
Week 4

COMP3231 Operating Systems

2005/S2

Slide 1

- ① Concurrency Control
 - Interprocess Communication (IPC)
 - Deadlocks
- ② Memory Management

IPC MECHANISMS

Several different mechanisms can be used for communication

- ① **Shared memory**: can be used to exchange information, but synchronisation issues remain
 - threads: run in common memory space
 - `mmap()` system call
- ② **File system**
 - normal files
 - pipes (FIFOs)
 - sockets
- ③ **Message passing**
 - more abstract communication mechanism

A CLOSER LOOK AT `mmap` AND PIPES

Memory mapped files:

- Processes can share files to communicate
- `mmap` maps by default a file into memory

Type of `mmap`:

Slide 3

```
void *mmap (void      *addr, /* dst address for map */
            size_t   len,  /* length of data to map */
            int      prot, /* protection           */
            int      flags, /* misc flags          */
            int      fildes, /* file descriptor     */
            off_t    offset) /* offset into file    */
```

By using `mmap` can be used to simulate "shared memory".

Example: simplified, no error checking!

```
int main () {
    char *data, *fname = "foo";
    int  fd;
    struct stat sbuf;
    fd = open (fname, O_RDONLY);
    stat (fname, &sbuf);
```

Slide 4

```
    data = mmap ((caddr_t) 0, /* let the system choose the dest.
                               will be rounded to pagesize
                               mem. area will be read only
                               shared: changes visible to all proc
                               file desc. of previously opened file
                               offset
                               */
                 sbuf.st_size,
                 PROT_READ,
                 MAP_SHARED,
                 fd,
                 0);

    printf ("Test mmap:...%c", data[0]);}
```


PIPES

What is a file descriptor?

- `fopen`, `fwrite`, etc are implemented in terms of systems calls like `open`, `write`
- `fopen` returns a FILE pointer
- `open` returns a file descriptor (int)
- `stdin`, `stdout`, `stderr` are file descriptors

Slide 5

`pipe()` returns a pair of file descriptors:



Example: (no error checking)

```
int fd[2];
char buf[20];

pipe (fd);

write (fd[1], "test1", 6);
write (fd[1], "test2", 6);
read (fd[0], buf, 6);
printf ("read %s from pipe \n", buf);
```

Slide 6

Usage:

- pipes can be used in combination with `fork`
- named pipes: `mknod`
`mknod ("testFIFO", S_FIFO | 0644, 0);`

ASSIGNMENT 1

Implementing (a sort of) Pipes

```
void pipe_create (pipe_in **p_in, pipe_out **p_out);
void pipe_destroy_in (pipe_in *p);
void pipe_destroy_out (pipe_out *p);
void pipe_read (pipe_out *p, void *dest, int n_bytes);
void pipe_write (pipe_in *p, void *src, int n_bytes);
```

Slide 7

Assignment 1:

- Will be out in the next few days
- Anybody not in a group on Wednesday will automatically have a partner assigned

SOCKETS

What is a socket?

- two-way communication pipe
- can be used to communicate in a wide variety of domains (e.g., internet)

Slide 8

Communication between Processes:

Sockets are also a file in the Unix file system, but offer a different interface

- `socket()`, `bind()`, `receive()` instead of `open()`, `read()` and `write()` (`read()` and `write()` are actually also available on sockets)
- typically used in client/server style programs

Server:

- ① create socket:
`s = socket(AF_UNIX, SOCK_STREAM)`
- ② bind socket to local address:
`bind(s, <socket name>)`
- ③ listen for incoming connections:
`listen(s, <max size of incoming connection queue>)`
- ④ main server loop:
 - ① accept connection
`s2 = accept(s, &<remote socket name>)`
 - ② receive/send
`recv(s2, &<request>)`
`send(s2, <answer>)`
 - ③ close
`close(s2)`

Slide 9

Client:

- ① create socket:
`s = socket(AF_UNIX, SOCK_STREAM)`
- ② connect:
`connect(s, <socket name>)`
- ③ send/receive:
`send(s, <request>)`
`recv(s, &<answer>)`
- ④ close:
`close(s)`

Slide 10

MESSAGE-PASSING IPC

Primitives:

- Sending a message: `send(dest, msg)`
- Receiving a message: `receive(source, &msg)`

Different message passing styles:

- **synchronisation**: blocking (synchronous) vs. non-blocking (asynchronous)
 - blocking send, blocking receive
 - non-blocking send, blocking receive
 - non-blocking send, non-blocking receive
- **addressing**: direct vs. indirect
 - identifier of destination process
 - message to shared data structure (mailbox, port), one-to-one, one-to-many, many-to-many
- **message format**: depends on objectives, single computer vs. distributed system, fixed vs. variable-length messages

Slide 11

IPC Implementation Issues:

- **Security & Safety**
 - messages may be lost
 - authentication
- How are links established?
 - automatically
 - have to be set up explicitly
- What is the **capacity** of a link?
- Is the message format fixed or variable?
- Is a link uni-directional or bi-directional?

Slide 12

IPC: DIRECT COMMUNICATION

Processes must name each other explicitly

- `send(pid, &msg)`
- `receive(pid, &msg)` — sometimes id of sender cannot be anticipated

Slide 13

Properties of communication link :

- links established automatically
- link associated with pair of processes
- exactly one link between each pair
- link may be uni- or bi-directional

INDIRECT COMMUNICATION

→ Messages go via mailboxes (aka. **ports**)

- each port has unique ID
- communication requires sharing of a port

→ Properties of communication link:

- links established if processes share a port
- link may be associated with many processes
- each pair may share many links
- link (port) may be uni- or bi-directional

→ Operations: create, delete, send, receive


Slide 14

GENERAL MESSAGE FORMAT

The format of a message depends on

- objectives of message facility
- local or distributed
- security and safety requirements

Slide 15



Message Buffering:

- Associate message buffer with link:
 - **Zero capacity**: 0 messages
 - sender blocks until receive (rendezvous)
 - **Bounded capacity**: finite # messages
 - if full sender blocks or fails
 - **Unbounded capacity**: infinite # messages
 - sender never blocks

Slide 16

IPC Exception Conditions:

- Partner process terminated
- Partner uncommunicative (protocol failure)
- Message buffer overflow
- Message lost
- Message scrambled