
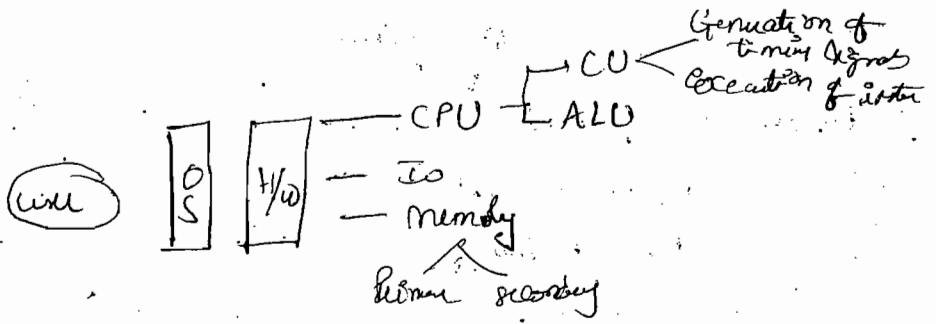


Operating systems:

- Process management
 - Process concepts
 - CPU scheduling
 - Synchronization and Concurrency
 - Deadlocks
 - Threads
 - memory management
 - Concepts
 - Techniques
 - virtual memory.
 - File system and device management
 - Interfaces
 - Implementation issues
 - Protection mechanisms
- what is an OS?
- Interface b/w user and h/w.
 - Control program
 - Resource manager (allocator)
 - Set of utilities to simplify application development.



Goals of an OS:

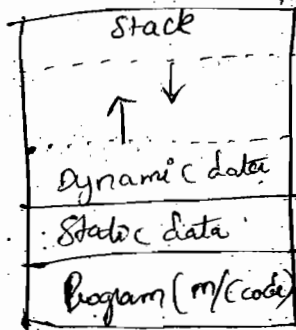
- user friendly (convenience)
- Reliable
- Efficiency (in utilization of resources)
- Portability (Shd be able to load and use across different environments)
- Scalability (ability to evolve)

Functions of OS:-

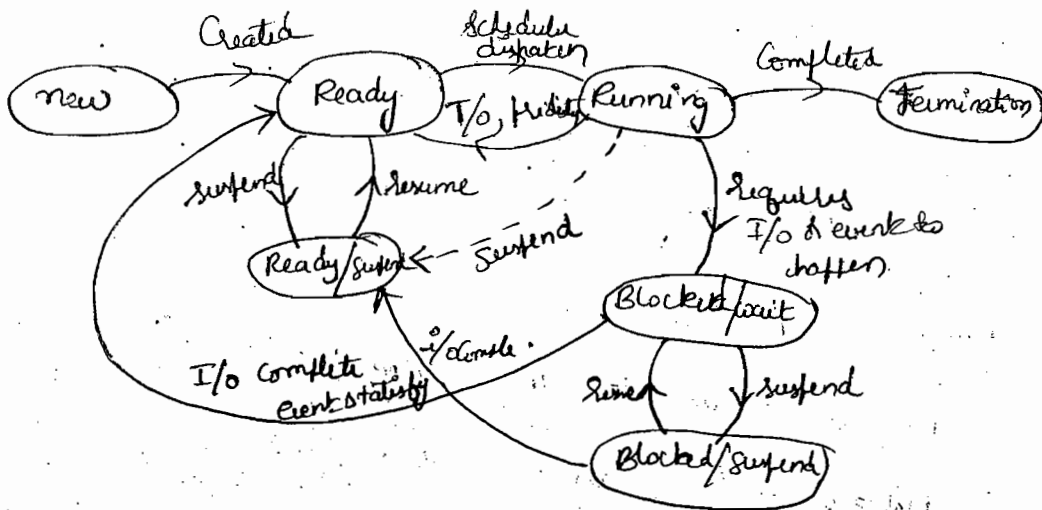
- Security
- CPU scheduling
- mem management
- Deadlock handling.

Process Concepts

- Program under execution (exists in CPU & I/O mem).
- Programs not under execution (not using CPU) are also called processes, when they are waiting in memory (RAM) for CPU, i.e. using a program is using a resource.
- The resources allocated to the process are CPU time, memory, files and I/O devices.



P.S.D. Process state transition diagram:



Process control block :-

Each process is represented in the OS by a PCB also called task control block.

Pointer	Process state
	Process number
	Program Counter
	Registers
	memory limits
	list of open files

PCB Contains many fields of information associated with a specific process, including the following -

Process state: The state may be new, ready, running etc.

Program's Counter: The address of the next instruction to be executed for this process.

CPU registers: The registers vary in number and type, depending on the computer architecture.

Along with the program counter, this state information must be saved when an interrupt occurs, to allow the process to be continued correctly afterward.

CPU-scheduling information: This information includes a process priority, pointer to scheduling queues, and any other scheduling parameters.

memory-management information: This includes values of base and limit registers, page tables, & segment tables depending on the memory system used by the OS.

Accounting information: This i/f includes amount of CPU and real time used, time limits, account numbers, job & process numbers and so on.

I/O status information: This i/f includes the list of I/O devices allocated to this process, a list of open files and so on.

→ Consider a computer with n -CPU's and m processes. ($n \geq 1$ & $m > n$) what is the lower limit and UL on the no. of processes which can be in the ready state, running state and blocked state.

	lower bound	upper bound
Ready	0	m
Running	0	n
Blocked	0	m

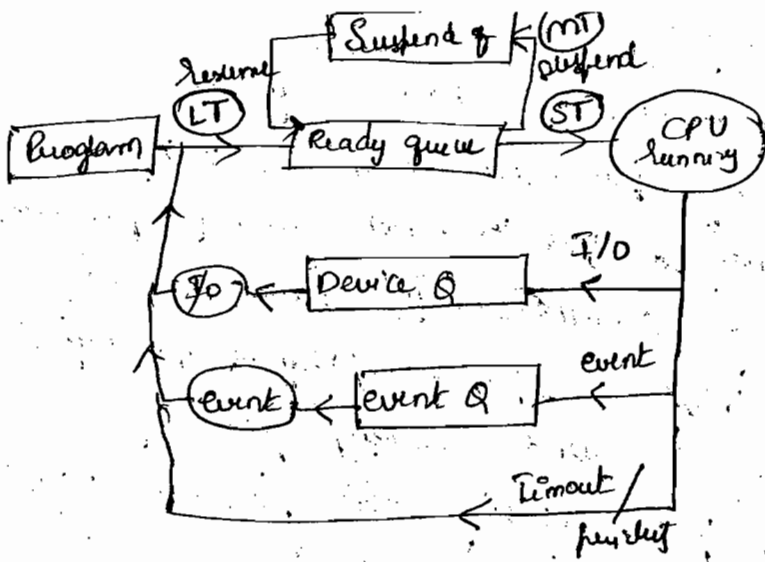
Multiprogramming :-

A uniprocessor system can have only one ~~two~~ running process at all times. But, to ^{make} ~~give~~ the user a feel as that many processes are running simultaneously, the processor switches b/w the processes very frequently. This concept is called multiprogramming. The process of switching the processor b/w processes is called time-sharing.

Multi-processing :-

A multi processor system can have more than one process running at all ~~times~~ given time. The process of running many processes ~~in~~ simultaneously is called multiprocessing.

Degree of multiprogramming :- The number of processes in the memory is called degree of multiprogramming.



Queuing diagram representation of process scheduling.

Long-term scheduler: The long-term scheduler or job scheduler, selects processes from secondary memory (disk) and schedules loads them into memory for execution.

Short-term scheduler or CPU scheduler: select among the processes that are ready to execute and allocates CPU to one of them.

MT scheduler: Removes processes from memory

and thus reduces degree of multi-programming. At some ~~time~~ later time, the process will be reintroduced into memory and its execution can be continued where it left off. This is called swapping.

Processes are generally of two types:

- a) I/O-bound process: It spends more of its time doing I/O than it spends doing Computations.
- b) CPU-bound process: It generates I/O requests infrequently, using more of its time doing I/O.

Conclusion

- The long-term scheduler should select a good process mix of I/O-bound and CPU-bound processes.
- If all processes are I/O bound, then CPU will be idle most of time and ready queues will be almost empty, which leaves short-term scheduler with little work to do.
- If all processes are CPU bound, then I/O waiting queues will be almost empty, devices will go around.

Context switching :-

Switching the CPU to another process & requires saving the state of the old process and loading the saved state for the new process.

This task is known as context switch.

The context of a process can be found in PCB.

Dispatcher :-

It is a part of short term scheduler.

When a ST scheduler, decides to ~~switch~~

move a process from running to ready and

another process from ready to running, context switching is needed. All the code needed

for this process will be in a module called

dispatcher.

The time spent on this process of context switching is called switching time or Dispatch latency.

It is always desirable to reduce CST as much as possible.

Context switching time & Contents of size of PCB

CPU scheduling techniques & short term scheduler & CPU scheduler

Functions of CPU scheduler:

making a decision on which process to schedule next.

Goals of CPU scheduler:

- maximize CPU utilization
- Be fair to all processes (give chance to all)
- Increase the throughput (number of jobs & processes completed per unit of time)
- minimize waiting time of processes, TAT, and Response time

Process times: t_{arr}

Arrival time (AT): AT is submission time in the point of time instance when the process is submitted.

Burst time (BT): BT & service time is the period of time needed by a process to complete its execution.

waiting time (WT): Time spent by the period of time spent by the process waiting for CPU in the ready queue is called waiting time (WT).

Completion time:

The point of time when the process completes its execution is called completion time.

Turn around time (TAT): The total period of time taken by a process from its arrival till completion is called

turn around time.

Deadline: The point of time by which the process need to complete is called its deadline.

of Response time: The period of time from the point of submission of a request to the point at which we get first response is called response time.

→ let us assume we have 'n' processes

P_1, P_2, \dots, P_n

→ let AT of $P_i \rightarrow A_i$

BT of $P_i \rightarrow B_i$

CT of $P_i \rightarrow C_i$

Deadline of $P_i \rightarrow D_i$

$$TAT(P_i) = C_i - A_i$$

$$\text{weighted TAT}(P_i) = \frac{C_i - A_i}{B_i}$$

$$\text{Avg TAT} = \frac{1}{n} \sum_{i=1}^n C_i - A_i$$

$$\text{Avg wt. TAT} = \frac{1}{n} \sum_{i=1}^n \frac{C_i - A_i}{B_i}$$

$$\text{weighting time}(P_i) = TAT - B_i$$

$$= C_i - A_i - B_i$$

Schedule length (L) = $\max(C_i) - \min(A_i)$
(Time to complete all n processes).

→ CPU scheduling techniques :

These techniques can be classified as

(i) Preemptive: If a process is deallocated a resource ~~temporarily~~ ^{temporarily} before its completion of the resource usage, then such a technique is called preemptive technique.

In scheduling context, the resource is CPU and techniques are scheduling techniques.

(ii) Non-preemptive: If a process is allowed to ^{allocated} use the resource only after it finishes its usage the resource usage, then such technique is called non-preemptive technique.

→ If there are 'n' processes, then how possible schedules are possible with a non-preemptive scheduling technique.

n!

→ Through put (T) = $\frac{\text{no of processes}}{\text{Total time}}$

→ Deadline overrun (P_i) = $C_i - D_i$

$C_i - D_i = +ve$ - Crossed deadline

$= -ve$ - before time

$= 0$ - on time

FCFS: It is based on AT, It is non-preemptive technique.

P.NO	AT	BT
1	0	4
2	1	5
3	2	6
4	3	7
5	4	3

4	5	6	7	3	
0	4	9	15	22	25


P. NO	AT	BT	CT	TAT	WT
1	0	4	4	4	0
2	1	5	9	8	3
3	2	6	15	13	7
4	3	7	22	19	12
5	4	3	25	21	18

$$L = 25, \quad M = \frac{5}{25} = \frac{1}{5}$$


$$\text{Avg TAT} = \frac{65}{5} = 13 \quad \text{Avg WT} = \frac{40}{5} = 8$$

Ex

Ped	AT	BT	CT	TAT	WT
1	4	4			
2	5	5			
3	2	2			
4	1	1			
5	7	2			
6	0	3			



Pid	AT	BT	CT	PT	WT
1	2	3	5	3	0
2	8	4	12	4	0
3	15	2	17	2	0
4	18	1	19	1	0
5	25	3	28	3	0



$$L = 28$$


$$\mu = \frac{5}{28}$$

There is a Convoy effect, as all the other processes wait for the one big process to get off the CPU. To overcome this, we go for SJF & SJN.


Shohitak Jor Park (SJP) & Shohitak-focus week (SPN)

This is based on BT. mode: ~~absorptive~~
mode: non-fluoride.

PNO	AT	BT
1	0	7
2	0	6
3	2	1
4	3	2
5	4	3
6	5	1



PNO	AT	BT
1	5	1
✓ 2	4	2
✓ 3	3	1
4	1	4
✓ 5	2	3
✓ 6	1	2




Shortest Remaining time first (SRTF)

It is preemptive version of SJF


Critical: BT, mode: P.

Ex:

PNO	AT	BT
1	0	7
2	1	5
✓3	2	3
✓4	3	1
✓5	4	2
✓6	5	3




NP-SJF



→

PNO	AT	BT
✓1	4	8
✓2	6	10
✗3	5	7
✗4	2	1
✗5	7	3
✗6	3	5



- SJF (FC and non-FC) favours shorter jobs and causes starvation to longer jobs.
- It gives high throughput at any point of time compared to any scheduling technique. So, it is optimal scheduling algorithm.
- ~~It~~ This technique is and minimum average waiting time for a given set of processes.
- Though SJF is optimal, it cannot be implemented at the level of ^{OS} CPU scheduling as we cannot know the ^{exact} CT of a process exactly before running it.
- One approach is to approximate SJF scheduling, by predicting the burst time of a process.
- We expect that the ~~next~~ CPU burst of the next process will be similar in length to the CPU burst of the previous ones.


