

# CS 370: OPERATING SYSTEMS

## [INTER PROCESS COMMUNICATIONS]

Shrideep Pallickara  
Computer Science  
Colorado State University

September 4, 2018

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.1

## Frequently asked questions from the previous class survey

- When you fork() are objects and data of the process shared or is a new copy of the heap created?
  - ▣ Everything is copied
- Why is wait() called in the parent and exec() in the child?
  - ▣ Can you wait for multiple children?
- When you call exec() on child, is the parent affected?
  - ▣ What does exec() destroy? COPY of the memory image of the parent
- Zombies and Orphans
  - ▣ What happens after adoption by init?

September 4, 2018  
Professor: SHRIDEEP PALICKARA

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.2

## Frequently asked questions from the previous class survey

- Why would you ever make copies of programs like we did in the code snippets?
- As you fork processes, upon completion of the process creation are they considered ready for scheduling by the kernel?
- Automatic variables? What are they?
- Kernel strategies for preventing some of the attacks?
  - ▣ ASLR: Address space layout randomization
  - ▣ Non-executable stack

September 4, 2018  
Professor: SHRIDEEP PALLICKARA

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.3

## Topics covered in this lecture

- Shells and Daemons
- POSIX
- Inter Process Communications

September 4, 2018  
Professor: SHRIDEEP PALLICKARA

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.4

## SHELLS AND DAEMONS

September 4, 2018

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.5

### Shell: Command interpreter

- Prompts for commands
- Reads commands from standard input
- Forks children to execute commands
- Waits for children to finish
- When standard I/O comes from terminal
  - ▣ Terminate command with the interrupt character
    - Default Ctrl-C

September 4, 2018  
Professor: SHRIDEEP PALICKARA

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.6

## Background processes and daemons

- Shell interprets commands ending with **&** as a background process
  - No waiting for process to complete
  - Issue prompt immediately
    - Accept new commands
  - Ctrl-C has no effect
  
- **Daemon** is a background process
  - Runs indefinitely

September 4, 2018  
Professor: SHRIDEEP PALICKARA

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.7

# POSIX

September 4, 2018

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.8

## Portable Operating Systems Interface for UNIX (POSIX)

- 2 **distinct, incompatible** flavors of UNIX existed
  - System V from AT&T
  - BSD UNIX from Berkeley
- Programs written from one type of UNIX
  - Did not run correctly (sometimes even compile) on UNIX from another vendor
- Pronounced *pahz-icks*

September 4, 2018  
Professor: SHRIDEEP PALICKARA

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.9

## IEEE attempt to develop a standard for UNIX libraries

- **POSIX.1** published in 1988
  - Covered a small subset of UNIX
- In 1994, X/Open Foundation had a much more comprehensive effort
  - Called **Spec 1170**
  - Based on System V
- Inconsistencies between POSIX.1 and Spec 1170

September 4, 2018  
Professor: SHRIDEEP PALICKARA

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.10

## The path to the final POSIX standard

### □ 1998

- Another version of the X/Open standard
- Many additions to POSIX.1
- **Austin Group** formed
  - Open Group, IEEE POSIX, and ISO/IEC tech committee
    - International Standards Organization (ISO)
    - International Electrotechnical Commission (IEC)
  - Revise, combine and update standards

September 4, 2018  
Professor: SHRIDEEP PALICKARA

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.11

## The path to the final POSIX standard: Joint document

- Approved by IEEE & Open Group
  - End of 2001
- ISO/IEC approved it in November 2002
- Single UNIX spec
  - Version 3, IEEE Standard 1003.1-2001
  - **POSIX**

September 4, 2018  
Professor: SHRIDEEP PALICKARA

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.12

## If you write for POSIX-compliant systems

- No need to contend with small, but critical variations in library functions
  - Across platforms

September 4, 2018  
Professor: SHRIDEEP PALICKARA

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.13

## INTER PROCESS COMMUNICATIONS (IPC)

September 4, 2018

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.14

## Independent and Cooperating processes

- Independent: **CANNOT** affect or be affected by other processes
- Cooperating: **CAN** affect or be affected by other processes

