stream)**
 - Write X then write Y => read gets X then read gets Y
- **When ready?**
 - File read gets whatever is there at the time.
 - Assumes writing already took place
 - Blocks if nothing has arrived yet
 - Like pipes!

9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.39

Socket Creation

- File systems provide a collection of permanent objects in a structured name space:
 - Processes open, read/write/close them
 - Files exist independently of processes
 - Easy to name what file to open()
- Pipes: one-way communication between processes on same (physical) machine
 - Single queue
 - Created transiently by a call to pipe()
 - Passed from parent to children (descriptors inherited from parent process)
- Sockets: two-way communication between processes on same or different machine
 - Two queues (one in each direction)
 - Processes can be on separate machines: no common ancestor
 - How do we *name* the objects we are opening?
 - How do these completely independent programs know that the other wants to “talk” to them?

9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.40

Namespaces for Communication over IP


- Hostname
 - www.eecs.berkeley.edu
- IP address
 - 128.32.244.172 (IPv4, 32-bit Integer)
 - 2607:f140:0:81::f (IPv6, 128-bit Integer)
- Port Number
 - 0-1023 are “well known” or “system” ports
 - » Superuser privileges to bind to one
 - 1024 – 49151 are “registered” ports (**registry**)
 - » Assigned by IANA for specific services
 - 49152–65535 ($2^{15}+2^{14}$ to $2^{16}-1$) are “dynamic” or “private”
 - » Automatically allocated as “ephemeral ports”

9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.41

Connection Setup over TCP/IP




- Special kind of socket: **server socket**
 - Has file descriptor
 - Can't read or write
- Two operations:
 1. **listen():** Start allowing clients to connect
 2. **accept():** Create a *new socket* for a *particular* client

9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.42

Connection Setup over TCP/IP



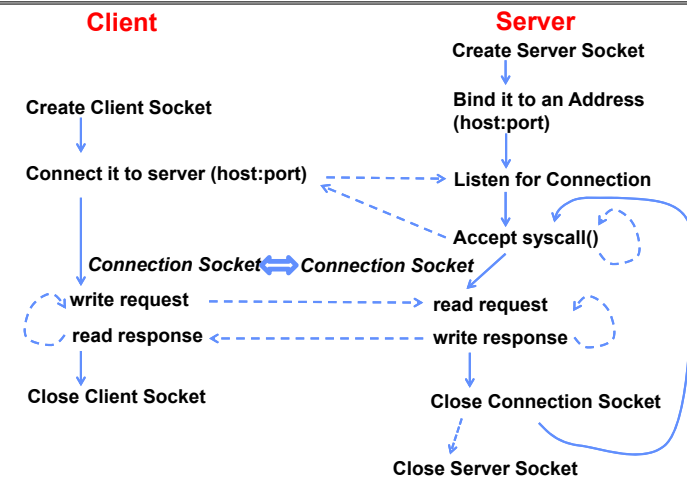
- 5-Tuple identifies each connection:
 1. Source IP Address
 2. Destination IP Address
 3. Source Port Number
 4. Destination Port Number
 5. Protocol (always TCP here)
- Often, Client Port “randomly” assigned
 - Done by OS during client socket setup
- Server Port often “well known”
 - 80 (web), 443 (secure web), 25 (sendmail), etc
 - Well-known ports from 0—1023

9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.43

Sockets in concept



9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.44

Client Protocol

```
char *host_name, *port_name;

// Create a socket
struct addrinfo *server = lookup_host(host_name, port_name);
int sock_fd = socket(server->ai_family, server->ai_socktype,
                    server->ai_protocol);

// Connect to specified host and port
connect(sock_fd, server->ai_addr, server->ai_addrlen);

// Carry out Client-Server protocol
run_client(sock_fd);

/* Clean up on termination */
close(sock_fd);
```

9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.45

Server Protocol (v1)

```
// Create socket to listen for client connections
char *port_name;
struct addrinfo *server = setup_address(port_name);
int server_socket = socket(server->ai_family,
                          server->ai_socktype, server->ai_protocol);
// Bind socket to specific port
bind(server_socket, server->ai_addr, server->ai_addrlen);
// Start listening for new client connections
listen(server_socket, MAX_QUEUE);

while (1) {
    // Accept a new client connection, obtaining a new socket
    int conn_socket = accept(server_socket, NULL, NULL);
    serve_client(conn_socket);
    close(conn_socket);
}
close(server_socket);
```

9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.46

How Could the Server Protect Itself?

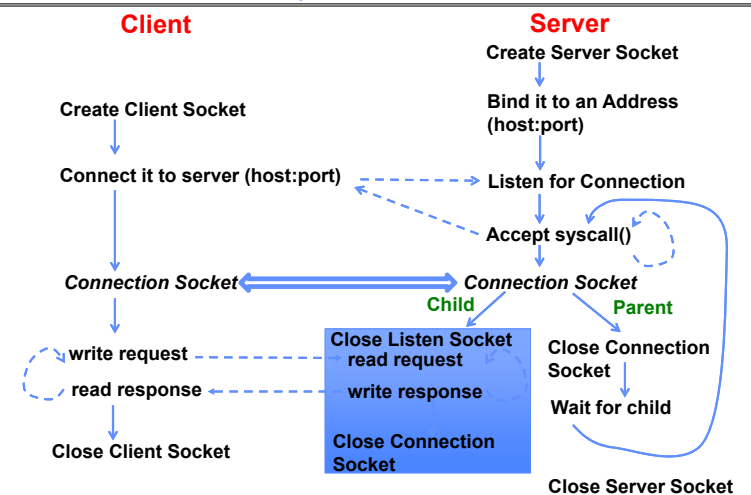
- Handle each connection in a separate process

9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.47

Sockets With Protection (each connection has own process)



9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.48

Server Protocol (v2)

```
// Socket setup code elided...
while (1) {
  // Accept a new client connection, obtaining a new socket
  int conn_socket = accept(server_socket, NULL, NULL);
  pid_t pid = fork();
  if (pid == 0) {
    close(server_socket);
    serve_client(conn_socket);
    close(conn_socket);
    exit(0);
  } else {
    close(conn_socket);
    wait(NULL);
  }
}
close(server_socket);
```

9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.49

Concurrent Server

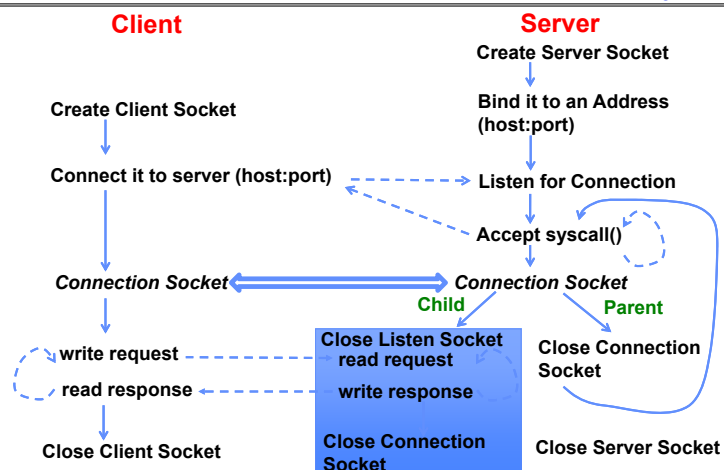
- So far, in the server:
 - Listen will queue requests
 - Buffering present elsewhere
 - But server waits for each connection to terminate before servicing the next
- A concurrent server can handle and service a new connection before the previous client disconnects

9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.50

Sockets With Protection and Concurrency



9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.51

Server Protocol (v3)

```
// Socket setup code elided...
while (1) {
  // Accept a new client connection, obtaining a new socket
  int conn_socket = accept(server_socket, NULL, NULL);
  pid_t pid = fork();
  if (pid == 0) {
    close(server_socket);
    serve_client(conn_socket);
    close(conn_socket);
    exit(0);
  } else {
    close(conn_socket);
    //wait(NULL);
  }
}
close(server_socket);
```

9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.52

Server Address: Itself

```
struct addrinfo *setup_address(char *port) {
    struct addrinfo *server;
    struct addrinfo hints;
    memset(&hints, 0, sizeof(hints));
    hints.ai_family = AF_UNSPEC;
    hints.ai_socktype = SOCK_STREAM;
    hints.ai_flags = AI_PASSIVE;
    getaddrinfo(NULL, port, &hints, &server);
    return server;
}
```

- Accepts any connections on the specified port

9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.53

Client: Getting the Server Address

```
struct addrinfo *lookup_host(char *host_name, char *port) {
    struct addrinfo *server;
    struct addrinfo hints;
    memset(&hints, 0, sizeof(hints));
    hints.ai_family = AF_UNSPEC;
    hints.ai_socktype = SOCK_STREAM;

    int rv = getaddrinfo(host_name, port_name,
                        &hints, &server);

    if (rv != 0) {
        printf("getaddrinfo failed: %s\n", gai_strerror(rv));
        return NULL;
    }
    return server;
}
```

9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.54

Concurrent Server without Protection

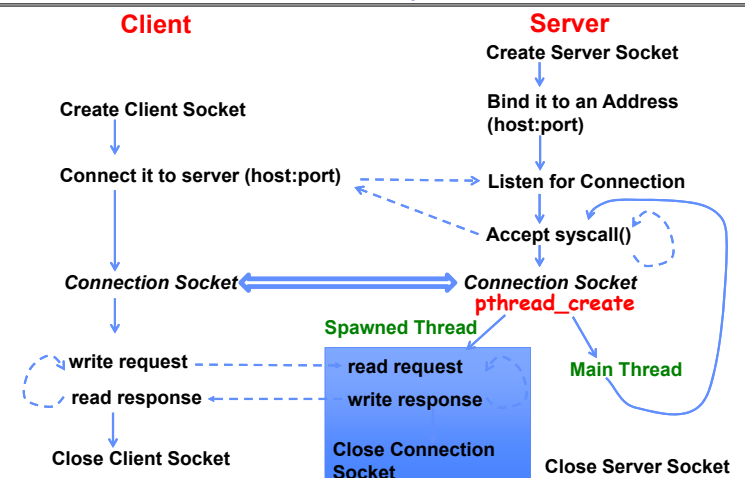
- Spawn a new thread to handle each connection
- Main thread initiates new client connections without waiting for previously spawned threads
- Why give up the protection of separate processes?
 - More efficient to create new threads
 - More efficient to switch between threads

9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.55

Sockets with Concurrency, without Protection




9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.56

Thread Pools

- Problem with previous version: Unbounded Threads
 - When web-site becomes too popular – throughput sinks
- Instead, allocate a bounded “pool” of worker threads, representing the maximum level of multiprogramming



```
master() {
    allocThreads(worker, queue);
    while(TRUE) {
        con=AcceptCon();
        Enqueue(queue, con);
        wakeUp(queue);
    }
}

worker(queue) {
    while(TRUE) {
        con=Dequeue(queue);
        if (con==null)
            sleepOn(queue);
        else
            ServiceWebPage(con);
    }
}
```

9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.57

Conclusion

- Interprocess Communication (IPC)
 - Communication facility between protected environments (i.e. processes)
- Pipes are an abstraction of a single queue
 - One end write-only, another end read-only
 - Used for communication between multiple processes on one machine
 - File descriptors obtained via inheritance
- Sockets are an abstraction of two queues, one in each direction
 - Can read or write to either end
 - Used for communication between multiple processes on different machines
 - File descriptors obtained via socket/bind/connect/listen/accept
 - Inheritance of file descriptors on fork() facilitates handling each connection in a separate process
- Both support read/write system calls, just like File I/O

9/14/20

Kubiatowicz CS162 © UCB Fall 2020

Lec 5.58