Midterm Class Statistics

Number of submitted grades: 165 / 171

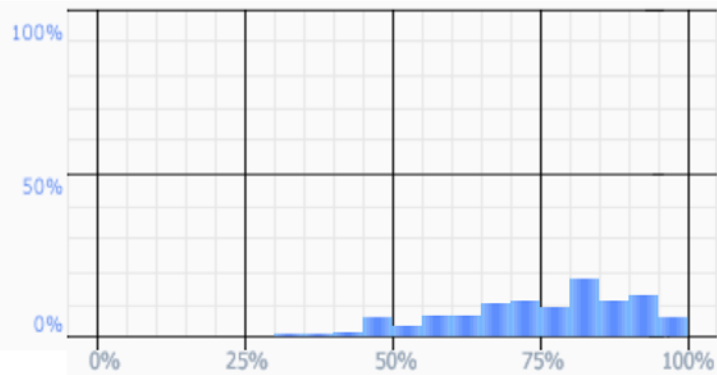|		
Minimum:	🟥	33.33 %
Maximum:	🟦	100 %
Average:	🟨	74.81 %

Mode: 83.33 %, 90 %

Median: 76.67 %

Standard Deviation: 14.7 % ❓

Grade Distribution

Number of Users (%)

Multiple Choice Questions

1. Which of the following about real time systems is not correct?

a) Classification of real time systems depends on how to define cost of missing deadline.

b) It requires a privileged instruction to run insmod.

c) Soft real-time tasks are those which have elastic deadlines.

d) There are three types of real-time applications. They are hard, firm and soft real-time applications.

C

2. What of the following is shared by multiple threads in the same process?

(a) program counter (PC)

(b) stack

(c) code

(d) Register

C

3. Which of the following statements is NOT correct?

a) Using threads to implement multiple tasks is more efficient than using processes. One of the reasons is that it is much quicker to create a thread in a process.

b) The principal objective of batch processing systems is to minimize response time.

c) Code written for the Linux kernel space should not have a main() function.

d) Running insmod requires superuser privileges

B

4. Which of the following is NOT a reason that a user program rarely makes a raw system call?

a) APIs provided in library achieve better security protection than raw system calls.

b) APIs provided in library is easy to use compared to raw system calls.

c) APIs provided in library has better portability than raw system calls

d) System calls are only available in kernel mode, so a user program cannot access raw system calls.

D

5. Given the following three task representations:

T1 (0; 10; 3; 10);

T2 (10; 3; 10);

T3 (10; 3).

Which of the following is correct?

a) They all represent the task with same parameters.

b) T1 is a special case of T2, and T2 is a special case of T3.

c) T2 represents a task which has an arbitrary phase.

d) None of the above.

A

6. In Unix, which system call creates a new process?

a) fork

b) create

c) new

d) none of the mentioned

A

7. What is the output of g value if the following code is given:

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

int g = 0;

void *aThread()
{
    g++;

```
    pthread_exit(NULL);
}


int main (int argc, char *argv[])
{
    int i;
    pthread_t thread[3];
    for (i=0; i<3; i++)
    {
        if( pthread_create( thread+i, NULL, aThread, NULL) )
        {
            printf("ERROR; return code from pthread_create()\n");
            return -1;
        }
    }
    printf("The value of g is %d\n", g);


    return 0;
}
```

a)   The value of g is 1

b)   The value of g is 2

c)   The value of g is 3

d)   There is a race condition, so the value of g cannot be determined.

D

8. Which of the following is a drawback of thread programming

a) Context switch time among different threads can be large.

b) Tasks implemented by threads are less efficient than tasks implemented by processes.

c) Without a synchronization, race condition on shared variables can be disastrous.

d) It does not share code among threads in the same process.

C

9. In UNIX, the _____ system call is used to send a signal.

a) signal

b) send

c) kill

d) sigsend

C

10. Which of the following output is not possible?

```
#include <pthread.h>

#include <stdio.h>

#include <stdlib.h>


int MAX = 10;

int count = 0;


void *even(void *arg)

{

   printf("This is even thread()\n");

   while(count < MAX)

     if(count % 2 == 0)

        printf("%d ", count++);

   pthread_exit(0);

}


void *odd(void *arg)
```

```c
{
    printf("This is odd thread()\n");

    while(count < MAX)

        if(count % 2 == 1)

            printf("%d ", count++);

    pthread_exit(0);

}


int main()

{

    pthread_t t1;

    pthread_t t2;


    pthread_create(&t1;, 0, &even;, NULL);

    pthread_create(&t2;, 0, &odd;, NULL);


    pthread_join(t1, 0);

    pthread_join(t2, 0);


    return 0;

}
```

