

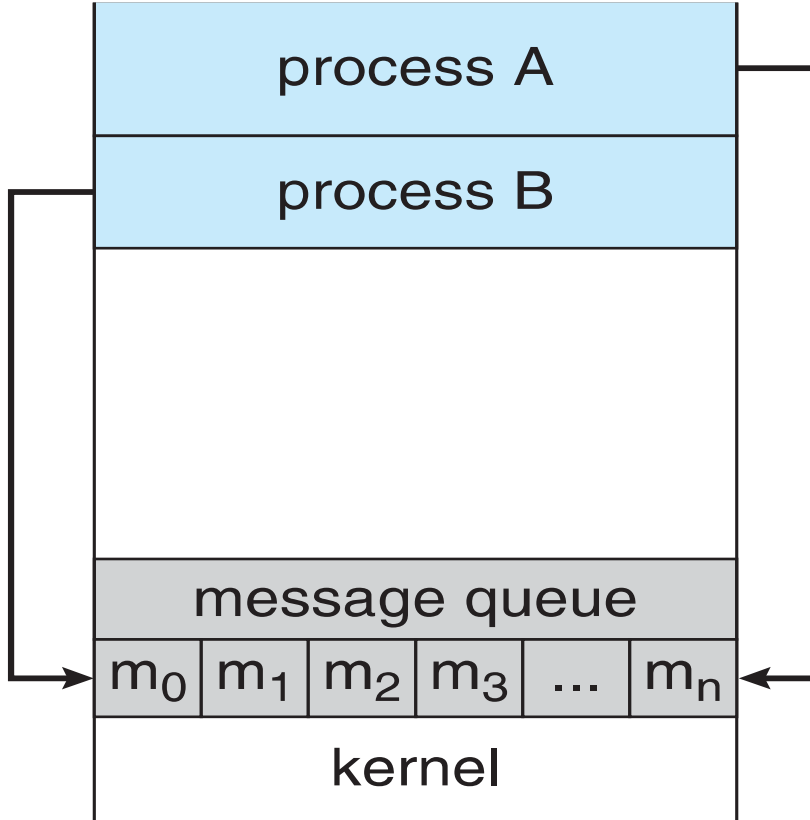
# Interprocess Communication

- Processes within a system may be *independent* or *cooperating*
- Cooperating process can affect or be affected by other processes, including sharing data
- Reasons for cooperating processes:
  - Information sharing
  - Computation speedup
  - Modularity
  - Convenience
- Cooperating processes need **interprocess communication (IPC)**
- Two models of IPC
  - **Shared memory**
  - **Message passing**




# Communications Models

(a) Message passing.



(a)

(b) shared memory.



(b)



**sns**  
INSTITUTIONS

# Cooperating Processes

- ***Independent*** process cannot affect or be affected by the execution of another process
- ***Cooperating*** process can affect or be affected by the execution of another process
- Advantages of process cooperation
  - Information sharing
  - Computation speed-up
  - Modularity
  - Convenience



**sns**  
INSTITUTIONS

# Producer-Consumer Problem

- Paradigm for cooperating processes, *producer* process produces information that is consumed by a *consumer* process
  - **unbounded-buffer** places no practical limit on the size of the buffer
  - **bounded-buffer** assumes that there is a fixed buffer size



# Bounded-Buffer – Shared-Memory Solution

- Shared data

```
#define BUFFER_SIZE 10

typedef struct {

    ...

} item;

item buffer[BUFFER_SIZE];

int in = 0;

int out = 0;
```

- Solution is correct, but can only use BUFFER\_SIZE-1 elements



**sns**  
INSTITUTIONS

# Bounded-Buffer – Producer

```
item next_produced;
```

```
while (true) {
```

```
    /* produce an item in next produced */
```

```
    while (((in + 1) % BUFFER_SIZE) == out)
```

```
        ; /* do nothing */
```

```
    buffer[in] = next_produced;
```

```
    in = (in + 1) % BUFFER_SIZE;
```

```
}
```



**sns**  
INSTITUTIONS

# Bounded Buffer – Consumer

```
item next_consumed;
```

```
while (true) {
```

```
    while (in == out)
```

```
        ; /* do nothing */
```

```
    next_consumed = buffer[out];
```

```
    out = (out + 1) % BUFFER_SIZE;
```

```
    /* consume the item in next consumed */
```

```
}
```





**sns**  
INSTITUTIONS

# Interprocess Communication – Shared Memory

- An **area of memory shared** among the processes that wish to communicate
- The communication is **under the control of the users processes** not the operating system.
- Major issues is to provide mechanism that will allow the user processes to synchronize their actions when they access shared memory.





**sns**  
INSTITUTIONS

# Interprocess Communication – Message Passing

- Mechanism for processes to communicate and to synchronize their actions
- Message system – processes communicate with each other without resorting to shared variables
- IPC facility provides two operations:
  - **send**(*message*)
  - **receive**(*message*)
- The *message* size is **either fixed or variable**



## Message Passing (Cont.)

- If processes  $P$  and  $Q$  wish to communicate, they need to:
  - Establish a ***communication link*** between them
  - Exchange messages via send/receive
- **Implementation issues:**
  - How are links established?
  - Can a link be associated with more than two processes?
  - How many links can there be between every pair of communicating processes?
  - What is the capacity of a link?
  - Is the size of a message that the link can accommodate fixed or variable?
  - Is a link unidirectional or bi-directional?



**sns**  
INSTITUTIONS

# Message Passing (Cont.)

Implementation of communication link

- **Physical:**
  - Shared memory
  - Hardware bus
  - Network
- **Logical:**
  - Direct or indirect
  - Synchronous or asynchronous
  - Automatic or explicit buffering



**sns**  
INSTITUTIONS

# Direct Communication

- Processes must name each other explicitly:
  - **send** ( $P$ , *message*) – send a message to process  $P$
  - **receive**( $Q$ , *message*) – receive a message from process  $Q$
- **Properties of communication link**
  - Links are established automatically
  - A link is associated with exactly one pair of communicating processes
  - Between each pair there exists exactly one link
  - The link may be unidirectional, but is usually bi-directional



**sns**  
INSTITUTIONS

# Indirect Communication

- Messages are directed and received **from mailboxes** (also referred to as ports)
  - Each mailbox has a unique id
  - Processes can communicate only if they share a mailbox
- **Properties of communication link**
  - Link established only if processes share a common mailbox
  - A link may be associated with many processes
  - Each pair of processes may share several communication links
  - Link may be unidirectional or bi-directional



**sns**  
INSTITUTIONS

# Indirect Communication

- Operations

- create a new mailbox (port)
- send and receive messages through mailbox
- destroy a mailbox

- Primitives are defined as:

**send**( $A, message$ ) – send a message to mailbox  $A$

**receive**( $A, message$ ) – receive a message from mailbox  $A$



**sns**  
INSTITUTIONS

# Indirect Communication

- Mailbox sharing

- $P_1$ ,  $P_2$ , and  $P_3$  share mailbox A
- $P_1$  sends;  $P_2$  and  $P_3$  receive
- Who gets the message?

- Solutions

- Allow a link to be associated with at most two processes
- Allow only one process at a time to execute a receive operation
- Allow the system to select arbitrarily the receiver. Sender is notified who the receiver was.





**sns**  
INSTITUTIONS

# Synchronization

- Message passing may be either blocking or non-blocking
- **Blocking** is considered **synchronous**
  - **Blocking send** -- the sender is blocked until the message is received
  - **Blocking receive** -- the receiver is blocked until a message is available
- **Non-blocking** is considered **asynchronous**
  - **Non-blocking send** -- the sender sends the message and continue
  - **Non-blocking receive** -- the receiver receives:
    - A valid message, or Null message



**sns**  
INSTITUTIONS

# Synchronization (Cont.)

Producer-consumer becomes trivial

```
message next_produced;
while (true) {
    /* produce an item in next produced */
    send(next_produced);
}
message next_consumed;
while (true) {
    receive(next_consumed);

    /* consume the item in next consumed */
}
```



**sns**  
INSTITUTIONS

# Buffering

- Queue of messages attached to the link.
- implemented in one of three ways
  1. Zero capacity – no messages are queued on a link.  
Sender must wait for receiver (rendezvous)
  2. Bounded capacity – finite length of  $n$  messages  
Sender must wait if link full
  3. Unbounded capacity – infinite length  
Sender never waits



**sns**  
INSTITUTIONS

# Examples of IPC Systems - POSIX

## POSIX Shared Memory

- Process first creates shared memory segment

```
shm_fd = shm_open(name, O_CREAT | O_RDWR, 0666);
```

- Also used to open an existing segment to share it
- Set the size of the object

```
ftruncate(shm_fd, 4096);
```

- Now the process could write to the shared memory
- `sprintf(shared memory, "Writing to shared memory");`



**sns**  
INSTITUTIONS

# Examples of IPC Systems - Mach

- Mach communication is message based
  - Even system calls are messages
  - Each task gets two mailboxes at creation- Kernel and Notify
  - Only three system calls needed for message transfer

**`msg_send()` , `msg_receive()` , `msg_rpc()`**

- Mailboxes needed for communication, created via **`port_allocate()`**
- Send and receive are flexible, for example four options if mailbox full:
  - Wait indefinitely
  - Wait at most n milliseconds
  - Return immediately
  - Temporarily cache a message