September 4, 2018  
Professor: SHRIDEEP PALICKARA

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.15

## Why have cooperating processes?

- Information sharing: shared files
- Computational speedup
  - ▣ Sub tasks for concurrency
- Modularity
- Convenience: Do multiple things in parallel
- Privilege separation

September 4, 2018  
Professor: SHRIDEEP PALICKARA

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.16

## Cooperating processes need IPC to exchange data and information

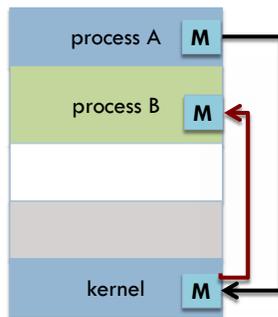
- **Shared memory**
  - Establish memory region to be shared
  - Read and write to the shared region
  
- **Message passing**
  - Communications through message exchange

September 4, 2018  
Professor: SHRIDEEP PALICKARA

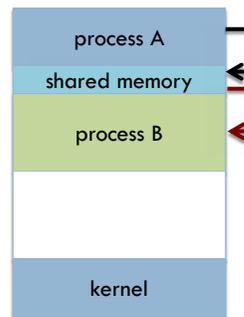
CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.17

## Contrasting the two IPC approaches



**Easier** to implement  
Best for **small** amounts of data  
**Kernel intervention** for communications



Maximum **speed**  
System calls to **establish** shared memory

September 4, 2018  
Professor: SHRIDEEP PALICKARA

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.18

## Shared memory systems

- Shared memory resides **in** the address space of process creating it
- Other processes must **attach** segment to their address space

September 4, 2018  
Professor: SHRIDEEP PALICKARA

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.19

## Using shared memory

- But the OS typically **prevents** processes from accessing each other's memory, so ...
  - ① Processes must agree to **remove** this **restriction**
  - ② Processes also **coordinate** access to this region

September 4, 2018  
Professor: SHRIDEEP PALICKARA

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.20

## Let's look a little closer at cooperating processes

- **Producer-consumer** problem is a good exemplar of such cooperation
- Producer process *produces* information
- Consumer process *consumes* this information

September 4, 2018  
Professor: SHRIDEEP PALICKARA

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.21

## One solution to the producer-consumer problem uses *shared-memory*

- Buffer is a shared-memory region for the 2 processes
- Buffer needed to allow producer & consumer to run **concurrently**
  - Producer fills it
  - Consumer empties it

September 4, 2018  
Professor: SHRIDEEP PALICKARA

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.22

## Buffers and sizes

- Bounded: Assume **fixed** size
  - Consumer waits if buffer is empty
  - Producer waits if buffer is full
- Unbounded: **Unlimited** number of entries
  - Only the consumer waits WHEN buffer is empty

September 4, 2018  
Professor: SHRIDEEP PALICKARA

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

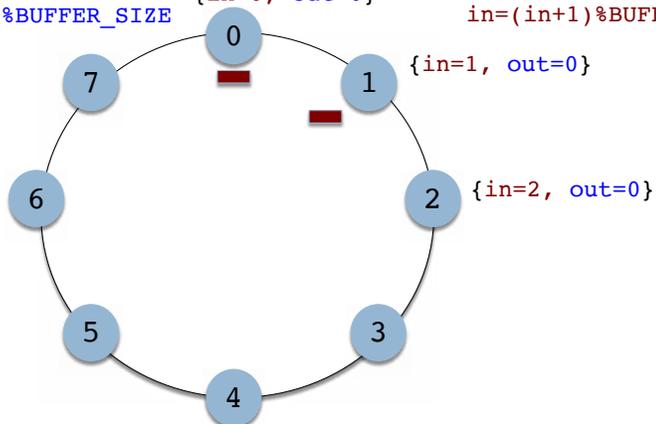
L5.23

## Circular buffer: Bounded

After consuming:  
 $out=(out+1)\%BUFFER\_SIZE$

{in=0, out=0}

After producing:  
 $in=(in+1)\%BUFFER\_SIZE$



in: next free position (producer)  
out: first full position (consumer)