A. 1 0 2 3 4 5 6 7 9 8

B. 1 0 2 4 3 5 6 7 8 9

C. 0 1 2 3 4 6 5 7 8 9

D. 1 0 2 4 5 3 6 7 9 8

D

11. What is the Output of the following program? Note that the SIGCHLD signal is sent to the parent of a child process when it exits, is interrupted, or resumes after being interrupted.

```c
#include<stdio.h>

#include<signal.h>

#include<sys/wait.h>


#define N 4

int val = 9;

void handler(sig) {

  val += 3;

  return;

}


int main() {

 pid_t pid;

 int i;

 signal(SIGCHLD,handler);

 for (i=0;i<N;i++) {

  if ((pid =fork()) == 0) {

    val -= 3;

    exit(0);

  }

 }

 for (i=0;i<N;i++) {

  waitpid(-1,NULL,0); //suspends execution of the calling process until any child has changed state.

 }

 printf("val = %d\n",val);

}
```

A. 9

B. 12

C. 18

D. 21

E. None of the above

D

12. Given the following three real-time tasks: T1(1, 6, 2, 8), T2(2, 10, 2, 10), and T3(12, 2, 10) in a system. What is the CPU utilization of the system if no other tasks exist?

A. 65%

B. 69.3%

C. 61.7%

D. 70%

D.

13. In cyclic executive scheduling, which of the following statement is correct.

A. CE scheduling algorithm is based on static priority of tasks. A task's priority is inverse proportional to its period.

B. A drawback of CE scheduling algorithm is that you have to compute an offline schedule which can be arbitarily long.

C. The framesize of the CE algorithm cannot be to small since we want an instance of a task is completed within a single frame.

D. The framesize of the CE algorithm cannot be to large, otherwise the computational complexity is too high.

C.

14. Given the task set: T1(4, 1); T2(6, 1); T3(12, 2). What is the largest suitable frame size f?

A. f=3

B. f=4

C. f=6

D. f=12

B.

15. Given the task set: T1(4, 1); T2(5, 1); T3(10, 2), and you are asked to find a Cyclic Executive (CE) schedule using flow graph. Assume that the frame size is 2. Which of the following is true?

A. There are 11 job nodes and 10 frame nodes in the flow graph.

B. There are 14 job nodes and 11 frame nodes in the flow graph.

C. Using the network flow to model CE scheduling cannot handle the task split in CE.

D. The algorithms finding the solution to network flow problem is NP-complete.

A

16. A Task set consists of 5 pre-emptive and periodic tasks. If the task set is NOT RM schedulable which of the following can be inferred?

A. The CPU utilization is over 0.743

B. The CPU utilization is over 1

C. The CPU utilization is less than 0.693

D. None of the above is correct

A

17. Given 3 periodic tasks T1, T2, and T3. They have the same execution time, but different periods. If the periods of the tasks are 4, 8, and 16. What is the maximum execution time so that the 3 tasks are RM schedulable?

A. 2

B. 2.28

C. 2.75

D. 3

B

18. Which of the following statements about RM is correct.

A. RM algorithm applies only when the periods of tasks are all integers.

B. Sufficient RM schedulability condition is based on processor utilization. It can be used at run-time to predict the scheduability, but may lead to poor processor utilization.

C. If the overall utilization of 3 periodic real time tasks is larger than 78%, the 3 tasks are not RM schedulable.

D. RM is a dynamic priority algorithm.

B

19. Consider the following set of independent, preemptable periodic tasks: T1(50; 12); T2(40; 10); T3(30; 10). Using Rate Monotonic(RM) scheduling algorithm, which of the following is ture.

A. The 3 tasks are schedulable, since the utilization is less than 100% according to scheduability test 1.

B. The 3 tasks are not schedulable, since the utilization is over 78% according to scheduability test 2.

C. The 3 tasks are schedulable according to the scheduability test 3.

D. The 3 tasks are not schedulable according to the scheduability test 3.

D.