→ Let t_n be the length of the n^{th} CPU burst
and let \hat{T}_{n+1} be our predicted value of the
next CPU burst. Then, for α , $0 \leq \alpha \leq 1$,
define

$$\hat{T}_{n+1} = \alpha t_n + (1-\alpha)\hat{T}_n.$$

- This formula defines an exponential average.
- t_n contains more recent information and
 \hat{T}_n stores the past history.
- The parameter α controls the relative weight of
recent and past history in our prediction.
- If $\alpha = 0$, $\hat{T}_{n+1} = \hat{T}_n$ and so recent
information has no effect.
- If $\alpha = 1$, then $\hat{T}_{n+1} = t_n$ and only the
most recent CPU burst matters.
- If $\alpha = 1/2$, then $\hat{T}_{n+1} = \frac{1}{2}(t_n + \hat{T}_n)$, i.e.
equal weightage is given to both a history's

Priority scheduling: (Based on priority).

→ The SJF algo is a special case of the general priority-scheduling algorithm.



→ It can operate in preemptive & non-preemptive mode.

~~NP priority~~ or


NP - priority scheduling:

It works exactly out of all ~~other~~ processes submitted this technique will select a process with highest priority and executes it first.


→ highest number may indicate highest priority & lowest number may indicate highest priority.

→ here let us assume highest number ~~is~~ ~~is~~ indicates highest priority.

<u>Priority</u>	<u>P.NO</u>	<u>AT</u>	<u>BT</u>
4	1	0	4
5	2	1	5
6	3	2	9
3	4	4	2
7	5	5	7




<u>Priority</u>	<u>P.NO</u>	<u>AT</u>	<u>BT</u>
✓ 4	1	6	4
✓ 3	2	10	5
✓ 2	3	8	6
✓ 2	4	4	2
✓ 1	5	2	3
1	6	12	1



b) Preemptive - priority:

This technique continues execution of a process till the arrival of next process, and, if the next process has higher priority, then preempt the current process.

	<u>Pi</u>	<u>P.NO</u>	<u>AT</u>	<u>BT</u>
X	4	1	0	4 3
X	5	2	1	5 4
X	6	3	2	6 4
✓	2	4	3	8
X	8	5	4	2 1
X	7	6	5	1



Round-Robin Scheduling:

The RR scheduling algo is designed especially for time sharing systems.

- It is ^{useful for} multi-programmed, multi-user and interactive systems like *unix*.

→ It is similar to FCFS scheduling, but preemption is added.

→ Criteria - AT + time quantum


→ Δ small / q

→ Time quantum (or) time slice is small unit of time.

→ To implement RR scheduling, we keep the ready queue as a FIFO queue of processes. New processes are added to the tail of the ready queue.

PNO	AT	BT
1	0	4 20
2	1	8 31
3	2	3 1
4	3	1
5	4	6 12
6	5	8 31

~~P₁ P₂ P₃ P₄ P₅ P₆ P₃ P₅~~
~~P₅ P₂ P₅ P₅ P₆~~




>

PNO	AT	BT
1	1	4/10
2	3	2/0
3	5	4/10
4	6	8/20
5	8	6/30
6	9	7/4

TQ = 3

~~P1~~, ~~P2~~, ~~P1~~, ~~P3~~, ~~P4~~, ~~P5~~, ~~P6~~, ~~P3~~, ~~P4~~, ~~P5~~




>

P.NO	AT	BT
1	6	4
2	2	5/3
3	3	6/4
4	1	8/6
5	4	10
6	5	3

TQ = 2

OUT

~~P4~~, ~~P2~~, ~~P3~~, ~~P4~~, ~~P5~~, ~~P6~~, ~~P2~~



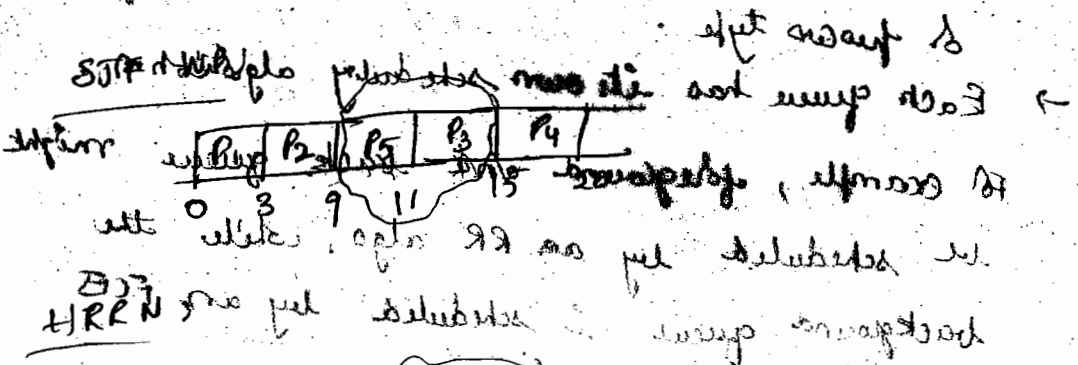
Highest Response Ratio rule: (HRRN)

RT or Response Ratio = $\frac{W+S}{S}$


- The scheduling is based on response ratio.
- It is a non-preemptive scheduling technique.
- Compute the RR of each ^{remaining} process after execution of every process.

Example

Process	AT	BT	WT
P1	0	3	2
P2	2	5	3
P3	6	8	5



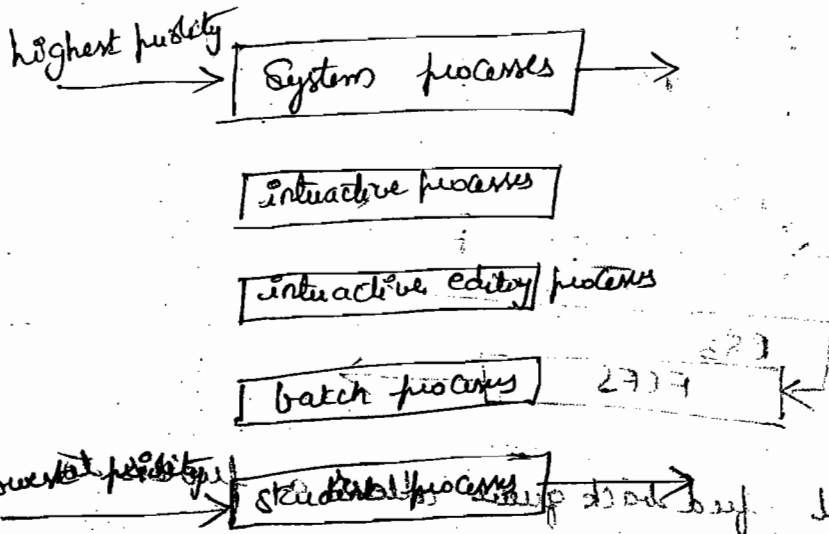
HRRN



Multilevel queue scheduling:

- Processes can be classified into different groups.
 - A common way to classify processes is division is made b/w foreground (or interactive) processes and background (or batch) processes.
 - A multilevel queue-scheduling algorithm partitions the ready queue into several separate queues.
 - The processes are permanently assigned to one queue, generally based on some property of the process, such as priority, process type, or process type.
 - Each queue has its own scheduling algorithm.
- For example, foreground and background queue might be scheduled by an RR algo, while the background queue is scheduled by an FCFS algo.

→ In addition, there must be scheduling among the queues, which is commonly implemented as fixed-priority preemptive scheduling.



Address ~~...~~ needs to be joined to a queue. When a process comes, it is scheduled based on its type and subpriority. At a certain queue, it can be operated in a different queue.

Disadvantage: If a queue of high priority is empty, then only lower priority processes can be operated. So processes in lower priority suffer from starvation.

and then it will be sent back to PO_2 , it still needs more time. Here

→ on PO_3 , the process will get executed till it completes.

→ The number of queues used and the scheduling algorithms within each queue are system dependent.

→ The scheduling algo gives highest priority to any process that is CPU bound & multi-processor.

→ Next priority will be to process with 8-24.

Complex cases

→ Long processes automatically sent to PO_3 and

gets served by FCFS order (schedules threads).

→ Solaris 2 uses multi-level queue scheduling.

contains the following queues

- (i) Real time
- (ii) system
- (iii) Time sharing and interactive


Process synchronization:

IPC and process synchronization:


- The concurrent processes executing in the OS may be either independent or cooperating processes.
- A process is independent if it cannot affect or be affected by the other processes executing in the system.
- A process is cooperating if it can affect or be affected by the other processes executing in the system.

The reasons for using cooperating processes are:

(i) Information sharing:



(ii) Computation speedup:



modularity: like many ~~part~~


Convenience:

→ These cooperating processes can communicate with each other via interprocess communication facilities like shared memory and message passing.

RAC Conditions:

Let $buff$ is a shared variable in process
Process 1: $a++$; process 2: $a--$

Process 1: $1. \text{mov } a, \text{reg1};$
 $2. \text{Add reg1, 1};$
 $3. \text{move reg1, a};$



$a - -$

$1. \text{move } a, \text{Reg1};$

~~2. Add~~

$2. \text{Sub reg1, 1};$

$2. \text{move reg1, a};$

reg: 4



P_1 : 1, 2 P_2 : 1, 2

P_1 : 1, 2 P_2 : 1, 2

P_1 : 1, 2	P_2 : 1, 2	P_1 : 3	P_2 : 3
--------------	--------------	-----------	-----------

P_1 : 1, 2	P_2 : 1, 2	P_1 : 3	P_2 : 3
--------------	--------------	-----------	-----------


Situations where two or more processes are reading or writing some shared data and the final result depends on who runs precisely first, are called race conditions.

Example: Print spooler:

Ex: Print spooler:

- when a process wants to print a file, it enters the file name in a special spooler directory.
- Another process, the printer daemon, periodically checks to see if there are any files to be printed, and if there are, it prints them and then removes their names.

Spades directory



P_1, P_2 both wants to queue a file for printing.

if P_1 has some data to print and P_2 is waiting for its turn, then P_1 will print first.

* next free slot = in ; will hold

* spades [~~in~~ next free slot] = abc

* $in++$; $in = 3$; $in = 4$; $in = 5$; $in = 6$

when a process wants to print a file, it will check the spades directory.

$P_1: x$	$P_2: x$	$P_1: y$	$P_2: y$	$P_1: z$	$P_2: z$
----------	----------	----------	----------	----------	----------

if a process wants to print a file, it will check the spades directory. if the slot is free, it will print. if the slot is occupied, it will wait.

Critical region of critical section!

→ The key to avoid race conditions is to find some way to prohibit more than one process from reading and writing to shared data.

→ In other words, we need what we need is mutual exclusion, that is, some way of making sure that if one process is using a shared variable of file, the other process is excluded from doing that task.

→ A process will be at such a point where it will be busy doing internal computations that do not lead to race conditions and remaining time, or it accesses the shared memory.

→ The part of that program where the critical memory is accessed is called the critical region or critical section.

them just before leaving it.

→ with interrupts disabled, no clock interrupts can occur.


→ The CPU is only switched from process to process as a result of clock or other interrupts and with interrupts turned off, the CPU will not be switched to another process.

ME ✓, no assumptions X, process ✓
BWX 17

Bounded waiting: There exists a bound on the number of times that other processes are allowed to enter their critical sections after a process has made a request to enter its critical section and a type that requests are granted.

Disadvantage: The control over interrupts should never be given to users (what if an user disables an interrupt and never enables it).

Lock variables & (S/w solution), 2 process, b/w.



```

if (lock == 0)
while (lock == 1);
lock = 1;
CS
lock = 0;
  
```

ME X, BWX, no AV, progress ✓

Stack allocation ✓ Sw, 2 process, b/w; 3m

P0

P1

go bound to this with: mutex between
 while (lock == 1) {
 CS
 NCS
 }
 while (lock == 0) {
 CS
 NCS
 }
 lock

- Conditional testing a variable which is being updated
- value appears in called by at language level
- busy waiting wastes CPU time

mE ✓, P X, No assum ✓, BQ ✓

The TSL instruction: How so and sw solution

Test & Lock

TSL RX, LOCK

- ~~It~~ The above command reads the contents of the memory word 'lock' into the register RX and then stores a nonzero value at the memory address 'lock'.
- The operations of reading the word and storing into it are guaranteed to be indivisible.
- No other process can access the memory word until the instruction is finished.
- The CPU executing the TSL instruction locks the memory bus to prohibit other CPUs from accessing memory until it is done.

enter-region;
Entry:

(80,000)

```
TSL REGISTER, LOCK // copy lock to register and it  
CMP REGISTER, #0 // was lock zero?  
JNE enter-region // if not equal non zero, lock nonzero  
RET // return to caller; critical region entered.
```

leave-region:
also program with
subroutine
MOVE LOCK, #0
first RET here // return to caller
distribution of copies distributed up to done
Peterson's solution
at Conferences
and learning
lock variables
processes are
of lock program
program process

```
lock to);  
# define FALSE 0  
# define TRUE 1  
# define N 2
```

```
int turn;  
int interested [N];
```

```
void enter-region (int process) /* process is 0 or 1 */
```

```
{  
1 int other; /* number of the other process */
```

```
2 other = 1 - process; /* the opposite of process */
```

```
3 interested [process] = TRUE; /* show that you  
are interested */
```

```
4 turn = process; /* set flag */
```

```
5 while (turn == process && interested [other] == TRUE)
```

```
}
```

```
void leave-region (int process) /* process: who's  
leaving */
```

```
{  
interested [process] = FALSE; /* indicate
```



```
}
```

departure from critical
region */




- ~~First~~. A process before entering critical section will call `enter_region` with its own process number, 0 or 1, as parameter.
- After it has finished with the shared variables, the process calls `leave_region` to indicate that it is done and to allow the other process to enter, if it so desires.

Analysis:

- Initially neither process is in its critical region. Now process 0 calls `enter_region`. It indicates its interest by setting its array element and sets `turn` to '0'. Since process 1 is not interested, `enter_region` returns immediately.
- If process 1 now calls `enter_region`, it will have the critical interested [0] goes to FALSE, an event that only happens when process '0' calls "leave_region" to exit the critical region.
- Now consider the case that both processes call `enter_region` at almost simultaneously.



Both will store their process number in 'turn'. whichever store is done last is the one that counts; the first one is overwritten and lost.



Suppose that process 1 stores last, so turn is 1. when both processes come to while statement, process '0' executes it zero times and enters its critical region. process '1' looks and doesn't enter its critical region until process 0 exits its critical region.

ME ✓, Proper ✓, BW ✓, no assertions.

Notes

Both Peterson's and the solution using TSL are correct but both have a defect of requiring busy waiting and priority inversion problem.

Consider a computer with two processes, 'H', with high priority and 'L', with low priority. The scheduling rules are such that 'H' is run whenever it is in ready state. At a certain moment, both L is only in the system and it enters critical region, now 'H' becomes ready to run and enters the ready queue. 'H' now begins busy waiting, but since 'L' is never scheduled and while H is running L never gets the chance to leave its critical region. So H loops forever. This situation is referred to as the priority inversion problem.

Semaphores


A semaphore 'S' is an integer variable, which is accessed only through two standard atomic operations; "wait" and "signal".

- one of the simplest ways to avoid busy waiting is to use a pair "sleep" and "wakeup".
- "sleep" is a system call that causes the caller to block, that is, be suspended until another process wakes it up. The wakeup call has one parameter, the process to be awakened.

~~wait~~ is a generalization of "sleep".
~~signal~~ is a generalization of "wakeup".

```
struct semaphore  
{  
    int value;  
int type; struct process *L;  
}
```

→ Semaphore is an OS resource - Executed in OS Kernel atomically.



Counting Semaphore

Down (decr semaphore S)

{
S-value = S-value - 1;

if (S-value < 0)

{

~~Block the process and put it in~~
add this process to
in S.L; block();

}

}

UP (struct semaphore S)

{ S.value = S.value + 1;

if (S.value <= 0)

{ ~~select a blocked process from~~

~~S.L() and wakeup()~~;

} Remove a process P from S.L;

wakeup(P);

}

→ Positive value of 'S' indicates no of successful down operations that can be performed.

→ -ve value indicates no of blocked processes.

→ A process may get blocked only while performing down operation.

Ex: In a certain computation, the value of semaphore S is initialized to S=13. Then

The following operations were performed on the semaphore in the given order. what is the status of it after these

15P, 12V, 10P, 8V, 8P, 7V, 6P, 1V

P - down V - up

-2, 10, 0, 8, 0, 7, 1, 2

→ S=11

10P, 4V, 18P, 6V, 8P, 12V, 3P, 1V

→ Binary Semaphores (mutex)

In general they are used for mutual exclusion hence called (mutex).

It takes only 0 or 1
↓ ↓
busy free

struct B Semaphore

{ Enum value (0, 1);

Queue type t;

BDown (stuck semaphore s)

{ if (s.value == 1)

{ s.value = 0 ;

return ;

}

{ else

{ Block the process ; and add it
to the queue ;

}

}

→ The no of blocked processes is not known in binary semaphores.

UP!
→ when a process performs OP operation and if it finds \emptyset not empty, then $S=0$ only and it wakes up on process.

→ only if no process is waiting blocked, then it makes S to 1.

So for no dead lock, both should have same entry motion i.e

P_0

$P(S)$

$P(T)$

P_0

$P(S)$

$P(T)$

→ we say that a set of processes is in a deadlock state when every process in the set is waiting for an event that can be caused only by another process in the set.

→ Another process related to deadlock is indefinite blocking starvation, a situation where a process waits indefinitely for a resource held by another process. This is a form of starvation where a process is blocked from execution because of a resource held by another process.


another

09

Deadlocks:

⇒ Permanent blocking of set of processes is called Deadlock.

⇒ Primary requirement for set of DL is two or more processes.



⇒ System model consists of a set of processes and resources. The system is a set of processes and resources. Resources are divided into several classes. Each class is characterized by a number of identical instances. Conflicting processes.

⇒ Each resource may have several identical instances.

CPU - 2 instances.

Printer - 5 instances.

Process (running)

request

OS

not granted

cancel

granted

Blocked

Dead lock

run

request

finished

Release

Request and release of semaphore
operations on semaphore
acompilshed through semaphore
work (and signal)

Deadlock characterization.

Necessary conditions.

Resource allocation graph.

Necessary Conditions

→ A deadlock situation can arise if the following four conditions hold simultaneously in a system:

1. Mutual exclusion: There should be at least one resource such that only one process can use it at a time.
2. Hold and wait: A process must be holding at least one resource and waiting to acquire additional resources that are currently being held by other processes.
3. No preemption: Resources cannot be preempted; i.e., a resource can be released only voluntarily by a process holding it, after that process has completed its task.

34. Circular wait: A set $\{P_0, P_1, \dots, P_n\}$ of waiting processes must exist such that P_0 is waiting for a resource that is held by P_1 , P_1 is waiting for a resource that is held by P_2 , ..., P_{n-1} is waiting for a resource that is held by P_n and P_n is waiting for a resource that is held by P_0 .

Resource-allocation graph

→ Deadlocks can be described precisely in terms of directed graph.

System resource-allocation graph

→ The graph consists of a set of vertices V and

set of edges E .


→ The set of vertices V is partitioned into two different types of nodes $P = \{P_1, P_2, \dots, P_n\}$

the set consisting of all the active processes in the system, and $R = \{R_1, R_2, \dots, R_m\}$, the set consisting of all the resource type in the system.




→ A directed edge from process P_i to resource type R_j is denoted by $P_i \rightarrow R_j$; it signifies that process P_i requested an instance of resource type R_j and is currently waiting for that resource.

→ A directed edge from resource type R_j to process P_i is denoted $R_j \rightarrow P_i$; it signifies that an instance of resource type R_j has been allocated to process P_i .

→ A directed edge $P_i \rightarrow R_j$ is called a **request edge**, a directed edge $R_j \rightarrow P_i$ is called an **assignment edge**.



→ Presence of a cycle in a resource-allocation graph is a necessary condition for a deadlock.




→ In summary, if a resource-allocation graph does not have a cycle, then ~~the~~ the system may or may not be in a deadlock state. System is not in a deadlock state.

On the other hand, if there is a cycle, then the system may or may not be in a deadlock state.

Deadlock handling strategies

1. Ignore: we ignore the problem altogether, and pretend that deadlocks never occur in the system.
2. Deadlock prevention: It is a set of methods for ensuring that at least one of necessary conditions cannot hold.
3. Deadlock avoidance: It requires that the operating system be given in advance additional information concerning

→ Consider a system with 'n' processes and a single resource type 'R' with 5 instances. Each process P_i requires 2 copies of resource R to complete its execution. What is the maximum value of 'n' that can exist in system with a deadlock free situation.




$d = 9$ (circled)
 I process = 2 up to an amount
 $(n-1)$ process = 1

$$2 + (n-1) = 5$$


$$n + 1 = 5$$

$$\Rightarrow n = 4$$

2) Consider a system with 3 processes and a single resource type 'R'. Each process P_i requires 2 copies of 'R'. What is the minimum no. of instances of 'R' in ~~the~~ system with a deadlock free situation.



3) n processes each require 2 resources $R=6$.
 max no. of processes = 3 so that no deadlock.



4) $n=4$ and a resource type R with many instances.

$P_1 = 10$

$P_2 = 20$

$P_3 = 8$

$P_4 = 19$

$$\begin{array}{r} 9 \\ 19 \\ 7 \\ 18 \\ \hline 53 \end{array}$$

\rightarrow max no of copies
Cause deadlock.

what is the minimum number of instances of R for a deadlock free situation.

5) n processes and a resource R with each R 10 instances. Let each process requires 1 instance of resource R . What is the max number of processes that can be allowed to avoid a deadlock free situation?

$n=20$
no deadlock at all

Deadlock Prevention

Deadlock prevention:

→ For a deadlock to occur, the four necessary conditions (ME, Hold, No preemption, Circular wait) must hold. By ensuring that at least one of these conditions cannot hold, we can prevent the occurrence of deadlock.

Mutual Exclusion:

we cannot prevent deadlocks by denying

the mutual-exclusion condition: some

resources are intrinsically non-sharable.

Hold and wait

Two protocols can be used:

- a) one protocol, is to ask each process to request and be allocated all its resources before it begins execution.

- a) other protocol is to request a process to request resources only when the process has resources and use them. Before it can request any additional resources, however, it must release all the resources that it is currently allocated.
- b) other protocol is to allow a process to request resources only when the process has none. A process may request some resources and use them. Before it can request any additional resources, however, it must release all the resources that it is currently allocated.

These two protocols have two main disadvantages.

- a) Resource utilization may be low, since many of the resources may be allocated but around for a long time.
- b) Starvation is possible. A process that needs several popular resources may have to wait indefinitely, because at least one of the resources that it needs is always allocated to some other process.

No preemption

To ensure that this condition does not hold,
we can use the following protocol.

If a process is holding some resources and
requests another resource that cannot be
immediately allocated to it, then all
the resources currently being held are
preempted.

Circular wait

Let $R = \{R_1, R_2, \dots, R_m\}$ be the set of
resource types. We assign to each resource
a unique integer number, which
allows

write with pen

body

number

1
2
3

Abstract

form

3
2
1

of
H
F

1
1
1

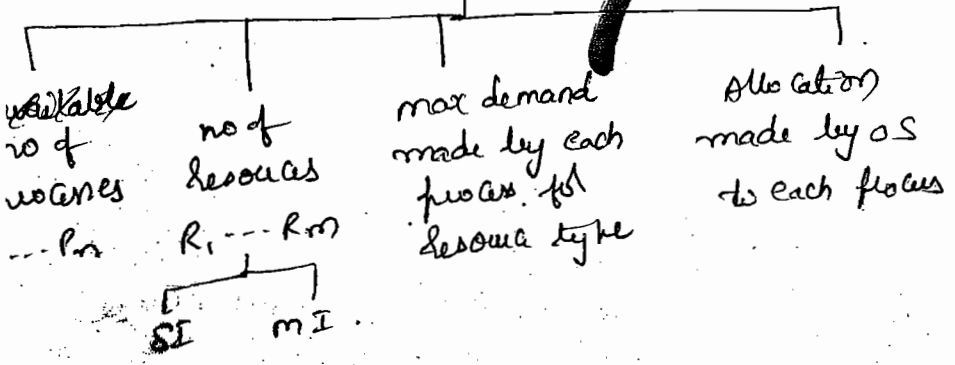
cm 12

Abstract form

Abstract form

Deadlock avoidance

resource allocation state of system \leftarrow available
 need of each process for every resource



Ex: no of processes = 3

no of resources = 1 with 12 tape drives

	<u>max</u>	<u>Allocation</u>	<u>Available</u>	<u>Need</u>
P ₀	10	5	3	5
P ₁	4	2		2
P ₂	9	2		7

$\langle P_1, P_0, P_2 \rangle \rightarrow$ safe sequence

is safe state

Safe state :- A safe state is safe, if the system can allocate resources to each process in some order and still avoid a deadlock.

Safe sequence :- A sequence of processes $\langle P_1, P_2, \dots, P_n \rangle$ is a safe sequence for the current allocation state, if for each P_i , the resources that P_i can still request can be satisfied by the currently available resources plus the resources held by all P_j with $j < i$.

Note :- A system is in a safe state, only if there exists a safe sequence.

Consider the following state of resources :-

no. of processes = 3, max Allocation

no of resources = 1 with 12 instances Available Need

	max	Allocation	Available	Need
P_0	10	5	2	5
P_1	4	2	7	2
P_2	9	3		7

no safe sequence, so it is not in safe state.

→ The algorithm's general algorithm used for deadlock avoidance is known as Banker's algorithm.

→ Data structures used for Banker's algorithm are:

Available: A vector of length 'm' indicates the number of available resources of each type.

Request: A vector of length 'm' indicates the number of resources of each type R_j available.

max: An $n \times m$ matrix defined the maximum demand of each process. If $\text{max}[i, j]$

$= k$, then process P_i may request at most k instances of resource type R_j .

Allocation: An $n \times m$ matrix defines the number of resources of each type currently allocated to each process.

If Allocation $[i, j] = k$, then process P_i is currently allocated k instances of resource type R_j .

Need: An $m \times m$ matrix indicates the remaining resource need for each process.

If $\text{Need}[i, j] = k$, then P_i may need ' k ' more instances of resource type R_j to complete its task.

Note: $\text{Need}[i, j] = \text{max}[i, j] - \text{Allocation}[i, j]$

Problem

Consider a system with 5 processes P_1 and three resource types A, B, C. Resource type A has 10 instances, resource type B has 5 instances, and resource type C has 7 instances. Suppose that, at time T_0 , the following snapshot of the system has been taken:

Allocation	max			Available		
	A	B	C	A	B	C
P ₀	0	1	0	7	5	3
P ₁	2	0	0	3	2	2
P ₂	3	0	2	9	0	2
P ₃	1	1	1	2	2	2
P ₄	0	0	2	4	3	3

the need as follows:

Need of A request

Request: $\langle P_1, P_3, P_4, P_0 \rangle$

Safe sequence: $\langle P_1, P_3, P_4, P_0 \rangle$

P₀ 0 1 0
 P₁ 2 0 0
 P₂ 3 0 2
 P₃ 1 1 1
 P₄ 0 0 2

now: P₁ request (1, 2). Can this request be granted? will this lead to an unsafe state?

To check this, let us assume that we have granted the request, then the system will change to another state.

Process	Allocation			Need			available		
	A	B	C	A	B	C	A	B	C
P ₀	0	1	2	7	4	3	2	3	0
P ₁	3	0	2	0	2	0			
P ₂	3	0	2	6	0	0			
P ₃	2	1	1	1	1	0			
P ₄	0	0	2	4	3	1			

new check if the state is safe or not. We found a safe sequence $\langle P_3, P_4, P_0, P_2 \rangle$ so this state is safe. Now a new request for (0, 2, 0) cannot be granted. Check if this can be granted. It cannot be granted because there is no safe sequence.

P _i	Allocation				Requirement				Available			
	A	B	C	D	A	B	C	D	A	B	C	D
P ₁	1	2	2	1	3	3	2	2	3	1	1	2
P ₂	0	0	3	3	1	2	3	4				
P ₃	1	1	0	0	1	1	5	0				

Need


P _i	A	B	C	D
P ₁	1	1	0	1
P ₂	0	0	0	1
P ₃	0	0	0	0

$\langle P_1, P_2, P_3 \rangle$ is a safe sequence.
 so this is safe state.
 P₂ requests for (0, 1, 0). It can be granted.
 now P₃ requests for (0, 0, 2, 0). It can be granted.
 P₃ asks for (0, 0, 1, 0). Safe.
 (P₁, P₂, P₃)

→	Allocation			max			Available		
	A	B	C	A	B	C	A	B	C
P ₀	0	0	1	2	0	0	1	2	1
P ₁	1	0	0	0	7	5	0	1	5
P ₂	1	3	5	4	2	3	5	6	1
P ₃	0	6	3	2	0	6	5	2	3
P ₄	0	0	1	4	0	6	5	6	4

Need

	A	B	C	D
P ₀	0	0	0	0
P ₁	0	7	5	0
P ₂	1	0	0	2
P ₃	0	0	2	0
P ₄	0	6	4	2




If a request from process P₁ arrives for (0, 4, 3, 0)
 Can the request be granted immediately?


Deadlock detection and recovery:

- If a system does not employ a deadlock-free-
tion or a deadlock avoidance algorithm, then
a deadlock situation may occur.
- In such a situation, we need to detect
the deadlock and recover from it.

Case (i): when all resources are of single instance type

→ In this case, we use a variant of the
resource allocation graph called a wait-for
graph.





There will be a deadlock in the system if and only if the wait-for graph contains a cycle.

Case 2: Several instances of a resource type:

when each resource has several instances, then we have to check for all possible if there is safe sequence or not.

Handwritten text, possibly a signature or name, appearing as "M. J. ...".

Handwritten text, possibly a date or address, appearing as "1911 ...".

Recovery from Deadlock:

a) Process termination:

In this method, we eliminate deadlocks by aborting a process and reclaiming all resources allocated to the terminated process. We can follow two strategies here:

(i) Abort all deadlocked processes:

advan: Simple to implement.
disadvan: whatever process have completed till it aborts, will be lost.

(ii) Abort one process at a time until the deadlock cycle is eliminated.

advan: all the processes need not be aborted. So some computations will be saved.

disadv: Implementation is not simple, we have to abort each process and check if the deadlock is eliminated or not.

Resource Preemption:

To eliminate deadlocks, using resource preemption, we successively preempt some resources from processes and give these resources to other processes until the deadlock cycle is broken.

That ~~if~~ preemption is required to deal with deadlocks, then.

In this method, three issues need to be

addressed:

1) Selecting a victim:
Which resource and which process?

2) Roll back.

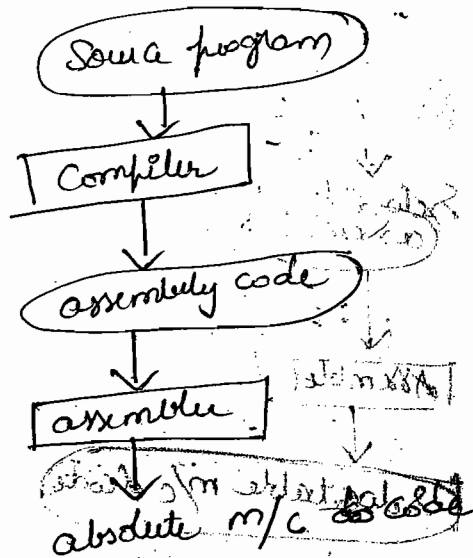
What should be done with that process?

3) Starvation:

How do we ensure that starvation will not occur?

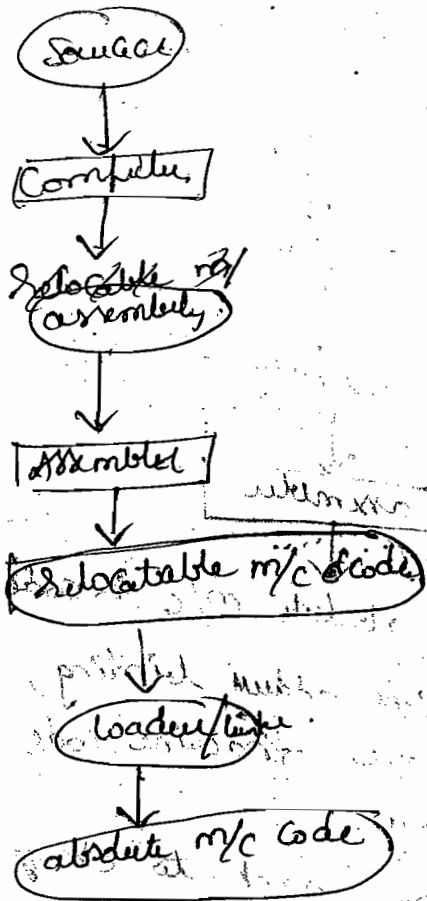
memory management:

Compile time address binding:



For compile-time address binding, we should know where the machine code will be loaded in the memory. Disadv: If we want to change the starting memory location, we will have to recompile the source code.

Load time address binding

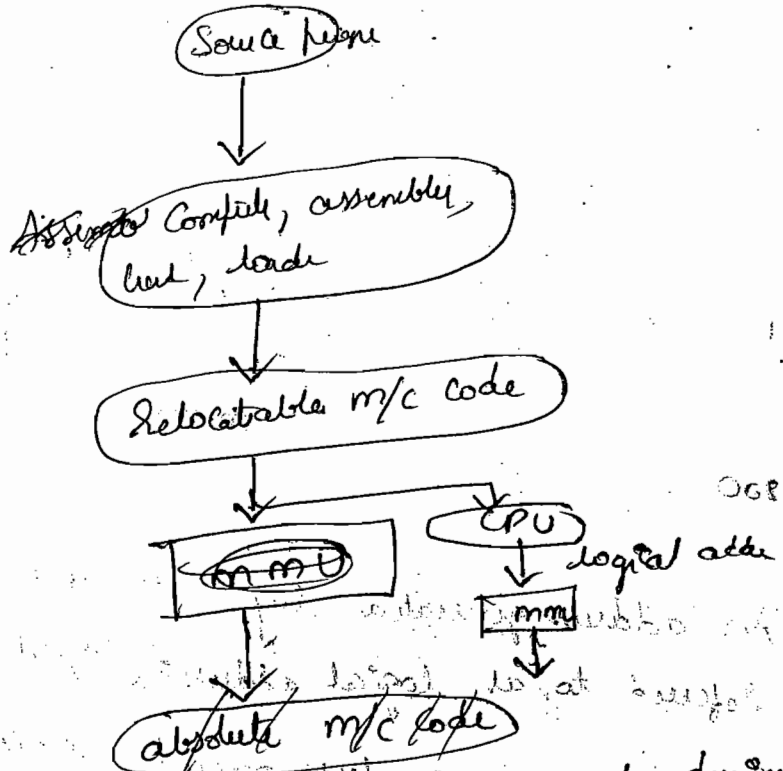


if the compiler stage

in load time binding, if the starting address changes, we need only to reload the user code to incorporate this changed value.

Execution time occurs binding or runtime occurs

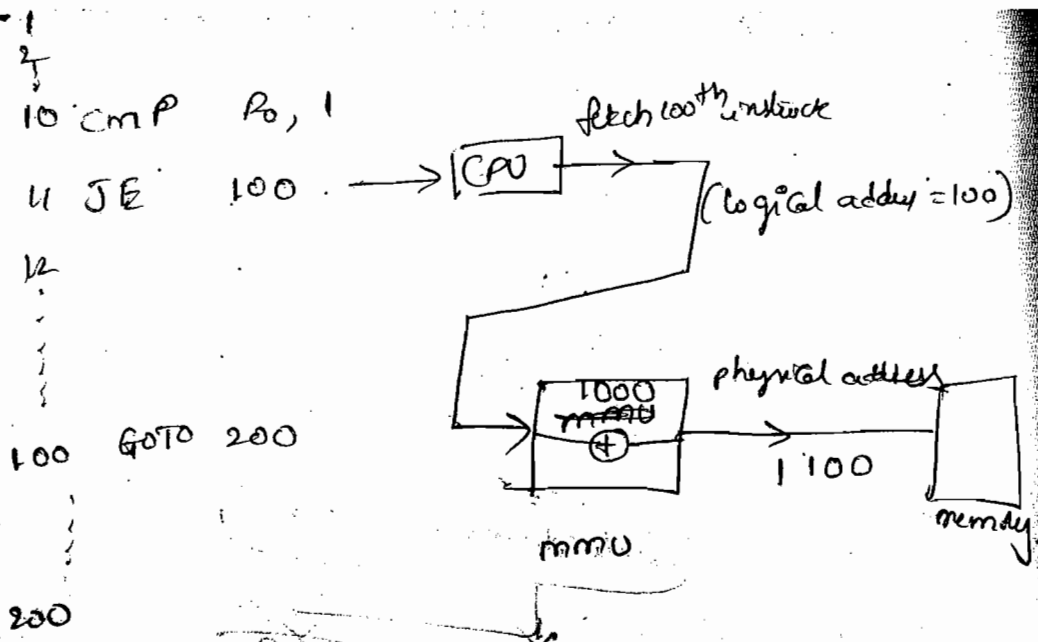
binding :



program
address
buffer

→ If the process has to be moved, during its execution from one memory segment to another, then binding must be delayed until run time.

level



- An address generated by the CPU is commonly referred to as logical address.
- An address seen by memory unit is referred to as physical address.
- In case of compile time and load time binding, logical address = physical address.
- The runtime mapping from virtual to physical addresses is done by a h/w device called

the memory-management unit (MMU).

Dynamic loading;

- A routine is loaded ^{into memory} only when it is called.
- All routines are kept on disk in a ~~rel~~ relocatable load format.
- The main program is loaded into memory and executed. When a routine calls another routine, then only relocatable linking loaded is called to load the desired routine.


Advantages

An unused routine is never loaded.

Note: Dynamic loading does not require special support from the operating system.

It is the responsibility of users to design their programs to take advantage of such a method. Operating systems may help the programmer, however, by providing library routines to implement dynamic loading.

Dynamic linking and shared libraries:



Here linking is postponed until run time.

- A stub is included in the m/c code for each library routine reference.
- This stub is a small piece of code that indicates how to locate the appropriate memory-resident library routine, & how to load the library if the routine is not already present.

- Stub replaces itself by the with the address of address of routine and executes the routine.
- under this scheme, all processes that use a language library execute only one copy of the library code.
- unlike dynamic loading, dynamic linking generally requires help from the OS.
- If the processes in memory are protected from one another, then OS is the only entity that can check to see whether the needed routine is in another process's memory space & that can allow multiple processes to access the same memory address.

→ overlays

- we use overlays to store a process that is larger than the amount of memory allocated to it.
- The idea of overlays is to keep in

memory only those instructions and data that are needed at any given time.

→ when other instructions are needed, they are loaded into space occupied previously by instructions that are no longer needed.

BC: provides a two pass assembler. During pass 1, it constructs a symbol table; then, during pass 2, it generates machine-language code, using symbol table. We can partition such an assembler into pass 1 code, pass 2 code, the symbol table and common symbols routines used by both pass 1 and pass 2. Let us assume that sizes of these components are as follows:

Pass 1	70 KB
Pass 2	80 KB
Symbol table	20 KB

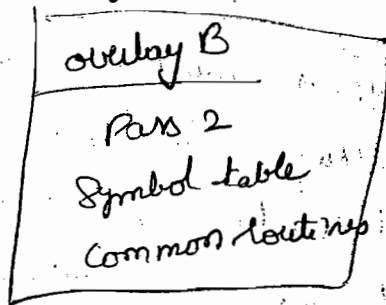
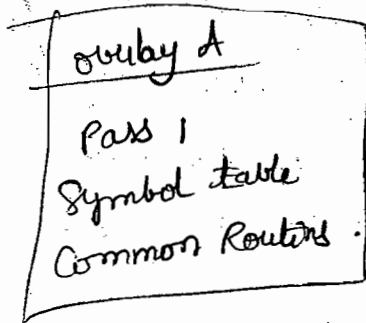
Common routines 30 KB.

To load everything at once, we would require 200 KB of memory. If only 150 KB of memory is available, we cannot run our program.


→ Hence we use an overlay device to load pass 1 code, then pass 2 code.

→ The code for

→ Thus we define two overlays:



→ code for overlay A and code for overlay B



are kept on disk as absolute memory images and are read by the overlay device as needed.


→ As in dynamic loading, overlays do not require any special support from the OS.

→ They can be implemented completely by user with simple file structures, reading from the files into memory and then jumping to that memory and executing the newly loaded instructions. The OS only notices only that there is more I/O than usual.

→ overlays is currently limited to mainframe computers and other systems that have limited amounts of physical memory and that lack the h/w support for more advanced techniques.

Swapping:

Swapping is moving a process b/w main memory and secondary memory.



→ If we use compile time load time binding, then process must be loaded into the same location; the process cannot be moved to different locations.

→ If we use execution time binding, then a process can be swapped into a different memory space, because the physical addresses are completed during execution time. Dispatcher is responsible for swapping in and swapping out.

→ Context switching time is a system which involves swapping is fairly high.

Ex: Let size of a process be 1MB.
and the disk has a transfer rate of 5MB/s.
Transfer of 1MB to & from memory
= $\frac{1}{5}$ seconds.
= 0.2 seconds.
= 200 milli seconds.

So both swap in and swap out takes 400 million words.

So for efficient CPU utilization, we want the execution time for each process to be long relative to the swap time.

memory manager

Functions

1. Allocation
2. Protection
3. File space management
4. De-allocation

Goals

1. Efficiency.
2. ability to run larger programs in small amounts of memory -

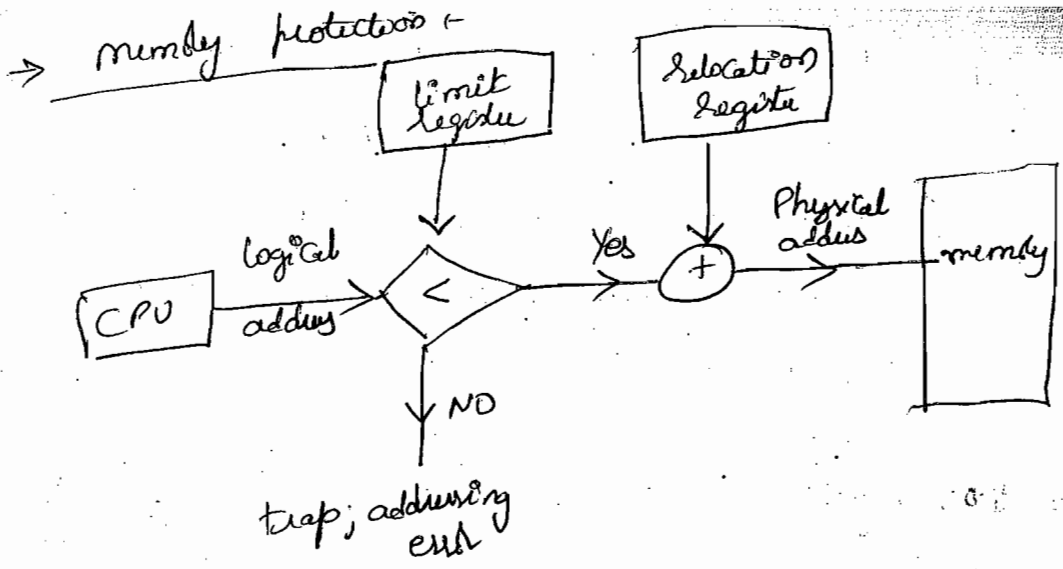
Techniques

Contiguous

- overlays
- Partitioning
 - Fixed
 - variable
- Buddy

Non-contiguous

- Paging
- Segmentation
- Virtual memory





Each logical add

- Protecting the OS from user programs and user programs from one another is needed.
- we can provide this protection by limit and relocation registers.
- Each logical address should be less than the limit register.
- The MMU converts the logical address to physical address dynamically by adding

the value in the relocation register.

→ when the CPU scheduler selects a process for execution, the dispatcher loads the relocation and limit registers with correct values as part of the context switch.

partitioning &



1) $P_i \rightarrow 14$

(i) Free - 22m

(ii) Best - 14m

(iii) Next -

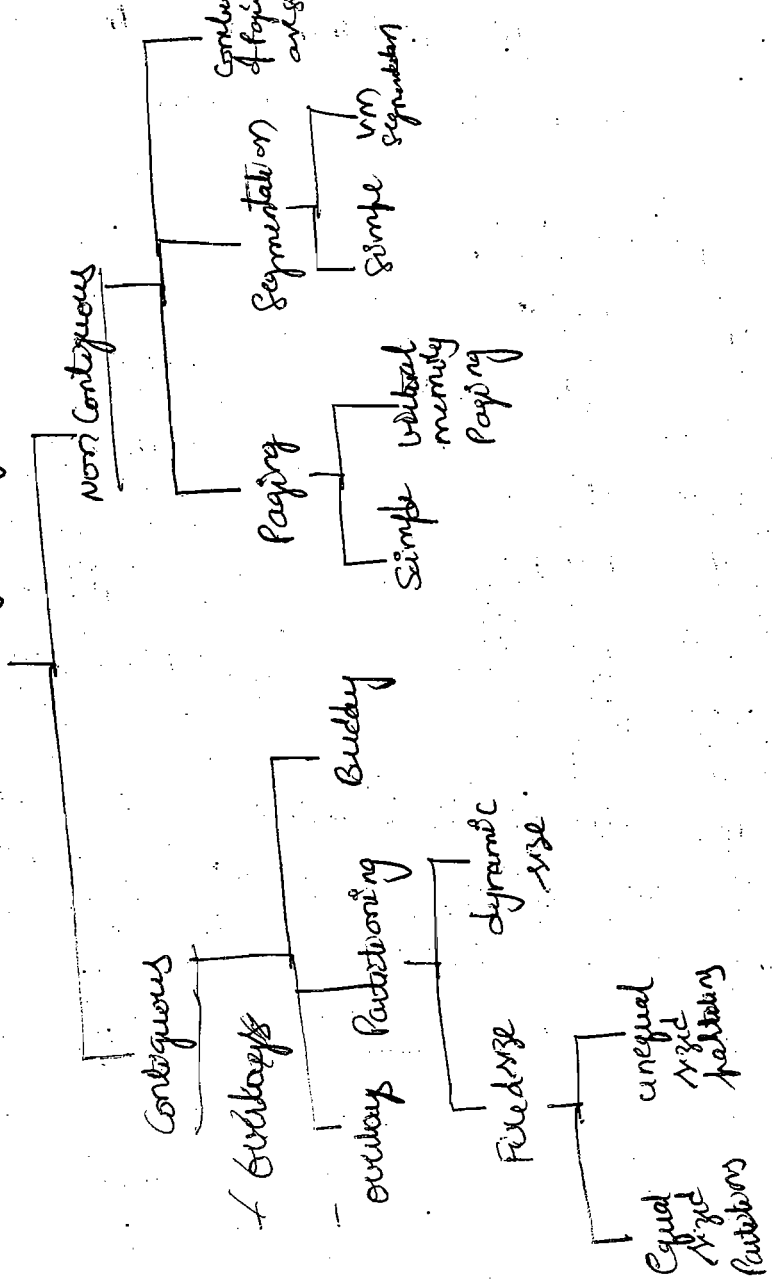
(iv) worst - 36m.

2) How many successive requests of size 5m could be satisfied 22.

memory management Techniques.



memory allocation



Partitioning :

→ core space can be partitioned ^{into} ~~into~~ ^{two} fixed size partitions and variable size partitions.

Fixed size partitioning : (MFT)


→ In this scheme, memory (user space) is partitioned into partitions, ~~whose size~~ and is fixed, the size of each partition does not change.

→ In order to run a process, it has to be placed in a partition which can hold it.

→ If the sizes of all partitions are equal, (decided at system generation time)

→ Since the size of each partition is fixed, and the size of memory is fixed, we can have only fixed number of processes at any time in the memory.

- Therefore, degree of multiprogramming is fixed in this scheme.
- So, it is also called multiprogramming with fixed number of tasks (MFT).
- If the sizes of all partitions are equal, then it is called, equal sized fixed partitioning.



- ~~If the prog. problem is too big to fit into a partition,~~
- Any process can be placed in any partition, since, all partitions are of equal size.

- If a process is too large to fit in a partition, then overlays should be used (overlays is ^{load on} outside).
- If the process is smaller than the size of a partition, then is some space wasted ^{due to} inside the partition, which cannot be used by any other process. This is called internal fragmentation.
- Both overlaying and internal fragmentation, can be decreased, though not solved, by using unequal size partitions.

0.8 MB
2 m
4 m
6 m
8 m
10 m
16 m
64 m
128 m

→ In unequal size partitioning, ~~the~~ ~~more~~ overlays will be used only if ~~the~~ ~~largest~~ largest partition also can't hold a process.

→ Internal fragmentation can be decreased by choosing a partition ^{empty} smallest, partition, which can hold a process. (Best fit)

This placement algorithm is called best fit algorithm. (If the process is very large and)

→ In some situations, if we add all the empty available partition sizes, and the memory ~~is~~ ^{is} available by ~~external~~ ^{internal} fragmentation, then also we can accommodate the process. But

Since this memory is scattered all around, we won't be able to accommodate all the process in memory. This phenomenon is called external fragmentation.

Disadvantages of fixed partitioning

- a) Internal fragmentation
- b) External fragmentation
- c) Degree of multiprogramming is fixed and is limited by the number of partitions in the memory.


→ To avoid the disadvantages (a) and (c) of the fixed partitioning, we go for dynamic & variable sized partitioning.

Dynamic & variable sized partitioning - (MVT)


→ In this partitioning, the no of partitions and size of each partition is not fixed at the system generation time.

→ Partitions are created dynamically, with the size of each process partition equal to the size of the process loaded. Therefore, degree of multiprogramming is variable. Hence it is also called multiprogramming with variable number of tasks (MVT).


Ex 1



↓ P1 finishes



⇒



→ In this scheme, at any time, there will be a set of holes, of various sizes scattered throughout memory.

→ whenever a request arises to load a process, which hole has to be selected is main issue here.

2. The allocation can follow four strategies,
- First fit
 - Best fit
 - Worst fit
 - Next fit

First fit: Search from starting of the list of holes, and allocate the first hole that is big enough.


Next fit: Same as first fit, but ^{instead of} searching from ~~first~~ starting of the list of holes, search the list from where you previous search ended.

Best fit: Allocate the smallest hole that is big enough to hold the process.


Worst fit: Allocate the largest hole.

Always

Consider a MVT system with following process & memory



→ ~~what is the available for C?~~




→ when a process request for 14m B, where will it be placed using

- (i) First fit - 22m
- (ii) Best fit - 14m
- (iii) Next fit
- (iv) worst fit - 36

→ How many successful requests of size 5m could be satisfied.

Ans: 2

→ Consider the MVT with following snapshots



300K, 25K,
125K, 50K

The request could be satisfied with which of the following:

- a) Both First and last b) neither First nor last
- c) First file but not last file
- d) last file but not first file

Advantage of MVT:

- a) degree of multiprogramming is not fixed
- b) no internal fragmentation

Disadvantage of MVT is external fragmentation.

external fragmentation

→ one solution to the problem of external fragmentation is compaction.

- Compaction is to shuffle the memory contents to place all free memory together in one large block.
- Compaction is possible only with dynamic address binding.
- Another possible solution to the external fragmentation problem is to permit the process to be non contiguous (Paging, segmentation and combination).

Paging :

<u>Bytes</u>	$2^5 = 32$	2^{50} - Peta
	$2^6 = 64$	2^{60} - Exa
	$2^7 = 128$	2^{70} - Zetta
	$2^8 = 256$	2^80 - Yotta
	$2^9 = 512$	
	$2^{10} = 1024 = K$	

$$2^{20} = M, \quad 2^{30} = G, \quad 2^{40} = T$$

$$2^{32} = 2^{30} \times 2^2 = 4 \times 2^{30} = 4 \text{ GB}$$

$$2^{45} = 2^5 \times 2^{40} = 32 \times 2^{40} = 32 \text{ TB}$$

$$2^{16} = 2^6 \times 2^{10} = 64 \text{ KB}$$

→ If you want to address two words then
no. of bits needed.

→ word: Smallest addressable unit in memory.
is called a word.


A word can be of one byte, two bytes or

word can be a collection of one or more bytes.


→ The collection depends

→ The size of a word generally depends on the
architecture of the computer.

→ How many bits are need to address words.



→ How many bits are needed to address 4 words



□ □ - 2 bits

→ ~~16~~ - 8 words - 3 bit address
 16 words - 4 bit address
 32 " - 5 bit address

of words - log
 33 words - 5x
 6 address

6 address - $2^6 = 64$
 $\frac{64}{33} = 31$ - wasted

→ 64 K words - $2^6 \times 2^{10} = 2^{16}$ KMGTP
 16 bits are required
(2 20 30 40 50)




128 MB = $2^7 \times 2^8 = 2^{15}$ 27 bits




Paging:

→ Paging is a memory management scheme that permits the physical address space of a process to be noncontiguous.

→ In this scheme, process (logical address space) is divided in pages and memory is divided into frames.


→ Size of a frame is equal to the size of a page.





> Advantages :

- > No external fragmentation, No need of
- > Before executing a process, all its pages will be loaded into any available ~~mem~~ memory frames.
- > Logical address generated by CPU, will be divided into two parts
a) Page number b) offset.



→ Using page number, page table is consulted and frame number is obtained. Now, frame number and offset forms PA.

→ Let size of a process be 2^n ~~bytes~~ words.

size of LA = n bits.

size of ~~one~~ a page be 2^k words.

offset contains = k bits.

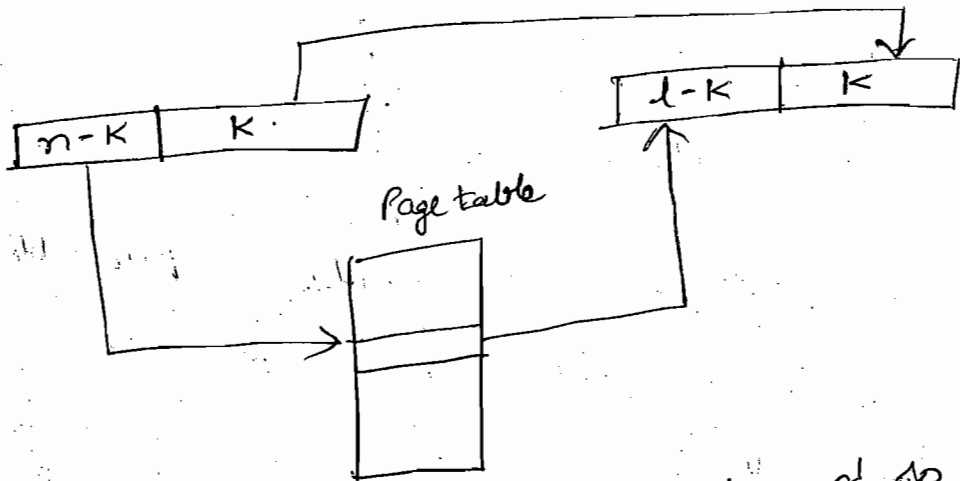
no of pages in process = $\frac{2^n}{2^k} = 2^{n-k}$

Page number field contains $n-k$ bits

→ Let size of main memory be 2^l words.

PA = l bits

no of frames in memory $\frac{2^l}{2^k} = 2^{l-k}$



→ There is no external fragmentation and so no need of compaction.

→ There may be internal fragmentation, if the size of a process is not a power of 2.

→ Average internal fragmentation = $\frac{(\text{Page size}) \cdot p}{2}$

→ Def average wastage per process = $\frac{\text{page size}}{2}$

- so smaller the page size, lesser the internal fragmentation.
- Page table size \propto ^{of process} no of pages of the process.
- ~~larger the~~ no of pages/process
- no of pages of process \propto $\frac{1}{\text{Page size}}$

$$\therefore \text{Page table size} \propto \frac{1}{\text{Page size}}$$

- larger the page size, smaller the page table size
- so page size should not be too large nor

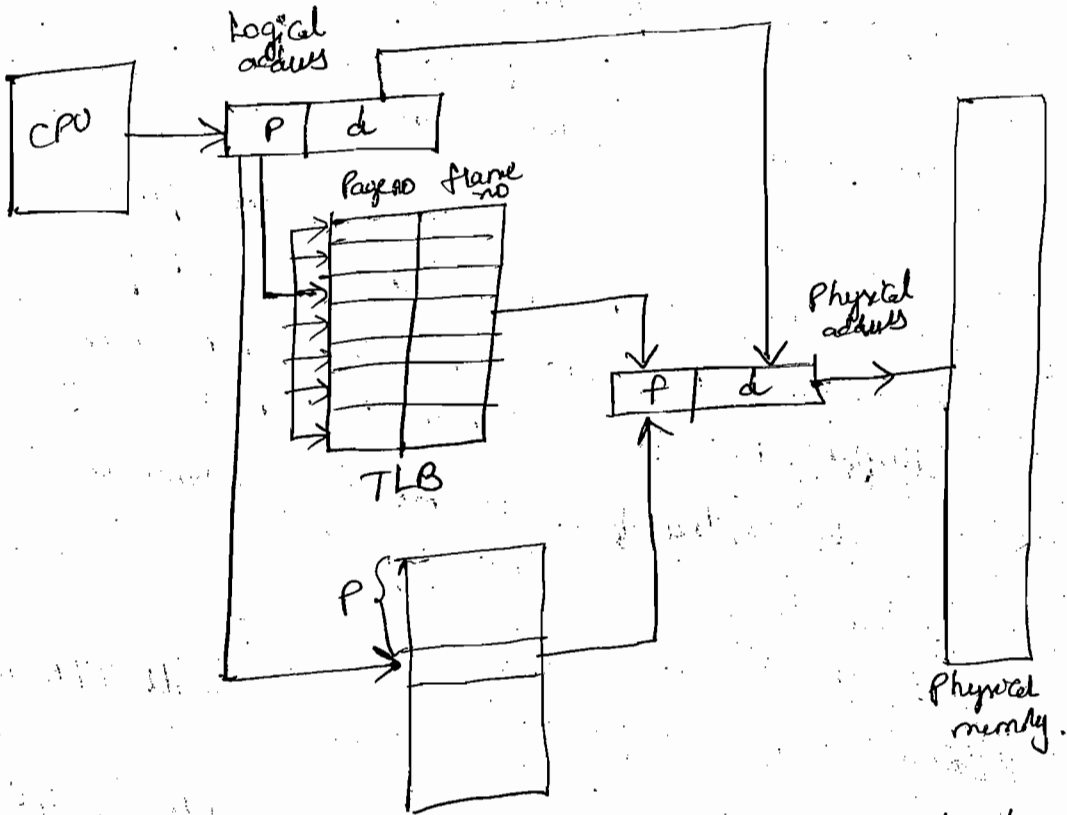
too small.

- The operating system will maintain a frame table. The frame table has one entry for frame. This ~~that~~ gives the information like which frames are allocated to which process & if the frame is free etc.

- Paging incurs context-switching time, because the page table has to be loaded into memory.

Hardware support for paging

- Page table can be placed in a set of registers.
- This is preferable if the page table size is ~~too small~~ reasonably small.
- The ~~best~~ standard solution to this problem is to use a special, small, fast lookup hardware cache, called translation look-aside buffer (TLB). (associative, high-speed memory).
- Each entry of a TLB contains a ~~tag~~ tag and a value.
- ~~But the size of~~
- But TLB is expensive and so very small TLB is used in many systems.
- So only a part of TLB of page table is kept in TLB. Each entry of TLB contains some page number as tag and frame number as value.
- When the LA is generated by CPU, its page number is presented to the TLB.



- If page number is not in TLB found, its frame number is immediately available and is used to access memory.
- If page number is not in TLB (known as a TLB miss), a memory reference ^{to} page table must be made in the memory will be consulted.

→ The page table entry which caused the miss will be again added to TLB, so that they will be found quickly on the next reference.

Hit ratio :-

→ The percentage of times that a particular page number is found in the TLB is called the hit ratio.

→ If hit ratio = 70, it means out of hundred times, we found a page 70 times.

Problem :-

Consider a system with in which it takes 20 nanoseconds to search TLB, and 100 nanoseconds to access memory. TLB has a hit rate of 80%. What is the effective memory access time.

Sol: Here we have two cases, if

(i) TLB HIT

(ii) TLB MISS

$$\begin{aligned}
 \text{TLB hit} &: 20 \text{ ns to search TLB} \\
 &+ \\
 &100 \text{ ns to fetch desired word.} \\
 &= 120 \text{ ns}
 \end{aligned}$$

$$\begin{aligned}
 \text{TLB miss} &: 20 \text{ ns to search TLB} \\
 &+ \\
 &100 \text{ ns to access page table in} \\
 &\quad \text{memory + frame number} \\
 &+ \\
 &100 \text{ ns to fetch the desired word.} \\
 &= 220 \text{ ns}
 \end{aligned}$$

To find effective memory access time, we must weigh each case by its probability.

$$\begin{aligned}
 \text{effective access time} &= 0.80 \times 120 + 0.20 \times 220 \\
 &= 140 \text{ nanoseconds}
 \end{aligned}$$

if hit ratio is 90%

$$0.90 \times 120 + 0.10 \times 220$$

98%

$$0.98 \times 120 + 0.02 \times 220$$

Problems:

→ If LAS = 8KB, PAS = 4KB and page size = 1KB.

- (i) size of logical address ^{and} no of bits in each field.
- (ii) size of physical address ^{and} no of bits in each field.
- (iii) If each entry size of each entry of page table is 'e' bytes, then the PT size = ~~N x e~~ bytes.

→ If KA = 29 size of logical address is 29 bits, no of pages = 2K ~~and~~ no of frames is 256, ~~Calculate~~ PA, and each word contains 4 bytes, then calculate

a) ^{LA and} LAS in bytes b) ^{PA} PAS in bytes

c) If each PTE is 4 bytes, then calculate PTS in bytes.

$$LAS = 2^9 \text{ words}$$

$$\text{So } LAS = 2^{29} \text{ words}$$

$$= 2^{29} \times 4 \text{ bytes}$$

$$PAS = \frac{\text{no of frames} \times \text{frame size}}{\text{Page size}}$$

$$LAS = \text{no of pages} \times \text{page size}$$

$$\Rightarrow \text{no of pages} =$$

$$\Rightarrow \text{Page size} = \frac{LAS}{\text{no of pages}}$$

$$= \frac{2^{29}}{2^k} = \frac{2^{29}}{2^{11}} = 2^{18} \text{ words}$$

$$PAS = \text{no} = 256 \times 2^{18} \text{ words}$$

$$= 2^8 \times 2^{18} \text{ words}$$

$$= 2^{26} \text{ words} = 2^{26} \times 4 \text{ bytes}$$

$$= 2^{28}$$

Page table size = no of entries in page table
 \times size of each entry.

$$= \text{no of pages} \times \text{size of each entry}$$

$$= 2K \times 4 \text{ Bytes}$$

$$= 2^4 \times 2^2 \text{ Bytes}$$

$$= 28 \text{ KB}$$

\rightarrow If LAS = 8GB, PAS = 64MB and PS = 16KB.

Calculate a) PA size. b) LA size c) PT size.

$$\text{No of Pages} = \frac{64 \text{ MB}}{16 \text{ KB}} = 4K$$

$$\text{No of frames} = 64M$$

$$\text{No of pages} = \frac{8 \text{ GB}}{16 \text{ KB}} = \frac{2^3 \times 2^{30}}{2^4 \times 2^{10}} = \frac{2^{33}}{2^{14}}$$

$$= 2^{19}$$

$$\text{No of frames} = \frac{64M}{16K} = \frac{2^6 \times 2^{20}}{2^4 \times 2^{10}}$$

$$= 2^{12}$$

Frame number field = 12 bits

PT entry size = 12 bits

PT size = no of ^{entries} pages \times size of each entry

= no of pages \times size of each bit

= $2^{19} \times 12 = 12 \times 2^{19}$ bits

= 12×2^{16} bytes

\rightarrow LA = 32 bits, PS = 4K, PAS = 64MB

then PT is bytes

Restrictions:

\rightarrow Some bits are stored in page table entry along with frame number, which the page

Information

\rightarrow read-write or read-only bit indicates if the page can be used only for reading & it can

be written also

\rightarrow A valid-invalid bit says if the address generated by CPU is valid & not-

Hierarchical paging:


- If the page table of a process is too large, we would not be able to allocate the page table contiguously in main memory.
- In such case, the page table will be divided into 2^k blocks of equal size.

Ex:

~~32~~ 32



Consider a system with 32 bit logical address space. Page size is 4KB. Then

Page table will have

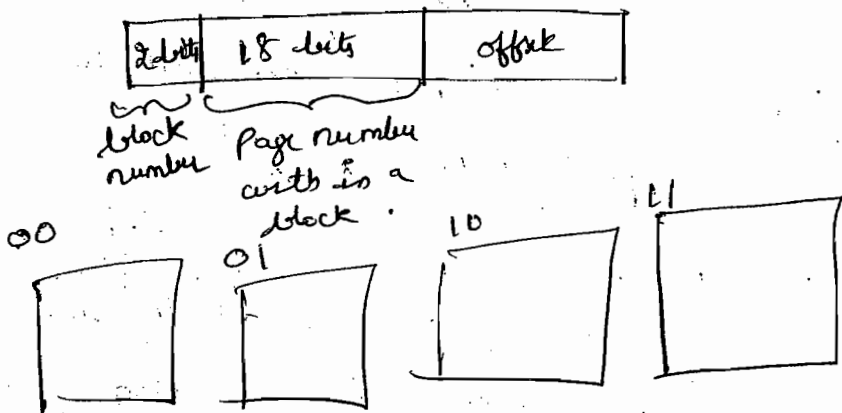


Page table will have 2^{20} entries.

If we divide it into two parts,



Divide the page table into 4 blocks, then
 LA can be divided as



→ If each page table entry is 4 bytes, then

$$PT \text{ size} = 2^{20} \times 2^2 = 2^{22} \text{ bytes}$$



→ If we split the page table in such a way that each block fits in a page, then we

will get $\frac{2^{22}}{2^{12}} = 2^{10}$ pages in the page table

→ Then we have a page table called
 two page tables called outer page table and
 inner page table.



→ now

LA looks like this



outer page table


Page of page table




→ This is called two level paging scheme.

→ we can go upto any levels of paging


Ex 1




→ Consider a system using two level paging in which inner page table is divided into 512 pages each of size 512 words. If the length of logical address is 34 bits, calculate the size of PAB if there are 256 frames in it.



Hashed page tables



Inverted page tables



one drawback with page table is that, each process will have a page table and each page table will contain millions of entries, which requires large amounts of PM.

→ To solve this problem, we use an inverted page table.

→ An inverted page table has one entry for each deal page (or frame) of memory.

→ Each entry consists of the virtual address of page table in that deal memory location, with information about the process that owns that page.

→ Thus, only one page table is in the system and it has only one entry for each page of PM.






→ Shared pages:

→ Another advantage of shared paging is the possibility of sharing common code.

→ Reentrant Code (or pure Code) is a non-self-modifying code.

→ If the code is reentrant, then it never changes during execution.

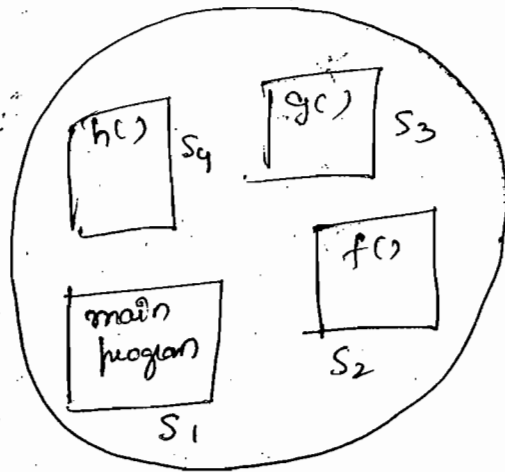
→ Thus, two or more processes can execute the same code at same time.



→ Examples of heavily used programs that are commonly used should include compilers, window systems, runtime libraries, database systems and so on.

Segmentation!

Paging does not preserve users view of a program.
User ~~uses~~ views a program as a collection of functions

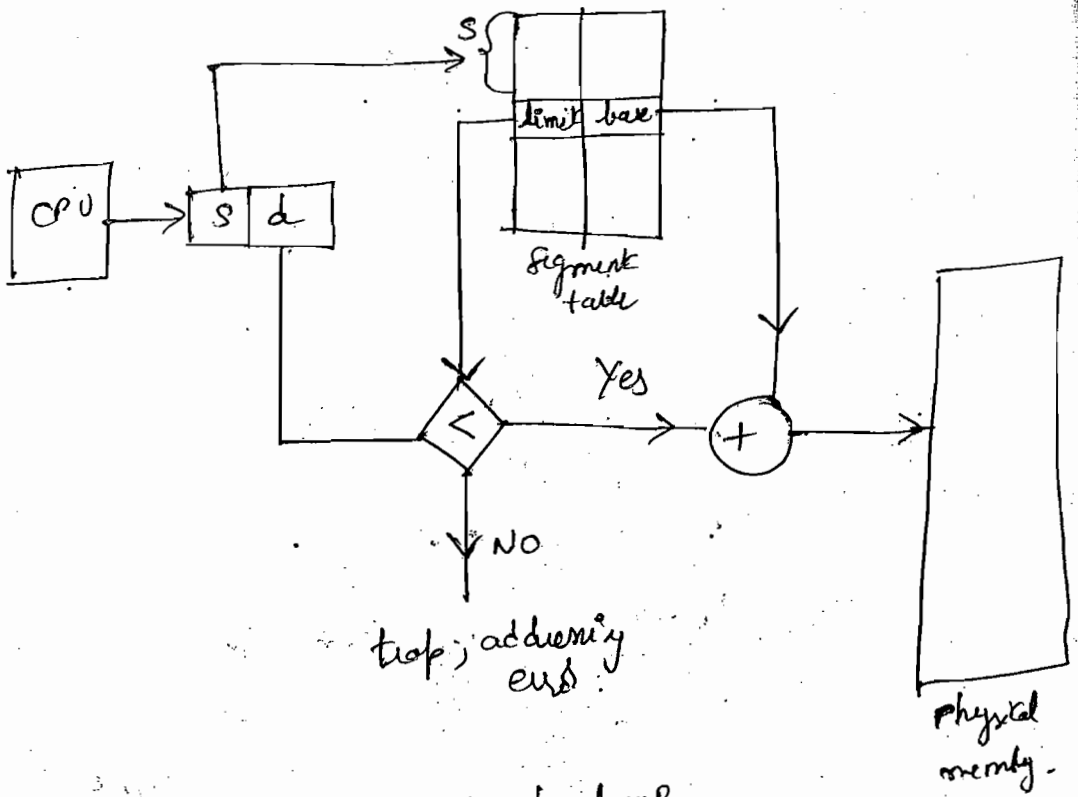


Segmentation is a memory management scheme that supports this user view of memory.

A logical-address space is a collection of segments according to segmentation.

This logical address consists of two tuple:


$\langle \text{Segment-number, offset} \rangle$



Disadvantage of segmentation:


- Link in ~~variable~~ length dynamic size partitioning, pure segmentation suffers from external fragmentation.
- This problem arises because, we wanted a segment to be contiguously allocated memory.

→ The solution to the above problem is to
page to divide each segment into pages and
memory into frames. This is called
→ the segmentation with paging.



Consider a system with segmented paging
 architecture. LAS = PAS = 2^{16} bytes. The LAS is
 divided into 8 equal size non overlapping
 segments. The segment is divided into pages
 and page tables of segments are kept in
 memory. The page table entry size is 2 bytes
 that must be the page size of segment so
 that the page table of ~~the~~ a segment
 exactly fits in 1 page. (assume that
 page size is a power of 2).

id:



Let page size = 2^K B

Page table size = no of pages \times size of each page entry.

$$= 2^{13-K} \times 2 = 2^{14-K} \text{ Bytes}$$

$$2^{14-K} = 2^K$$

we need to fit a page table in one page.

$$\therefore 2^{14-K} = 2^K$$

$$\Rightarrow \frac{2^{14-K}}{2^K} = 1$$

$$\Rightarrow 2^{14-2K} = 2^0$$

$$\Rightarrow 2K = 14 \Rightarrow K = 7$$

$$\therefore \text{page size} = 2^K \text{ B}$$

$$= 128 \text{ B}$$

Virtual memory :-

→ virtual memory gives illustration to the programmer that a huge amount of memory is available at his disposal for writing programs larger than the size of available physical memory.

- > virtual memory is commonly implemented by demand paging.
- > Demand paging is rather than similar to rather than loading an entire process into memory, we load only those pages which are needed.
- > A page that is not needed will be ~~kept~~ ^{remain} on disk until it is needed.

Frame allocation policies

Let total frame = m

no. of processes = n

Demand of each process i = s_i

$$\sum s_i \leq S$$

Let frames allocated for each process = a_i

Frames can be allocated to processes using:

- (a) Equal allocation (b) proportionate allocation.

Equal allocation:-

$$a_i = \frac{m}{n}$$

allocate frames to all processes equally.

$$a_i = m/n$$

It is useful only when all programs are equal sized.

Proportional allocation:-

$$a_i = \frac{S_i}{S} \times M$$

Page replacement algorithms:-

Reference string: set of successively unique pages referred in the given list of LA's, is called reference string.

Ex: Let size of a page be 100 words, and a process generates the LA's as follows:
122, 142, 148, 150, 172, 112, 754, 787,
756, 896, 012, 064, 074, 158, 172,
666

new reference string = (1, 7, 8, 0, 1, 6)

length of ref string = 6

unique pages referred = $n = 5$

Demand paging

Pure demand paging

Pre demand paging

Pure demand paging


we start with empty frames

Pre demand paging

we start with frames filled

FIFO: refs 1, 2, 3, 4, 1, 2, 5, 1, 2, 3, 4, 5

3)



PF = 9



PF = 10

5F - 5

6F - 5

7F - 5

Ref: 5, 4, 3, 2, 1, 4, 3, 5, 4, 3, 2, 1, 5

2) 3 frames = 210 4 frames = 11

~~A B C D A B 5 2 1~~

~~Belady's anomaly~~ ~~A B C D A B 3 2 1 5~~

Ref: 7, 0, 1, 2, 0, 3, 0, 4, 2, 3, 0, 3, 2, 1, 2,

1) 0, 1, 7, 0, 1.

~~A B C D A B C D A B C D 7 0 1~~

3 frames = 15 faults

4 frames = 10 faults

~~A B C D A 0 1 2 7~~

- For some page replacement algorithms, the page fault rate may increase as the number of allocated frames increases.
- This most unexpected result is known as Belady's anomaly.


Optimal page replacement

→ To avoid Belady's anomaly, optimal page replacement algorithm was proposed.

→ ~~the~~ optimal page replacement is to replace a page that will not be used for the longest period of time.

$Z = 7, 0, 1, 2, 0, 3, 0, 4, 2, 3, 0, 3, 2, 1, 2, 0, 1, 7, 0, 1.$

$R = 1, 0, 7, 1, 0, 2, 1, 2, \dots$



3 frames - 9 PF.

4 frames -

$= 1, 2, 3, 4, 1, 2, 5, 1, 2, 3, 4, 5. \quad S^R_2$

3 frames -

4 frames -

$\cdot 5, 4, 3, 2, 1, 4, 3, 5, 4, 3, 2, 1, 5 \quad S^R_2$

3 frames

or frames

→ optimal page replacement algorithm is difficult to implement, because we need to know the future ^{values} knowledge of reference string to take current decision.

→ since no other page replacement algorithm can outperform optimal page replacement algorithm, this is ~~used as~~ mainly used for comparison studies.

LRU Page replacement

LRU page replacement ^{algo} will replace the page that has not been used for the longest period of time.

7, 0, 1, 2, 0, 3, 0, 4, 2, 3, 0, 3, 2, 1, 2, 0, 1, 7, 0, 1.

A 0 1 2 3 4 2 3 0 1 0 7

3 frames = 12 PFS.

4 frames = ?

SR

2?

LRU Implementation

a) counter: we maintain a register called clock register which is incremented for every memory page reference.

we maintain a time of use field in every entry of page table.

with every page reference / memory reference, the value of the clock register is copied to the time of use field in the page table.

b) stack: we maintain a stack of page numbers. whenever a page is referenced, it is removed from the stack and puts on the top.

In this way, the top of stack is always the most recently used page and the bottom of

is LRU page.

This stack is implemented as a double linked list.

Stack algorithms :-

- A stack page replacement algorithm is called a stack algorithm, if it can be shown that the set of pages in memory for n frames is always a subset of the set of pages in that would be in memory with $n+1$ frames.
- Stack algorithms never exhibit Belady's anomaly.

LRU approximation page replacement :-

Reference bit

- Each entry in page table will have a reference bit, associated with every page.
- The reference bit for a page is set, whenever the page is referenced.

Additional-reference-bits algorithm :-

- we keep an 8-bit byte for each page in a table in memory.
- At regular intervals, OS ~~will~~ shifts the

reference bit for each page into high-order
bit of its 8-bit shift register byte,
shifting the other bits right 1-bit, discarding
the low-order bit.

→ These 8-bit shift registers contain the history
of page use for the last eight time periods.

Ex) $\begin{matrix} 11000100 \\ 01110111 \end{matrix}$ → more recently used
among the two.

Second chance algorithm

→ It is a variant FIFO replacement algorithm.

→ when a page has to be selected for
replacement, we inspect its reference bit. If the

value is 0, we proceed to replace this page.

If it is 1, we give that page a second

chance and move on to select the next FIFO
page.

when a page gets a second chance, its
reference bit is cleared and its arrival
time is reset to the current time.

Enhanced second channel

we use both reference and modified link as an ideal pair.

(0,0) might be

- best for reference

(0,1)

- not quite good.

(1,0)

- it probably be used again.

(1,1)

- it probably will be used again soon

Counting-based page replacement:


we could keep a counter of the number of references that have been made to each page. we develop the following two schemes.

at least frequently used page replacement algo:

As this requires that pages with smallest count be replaced. The reason for this selection is that an actively used page should have a large reference count.

Disadvantages

1, 2, 3, 2, 3, ~~2~~ 4, 1, 4, 1, 4, 1, ...



A page which is used heavily during the initial phase of a process, but then is never used again will always ~~stay~~ remain in memory.

The solution is to right shift the count by 1 bit at regular intervals.

→ most frequently used (MFU) Page-replacement algo

This is based on argument that the page with the smallest count was probably just brought in and has yet to be used.

→ Both MFU and LFU ~~are~~ implementations are expensive and they do not approximate OPT replacement well.

Thrashing:

1, 2, 3, 4, 1, 2, 3, 4

- If no of frames allocated ~~are~~ ^{is} less than the number of pages in active set, ~~is~~ ^{is} ~~repeated~~ ^{repeated} page faults occur very often.
- Page faults will have significant effect on the performance of a computer system.
- If there are no page faults, effective memory access time = main memory access time.
- If there is page faults, effective memory access time = Time taken to fetch read the relevant page from disk, and then access the desired word.
- Let us call ^{the} ~~the~~ time taken to process ~~the~~ page fault service time as page fault service time.

Let p be the probability of page fault. Let

Then

effective memory access time

$$= (1-p) \times \text{ma} + p \times \text{page fault service time}$$

If page fault service time = 25 milliseconds

and memory access time = 100 nano seconds

Then effective access time in nano seconds

$$= (1-p) \times 100 + p (25 \text{ milliseconds})$$

$$= (1-p) \times 100 + p \times 25,000,000$$

$$= 100 + 24,999,900 \times p$$

↳ Memory

Effective memory access time $\propto p$

→ A high paging activity is called thrashing.

A program is thrashing, if it is spending more time in paging than executing.

→ If trashing occurs, then ~~decrease~~^{increase} the degree of multiprogramming to λ .

→ Program structure.

Assume pages are 128 words in size.

Consider a Java program whose function is to initialize to '0' each element of a 128 by 128 array. The following code is typical:

```
int A[][] = new int[128][128];
```

```
for (int j=0; j < A.length; j++)
```

```
for (int i=0; i < A.length; i++)
```

```
    A[i][j] = 0;
```

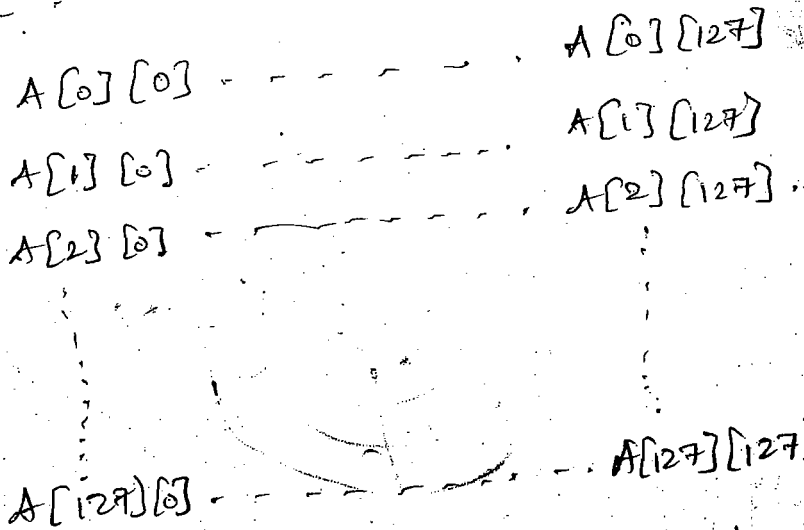
note that an array is stored in row-major order.

we.

$A[0][0]$ $A[0][1]$... $A[0][127]$; $A[1][0]$,

$A[1][1]$... $A[127][127]$.

Let size of each element of array be o
 then each row will be present in one
 page.



If OS allocates less than 128 frames to this
 process, then each page reference will produce
 a page fault. So no. of page faults
 $= 128 \times 128 = 16,384$ faults.

now change the code to as follows:


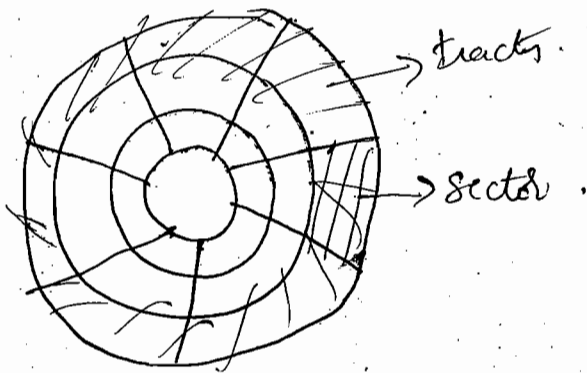

```

int A[128][128];
for (int i=0; i < A.length; i++)
    for (int j=0; j < A.length; j++)
        A[i][j] = 0;
    
```

This modules only for tape faults.

~~ditto~~

Diagram 1-



$$\text{Capacity} = \text{no of tracks surfaces / disk} \times \text{no of tracks / surface} \\ \times \text{no of sectors / track} \times \text{no of bytes / sector}$$

- Given that no of surfaces = 32, ~~sectors~~
 no of ~~the~~ tracks/surface = 64
 no of sectors / track = 128
 sector size = 4096 byte

what is the capacity of disk.

- writing on disk is called recording. There are two types of recording
- variable recording density
 - Fixed recording density.

→ Density of a disk of a track is number of bytes stored per unit distance.

→ In Fixed recording density, the no of bytes/unit distance is constant and so is the distance b/w the bits in all tracks and so ~~no~~ the no of bits/track will vary and so ~~inner~~ the track capacity is less the capacity of inner track is less than that of outer track.

→ Linear velocity is used to read/write from disk.

variable recording density

- In variable recording density, the no of bits per unit distance is not constant.
- The Capacity of every track is same.
- So the bits in the inner tracks are densely packed and the ^{bits in} outer track are sparsely packed.
- Angular velocity is used to read/write

Seek time (T_{seek})

Seek time is the time taken for the disk arm to move the heads to the cylinder containing the desired sector.

The Rotational

Rotational latency - It is the ~~rotational~~ ^{rotational} latency additional time waiting for the disk to rotate the desired sector to the disk head.

Transfer time - The time taken to transfer a ~~data~~ block of sector is called transfer time.