September 4, 2018  
Professor: SHRIDEEP PALICKARA

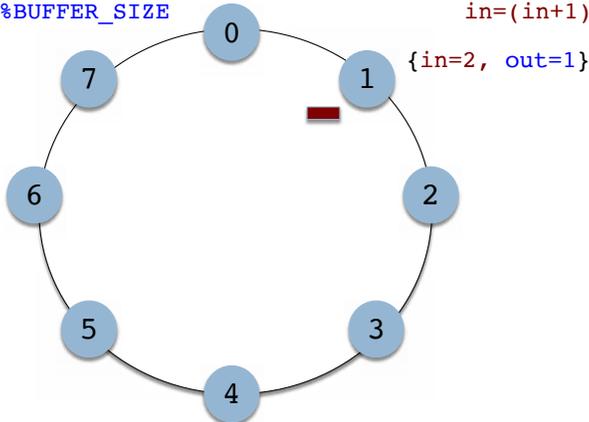
CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.24

## Circular buffer: Bounded

After consuming:  
`out=(out+1)%BUFFER_SIZE`

After producing:  
`in=(in+1)%BUFFER_SIZE`



`in`: next free position (producer)  
`out`: first full position (consumer)

September 4, 2018  
Professor: SHRIDEEP PALICKARA

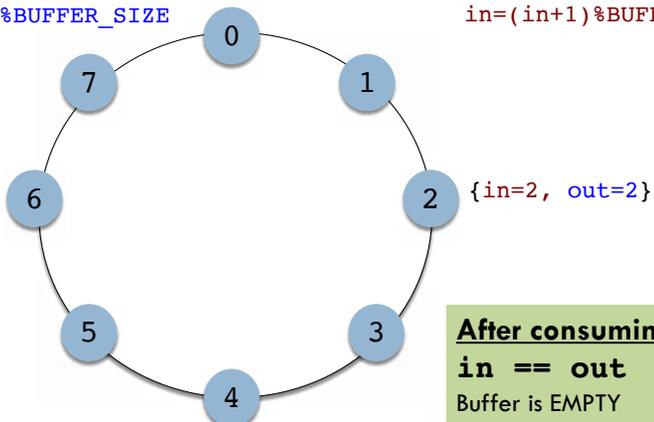
CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.25

## Circular buffer: Bounded

After consuming:  
`out=(out+1)%BUFFER_SIZE`

After producing:  
`in=(in+1)%BUFFER_SIZE`



`in`: next free position (producer)  
`out`: first full position (consumer)

September 4, 2018  
Professor: SHRIDEEP PALICKARA

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.26

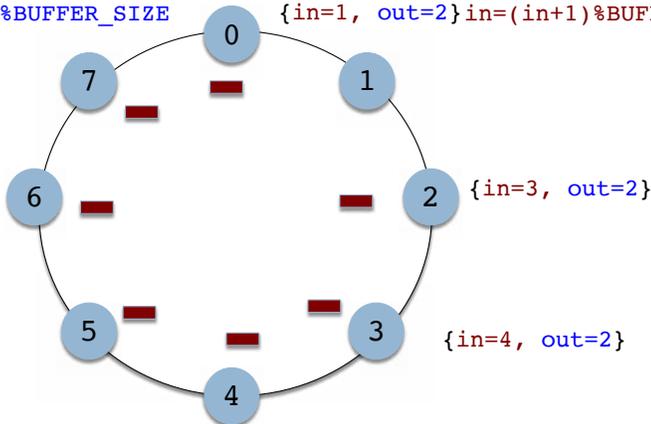
## Circular buffer: Bounded

After consuming:

`out=(out+1)%BUFFER_SIZE`

After producing:

`{in=1, out=2} in=(in+1)%BUFFER_SIZE`



`in`: next free position (producer)  
`out`: first full position (consumer)

September 4, 2018  
 Professor: SHRIDEEP PALICKARA

CS370: Operating Systems [Fall 2018]  
 Dept. Of Computer Science, Colorado State University

L5.27

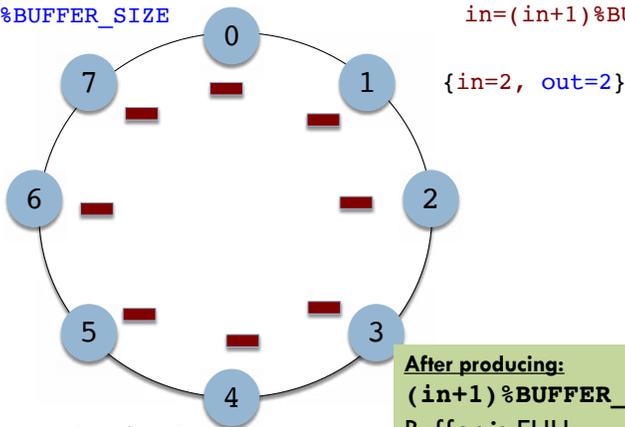
## Circular buffer: Bounded

After consuming:

`out=(out+1)%BUFFER_SIZE`

After producing:

`in=(in+1)%BUFFER_SIZE`



`in`: next free position (producer)  
`out`: first full position (consumer)

September 4, 2018  
 Professor: SHRIDEEP PALICKARA

CS370: Operating Systems [Fall 2018]  
 Dept. Of Computer Science, Colorado State University

L5.28

## INTER PROCESS COMMUNICATIONS

### SHARED MEMORY

September 4, 2018

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.29

## POSIX IPC: Shared Memory

### Creating a memory segment to share

- First **create** shared memory segment `shmget ( )`
  - `shmget (IPC_PRIVATE, size, S_IRUSR | S_IWUSR)`
    - `IPC_PRIVATE`: key for the segment
    - `size`: size of the shared memory
    - `S_IRUSR | S_IWUSR`: Mode of access (read, write)
- Successful invocation of `shmget ( )`
  - Returns integer ID of shared segment
    - Needed by other processes that want to use region

September 4, 2018  
Professor: SHRIDEEP PALICKARA

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.30

## Processes wishing to use shared memory must first attach it to their address space

- Done using `shmat()`: SHared Memory ATtach
  - Returns pointer to beginning location in memory
  
- **(void \*) shmat(id, asmP, mode)**
  - `id`: Integer ID of memory segment being attached
  - `asmP`: Pointer location to attach shared memory
    - `NULL` allows OS to select location for you
  - Mode indicating read-only or read-write
    - `0`: reads and writes to shared memory

September 4, 2018  
Professor: SHRIDEEP PALICKARA

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.31

## IPC: Use of the created shared memory

- Once shared memory is attached to the process's address space
  - Routine memory accesses using `*` from `shmat()`
    - Write to it
      - `sprintf(shared_memory, "Hello");`
    - Print string from memory
      - `printf("%s\n", shared_memory);`
  
- **RULE**: First attach, and then access

September 4, 2018  
Professor: SHRIDEEP PALICKARA

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.32

## IPC Shared Memory: What to do when you are done

- ① **Detach** from the address space.
  - `shmdt ( )` : SHared Memory DeTtach
  - `shmdt (shared_memory)`
  
- ② To **remove** a shared memory segment
  - `shmctl ( )` : SHared Memory ConTrol operation
    - Specify the segment ID to be removed
    - Specify operation to be performed: `IPC_RMID`
    - Pointer to the shared memory region

## INTER PROCESS COMMUNICATIONS MESSAGE PASSING

## Communicate and synchronize actions without sharing the same address space

- Two main operations
  - send(message)
  - receive(message)
  
- Message sizes can be:
  - Fixed: Easy
  - Variable: Little more effort

September 4, 2018  
Professor: SHRIDEEP PALICKARA

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.35

## Communications between processes

- There needs to be a communication link
  
- Underlying physical implementation
  - Shared memory
  - Hardware bus
  - Network

September 4, 2018  
Professor: SHRIDEEP PALICKARA

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.36

## Aspects to consider for IPC

### ① Communications

- Direct or indirect

### ② Synchronization

- Synchronous or asynchronous

### ③ Buffering

- Automatic or explicit buffering

## Communications: Naming allows processes to refer to each other

- Processes use each other's identity to communicate
- Communications can be
  - Direct
  - Indirect

## Direct communications

- Explicitly name recipient or sender
- Link is established automatically
  - ▣ Exactly one link between the 2 processes
- Addressing
  - ▣ Symmetric
  - ▣ Asymmetric

September 4, 2018  
Professor: SHRIDEEP PALICKARA

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.39

## Direct Communications: Addressing

- Symmetric addressing  Explicitly name recipient and sender of message
  - `send(P, message)`
  - `receive(Q, message)`
- Asymmetric addressing  Only sender names recipient  
Recipient does not
  - `send(P, message)`
  - `receive(id, message)`
    - Variable `id` set to name of the sending process

September 4, 2018  
Professor: SHRIDEEP PALICKARA

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.40

## Direct Communications: Disadvantages

- **Limited modularity** of process definitions
- **Cascading effects** of changing the identifier of process
  - Examine *all* other process identifiers

September 4, 2018  
Professor: SHRIDEEP PALICKARA

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.41

## Indirect communications: Message sent and received from mailboxes (ports)

- Each **mailbox** has a unique identification & owner
  - POSIX message queues use `integers` to identify mailboxes
- Processes communicate *only* if they have **shared mailbox**
  - `send(A, message)`
  - `receive(A, message)`

September 4, 2018  
Professor: SHRIDEEP PALICKARA

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.42

## Indirect communications: Link properties

- Link established only if both members share mailbox
- Link may be associated with more than two processes

September 4, 2018  
Professor: SHRIDEEP PALICKARA

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.43

## Indirect communications

- Processes P1, P2 and P3 share mailbox A
  - ▣ P1 sends a message to A
  - ▣ P2, P3 execute a `receive()` from A
- Possibilities? Allow ...
  - ① Link to be associated with at most 2 processes
  - ② At most 1 process to execute `receive()` at a time
  - ③ System to arbitrarily select who gets message

September 4, 2018  
Professor: SHRIDEEP PALICKARA

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.44

## Mailbox ownership issues

- Owned by process
- Owned by the OS

September 4, 2018  
Professor: SHRIDEEP PALICKARA

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.45

## Mailbox ownership issues: Owned by process

- Mailbox is part of the **process's address space**
  - Owner: Can *only receive* messages on mailbox
  - User: Can *only send* messages to mailbox
- When process terminates?
  - Mailbox disappears

September 4, 2018  
Professor: SHRIDEEP PALICKARA

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.46

## Mailbox ownership issues: Owned by OS

- Mailbox has its own existence
- Mailbox is **independent**
  - ▣ Not attached to any process
- OS must allow processes to
  - ▣ Create mailbox
  - ▣ Send and receive *through* the mailbox
  - ▣ Delete mailbox

September 4, 2018  
Professor: SHRIDEEP PALICKARA

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.47

## Message passing: Synchronization issues Options for implementing primitives

- Blocking send
  - ▣ Block *until* received by process or mailbox
- Nonblocking send
  - ▣ Send and *promptly resume* other operations
- Blocking receive
  - ▣ Block *until* message available
- Nonblocking receive
  - ▣ Retrieve *valid* message or *null*
- Producer-Consumer problem: Easy with blocking

September 4, 2018  
Professor: SHRIDEEP PALICKARA

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.48

## Message Passing: Buffering

- Messages exchanged by communicating processes reside in a **temporary** queue
- Implementation schemes for queues
  - ZERO Capacity
  - Bounded
  - Unbounded

September 4, 2018  
Professor: SHRIDEEP PALICKARA

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.49

## Message Passing Buffer: Consumer always has to wait for message

- ZERO capacity: No messages can reside in queue
  - Sender **must block** till recipient receives
- BOUNDED: At most  $n$  messages can reside in queue
  - Sender **blocks only if queue is full**
- UNBOUNDED: Queue length potentially infinite
  - Sender **never blocks**

September 4, 2018  
Professor: SHRIDEEP PALICKARA

CS370: Operating Systems [Fall 2018]  
Dept. Of Computer Science, Colorado State University

L5.50

## The contents of this slide-set are based on the following references

- *Avi Silberschatz, Peter Galvin, Greg Gagne. Operating Systems Concepts, 9<sup>th</sup> edition. John Wiley & Sons, Inc. ISBN-13: 978-1118063330. [Chapter 3]*
- *Kay Robbins & Steve Robbins. Unix Systems Programming, 2nd edition, Prentice Hall ISBN-13: 978-0-13-042411-2. [Chapter 2, 3]*
- *Andrew S Tanenbaum. Modern Operating Systems. 4<sup>th</sup> Edition, 2014. Prentice Hall. ISBN: 013359162X/ 978-0133591620. [Chapter 2]*