20. Consider the task set T = {(8, 4), (10, 2), (12, 2)}, which of the following is true:

A. The task set is not RM schedulable

B. The task set is RM schedulable

B

Written Response Questions

21. Given real time tasks T1(4, 2) and T2(8, 3), reduce the cyclic executive (CE) scheduling problem as a network flow problem, assuming the frame size is 4. Then get a maximum attainable flow on the network, and give the corresponding CE schedule based on the directed graph, using "I" represents "idle" (4 points).

22. T1(5.5; 20; 4; 20[R3; 2]); T2(2; 22; 5; 22[R2; 3[R1; 1]]); T3(0; 24; 7; 24[R1; 5.5[R2; 2.5]]) where the information in square brackets shows the resource requirements of different processes. [3+3 points]

Graphically show a schedule for this task set based on RM algorithm and using priority ceiling protocol for resource access control. Given that: Task T3 requests a lock on resource R1, one time unit after start of its execution. It needs R2 after using R1 for 2 units of time during which it holds R1 and finally uses R1 for one time unit again (after finishing with R2). Thus R1 is held by this task for a total of 5.5 units of time.

Task T2 executes a lock on R2, 1.5 time units after start of its execution, uses it for 0.5 units of time, then needs R1 for one time unit while holding R2 which is used again for 1.5 time units before it is released. Thus this task holds R2 for a total of 3 units of time.

Task T1 executes a lock of resource R3, one time after start of its execution.
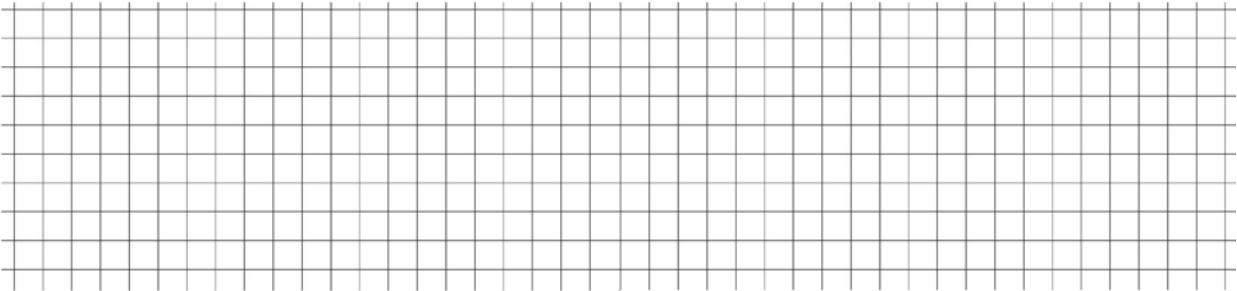
Ans:

$$\pi(R_1) = \pi(R_2) = 2, \pi(R_3) = 1$$
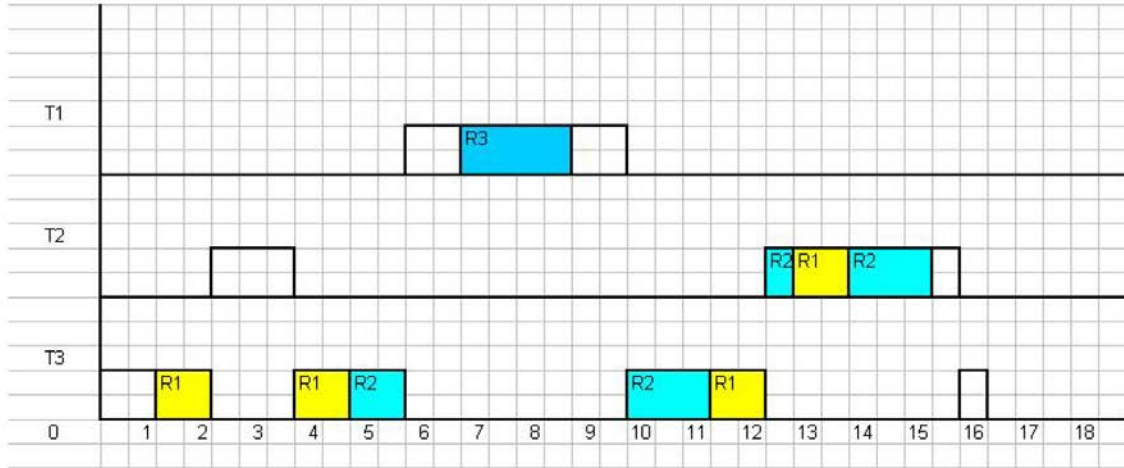
The schedule is shown in figure 3

Figure 3: