

**CMP426-697: Operating System**  
***Mid-term Examination (take-home test)***  
Department of Computer Science, Lehman College

Submit your exam report including answers to the questions via email with subject OSMidExam  
to: [gwang.jung@lehman.cuny.edu](mailto:gwang.jung@lehman.cuny.edu)

**Due by April 7, 2017**

**[Q.1] Answer the following questions**

1. Which of the following items are shared across multiple threads belonging to the same process?

- A) code, data, files
- B) registers
- C) stack
- D) all of the above

2. In a multithreaded server architecture, which of the following is used to service a new user request?

- A) a new created thread
- B) a new created process
- C) the same process for prior users
- D) none of the above

3. Which of the following benefits go to multithreaded programming?

- A) responsiveness
- B) resource sharing
- C) economy
- D) scalability
- E) all of the above

4. Which of the following refers to the capability to allow multiple tasks make progress on a single processor system?

- A) concurrency
- B) parallelism
- C) data parallelism
- D) task parallelism

5. \_\_\_\_\_ is a formula that identifies potential performance gains from adding additional computing cores to an application that has a parallel and serial component.

- A) Task parallelism
- B) Data parallelism
- C) Data splitting
- D) Amdahl's Law

6. \_\_\_\_\_ is not considered a challenge when designing applications for multicore systems.
- A) Deciding which activities can be run in parallel
  - B) Ensuring there is a sufficient number of cores
  - C) Determining if data can be separated so that it is accessed on separate cores
  - D) Identifying data dependencies between tasks.
7. Which of the following models are possible for the relationship between the user threads and kernel threads?
- A) many-to-one model
  - B) one-to-one model
  - C) many-to-many model
  - D) two-level model
  - E) all of the above
8. \_\_\_\_\_ is a thread library for Solaris that maps many user-level threads to one kernel thread.
- A) Pthreads
  - B) Green threads
  - C) Sthreads
  - D) Java threads
9. The \_\_\_\_\_ model maps each user-level thread to one kernel thread.
- A) many-to-many
  - B) two-level
  - C) one-to-one
  - D) many-to-one
10. The \_\_\_\_\_ model multiplexes many user-level threads to a smaller or equal number of kernel threads.
- A) many-to-many
  - B) two-level
  - C) one-to-one
  - D) many-to-one
11. Which of the following is a function that can be provided by Pthreads API for constructing a multithreaded program?
- A) pthread\_attr\_t
  - B) pthread\_create
  - C) pthread\_join
  - D) all of the above
12. Which of the following statements regarding threads is false?
- A) Sharing is automatically provided in Java threads.
  - B) Both Pthreads and WinAPI threads share global data.
  - C) The start() method actually creates a thread in the Java virtual machine.
  - D) The Java method join() provides similar functionality as the WaitForSingleObject in Win32.

13. The most common technique for writing multithreaded Java programs is \_\_\_\_\_.  
A) extending the Thread class and overriding the run() method  
B) implementing the Runnable interface and defining its run() method  
C) designing your own Thread class  
D) using the CreateThread() function
14. Which of the following is a method for implicit threading?  
A) thread pools  
B) OpenMP  
C) all of the above
15. Which of the following options exist to deliver signals in multithreaded program?  
A) deliver the signal to the thread to which the signal applies  
B) deliver the signal to every thread in the process  
C) deliver the signal to certain threads in the process  
D) assign a specific thread to receive all signals for the process  
E) all of the above
16. Which of the following options to deliver signals in multithreaded program should be applied to an asynchronous signal that terminates a process (<control><C>, for example)?  
A) deliver the signal to the thread to which the signal applies  
B) deliver the signal to every thread in the process  
C) deliver the signal to certain threads in the process  
D) assign a specific thread to receive all signals for the process  
E) all of the above
17. Which of the following options to deliver signals in multithreaded program should be applied to a synchronous signal?  
A) deliver the signal to the thread to which the signal applies  
B) deliver the signal to every thread in the process  
C) deliver the signal to certain threads in the process  
D) assign a specific thread to receive all signals for the process  
E) all of the above
18. To associate each thread created using an implicit technique such as a thread pool, with its unique transaction identifier, we could use \_\_\_\_?  
A) global variable  
B) local variable  
C) static data  
D) thread-local storage
19. LWP is \_\_\_\_\_.  
A) short for lightweight processor  
B) placed between system and kernel threads

- C) placed between user and kernel threads
  - D) common in systems implementing one-to-one multithreading models
20. Which are included in the context of a thread?
- A) register set
  - B) stacks
  - C) private storage area
  - D) all of the above
21. Which of the following information is shared when the flag CLONE\_VM in Linux clone() system call is set up?
- A) file-system information
  - B) memory space
  - C) signal handlers
  - D) set of open files
22. The ready queue can be implemented as a \_\_\_\_\_.
- A) FIFO queue
  - B) priority queue
  - C) tree
  - D) unordered linked list
  - E) all of the above
23. Which of the following circumstances can cooperative scheduling take place?
- A) when a process switches from the running state to the waiting state
  - B) when a process switches from the running state to the ready state
  - C) when a process switches from the waiting state to the ready state
  - D) none of the above
24. Which of the following circumstances can preemptive scheduling take place?
- A) when a process switches from the running state to the waiting state
  - B) when a process switches from the waiting state to the ready state
  - C) when a process terminates
  - D) none of the above
25. Which of the following items does not belong to the function of a dispatcher?
- A) switching context from one process to another
  - B) selecting a process among the available ones in the ready queue
  - C) switching to user mode
  - D) jumping to the proper location in the user program to resume that program
26. Assume process  $P_0$  and  $P_1$  are the process before and after a context switch, and  $PCB_0$  and  $PCB_1$  are respectively their process control block. Which of the following time units are included inside the dispatch latency?
- A)  $P_0$  executing

- B) save state into PCB<sub>0</sub>, and restore state from PCB<sub>1</sub>
- C) P<sub>1</sub> executing
- D) all of the above

27. Which of the following criteria is more important for an interactive system?

- A) CPU utilization
- B) Response time
- C) Turnaround time
- D) Throughput

28. Which of the following criteria is more important from the point of view of a particular process?

- A) CPU utilization
- B) Response time
- C) Turnaround time
- D) Throughput

29. For interactive systems, it is more important to minimize \_\_\_\_\_.

- A) the average response time
- B) the average turnaround time
- C) the variance in response time
- D) the waiting time

30. Which of the following scheduling algorithm may suffer from convoy effect?

- A) SJF
- B) FCFS
- C) RR
- D) Multilevel queue

31. Which of the following scheduling algorithms must be nonpreemptive?

- A) SJF
- B) RR
- C) FCFS
- D) priority algorithms

32. I/O-bound program typically has many short \_\_\_\_\_ and a CPU-bound program might have a few long \_\_\_\_\_.

- A) I/O burst, CPU burst
- B) I/O burst, I/O burst
- C) CPU burst, CPU burst
- D) CPU burst, I/O burst

33. The \_\_\_\_\_ occurs in first-come-first-served scheduling when a process with a long CPU burst occupies the CPU.

- A) dispatch latency

- B) waiting time
- C) convoy effect
- D) system-contention scope

34. Which of the following scheduling algorithms gives the minimum average waiting time for a given set of processes?

- A) SJF
- B) FCFS
- C) RR
- D) Multilevel queue

35. Shortest-remaining-time-first scheduling is the preemptive version of \_\_\_\_\_?

- A) SJF
- B) FCFS
- C) RR
- D) Multilevel queue

36. Which of the following scheduling algorithms gives the minimum average response time?

- A) SJF
- B) FCFS
- C) RR
- D) Multilevel queue

37. A significant problem with priority scheduling algorithms is \_\_\_\_\_.

- A) complexity
- B) starvation
- C) determining the length of the next CPU burst
- D) determining the length of the time quantum

38. If the time quantum get too large, RR scheduling degenerates to \_\_\_\_\_?

- A) SJF
- B) FCFS
- C) Shortest-remaining-time-first
- D) Multilevel queue

39. Which of the following can be a solution to the problem of indefinite blockage of low-priority processes?

- A) Aging
- B) Starvation
- C) Multilevel queue
- D) All of the above

40. Which of the following processes usually have highest priority?

- A) real-time processes
- B) system processes

- C) interactive processes
- D) batch processes

41. Which of the following is allowed on Linux and Mac OS X systems?

- A) only PTHREAD\_SCOPE\_PROCESS
- B) only PTHREAD\_SCOPE\_SYSTEM
- C) Both PTHREAD\_SCOPE\_PROCESS and PTHREAD\_SCOPE\_SYSTEM
- D) none of the above

42. Which of the following system architectures involves multiprocessor scheduling?

- A) multicore CPUs
- B) multithreaded cores
- C) NUMA systems
- D) heterogeneous multiprocessing
- E) all of the above

43. Which of the following is preemptive?

- A) rate-monotonic scheduling
- B) earliest-deadline-first scheduling
- C) both of the above
- D) none of the above

44. Which of the following POSIX API sets the scheduling policy?

- A) pthread\_attr\_getsched\_policy
- B) pthread\_attr\_setsched\_policy
- C) pthread\_attr\_getscope
- D) pthread\_attr\_setscope

**[Q.2] Answer the following questions**

- (a) Why should a web server not run as a single-threaded process?
- (b) List the four major categories of the benefits of multithreaded programming. Briefly explain each.
- (c) Distinguish between parallelism and concurrency.
- (d) Multicore systems present certain challenges for multithreaded programming. Briefly describe these challenges.
- (e) Distinguish between data and task parallelism.
- (f) What are the two different ways in which a thread library could be implemented?
- (g) Explain the concept of a CPU-I/O burst cycle.

- (h) What role does the dispatcher play in CPU scheduling?
- (i) Explain the difference between response time and turnaround time. These times are both used to measure the effectiveness of scheduling schemes.
- (j) Explain the process of starvation and how aging can be used to prevent it.
- (k) What effect does the size of the time quantum have on the performance of an RR algorithm?
- (l) Explain the fundamental difference between asymmetric and symmetric multiprocessing.
- (m) Describe two general approaches to load balancing.

**[Q.3] Answer the following questions**

Process	Burst Time	Priority
$P_1$	8	3
$P_2$	4	1
$P_3$	2	5
$P_4$	3	4
$P_5$	5	2

- (a) Consider the set of processes shown in the table above, with the length of the CPU-burst time given in milliseconds. The processes are assumed to have arrived in the order  $P_1, P_2, P_3, P_4,$  and  $P_5,$  all approximately at time 0. Draw three Gantt charts illustrating the execution of these processes using FCFS, round-robin (quantum size = 5), and a non-preemptive priority (a smaller priority number implies a higher priority) scheduling. Calculate average waiting time of each scheduling.
- (b) What is convoy effect?
- (c) Explain the difference multilevel queue scheduling and multilevel feedback queue scheduling