

Operating System

Done By:
Sarah Alkasasbeh



POWERUNIT

slides (4-8)

Lecture 1

• The main component of the operating system is -
Kernel

• The operating system abstracts (simplifies) the complexity of the hardware, and the RAM size ... etc. → (Abstracts)

• The operating system controls the hardware of the computer (ex: CPU, memory, disk), it works to ensure a fair share of resources among different applications and users. (resources: CPU, memory, disk)
↳ (Arbitrates)

• The hard disk is I/O device, so the main components to operate the computer are the CPU and the memory.

• NIC (network interface card) is input and output device. "receives and sends packets"

• It is possible that there is a system that doesn't use an operating system, if there is only one application we want to execute on that hardware.

• The operating system is stored on the hard disk, or on external storage like flash.

• The operating system deals ~~to~~ with the controllers of the devices through the driver of the device, each component has different device driver.

• The alternative of the device driver is the ability of the operating system to deal with all hardware or I/O devices. (not practical).

• Note:- controller :- hardware / Driver :- Piece of software

• every device controller has local buffer, to store enough data in it before it is used by the CPU, in order to avoid the speed mismatch between the I/O devices and the CPU.

• The interrupt temporarily terminates the task of the CPU to do something else. ~~When~~ ~~data~~ dealing with interrupts is according to priority of each one.

• hardware interrupt :- interrupt caused by wires in the CPU.

• software interrupt :- no wires in the CPU determine if there is interrupt. (ex:- I/O operation).

• interrupt is external, ~~and~~ ^{but} the exception is internal error.

• Two methods to handle interrupts :-

1- Blocking I/O → CPU becomes idle until the interrupts end.

2- Non-blocking I/O → The CPU completes executing the instructions. (it helps avoiding speed mismatch)
~~and it helps avoiding speed mismatch~~

• BIOS (Basic input output system) → makes sure that ~~that~~ there is not any problem in the main components, and it is the first one operate once turning on the device, then the bootstrap program runs.

• Bootstrap stored in a permanent storage, it loads the operating system from the disk to the memory.

~~First program runs~~ First program runs.

• init's first service in the operating system, it initializes the system, operates memory management service, operates CPU scheduling service...etc.

• Bootstrap program is typically stored in ROM or EPROM, generally known as firmware. (firmware ex :- bootstrap,
 ↓ BIOS)

any software that controls hardware.

slides (25-37) lecture 3

- DMA (Direct memory access), works on transferring the data between I/O devices and main memory.
- before the DMA, transferring the data between I/O devices and memory was through the CPU, which affected the CPU time.

• Kernel is basic service in the operating system, controls all other services provided by the operating system.

- Hardware interrupt on kernel → by one of the I/O devices.
(OS elements) →

• one of the OS elements is the abstraction (like: process, thread, socket... etc), and to deal with these abstractions we need mechanisms (like create, open, write... etc), so we need policies to avoid violation (like LRU, EDF).

- ~~When~~ the OS ^{allocate} ~~reserves~~ a place in the memory for each process (application) called memory page, it equals 4KB.
ex: (memory management) →
- LRU: swap between the "least recently used" page and other page the process needs. (LRU page from the memory to the disk).

• modes of operation :- 1- privileged mode (kernel mode) →

means that the instructions are operating system's instruction (specifically related to the kernel).

2- unprivileged mode (user mode) → executed instructions are user application instructions.

• some instruction are not allowed to be executed on the user mode like halt instruction and timer instructions. (halt instruction → shut down the system)

• System call → when the user program tries to execute a privileged instruction, or to deal with an I/O device. → trap mode bit becomes zero → after finishing the system call and execute the instructions the trap bit becomes one, and the control returns to the user application.

• Program is a passive entity (before execution)

• Process is an active entity (~~the~~ program in execution)

Mode bit = 0 → kernel mode

Mode bit = 1 → user mode

note

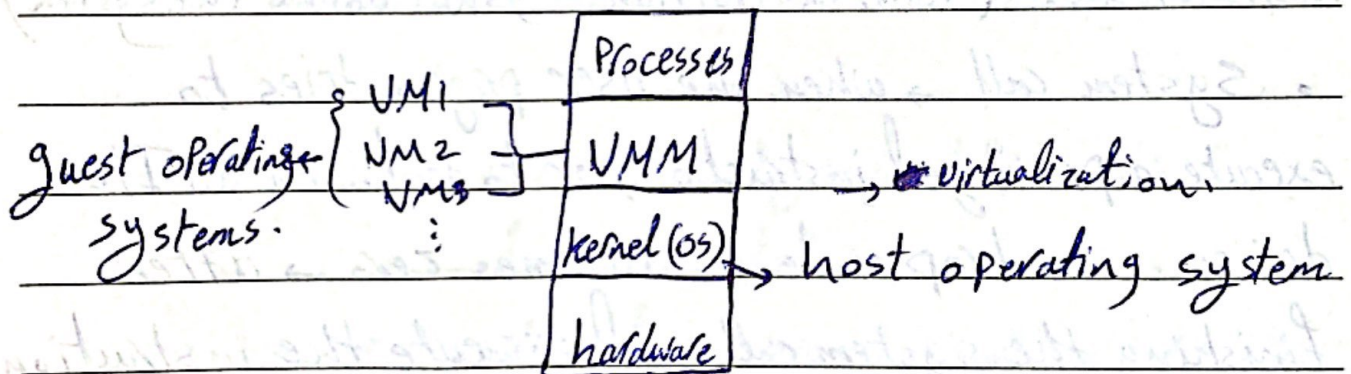
slides (38-55)

Lecture 4

- ~~software~~ interrupt number → interrupt vector → determine the interrupt service routine

* (difference between emulation, virtualization)

- emulation :- Computer program's ability in an device to emulate the behavior of another program or device. (slow)
- Virtualization :- The process of creating a virtual instance of computer hardware platforms.



slide 40

- Local Area network (LAN) → building
 - Wide Area network (WAN) → Countries
 - Metropolitan Area network (MAN) → City
 - Personal Area network (PAN) → one person
- ↳ ex:- health care applications.

Throughput :- number of tasks a system can process (finish) in a given amount of time.

- Fault Tolerance :- system's ability to continue operating uninterrupted despite the failure of one or more of its components. (the user won't notice the failure)
- Graceful degradation :- The ability of the system to maintain limited functionality (at reduced level) when portions of the system break down.
- System ~~reliability~~ reliability :- related to the frequency of failure, when the system is reliable, the frequency of failure is minimal.

- Mesh topology :- Network setup where each component is interconnected with one another.
 - advantages → high performance
 - drawback → complexity, high cost

- Client server → The client nodes request for services and server node responds with services.
- peer to peer ^{network} ~~network~~ (P2P) → each node can request for services and provide services. (needs more effort)

Cloud computing :- access different services through internet

Lecture 5

Open source operating system, operating system available in source-code format, it doesn't have to be free to use like the free operating system.

End of Chapter one.

slide(1-12)

Chapter 2

• System call, is a request from computer software to an operating system's kernel to perform a specific service.

→ Provide useful functions

• operating system services :- (user can notice)

1. User interface (UI) → The part of an operating system that allows a user to enter and receive information.
2. Program execution and identifying errors.
3. I/O operations (like dealing with files stored on the HD)
4. File-system manipulation → gives permission for operation on files.
5. Communications between different processes.
6. Error detection → The operating system keeps track of all the operations and processes and makes sure that they are running correctly without any errors, by using certain techniques.

→ just to make the system runs operating system services :- (user can't notice) efficiently

- 1- Resource allocation → the process of managing the different processes between hardware operational resources.
- 2- logging :- to keep track with each user and their actions on the system, especially when there are multiple users.
- 3- protection and security

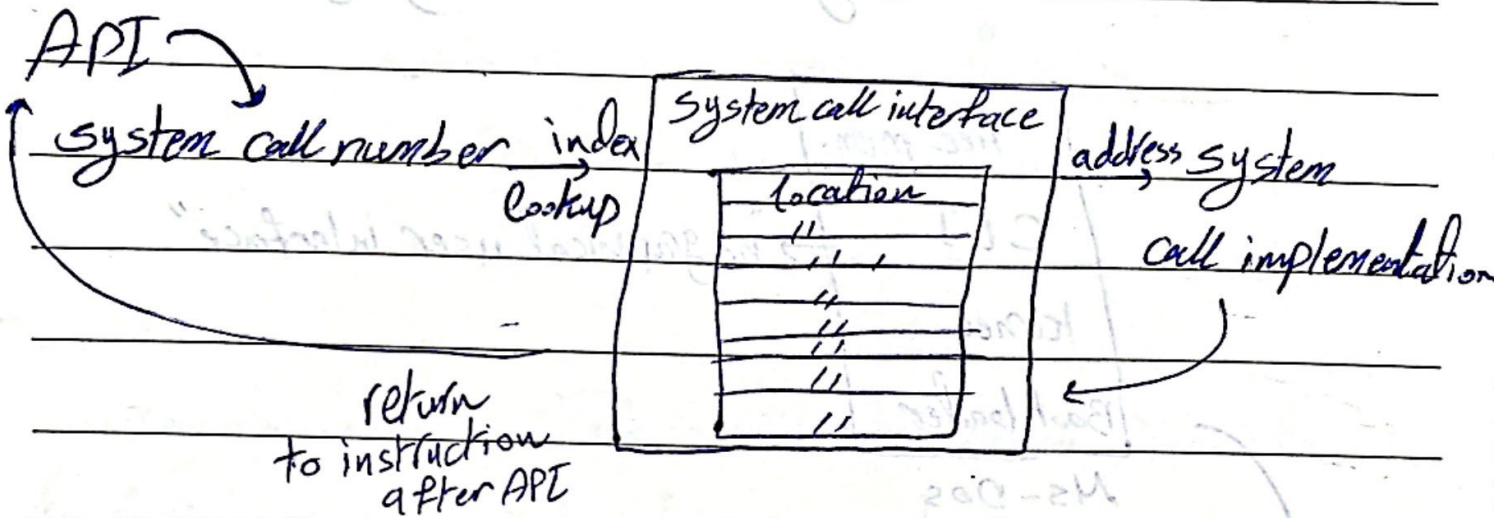
in System call → control is transferred from user mode to the kernel mode