→ Disk access time = Seek time + Rotational latency + data transfer time.

→ Rotational latency can be known if the current position of the head R/W head is given, else, we have to take rotational latency as $\frac{1}{2}$ (time taken for one rotation).

→ To have eff

→ To increase the efficiency, many disks are stacked together and it is called hard disk.

→ If each sector is 1KB.

→ If we have a file of 10KB, and sector size is 1KB then, instead of storing each 1KB on a single track, we will distribute it across the disks by a concept called cylinder.

→ If we have

> A disk system of 2GB

> Disk formatting :

Each sector of should contain some control information like sector number etc. Such information should be not accessible by users and that space should not be used for storing user information. This is called formatting.

> A disk system of 2GB is constructed with each disk having the following specifications :

(i) It will have 64 tracks/surface

(ii) It will have ~~100~~ 1024 sectors/track.

(iii) It will have 1024 Bytes/sector and the disk is moving at an angular velocity

6000 rpm.

The formatting overhead is 128 bytes/sector

Inner track

Capacity is 0.7cm - seek time

= 10 ms.

- Q: (i) what is the data transfer rate of each disk
 (ii) what is the capacity of formatted disk
 (iii) what is the average access time
 (iv) if the disk speed is doubled, what is the % of reduction in average access time.
 (v) what is the maximum recording density employed in the above recording.
 (vi) if each disk is having the data on both surfaces except the I and last one, for which one surface is used as a guard surface. what will be the no. of disks needed for the desired capacity.

Sol: (i) no of bytes transferred in 1 sec.

In one complete revolution
 in one rotation, it will cover 1 track
 Given speed is 6000 rpm.

1 min - 6000 rotations

$$1 \text{ rotation} = \frac{1}{6000} \text{ min} / 60 \text{ sec}$$

$$1 \text{ rotation} = 1 \text{ ms}$$

$$\text{Capacity of each track} = 1024 \times 1024 \\ = 1048576 \text{ KB}$$

$$\text{In } 10 \text{ ms} - 1 \text{ MB}$$

$$1 \text{ sec} - 100 \text{ MB}$$

$$\therefore \text{data transfer rate} = 100 \text{ MBPS}$$

$$\text{ii) Capacity of a disk} = 2 \times 64 \times 1024 \times \left(\frac{1024 - 1}{2} \right) \text{ bytes}$$

$$\text{iii) } T_{\text{ave}} = 10 \text{ ms} + \frac{1}{2} (10 \text{ ms}) + \frac{1024}{100 \text{ MBPS}} \\ = 15.01 \text{ msec}$$

iv) If speed of disk is double

$$T_{\text{ave}_1} \approx 15 \text{ ms}$$

$$T_{\text{ave}_2} \approx 12.5 \text{ ms}$$

$$\text{w.r.t } 15 \text{ ms} - 2.5 \text{ ms}$$

$$100 \text{ ms} - \frac{2.5}{15} \times 100 = 16.67$$

(v) Inner track diameter = 0.7 cm

$$\text{Circumference} = 2\pi r$$

$$= \pi (2r)$$

$$= \pi D$$

no of ~~bits~~ Bytes in $\pi(D)$ length = 1024 KB

$$\Rightarrow \frac{22}{7} \times 0.7 \text{ cm} = 1 \text{ mB}$$

$$\Rightarrow 2.2 \text{ cm} \approx 1 \text{ mB}$$

$$\Rightarrow 1 \text{ cm} = \frac{1}{2.2} \text{ mB/cm}$$

(vi)

$(2n-2) \times 2^x$

Let n be the no of disks

$$(2n-2) \times 2^{26} = 2^{31}$$

$$\Rightarrow 2n-2 = 32$$

$$\Rightarrow n = 17$$

→ A disk has 8 equidistant tracks, the diameter of inner track is 1 cm, while outermost track is 8 cm. The inner most track has a storage capacity of 10 MB. What is the amount of data that can be stored on disk if it is used with a drive that rotates (i) with constant linear velocity (ii) constant angular velocity.

a) 80 MB, 2040 MB b) 2040 MB, 80 MB

c) 80 MB, 360 MB d) 360 MB, 80 MB

In linear $10 + 20 + \dots + 80 = 360 \text{ MB}$

In angular, $10 + 10 + \dots + 10 = 80 \text{ MB}$

Resource-request algorithm

1. If $Request_i \leq need_i$, go to the step 2. otherwise, have an end condition, since the process has exceeded its maximum claim.

2. If $Request_i \leq Available$, go to step 3. otherwise P_i must wait, since the resources are not available.

3. Have the system pretend to have allocated the requested resources to P_i by modifying the state as follows:

$$Available = Available - Request_i;$$

$$Allocation = Allocation + Request_i;$$

$$Need_i = Need_i - Request_i;$$

If the resulting resource allocation state is safe, the transaction is completed and process P_i is allocated its resources.

Safety algorithm

Let "work" and "finish" be vectors of length m and n , respectively. Initialize work = Available and finish $[i] = \text{false}$ for $i = 1, 2, \dots, n$.

2. Find an i such that both

a. finish $[i] = \text{false}$

b. need $_i \leq \text{work}$.

if no such i exists, go to step 4.


3. work = work + allocation $_i$

finish $[i] = \text{true}$

goto step 2.



4. If finish $[i] = \text{true}$ for all i , then the system is in a safe state.

This algorithm requires an order of $m \times n$




Compile time
loading

absolut m/c code



load time



~~Exercise~~
Allocation

~~15~~

	A	B	C	D
P_1	1	2	2	1
P_2	1	0	3	3
P_3	1	1	1	0

available

A	B	C	D
3	1	1	2

maximum requirement

3	3	2	2
1	2	3	4
1	1	5	0

need

P_1	2	1	0	1
P_2	0	2	0	1
P_3	0	0	4	0

$\langle P_1, P_2, P_3 \rangle$ ✓

safe state

At the fork P_2 asks for 1B.

$$\text{request} = 0 \ 1 \ 0 \ 0$$

$$\text{because} = 3 \ 0 \ 1 \ 2$$

unsafe

~~request~~ P_3 asks for $(0, 0, 2, 0)$

allocation cannot be met

P_3 asks for $(0, 0, 1, 0)$

$\langle P_1, P_2, P_3 \rangle$ safe sequence

$$T_{n+1} = \alpha t_n + (1-\alpha) T_n$$

(45)

$$= \alpha t_n + (1-\alpha) (\alpha t_{n-1} + (1-\alpha) T_{n-1})$$

$$= \alpha t_n + (1-\alpha) \alpha t_{n-1} + (1-\alpha)^2 T_{n-1}$$

~~P1 P2 P3 P4 P5 P6~~

~~P1 P2 P3 P4~~

~~P1 P2 P3 P4 P5 P6 P1 P2 P3~~

P1	P2	P3	P4
----	----	----	----


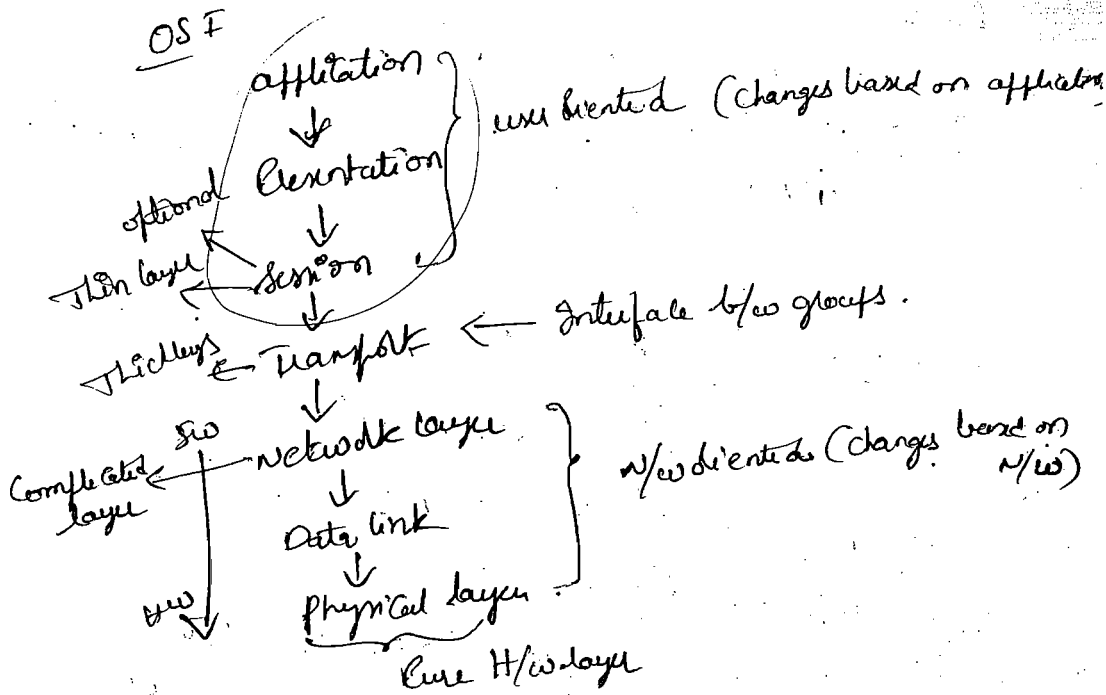
0 2 4 6 7

~~P1 P2 P3 P4 P5 P6 P3 P5 P6~~

	AT	BT	BT
1	0	4 2	4 2
2	1	5 3	5 3
3	2	3 1	3 1
4	3	4 0	4 0
5	4	6	6 4
6	5	5	5 3

P1	P2	P3	P1	P5	P6	
----	----	----	----	----	----	--

0 2 4 5 8 10 12



Process



→ application layer

Socket (It is also known as API b/w application and network layer).

→ Transport layer

Socket programming with TCP

Any n/w application consists of a pair of programs, a client program and a server program - residing in two different hosts.

When these two programs are first executed, a client and a server process are created and these two processes communicate with each other by reading from and writing to sockets.

→ To write any of n/w program, there will be set of rules called prototype specified by RFC.

→ we can directly implement RFC's.

Ex: RFC-959 client side of FTP.
RFC-959 server side of FTP.

network programming

follow RFC protocols

Ex) RFC-959
FTP server and client

Client can be created by a programmer and server by another programmer

need not follow RFC

Ex: Proprietary n/w application

client and server has to be created by same developer & development team

→ ~~Proprietary n/w applications~~

~~Step 1. decide whether the application is to run on TCP or UDP.~~

~~n/w applications~~

~~Socket programming with TCP~~

~~Socket programming with UDP~~

Socket programming with TCP:-

- The client initiates the contact with the server.
- For the server should always be listening.
- You can compare this as someone knowing Socket programming

Step 1: (Socket Creation)

Create a socket, which is done with the following

operation:

`int socket (int domain, int type, int protocol).`

domain: specifies protocol family that is going to be used

`PF_INET` denotes internet family.

`PF_UNIX` denotes unix pipe facility.

`PF_PACKET` denotes direct access to network interface.

Type: Type argument indicates the semantics of the communication.

`SOCK_STREAM` is used to denote byte stream.

`SOCK_DGRAM` is used to denote message oriented service.

Protocol: Protocol argument identifies the specific protocol that is going to be used.

→ The return value from the socket is a handle for the newly created socket, i.e., an identifier by which we can refer to socket in the future.

Step 2: ^(Connection establishment) This step depends on whether you are a client or a server. ☹

Server: on a server machine, the application process performs a passive open - the server says that it is prepared to accept connections, but it does not actually establish a connection. The server does this by invoking the following three operations:

`int bind (int socket, struct sockaddr * address, int addrlen)`

`int listen (int socket, int backlog)`

`int accept (int socket, struct sockaddr * address, int * addrlen)`

Bind: The bind operation, as its name suggests, binds the newly created socket to the specified address.

This is the n/w address of the local participant - the server. when used with internet protocols, "address" is a data structure that includes both IP address of the server and a TCP port number.

web servers - 80


listen: The listen operation then defines how many connections can be pending on specified socket.

→ accept: This operation carries out the passive open.

It is a blocking operation that does not return until a remote participant has established a connection, and when it does complete, it returns a new socket that corresponds to this just established connection, and the address argument contains the remote participant's address.

client process

server process



→ client machine:

on the client machine, the application process performs an active open, that is, it says who it wants to communicate with by invoking the following single operation:

```
int connect(int socket, struct sockaddr *address, int addrlen).
```

address contains the remote participant's address.

This is a blocking call, and does not return

until TCP has successfully established a connection.

step 3: (Communication)

once a connection is established, the application processes invoke the following two

operations to send and receive data

```
int send(int socket, char * message, int msg-len,  
int flags)
```


```
int recv(int sock, char * buff, int buf-len,  
int flags).
```

→ Here flags control details of operation.


→ `close(int socket)` is used to close a socket.

File systems


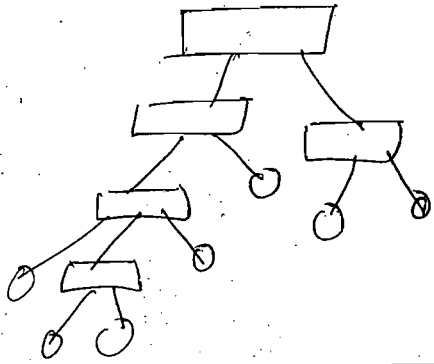
Single level



Two level




multiple hierarchy




Boot block	Super block	Free space mgmt	i-nodes	root dir	Files and directories
------------	-------------	-----------------	---------	----------	-----------------------

Contiguous allocation of files




Storing a file as a linked list of disk blocks

File A.



File B.



linked list

0	
1	
2	10
3	11
4	7
5	
6	3
7	2
8	
9	
10	12
11	14
12	-1
13	
14	-1
15	

disk allocation using a FAT in main memory


Example 1 - node

File Attributes	
Add of file blocks	
	2
	3
	4
	5
	6
	7
Address of blocks of points	

Disk block
 containing
 additional
 disk addresses

Implementing directories

f1	attributes
f2	attributes
f3	attributes
f4	attributes




a) Filled sse entries with disk address and attributes in directory entry.

data structure containing the attributes.

FFS

0 14 37 53 65 67 98 122 124 183 199.



SSTF:

