

RR

Interactive, preemptive

Like FCFS but preemptive so when time quantum lapses, task forced to give up CPU and placed to end of queue (after I/O)

Response time (guaranteed) to be upper bounded by $(n - 1)q$ with n tasks and time quantum q

Larger time quantum yields better CPU utilization but longer waiting time, shorter time quantum yields larger context switching overhead but shorter waiting time

Behaves like FCFS if job lengths are shorter than time quantum

Performs worse than FCFS if job lengths are all the same and greater than time quantum (higher average turnaround time) or when there are many jobs and job lengths exceed time quantum (reduced throughput due to increased overhead)

Priority Scheduling

Interactive, preemptive/non-preemptive

Assigns priority to processes and run highest priority first

Preemptive: Higher priority process can preempt running process with lower priority

Non-preemptive: Late coming waits for next round of scheduling

Low priority process can starve if high priority keeps hogging CPU (worse for preemptive); resolved by decreasing priority or current process after each time quantum or giving current process a time quantum

Priority inversion: lower priority preempts higher priority because resource is locked so higher priority is blocked

MLFQ

Interactive, preemptive

Multiple queues with different priority levels, minimizes response time for I/O bound processes and turnaround time for CPU bound processes

Highest priority queue must be empty before next queue is used

Rules:

- Priority(A) > Priority(B) -> A runs
- Priority(A) == Priority(B) -> A and B runs in RR
- New job -> highest priority
- Job fully used time quantum -> reduce priority
- Job gives up before time quantum -> retain priority

Change of heart: process with lengthy CPU phase right before I/O phase sinks to the lowest priority and gets starved; periodically move all tasks to highest priority to fix

Gaming the system: process repeatedly gives up CPU before time quantum lapses remains in high priority and starves other processes; accumulate total CPU use time across all quantum

Lottery Scheduling

Interactive, preemptive

Give out "lottery tickets" to processes for system resources, randomly choose lottery ticket, winner gains resource; lottery tickets can be distributed to children

Process holding X% of tickets has X% chance to win and use X% of resource

Responsive as new processes can participate

Threading

- Units of work within a process
- Shares memory context (text, data, heap) and OS context (process ID, files, etc.)
- Uniquely identified by thread ID, registers, and stack


Thread switching: change hardware context like registers and stack

Benefits:

- Economy: multiple threads in the same process requires less resources to manage
- Resource sharing: threads share most of the resources of a process
- Responsiveness
- Scalability: take advantage of multiple CPUs

Problems:

- System call concurrency: ensure that parallel system calls have no race condition
- Process behavior: what happens when a single thread calls exit() or exec()



User Thread

- Thread implemented as user library (process handles thread related operations)
- Kernel not aware of user threads
- User thread with no kernel thread still runs on a single thread

Pros:

- Can have multithreaded program on any OS
- Thread operations are just library calls
- More configurable and flexible

Cons:

- OS not aware of threads and scheduling is performed at process level
- If one user thread blocks, the entire process is blocked so all threads are blocked
- Cannot exploit multiple CPUs

Kernel Thread

- Thread is implemented in the OS handled as system calls
- Kernel schedule by threads, not processes; may make use of threads for its own execution

Pros:


- Kernel schedules on thread levels
- Multiple threads can run for the same process and be non-blocking

Cons:

- Thread operations are system calls that are slower and more resource intensive
- Less flexible since kernel threads used by all multithreading programs

Hybrid Thread

- User threads bind to a kernel thread



POSIX thread

- Implemented as either user or kernel threads
- pthread_t:** data type to represent thread id
- pthread_attr:** data type to represent attributes of thread

pthread_create:

```
int pthread_create(
    pthread_t *tidCreated,
    const pthread_attr_t *threadAttributes,
    void * (*startRoutine)(void*),
    void * argForStartRoutine);
Return 0 if successful, !0 if errors
```

tidCreated: thread id for the created thread

threadAttributes: control the behavior of the new thread

startRoutine: function pointer to the function to be executed by thread

argForStartRoutine: arguments for the startRoutine function

pthread_exit: pthread terminates automatically at the end of the startRoutine with return value defined by return statement

```
int pthread_exit(void * exitValue);
exitValue: value to be returned to whoever synchronize with this thread
```

pthread_join: wait for the termination of another pthread

```
int pthread_join(pthread_t threadID, void **status);
Return 0 if successful, !0 if errors
threadID: TID of the pthread to wait for
status: exit value returned by the target pthread
```

Shared Memory (IPC)

- Shared between processes
- Process P1 creates shared memory M (shmget) and P2 attaches M to its memory space (shmat)
- OS involved in creating and attaching memory region; easy to use as all operations are array operations
- Requires synchronization and implementation is harder
- Must detach M (shmdt) and destroy M (shmctl) after all processes detached

Message Passing (IPC)

- Process P1 sends message M to P2 and P2 receives M (all as system calls)
- Additional properties: identification of other party (naming scheme) and synchronization of send/receive operations
- Messages are stored in kernel memory space

Direct communication: sender/receiver of message explicitly name the other party

- One link per pair of communicating processes
- Identify of other party must be known

Indirect communication: messages are sent to/received from message storage (mailbox/port)

- One mailbox shared among multiple processes

Synchronization behaviors for send() and receive(): blocking primitives (synchronous) or non-blocking primitives (asynchronous)

Advantages: easier to implement on different processing environments and easier to synchronize

Disadvantages: requires OS intervention so inefficient and harder to use due to limits on message size/format

Unix Pipes (IPC)

- Share input/output between processes (producer/consumer)

Process communication channels: stdin, stdout

Piping (A | B): links the input/output channels of one process to another

Pipe in C: circular bounded byte buffer with implicit synchronization (writer waits when buffer is full, reader waits when buffer is empty)

Unidirectional (half-duplex): one write end, one read end

Bidirectional (full-duplex): any end for read/write

System calls:

```
int pipe(int fd[]);
close(int) -> side to close
write(int, data, length) -> side to write to
read(int, buffer, length) -> read from
0 for success, !0 for errors
fd[0] is reading, fd[1] is writing
```

- Must close the end not in use, otherwise, other processes might accidentally write/read to it

Redirecting communication channels to pipes: use dup() and dup2()

messages downloaded_by

- Sets new file descriptor to refer to old as well (i.e. old and new are the same now)

```
dup2(fd(0), STDIN_FILENO)
dup2(fd(1), STDOUT_FILENO)
dup2(file, STDOUT_FILENO)
```

Unix Signals (IPC)

- Asynchronous notification regarding an event sent to a process/thread
 - E.g. kill, stop, continue
- Recipient of the signal must handle the signal via default handlers or user supplied handlers (signal(type, handler))

Race Conditions

- Execution outcome depends on the order in which the shared resource is accessed/modified (non-deterministic)
- General modification flow: load memory value into register, update register value, load register to memory
 - Race conditions happen when loading happens at the wrong time
- Caused by unsynchronized access to shared modifiable resources

Solution: designate code segment with race condition as critical section, only one process can execute in CS (synchronization)

Critical Section

```
// Normal code
Enter CS
// Critical work
Exit CS
// Normal code
```

Mutual exclusion: if process is executing in critical section, all other processes are prevented from entering the section

Progress: if no process is in a critical section, one of the waiting processes should be granted access

Bounded wait: after a process requests to enter the critical section, there exists an upper bound on how many other processes can enter the section before the process

Independence: process not executing in critical section should never block other process

Incorrect Synchronization

- Deadlock: all processes blocked so no progress
- Livelock: related to deadlock avoidance mechanism where processes keep changing state to avoid deadlock and make no other progress; processes are not blocked
- Starvation: some processes never get to make progress in their execution as it is perpetually denied necessary resources

Test and Set

- TestAndSet(Register, MemoryLocation)
- Machine instruction provided by processors to aid synchronization
- Load the current content at MemoryLocation into Register and stores 1 into MemoryLocation (treated as lock)
- Performed as single machine operation (atomic)
- Used to create spinlocks

```
void EnterCS(int* Lock) {
    // Cannot enter if lock value is 1
    while (!TestAndSet(Lock) == 1);
}
void ExitCS(int* Lock) {
    *Lock = 0;
}
```

Busy waiting: keep checking the condition until it is safe to enter critical section which is wasteful use of processing power

Conditions: compare and exchange, atomic swap, load link/store conditional