• System Programs → Programs used by the kernel.
↳ in user mode.

shells examples :- Csh, Sh, Bourne

- The transition from user mode to kernel mode occurs when the application requests the help of operating system or an interrupt or a system call occurs.
- if the privileged instructions are executed in user mode, it is illegal and a trap is generated.
- The mode bit is set to 1 in the user mode - it is changed from 1 to 0 when switching from user mode to kernel.
- To copy the contents of one file to another file we need a lot of system calls, at every step. Even if any error occurs, this requires a lot of system calls, so any command needs thousands of system calls.
- (any simple operation needs thousands of system calls)
- Programmers deal with API not the system call because the system call is operating specific.
- Each API has a specific system call that we can access through it. (ex: ~~stdio.h~~) ~~include~~ `read()` ~~include~~
- The set of instructions to use or access a certain system call are in the kernel. ~~The system~~

- each system call has a number that represents an index in a table of system calls in the system call interface.



- Usually the names of the system calls aren't identical to the names of the APIs, but there is a great similarity between them.

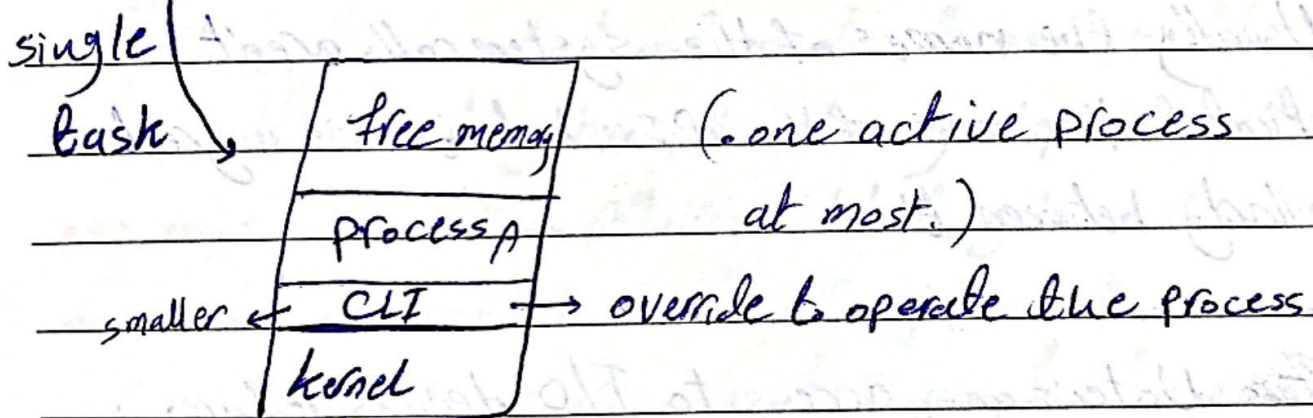
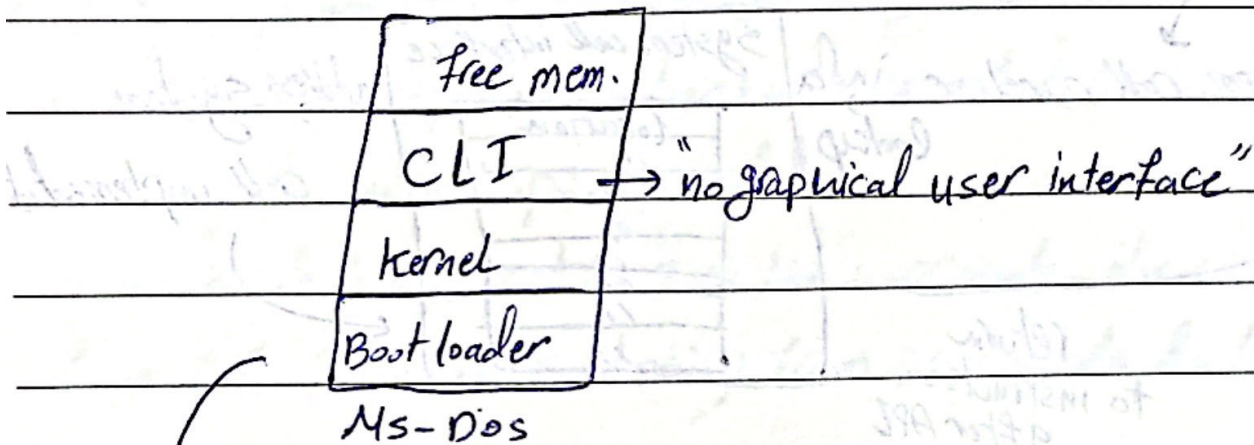
~~Then~~ Note:- any access to I/O devices requires a system call. (ex:- open file stored in HD)

API contains → return value / function name / parameters

- Note: Libraries that are included maps the API with the system call (read (API) → open (sc))
- Libraries have the implementation of each system call.

• Single task \equiv No operating system

• MS-DOS :- single task operating system



• Free BSD, multitasking operating system

• system calls → in the kernel

• system programs → services by the OS to provide a convenient environment for program development and execution → in the user mode

• most users' view of the operating system is defined by system programs and user application.

• Registry → manages resources and stores configuration settings for applications in the system (ex: windows registry)

• remote login :- The ability to access the data stored on a computer from a remote location.

↳ example of system services / programs

(application examples :- PuTTY, WinSCP

↓
transfer files



• The Role of linker / loader :-

{ linker → generate executable files


{ loader → load executable files to memory.

↳ system programs

• APPs compiled on one system usually not executable on other operating systems due to different ~~ISA~~ system calls, file formats ... etc.



linker :- system program, links two object files in one object file (library file and machine code file), gives executable file.



ex:- linker maps the abs value in the Program to the math library using ABI

Remember :- operating system elements :-

- 1- Abstraction
- 2- mechanism
- 3- Policy

Ex :- Policy :- LRU, LFU

mechanism :- Count the number of

times a page is used

the two policies may use the same

mechanism

P₁ | P₂ | P₃ | P₁ | P₂ LRU → P₂
LFU → P₃

• simple structure OS (like MS-DOS) → small kernel with few functionalities (limited), → fast boot

→ monolithic and windows

• More complex OS (like UNIX) → very large kernel with a lot of functions, and number of services in the user mode much less than services in the kernel. slower boot

→ system programs (very limited)

• layered operating system

use

or

layers ~~alternatively~~ modules (like linux)

→ more complex

→ The kernel executes a lot of

→ each layer contains

functions that are divided into modules

certain functions

install OS → kernel with basic modules

→ other modules are installed by the user.

- microkernel → (like MACOSX) → very small kernel but can be extended easily.

in microkernel →

- Two processes can communicate through the kernel, using interprocess communication.

- Context switch is the process of storing the state of a process, so that it can be restored and resume execution at a later point. This allows multiple processes to share a single CPU. (switching between more than one process (ex:- user process and system call))

- push and pop and adjusting the PC counter ~~and~~ reg.

Context switching take time, so the higher the number of context switching, the lower the performance.

- difference between BIOS and uefi

- ↳ The way they store information, configuration and setting

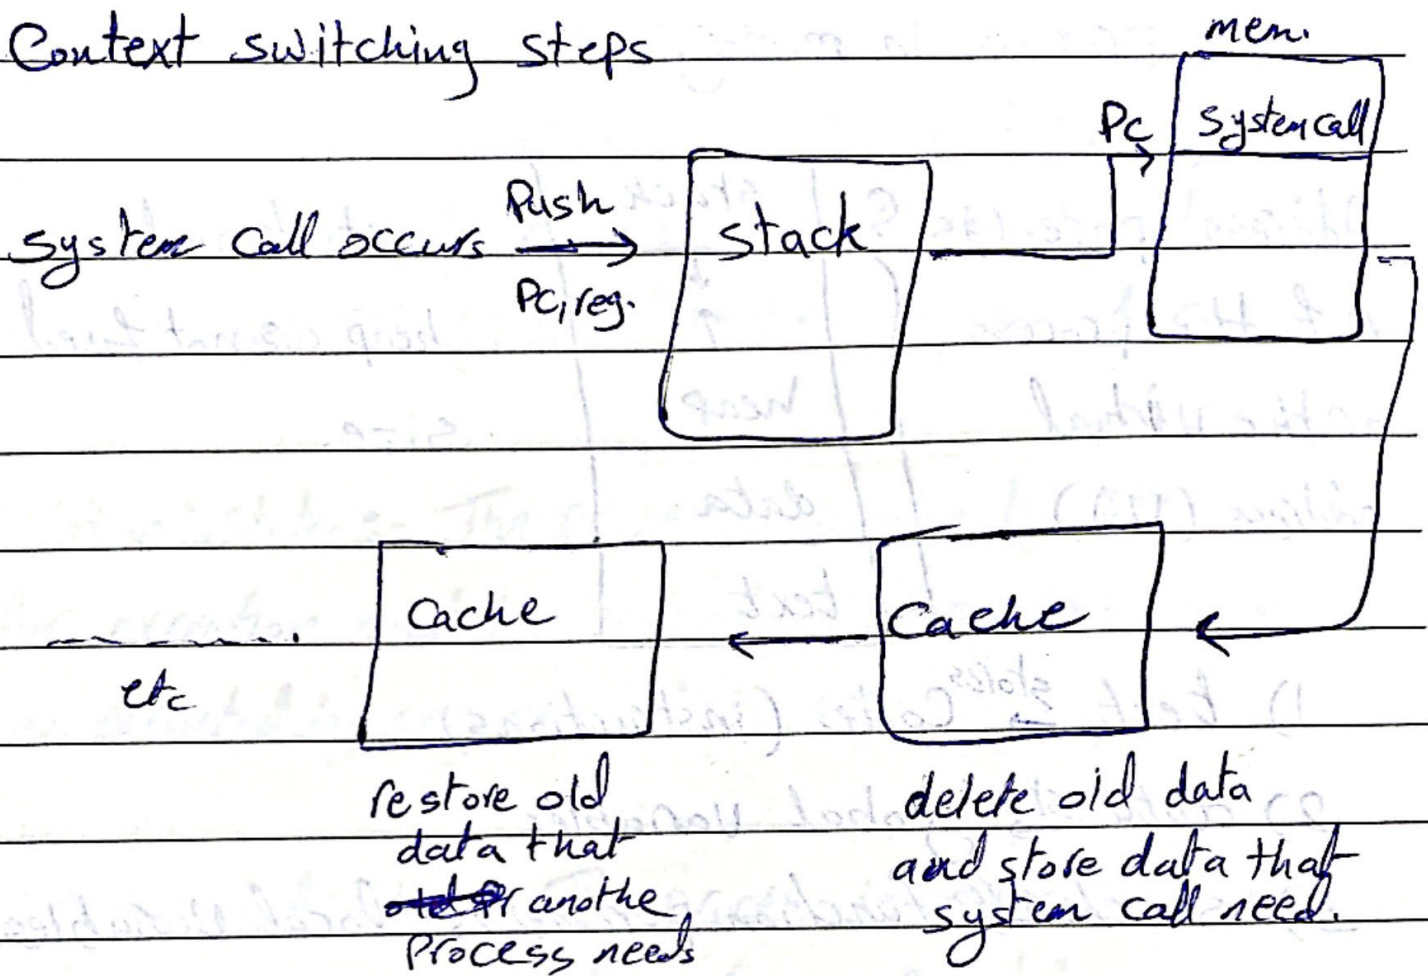
- ↳ Updating BIOS firmware is a bit difficult.

- Firmware → has direct control over low level hardware

- ↳ ≡ BIOS.

- tcpdump is a packet analyzer that is launched from the command line. (alternative - Wireshark)

• Context switching steps



cache hit → hit } affect performance.
cache cold → miss }

End of Chapter 2

Chapter 3

slides (1-9)

Lecture 10

- Process in memory

Address space (AS) of the process or the virtual address (VA)	stack	<ul style="list-style-type: none"> stack and heap are not fixed size. because the size depends on the application
	↓	
	↑	
	heap	
	data	
	text	

- 1) text ^{stores} → Code (instructions)
- 2) data ^{stores} → global variables
- 3) stack ^{stores} → Function parameters, local variables, return addresses
- 4) heap → Variables that are allocated at runtime.
(ex:- by malloc function)

• global variables :- are not defined inside any function and have a global scope.

- Process state

- 1) "New" state :- In this step, the process is about to be created but not yet created. (Command or double click)
- 2) "Ready" state :- The process is ready to run, after it has been approved by the operating system.
(according to security and resources)

Processes that are ready for execution ~~are~~ by the CPU are maintained in a queue for ready processes (run/ready queue), each process has its own process ID (PID) [→] unique integer number.

3) "Run" state :- The process is chosen by CPU for execution and the instructions within the process are executed by ~~any~~ one of the CPU cores.

4) "terminated" state :- The process is killed, (finished) as well as PCB is deleted. (use "exit" system call.)

• The short term scheduler or the dispatcher moves the process from ready state to running state.

• if the time ^{allotted} ~~limit~~ for the process ^{has} expired (run → ready) "interrupt".

• if the process needs any I/O operations during the running state (run → waiting queue) → (waiting queue → ready queue)

slides (10-29) Lecture 11

• Process Control block (PCB) :- Data structure used by the operating system to store all the information about a process.

• Process scheduling allows OS to allocate a time interval of CPU execution for each process. Also, the CPU scheduler selects a process among the processes that are ready to execute and allocates CPU to one of them. \rightarrow according to priority

• Child process :- process created by another process.

• Context switching time is affected by the PCB size.

\downarrow PCB \downarrow time (ct-sw) \downarrow
as well as it is affected by the operating system itself and its complexity.

• The context switching is useful when :-

($2 \text{ ct-time} < \text{idle time}$)

Note :- queues \rightarrow tail points to the last process
head " " " ~~process~~ process added first

• Operations on Processes :-

- Process Creation → using certain system call.

- Process termination

- Scheduling

- Communication

$Pid = Fork();$ → Process creation, returns the new process id. ~~it~~ (until terminates)


- if the parent process wait for the child process ↑ :-

• $exec()$ system call → The previous executable file is replaced and new file is executed / ~~replace the old~~ all code (text) and data in the process is lost and replaced with the the executable of the new program.

→ it makes the child process perform tasks different from the tasks of the parent process.

• $wait()$ system call → Parent process waiting for the child to terminate.

• Once the child process is terminated, it will return information to the "wait" system call (parent process) through the wait system call ~~it~~ using $exit()$ system call.



P_x takes the same code after `fork` in P_1 and only variables before `fork`.

- The parent process (`id = id` returned from `fork()`);
- child process (`pid = 0`)

Ex:-

```
cout << "Hello OS" << endl;
```

```
type - Pid_t pid = fork();
```

```
cout << "Bye, see you later!" << endl;
```

output:-

Hello OS	
Bye, see you later!	→ parent
Bye, see you later!	→ child

Note: `exec()` → take 2 Parameters, `Pid` of the process which perform the task, and the new executable file.

- if `fork()` returns a negative value → operation failed.

Ex:- `cout << "Hello os" << endl;`

`Pid t Pid1 = fork();`

`Pid t Pid2 = fork();`

} → 4 processes
2 Parents, 2 children

`cout << "Bye, see you later!" << endl;`

output:-

output after fork

of forks
= 2

Hello os

Bye, see you later!

3 forks → 8 times

4 times.

- A call to `wait()` blocks the calling process (parent) until its child process exits. After child process terminates, parent continues its execution.

- Once the parent process terminates, all of its children processes should terminate too. (best practice) else → if the children processes don't terminate with the parent process they become called orphan processes. (they keep the resources).


- Multiprocess Architecture → better performance than single process Architecture.

→ in Browsers → each tab is separate process!

Sandbox → improve security by isolation and restrictions

• Communications Models :-

a) shared memory



→ accessible by PA and PB, both of them can read from it, or write on it.

↳ it is located in the address

space (AS) of the process

that requested communication.


• The function of the kernel here is to allow one process to access the address space of the other process which has the shared memory.

b) message passing



• Communication in this model is through the kernel.

• The kernel here is involved in the communication process.



- Compiler/assembler → Compiler is producer, and assembler is consumer.
- assembler/linker → assembler is a producer, and linker is consumer.
- etc.

Two Variations (producer-consumer problem)

1. unbounded buffer : no practical limit on the size of the shared memory.
2. bounded buffer : shared memory with fixed size.

- Synchronization problem in the shared memory model:-
 - two processes need to access the same area at the same time (without any synchronization) (shared area)

~~buffer size shared memory size~~ → circular buffer
→ bounded buffer
In/out values in the buffer (shared memory)

- idx should be in range (0 to size - 1)
- * Counter variable = # of items in the buffer, (valid data)
 - ↳ 3 steps instruction.

Consumer / in/out شافيه التبين ال Producer وال Consumer.



Bounded-Buffer – Shared-Memory Solution

```
#define BUFFER_SIZE 10
```

```
typedef struct {
```

شبه ال class
بالجافا
مش مهم
التفصيل.

...
item;
بالباية in و out
حياتشوا بنفس المكان
يعني فاضي ال buffer

```
item buffer[BUFFER_SIZE];
```

→ bounded array (Circular Array)

```
int in = 0;
```

```
int out = 0;
```

- The shared buffer is implemented as a circular array with two logical pointers: *in* and *out*

*• *in* points to the next free position in the buffer → بالباية يكون 0

*• *out* points to the first full position in the buffer.

*• The buffer is empty when $in == out$

*• The buffer is full when $((in + 1) \% BUFFER_SIZE) == out$

- Solution is correct, but can only use $BUFFER_SIZE - 1$ elements

عملتوا عشان



أخليه Circular
بنعالج حالة لو كان ال in

عاشر Pointer فال out حيكون أول ال array لأنه Circular.

ما بقدر أستخدم كل العنصر
تاع ال buffer لأنه يفصل في
وحدة empty تباشر

in out



Producer Process – Shared Memory

local variable ←

item next_produced;

infinite
loop ←

```
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;
```

```
}
```

مطلبا ال buffer
full ما بقدر أحط
لحد ما ال Consumed
يفتني ويحرك ال out
ويجبر عني مساحة أحط.



* اللي بوقف ال Producer، انه يكون ال buffer (full) واللي بوقف ال Consumer يكون ال buffer (empty)



Consumer Process – Shared Memory

```
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 */
}
```

* بحالة ال Full وايضا في element مش متجيب ، ما بقدر امشي
ال in خطوة ويغير يساوي ال out .

* هي الحالة ما فيها Race conditions لانه ال Race ما يجبر بال
local variables بجبر بال Shared var .

Operating System Concepts – 10th Edition

3.43

Silberschatz, Galvin and Gagne ©2012

43

ال buffer وال in وال out هم ال Shared ،
- in/out بي لما يكون ال buffer empty فمافى ال buffer overwriting
- ال Producer بيمل update ال in وال Consumer ال out فمافى
بيتم (Race conditions)



Producer

```
while (true) {  
    /* produce an item in next produced */
```

```
    while (counter == BUFFER_SIZE)  
        ; /* do nothing */  
    buffer[in] = next_produced;  
    in = (in + 1) % BUFFER_SIZE;  
    counter++;  
}
```

حالة ال Full →
مينا نفند عال counter
لتحيد ال buffer
لو فاضي او مليان
مش عال in وال out.





Consumer

```
while (true) {  
    while (counter == 0) → Empty  
        ; /* do nothing */ → لحد ما يصير في  
    next_consumed = buffer[out]; ← Item جوا  
    out = (out + 1) % BUFFER_SIZE; buffer  
    counter--;  
    /* consume the item in next consumed */  
}
```

* هيك حلينا مشكلة، انه فيه element بضل فاضي بحال Full
بس صار فيه Race condition لأنه التين
يعملوا على counter



Race Condition

- `counter++` could be implemented as

load \leftarrow register1 = counter
Increment \leftarrow register1 = register1 + 1
Store \leftarrow counter = register1

- `counter--` could be implemented as

load \leftarrow register2 = counter
decrement \leftarrow register2 = register2 - 1
Store \leftarrow counter = register2

- Consider this execution interleaving with "count = 5" initially:

S0: producer execute	register1 = counter	{register1 = 5}
S1: producer execute	register1 = register1 + 1	{register1 = 6}
S2: consumer execute	register2 = counter	{register2 = 5}
S3: consumer execute	register2 = register2 - 1	{register2 = 4}
S4: producer execute	counter = register1	{counter = 6}
S5: consumer execute	counter = register2	{counter = 4}

النتيجة هي أن counter قبل ما نعدل
هو 5. incorrect لأنه كان لازم تفضل 5.



• If the "in" pointer points to the same location as the "out" pointer, we assume that the buffer is empty, or $\text{Counter} = 0$.

• Atomic operation:- operation that cannot be interrupted while being executed

ex:- $\text{register 1} = \text{Counter}$

• ~~no~~ no race condition in atomic operations

• $(\text{Counter}++)$ is not atomic because its 3 steps instruction

• Question - why was there no race condition in the solution (where at most $N-1$) ?


• because there aren't any shared variables and they aren't modifying the same variables.

• message passing model :-

• if the message size is fixed \rightarrow easier in system implementation.

• if the message size isn't fixed \rightarrow easier for the programmer.

\rightarrow direct communication



• Buffering

• we have a link between 2 processes,

this communication link uses a buffer which is a queue of messages, and attached to the link. It is ~~used~~ used to avoid losing messages in case of speed mismatch.

• This buffer ~~may~~ could be :-

1- zero capacity → if the sender send a message, the receiver should receive it directly, otherwise the message will be lost. (no buffer)

• The sender must remain blocked until the recipient receives the message.

• If the receiver uses the (receive system call) before the sender send a message, ~~so~~ the receiver will receive a Null message, so it must remain blocked too.
* in zero capacity → send, receive are ~~not~~ blocking system calls.

2- bounded capacity → (real life scenario)

3- Unbounded capacity

• Factors taken into ~~xxxxxx~~ consideration to determine the size of the buffer :-

- 1 - Communication rate
- 2 - speed mismatch between sender and receiver

• if producing speed is faster than the speed of consumption \rightarrow large buffer

• if producing speed is slower than the speed of consumption \rightarrow could be zero capacity buffer.

~~What is Remote Procedure Call~~


↳ ~~RPC~~ RPC (remote procedure calls) occurs when there is communication between two processes, each one is on different system, in (RPC) the socket is created. (socket is the end point of communication link) (socket is a combination of an IP address and a port number).

↳ Port \rightarrow represents the type of the service
Well known services (0 - 1024)

ex:- http in port 80 (web server)

Pipe:- implementation of the communication link between two processes which communicate through message passing.

- Ordinary Pipes → unidirectional



- ~~creates~~ parent-child relationship.

- Named pipes → unidirectional but ~~are~~ not Parent-child relationship

Note:- half-duplex → send or receive at a specific moment.

Full-duplex → Processes send and receive messages at the same time.

slides (55 - additional) lecture 16

- ordinary pipe (unnamed) pipe lasts only as long as the process. → created by the process
- named pipes can last as long as the system (kernel) is up, beyond the life of the process. → created by the kernel.

• fd (file descriptor) :- number that uniquely identifies an open file in a computer's OS. (id file)

std in :- standard input

↳ ex :- keyboard

std out :- default / standard out

↳ ex :- monitor

kernel	
fd table	
0	std in
1	std out
2	std error

reserved
fd

- std error :- where a process can write error message.

Parent

child

fd [0]
WR

fd [1]
read

{ int myFD[2]; → create 2 fd (for write and read)
Pipe (myFD); → API to create a pipe
pid_t x = fork(); → create a child (ordinary pipe)


```
if (x > 0) { → Process ID > 0 → Parent  
    write (myFD[1], "hello child", 30); → message size  
}
```

```
else { → child
```

```
int n = read (myFD[0], char msg); → returns # of  
} read on charac.
```

Note:- once the parent process is terminated,
the pipe will be deleted.

Note :- if we use (write (1, "hello child", 30);)
the message will appear
on the monitor because the FD of the (std out) is 1

• Bidirectional Pipes → only using 2 Pipes.

Chapter 4 slides (1-6)

Process :- Program in execution, and the execution style is "sequential" → (there is no parallelism in the process). if one instruction is blocked, then all the other instructions will be blocked too.

↳ Therefore, the main process is divided into sub processes

- if one of the sub processes is blocked, other sub processes won't be affected.

- each sub process is called "Thread"

- ~~is~~ multithreaded process → Parallel execution

- The code and Data segment is shared between all threads that are related to the same process, as well as files.

- each Thread has a unique ID, PC, stack and registers

- It is possible to split data between threads so that each thread performs the same instructions on a specific set of data, or to split the instructions so that each thread performs a different task.


slides (7-9)

type of context switching / lecture 17

• Thread switching → ~~process switching~~ From one thread to another thread in the same process.

Thread switching is very efficient and cheaper than process switching because it involves switching out the PC, registers and stack pointers. (lower overhead)

• MTP (T_1, T_2)



• any blocking or I/O operation on core 1 (T_1) won't affect the execution on Core 2.

→ multithreaded process takes advantage of multicore arch.

→ It is preferable to use a number of ^{Threads} ~~processes~~ equal to the number of cores. ~~(N)~~


→ multithreaded processes take advantage of single core arch. too.

• Parallel execution → ^{multi} different processes (task) on the same time. → multicore system. (2 cores → 2 threads max)

• Concurrent execution → only one thread (task) is executed. → in single core system. at a time

- Data Parallelism :- Distribute data to more than 1 thread while performing the same operations on it.
ex:- summation on elements in an array.

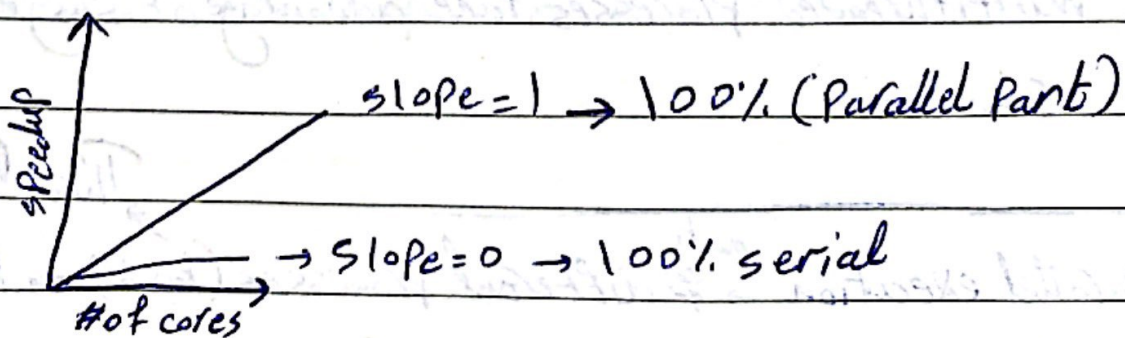
- Task Parallelism :- distributing thread across cores, each thread performing unique operation.



- ex:- Performing different operations on an array.

Note :- $\text{Parallel part} = (1 - \text{serial part})$.

- Inverse proportion between the serial portion and the speedup that can be obtained.



- We cannot have a line slope of more than 1.

• if the slope = 0.5, find the percentage of serial and parallel portion?

• Kernel is multithreaded, to serve the requests that are from the user mode.

• ~~Threads~~ Threads in user mode are ~~managed~~ managed by library in the user level.

• The threads that are in the kernel mode, are managed by the kernel.

• Multithreading models :- relationships between user threads and kernel threads.

1. many-to-one → no parallel execution. no limitations on # of user threads

2. one-to-one → kernel threads = user threads

→ There is parallel execution.

→ There is a limitation on the # of user threads to avoid overhead.

3. many-to-many → $UT \gg KT$ / no limitations / parallel execution



Amdahl's Law

18

- Identifies performance gains from adding additional cores to an application that has both serial and parallel components
- S is serial portion
- N processing cores

Parallel Part of the code ← $\text{speedup} \leq \frac{1}{S + \frac{(1-S)}{N}}$ → Parallel

- That is, if application is 75% parallel / 25% serial, moving from 1 to 2 cores results in speedup of 1.6 times
- As N approaches infinity, speedup approaches $1 / S$

if the execution time was 100s after speedup → new execution time = $\frac{100}{1.6}$

Serial portion of an application has disproportionate effect on performance gained by adding additional cores

- But does the law take into account contemporary multicore systems?



- any variable in the Data segment (global variable), are shared between all threads in the process.

- pthread_join (API used to prevent the main thread from terminating the process before the created threads have finished their tasks.) without this API the output won't be correct.

- In windows ~~to~~ when creating a thread we need to define ~~the~~ HANDLE first.

- same other steps are in all threading libraries.

↑ level of parallelism ↑ creating, managing threads

↓
to avoid this → implicit threading

Thread pools → pool of threads is created by the compiler, where they await work. (its size ^{is} determined by compiler)
→ faster than thread creation while execution, serving requests is faster.



Pthreads Example

```

#include <pthread.h>
#include <stdio.h>

#include <stdlib.h>

global variable shared between all threads
int sum; /* this data is shared by the thread(s) */
void *runner(void *param); /* threads call this function */

int main(int argc, char *argv[])
{
    data type
    pthread_t tid; /* the thread identifier */
    pthread_attr_t attr; /* set of thread attributes */

    /* set the default attributes of the thread */
    pthread_attr_init(&attr);
    /* create the thread */
    pthread_create(&tid, &attr, runner, argv[1]);
    /* wait for the thread to exit */
    pthread_join(tid, NULL);

    printf("sum = %d\n", sum);
}

```

function to execute (pointing to runner)

Parameters (pointing to argv[1])

waiting for the created thread to terminate first before the main thread. (pointing to pthread_join)



Pthreads Example (Cont.)

```

/* The thread will execute in this function */
void *runner(void *param)
{
    int i, upper = atoi(param);
    sum = 0;

    for (i = 1; i <= upper; i++)
        sum += i;

    pthread_exit(0);
}

```

sum numbers from 1 up to param. (pointing to the for loop)

termination, return to main function (pointing to pthread_exit)



Pthreads Code for Joining 10 Threads

```
#define NUM_THREADS 10
```

```
/* an array of threads to be joined upon */  
pthread_t workers[NUM_THREADS];
```

tid ⇒ workers[i]
⇒ array with size 10, each element is pthread type

```
for (int i = 0; i < NUM_THREADS; i++)  
    pthread_join(workers[i], NULL);
```



Silberschatz, Galvin and Gagne ©2018

- Performance in multithreaded programs depends on the hardware (# of cores), so we need to determine the best # of threads.

In "Implicit Threading", This task which consists of specifying the number of threads, and task of each thread becomes to the compiler instead of the programmer.

- In Implicit threading the compiler create the threads, and distribute the tasks ~~to~~ between them. (S)

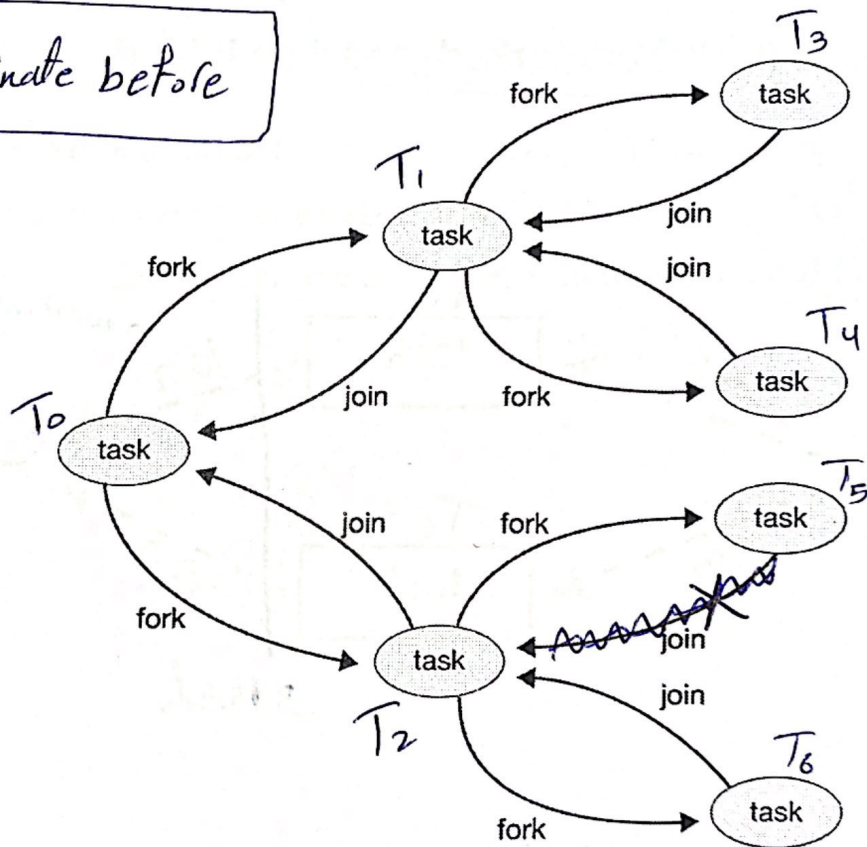
- Note :- In linux (Thread \equiv Task)

- The parent ~~and~~ thread cannot terminate its work before its child ~~thread~~ thread has ~~to~~ terminated. Unless the child thread doesn't (Join).

- In openMP we have to determine the parallel region.

- The number of times the instructions within the parallel ~~region~~ region are executed is equal to the # of ~~times~~ \rightarrow In openMP. Threads / cores

*T₀ can terminate before
T₅*



OpenMP

- Set of compiler directives and an API for C, C++, FORTRAN
- Provides support for parallel programming in shared-memory environments
- Identifies **parallel regions** – blocks of code that can run in parallel

#pragma omp parallel

Create as many threads as there are cores

compiler

Threads = Cores

```
#include <omp.h>
#include <stdio.h>

int main(int argc, char *argv[])
{
    /* sequential code */

    #pragma omp parallel
    {
        printf("I am a parallel region.");
    }

    /* sequential code */

    return 0;
}
```

if Dual core
this instruction
will be executed
two times: (T₁, T₂)

quad → 4 times



- overloading :- multiple functions with the same name, but different implementations.

~~a different~~

①

Pa

`fork();`
`exec();`

2 different forks in UNIX :-

~~fork~~ fork only duplicates the calling thread. ex: T_1, T_2, T_3

↓
duplicates calling thread only
→ T_4 → new code

②

Pb

`fork();`

→ fork duplicates all threads.


ex: T_1, T_2, T_3

↓
duplicates all ~~threads~~ threads.
→ T_4, T_5, T_6

- Synchronous signal :- signal generated by the process itself (ex: division by zero, access illegal memory location) → internal source

- ~~asynchronous~~ asynchronous signal :- from external source.
(ex: I/O operations)

- There is default handler ~~in~~ in the kernel, executed in case the user doesn't override the handler.



In single threaded process, the thread receive the signal and handle it.

In multi threaded process :-

- 1) Deliver the signal to the thread whose task caused it. (ex:- division by zero)
- 2) Deliver the signal to every thread in the process. (ex:- Termination process signal)
- 3) Deliver to a certain thread. (ex:- "kill" system call)
- 4) Deliver ~~the~~ all signals to a certain thread.

Asynchronous cancellation → The target thread should terminate its task immediately without completing any extra work after cancellation. → causes some errors like cancel before modifying certain variable.

- Threads allows threads to disable or enable cancellation. a thread cannot be canceled if cancellation is disabled, but cancellation requests remain pending, so the thread can later enable cancellation and respond to the request (mode: off)
- Asynchronous cancellation: one thread terminates immediately the target thread.
- Deferred cancellation: one thread indicates that a target thread is to be canceled, but cancellation occurs only after the target thread has checked a flag to determine whether or not it should be canceled.
- after cancellation, the "cleanup handler" invoked, to free the resources that a thread may hold at the time it terminates.

Thread-local storage: Allows a thread to store data that is visible to all the function invocation within that thread, but it is invisible to other threads.

The LWP appears to be a virtual processor on which the application can schedule a user thread to run, to the user thread library.

↑ LWP ↑ Overhead → more read and write operations

How many LWPs to create? according to the upcalls and upcalls ~~work~~ handler.

Note :-

CPU-bound Process :- Process that spends the majority of its time using the CPU (ex:- calculations)

I/O-bound Process :- Process that spends the majority of its time doing I/O operations. It needs more kernel level threads than the CPU-bound process due to a lot of blocking system calls.


"End of chapter 4"

Chapter 5

slides (1-10)

lecture 23

remember:-

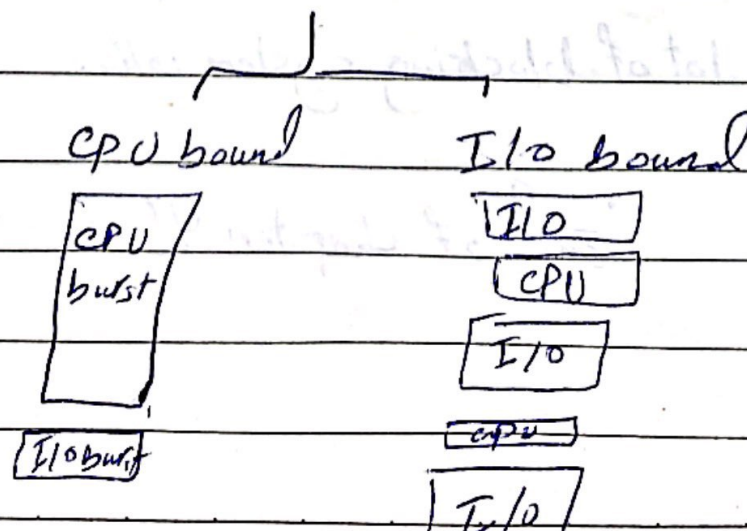


• if P_a needs any I/O operation, the CPU will start working on P_b .

• Profiler :- software that examine, analyze, and create useful summaries of the execution of a process.

• execution of any process is CPU burst and I/O burst. Their distribution affects CPU utilization.

• # of CPU, I/O bursts and their length depends on the nature of the process



• Switching from running to waiting state \rightarrow when the process needs any I/O device, then the scheduler will determine which process from the ready queue should allocate the CPU.

• Switching from running to ready state \rightarrow when the process time expires, " " " "

• Switching from waiting to ready \rightarrow when the process is finished with the input and output device.

• Resources utilization in preemptive scheduling is better, but can result in race conditions.

• Scheduler works in kernel mode

• Dispatcher works in kernel mode then switches to user mode.

Dispatch latency $>$ ~~switch~~ context sw. time.

~~DEFINITION~~ Scheduling Criteria \Rightarrow scheduling goals or objectives

Turnaround time is the completion time (CPU time / I/O time / waiting time)

• optimization in any system design is trade off, so we have to work on the balance of any improvement or optimization, so that we don't affect the rest of the tasks in a negative way.

scheduling algorithms:-

1) FCFS \rightarrow in the order they reach the ready queue.

(slide 12) \rightarrow CPU utilization = $\frac{30}{30} = 100\%$ \rightarrow CPU is busy
 \hookrightarrow total time

\hookrightarrow average Turnaround time (completion time) = $\frac{24+27+30}{3}$

\hookrightarrow Throughput = $\frac{\text{\# of completed tasks}}{\text{time}}$

$= \frac{3}{30} = 0.1 \text{ P/ms}$


2) SJF \rightarrow Associate with each process the length of its next CPU burst \rightarrow better average waiting time but we need to determine the next CPU burst length

- Scheduling Algorithms \rightarrow decide which and when the process should leave the CPU.

ex: non-preemptive SDF \rightarrow the process won't leave the CPU until it finishes its burst.

- preemptive SDF \rightarrow The process will be ~~preempted~~ preempted if there is another process with CPU burst $<$ remaining time.

ex slide 20



• wait time (P₂) = 8 - 1 = 7

• wait time (P₃) = 12 - 2 = 10


• wait time (P₄) = 21 - 3 = 18

• average turnaround time = $\frac{\text{wait time} + \text{Burst time}}{\text{\# of Processes}}$

$$= \frac{8 + 11 + 19 + 23}{4}$$

- if there are more than 1 wait times they are all included in turnaround time.

Algorithm \rightarrow SJF (non-preemptive)



• wait time (P_1) = 0

• wait time (P_2) = 7

• wait time (P_3) = 15


• wait time (P_4) = 9

- FCFS \rightarrow affects interactive processes if they're ~~not~~ arrive to the CPU late, after long processes.
- SJF \rightarrow improves performance

Round Robin ex:-

$q = 10$

4 Process with 25 burst time




Max response time $\rightarrow n-1 \cdot q = 3 \cdot 10 = 30$

• $q \downarrow \rightarrow \uparrow \text{context switching time} \uparrow \text{overhead}$

• $q \uparrow \rightarrow \text{FCFS}$

• best value for (q) $\rightarrow \begin{cases} q > 80\% \text{ of CPU burst} \\ \text{time} \end{cases}$



• Response time in RR is better than the response time in SJF.

• Average TAT time in RR is greater than the average TAT time in SJF. (SJF is better in TAT)

• $q > \text{context switching time}$.




Example of Shortest-remaining-time-first

- Now we add the concepts of varying arrival times and preemption to the analysis

Process	Arrival Time	Burst Time	WT
P_1	0	8	4
P_2	1	4	0
P_3	2	9	15
P_4	3	5	2

- Preemptive SJF Gantt Chart



- Average waiting time = $[(10-1)+(1-1)+(17-2)+(5-3)]/4 = 26/4 = 6.5$





Example of RR with Time Quantum = 4

FFS


Process

Burst Time

3 times → context switching

P_1	24	WT 10	34 TA
P_2	3	4	7
P_3	3	7	10

- The Gantt chart is:



- Typically, higher average turnaround than SJF, but better **response**
- q should be large compared to context switch time
 - q usually 10 milliseconds to 100 milliseconds,
 - Context switch < 10 microseconds

• response time $R_1 = 0$

• " " $P_2 = 4$

• " " $P_3 = 7$



slides (25-34)

lecture 26

priority scheduling

- each process has an integer number which represents its priority (smallest integer \rightarrow highest Priority)

- highest priority \rightarrow lower burst time
 \hookrightarrow so its type of SJF.

ex. slide 28

$$\text{wait time } (P_2) = 0 \quad (0-0)$$

$$\text{wait time } (P_5) = 1 \quad (1-0)$$

$$\text{wait time } (P_1) = 6 \quad (6-0)$$

$$\text{response time } (P_4) = 18$$

in Example slide 26 :-

10 non-Preemptive

					arrival time
P_1	P_2	P_5	P_3	P_4	0
					10
					4
					2
					3

- if there ~~is~~ ^{is} more than 1 process with the same priority level, we ~~can~~ ^{can} use priority scheduling with other scheduling algorithms like Round-Robin, FCFS, or SJF.

run-to-completion is similar to non-preemptive

- Searching in ready queue takes $O(n)$, where "n" is the number of processes in the queue.



So → use multilevel queue.

- In multilevel queue, we use round robin at the same queue, and priority scheduling between different queues.

Thread scheduling

- Scheduling in user level threads is done by Threading libraries. → many to many model / many-to-one

- Scheduling in kernel level threads is done by CPU scheduler.


- in one-to-one model → no need to ~~schedule~~ ^{schedule} in the user level threads.

• In windows and linux the scope is \rightarrow SCS,
because they are one-to-one model

• Multicore CPUs are faster, and less power consumption

• Hardware threading (2 register sets or more on the same core "2 threads or more on the same core") \rightarrow multithreaded cores.

• Homogeneous multiprocessing \rightarrow Identical cores



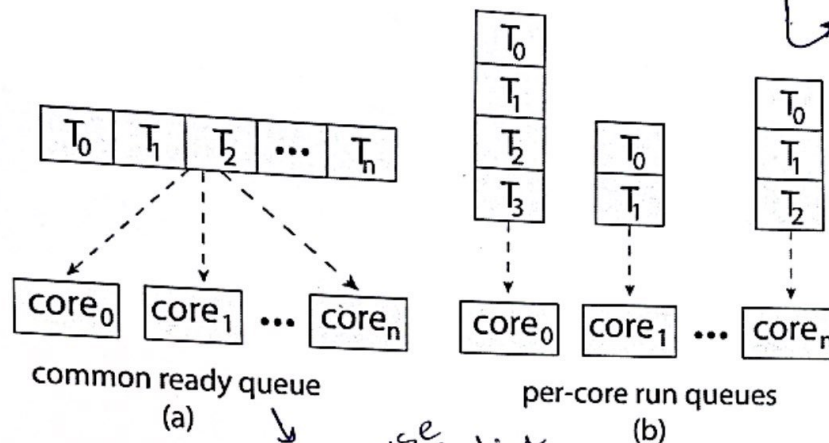
• Heterogeneous multiprocessing \rightarrow each core ~~was~~ has different tasks.

• Memory stall frequency depends on the code itself, size of the cache, and the ~~memory~~ locality.

• Memory stall time depends on the speed difference between CPU and RAM, and the RAM size.

Multiple-Processor Scheduling

- Symmetric multiprocessing (SMP) is where each processor is self scheduling.
- All threads may be in a common ready queue (a)
- Each processor may have its own private queue of threads (b)



Can cause
race condition

solution


should use
load balance.



Multicore Processors

- Recent trend to place multiple processor cores on same physical chip
- Faster and consumes less power
- Multiple threads per core also growing HW threading
 - Takes advantage of memory stall to make progress on another thread while memory retrieve happens

- Figure



T₀/T₁

↳ one of them to memory stall, and the other one continues its work, to avoid overhead.



Multithreaded Multicore System

- **Chip-multithreading (CMT)** assigns each core multiple hardware threads. (Intel refers to this as **hyperthreading**.)

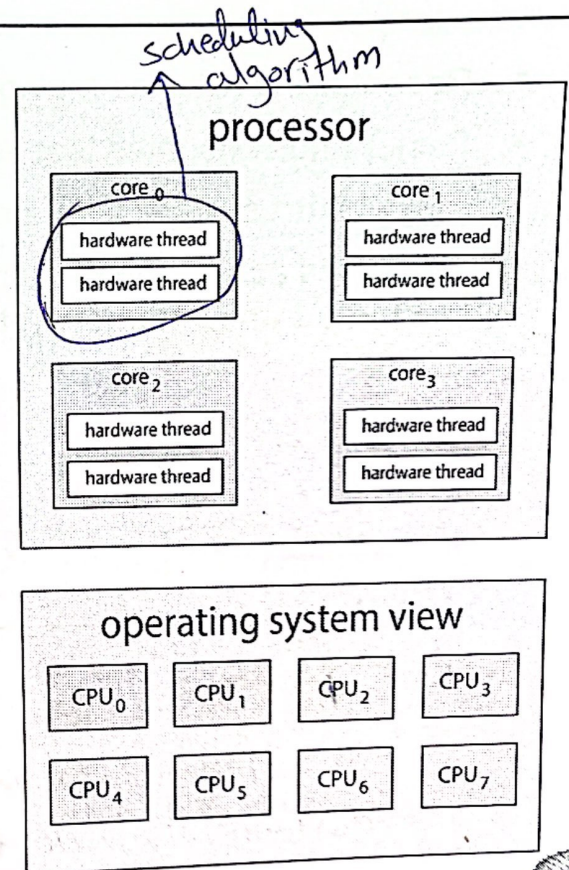
if the core supports 4 HW threads, then there are 4 Program counters.

- On a quad-core system with 2 hardware threads per core, the operating system sees 8 logical processors.

How many software threads can be scheduled on this Processor?

8 threads

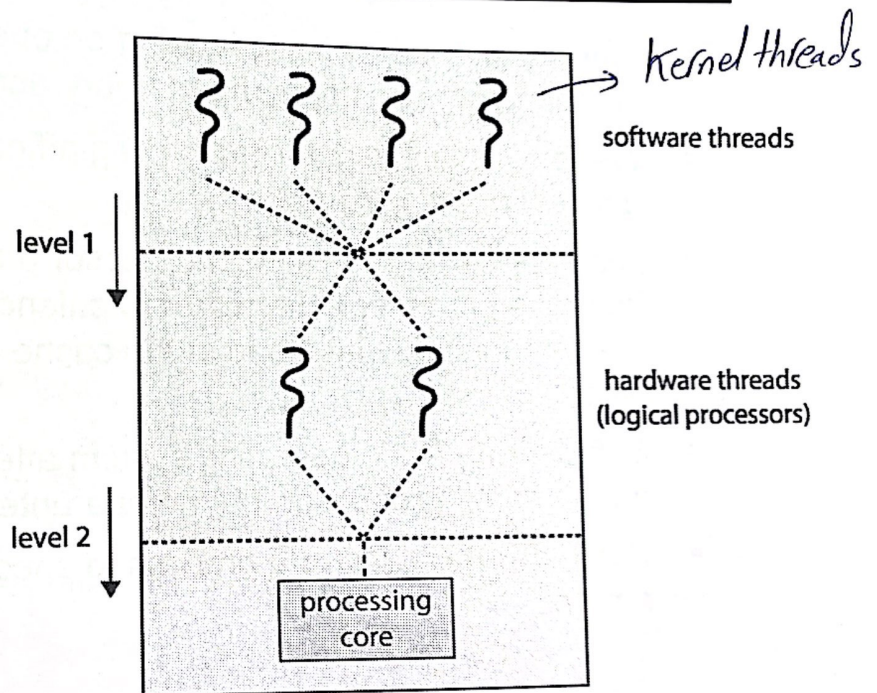
*↓
4 cores X 2 threads on each core.*



Multithreaded Multicore System


Two levels of scheduling:

1. The operating system deciding which software thread to run on a logical CPU
2. How each core decides which hardware thread to run on the physical core.



Hyper threaded
 • many-to-many and HW threading, how many scheduling decision?
3 levels.





- Cache 0 is warm cache for T_1 , and Cache 1 is cold for T_1
- ~~Affinity~~ affinity of core 0 \rightarrow which thread consider cache 0 as a warm cache. $\rightarrow T_1$
- load balancing affects Processor affinity negatively

"end of chapter 5"

Synchronization

Race Condition :- occurs when two threads or processes access a shared variable at the same time, and modify its value.

~~Atomic~~ instruction :- cannot be interrupted while being executed.

(increment, decrement, are not atomic instructions)

P₀) Pid & child = fork();

P₁) Pid & child = fork();

P₂) Pid & child = fork(); → next available id = 10

P₃) Pid & child = fork(); →

Q) What is the minimum number of processes that will take the Pid = 10?

Optimal ans :- 1 process. (Race condition may occur)

Q) What is the max. number of Processes that will take the Pid = 10?

ans :- 4 processes. (Worst case)

DATE

S M T W T F S

~~Critical~~

• Critical section :- Part of the program which tries to access shared resources.

• The critical section cannot be executed by more than one process at the same time.

• entry section → each process requires permission to enter the critical section.

• exit section → gives permission to the next process to enter the critical section.

Entry section (disable interrupts), so the process doesn't leave the critical section until it completes its task.

exit section (enable interrupts)

slides (9 - ²² ~~23~~)

lecture 29

critical section solutions:-

1- Interrupt-based solution, will cause some problems when the system is multicore, or the critical section code runs for long time.

2- Software solution, valid for two processes only.

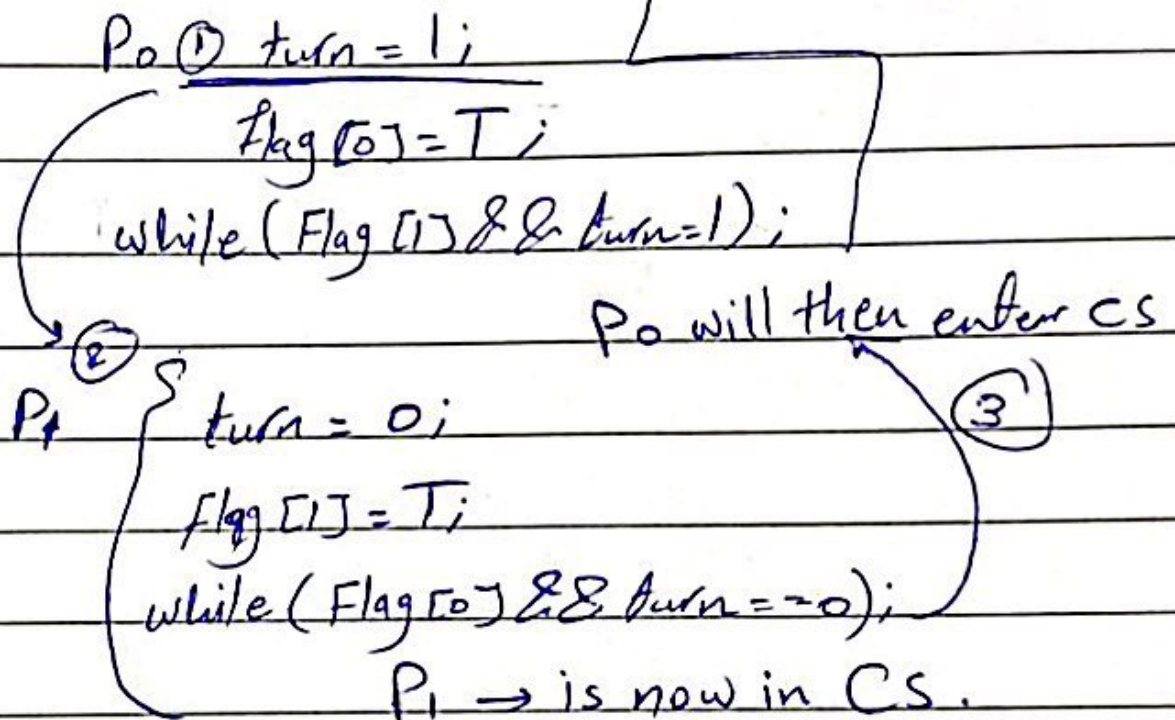
3- Peterson's solution, valid for two processes only. uses "flag" variable to indicate if a process is ready to enter the critical section.

Slide 17: \rightarrow It will stuck in the while statement until the flag becomes true. (wait) / (busy wait)
the expected output $\rightarrow x = 100$

In modern architecture there is reordering to memory accesses within the same thread (ex: inside 17 \rightarrow Print x before the while statement)
 \rightarrow

so the final expected output will be ($x=0$) if the execution started from Thread 1.

Example :- initially $\text{Flag}[0] = F$



The result of reordering \Rightarrow invalid situation

Hardware supported solution (Memory Barrier)


write back \rightarrow strongly ordered

write through \rightarrow weakly ordered.

• Mutex locks $\left\{ \begin{array}{l} 1 / \text{True} \\ 0 / \text{False} \end{array} \right\} \rightarrow \text{boolean variable}$

• All critical sections that access the same shared resource, have the same locks.

example:-



P_1, P_3 have the same mutex locks because they access the same shared resource

ex:- (mutex-lock m_1 ;

• When the lock equals 1 (ex:- $m_1 = 1$), this means that this critical section is available (shared resource is available).

• acquire lock $\rightarrow m_1 = 0 / \text{False}$

• release lock $\rightarrow m_1 = 1 / \text{True}$

acquire(m_1) {

while (! m_1); \rightarrow (when $m_1 = F$, busy wait).

$m_1 = \text{False}$; \rightarrow The lock is now unavailable

(critical section)

\rightarrow


```
release (m1) {
```

$m_1 = \text{True};$ → The lock is available now.

```
}
```

- Busy wait :- technique in which a process repeatedly checks to see if a condition is true. (CPU is busy without any progress or process).

So mutex locks cause busy wait situation.

- Semaphore → used to avoid busy wait.

↳ use wait(), signal() instead of acquire, release to avoid busy wait.

→ semaphore variable

```
• wait(s) {
```

while ($s \leq 0$); → the critical section is not available

wait queue ← // wait

until it becomes 1. $s--;$ (critical section) (if $s = 1$ it will be 0 to protect the CS)


```
}
```

```
• signal(s) {
```

$s++$ → s becomes 1 (from wait queue to ready queue)

```
}
```


- Counting semaphore: There are multiple instances of the shared resource (initial value for $S \neq 1$). The initial value for $S = \# \text{ of available instances from the shared resource}$.



- How the busy wait problem was solved by using semaphore?
 - with each semaphore there is an associated waiting queue

typedef struct {


int value; \rightarrow value of the semaphore

struct process * list; \rightarrow array of pointers, point the process.

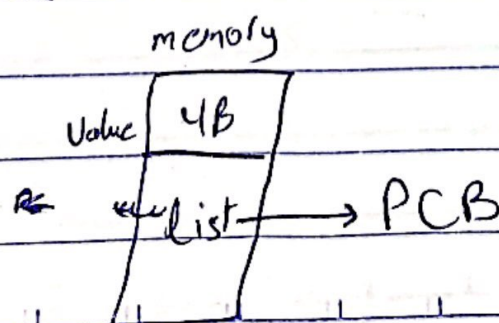
} semaphore

hs(struct name)

- Wait queue contains the PCBs of waiting processes \rightarrow Pointers to




Semaphore S_i



slides (34-44)

lecture 81



- Semaphore → when it is declared as variable not struct → it won't solve the busy wait problem

slide 34.

```
wait(s) {
    while(s <= 0);
```

```
    // busy wait
    (s--); }
may cause race condition.
    signal(s) {
    (s++);
    }
```

→ So, (s-- & s++) should be atomic instructions
(wait, signal) → atomic operations.

Struct :- is a composite data type declaration that defines a physically grouped list of variables.

DATE

S M T W T F S

Slide 40

assume initially $S=1$, $Value=1$ ~~Process execution~~ $P_1 \rightarrow \textcircled{1} Value = 0 \rightarrow \textcircled{2} CS$ $P_2 \rightarrow \textcircled{3} Value = -1 \rightarrow \textcircled{4} \text{add the process to the wait list} \rightarrow \textcircled{5} \text{block}$ $P_3 \rightarrow Value = -2 \rightarrow \text{wait queue} \rightarrow \text{block (delete from ready queue)}$

$\# \text{ of processes in the wait queue} = |Value|$
 $= 2$

 $P_1 \rightarrow \textcircled{1} \text{signal} \rightarrow \textcircled{2} Value = -1$ $P_2 \rightarrow \textcircled{3} \text{wakeup} \rightarrow \textcircled{4} \text{ready queue}$

• Condition Variables: statements only (don't have values)

↳ data type: Condition → ex: Condition X;

• only two operations are allowed to execute on Condition Variables

↳ 1) $x.wait()$ → suspended (wait) until another process
condition variable execute $x.signal()$

2) $x.signal()$ → resumes only one process from processes that invoked $x.wait()$

ex: P_2 P_1 → P_2 before P_1
 $(x.wait())$ $(x.signal())$ ↳ Then P_2 will wait until
 P_1 executes $x.signal()$, then
 P_2 will return to the ready queue.

• every $(x.signal())$ returns only one process from waiting queue to ready queue.

ex: $P_1: y.wait()$

$P_2: y.wait()$ How many processes are in the waiting queue

$P_3: y.wait()$ after executing $y.signal()$? 2 processes

$P_4: y.signal()$ How many processes are in the waiting queue?

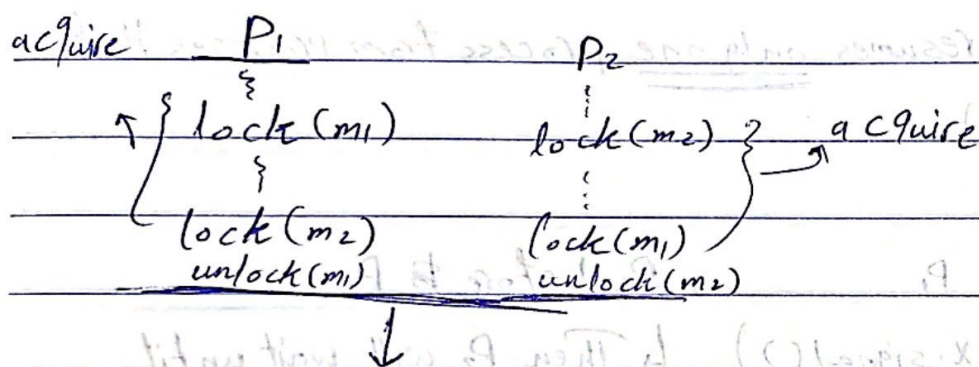
$P_5: x.signal();$ 2 processes.

• each condition variable has its own wait queue.

• **Deadlock** :- situation in which the process have to wait for long time (unlimited), while trying to acquire a synchronization tool.

ex:- critical section 1 \rightarrow with lock (m_1)

" " 2 \rightarrow with lock (m_2)



P_1 will not continue to execute after it acquired lock (m_1) because lock (m_2) is already acquired by P_2 , so P_1 will wait until P_2 release the lock. \rightarrow Here P_1 will wait indefinitely because P_2 needs to acquire lock (m_1) to release lock (m_2) which is already acquired by $P_1 \rightarrow$ so, P_1 and P_2 will wait indefinitely.

Deadlock, Liveness Failure.

• **Liveness Failure** another example \rightarrow infinite loop.

slides (1-6)

chapter 7

synchronization examples

Synchronization Problems → 1) bounded-buffer
2) readers-writers.

• Bounded-buffer → if there is at least one writer process at the buffer that will modify the original stored value, other writer and reader processes are not allowed to use the buffer. → using synchron. tools. → we need to create semaphore mutex (value = 1), and count semaphore called semaphore full (value = 0) which represents the number of full buffers. Also, semaphore empty which represents number of empty buffers (value = # of buffers).

slide (5) →

• Assumption → the producer will make the buffer full, and the consumer will empty it.

• when mutex = 1 → the process can enter the critical section.

↳ after executing (~~mutex~~ wait(mutex), mutex becomes zero.

↳ then signal(mutex) returns the value of the mutex to ~~become~~ = 1

↳ signal(full) → increment.



Bounded Buffer Problem (Cont.)

- The structure of the producer process

```
while (true) {  
    ...  
    /* produce an item in next_produced */  
    ...  
    wait(empty); → DECREMENT THE VALUE OF EMPTY  
                  WAIT ONLY WHEN EMPTY = 0 ( WHEN ALL BUFFERS ARE FULL )  
    wait(mutex); → THE PRODUCER WILL ENTER THE CRITICAL SECTION ( WRITING ON THE BUFFER ).  
    ...  
    /* add next produced to the buffer */  
    ...  
    signal(mutex);  MUTEX=1 AGAIN  
    signal(full);   INCREMENT ( FULL++)  
}
```





Bounded Buffer Problem (Cont.)

- The structure of the consumer process

```
while (true) {  
    wait(full);  DECREMENT (FULL = N-1)  
    wait(mutex); MUTEX BECOMES ZERO  
  
    ...  
    /* remove an item from buffer to next_consumed */  
    ...  
    signal(mutex); MUTEX = 1  
    signal(empty); INCREMENT EMPTY ++  
  
    ...  
    /* consume the item in next consumed */  
    ...  
}
```



~~slides (7-15)~~ slides (7-15)

• Readers - Writers Problem → only one single writer can access the shared data at the same time. (without any readers or writers)

• First reader writer ^{variation} → if we have ^{at least} one reader in the shared area, all incoming readers are allowed to use the shared area. → which cause starvation (writers)

• second reader writer variation → once a writer is ready, that writer performs its write as soon as possible (no new readers may start reading). even if there are readers in the shared area. → which cause starvation (readers)

slide 8)

read_count → represents # of readers in the shared area → if read_count ≥ 1 new coming readers can use the shared area without acquiring the lock. only the first reader acquire the lock to prevent the writers only from using the shared area.

• The last reader leaving the shared area release the lock to when (read_count = 0)

• read_count → shared variable (may cause race condition)

- difference between named and unnamed semaphores → declaration

Remember :- Condition variables are user-made objects that cannot be shared between processes → used to make sure that the shared resource is accessed by one process at a time → so mutex or binary semaphore is used with the condition variables.

Slide 18)

a, b → shared resources

~~pthread_mutex_unlock~~

pthread_cond_wait(&cond_var, &mutex);

↳ to protect a, b
(release the lock)


Chapter 8 | slides (1-6)

Remember, Deadlocks :- situation in which the process have to wait for long time, while trying to acquire a synchronization tool.

• CPU utilization in dead lock situation = 0%.


• Deadlock characterization → formal definition.

- Resource - Allocation Graph. \rightarrow group of nodes (vertices) and directed edges.



- request edge (from processes or threads to a resource)
- assignment edge (from a dot in the resource to a thread or process)
- if there is a cycle in the graph (The starting point is the same as the end point), There may be a deadlock situation or not, but if there isn't any cycle in the graph then there is no any deadlock situation.

slide 9




cycle, is there any dead lock?

- T3 acquires R3, and needs R2
- T2 acquires R2, and needs R3

Deadlock situation.

slide 9




- T_1 acquires R_2 , and needs R_1
- T_2 acquires R_1 , and needs R_3
- T_2 acquires R_2 , and needs R_3
- T_3 acquires R_3 , and needs R_2

↓

T_1, T_2, T_3 Deadlock

- If there is only one instance for resources that are involved within the cycle \rightarrow 100% there is a deadlock situation for all threads.

slide 10



- T_1 acquires R_2 , and needs R_1
- T_3 acquires R_1 , and needs R_2
- if T_4 release $R_2 \rightarrow$ no deadlock (break the cycle)

- Deadlock prevention \rightarrow By preventing the ~~occurrence~~ occurrence of the 4 deadlock ~~at~~ conditions at the same time.

- Deadlock avoidance :- when the kernel has previous knowledge about future requests of resources.

Chapter 9 | Slides (4 - 6)

- CPU can only access Main memory and registers
- read, write operations take 1 CPU cycle or less. (registers)
- read, write from/on the main memory is slower than accessing registers. (more cycles are needed), so cache is used and sits between main memory and CPU registers.

- base address \rightarrow The beginning of the ~~space~~ address space
- base + limit \rightarrow where the limit is the size of the address space, so (limit + ~~address~~_{base}) the end of the address space.

- The process itself can't change the base and the limit because they are ~~are~~ privileged.

base \leq logical address space of a process $<$ base + limit

Good Luck!

