

i About The Exam Questions

FINAL EXAM

Course: ITF23019 Parallel and Distributed Programming

Examination Date: 04.05.2023 14:00 2.25 hours

Lecturer: Nga Dinh

Permitted Aids: None

About the examination set

This exam contains 8 sections each of which consists of several questions. There are multiple-choice questions, True/False questions, and essay questions.

- For each multiple-choice question, you will answer by selecting one alternative from several alternatives. A correct answer is awarded 3 points or 2 points depending on particular section. A wrong answer will get 0 point.
- For each True/False question, a correct answer is awarded 2 points and a wrong answer will get 0 point.
- For essay questions, the total points in each section vary and are specified in that section. If a section contains several essay questions, the points for each question are divided fairly among them. Your answers to these questions must be written in the text box provided.

Most of the answers to the exam questions can be found from class materials. However, for some questions, you have to answer by your own understanding from this course. You have **2 hours and 15 minutes** to complete 100 points worth of the exam questions.

The difficulty of each question and, for essay questions, the expected length of responses are randomly distributed across the 8 sections. Therefore, you should allocate your time accordingly.

The results will be announced in Studentweb May 25, 2023.

Good luck!

1 Parallel Hardware

1. Cache coherence happens in which of the following architectures?

Select one alternative:

- Von Neumann Architecture
- Shared-memory system
- Distributed-memory system

2. Supposed that there is one program with five computations and one data input. The data input is divided into five smaller data items. Now that there are five cores to perform five computations on the five data items. Which of the following is task-parallelism?

Select one alternative

- Each core executes one computation on five data items.
- Each core executes five computations on one data item.

3. Supposed that there is one big problem which is divided into a number of smaller sub-problems. Which of the following statements is true about distributed programming?

Select one alternative

- Each sub-problem is executed by a different core in a computer and all sub-problems are executed simultaneously.
- Each sub-problem is executed alternatively in a time-sharing fashion.
- Each sub-problem is executed in a different computer in a network and all sub-problems are executed simultaneously.
- Each sub-problem is executed one after another.

4. Which of the following statements is NOT true about cache memory?

Select one alternative

- It helps to reduce the bottleneck problem in Von Neumann architecture.
- It has higher capacity than the main memory.
- An decision has to make to decide which data and instructions should be stored in the cache.
- It is a collection of memory locations that the CPU can access more quickly than it can access main memory.

5. Which of the following statements is NOT true about share-memory architecture?

Select one alternative

- Each core can read or write to the a shared variable.
- Using message passing is the only way for different cores to communicate with each other.
- We can coordinate cores by having them examine and update shared-memory location.
- Several cores share access the main memory.

6. Which of the following architectures can be used for MIMD systems?

Select one alternative

- Shared-memory architecture and distributed memory.
- Distributed-memory architecture.
- Shared-memory architecture.
- The Von Neumann Architecture.

7. Which of the following statements is correct about pipe-lining?


Select one alternative

- Pipe-lining makes a CPU run faster by overlapping the executions of different block of code.
- Parallelism is archived by having multiple instructions simultaneously initiated.
- An instruction is divided into several stages and different stages of different instructions can run simultaneously.
- It is not an instruction-level parallelism.

Maximum marks: 21

2 Foster Methodology

1. List the name of four steps in Forster's methodology to design a parallel program.
2. Explain the Communication step within 10 sentences.
3. One program is divided into the number of Tasks, each of which is assigned to a processor as shown in the figure below (where P stands for Processor). Tasks at the same level can run in parallel. Which processors need to communicate with each other to complete their work? Additionally, is this mapping scheme optimal?



Fill in your answer here

Format
B
I
U
x₂
x²
I_x
📄
📂
↶
↷
🔄
☰
☷
Ω
📊
✎
Σ

Words: 0

Maximum marks: 9

3 Distributed System - Fundamental Concepts

1. A port number is used to identify a unique service running in a computer.

Select one alternative:

False

True

2. Resource sharing is one advantage of a distributed system.

Select an alternative

True

False

3. There is a common clock to synchronize activities of different processes running on different computers in a distributed system.

Select an alternative

True

False

4. Share variables are used to synchronize activities of different processes belonging to a program in a distributed system.

Select an alternative

True

False

5. In a distributed system, different tasks at different computers cannot be executed at the same time.

Select an alternative

True

False

6. Transport control protocol (TCP) and User Datagram Protocol (UDP) are in the Application layer of network software stack.

Select an alternative

- True
- False

7. User Datagram Protocol (UDP) is connection-oriented.

Select an alternative

- True
- False

8. In a synchronous communication, both sending and receiving operations are blocking operations.

Select an alternative

- False
- True

9. In a distributed system, a message is divided into a number of packets before being transmitted.

Select an alternative

- False
- True

10. Network software is arranged in a hierarchy of layers.

Select an alternative











- True
- False


Maximum marks: 20

4 Interprocess Communication.

1. Explain about Transport Control Protocol (TCP).
2. What are the advantages and disadvantages of the TCP?
3. What types of applications should TCP be used for?

Fill in your answer here

Format ▾ | **B** *I* U x_2 x^2 | I_x |   |   |   |   |  |  |



Words: 0

Maximum marks: 10

5 Mobile Computing

1. Which of the following statements is NOT true about Mobile Computing?

Select one alternative:

- Users have to connect to the network using a network cable.
- Portability is one aspect of Mobile Computing.
- Users are able to use a computing device while moving.
- Users can connect to the network while moving.

2. Which of the following statements is NOT true about Mobile Computing?

Select one alternative

- Mobile Computing refers to the use of portable computing devices such as smartphones, tablets, and wearables, to access information and services.
- It exploits the development of wireless technology and hardware miniaturization.
- It allows the transmission of multimedia data via a wireless device.
- A device can connect to the network using a wireless connection but it has to be fixed in location.

3. Which of the following forms is NOT a relevant form of a volatile system?

Select one alternative

- Users use their fitness tracker to monitor their daily physical activity and sleep patterns.
- Changes in the bandwidth of wireless connections in the system.
- A new user joins the system.
- Failures of devices in the system.

4. Which of the following statements is NOT true about context-aware computing?

Select one alternative

- The system uses sensors to collect data about the environment.
- The system changes its behavior based on the context.
- The system cannot combine data from several sensors.
- The system analyzes the collected data to understand the context.

5. Which of the following is NOT a sub-field of Mobile and Ubiquitous Computing?

Select one alternative

- Context-aware Computing
- Ubiquitous Computing
- Parallel Computing
- Wearable Computing

Maximum marks: 10

6 Cloud Computing

1. Which of the following statement is NOT correct about Cloud Computing?

Select one alternative:

- It can reduce operation cost for a company.
- With Cloud Computing services, users can store their data in the Internet.
- It provides ability to users to implement services without having resources installed on premise.
- It does not allows users to access their data in the Internet.

2. When using Cloud Computing services from a service provide, users have to pay for

Select one alternative

- All other options
- Network usage
- Dat storage
- Data processing

3. Which of the following types allows Cloud Computing systems to be accessible by only one organization?

Select one alternative

- Community Cloud
- Hybrid Cloud
- Private Cloud
- Public Cloud

4. Which of the following is the pricing model for Cloud Computing services?

Select one alternative

- Fixed amount per month
- Prepaid
- Free
- Pay- per- use

5. Which of the following services does not belong to SaaS?

Select one alternative

- Dropbox
- Email
- Zoom
- Operating System

Maximum marks: 10

7 Programming - Part 1

1. Which of the following statements is correct about ReadWriteLock?

Select one alternative:

- One write operation and multiple read operations can be performed simultaneously.
- One read operation and multiple write operations can be performed simultaneously.
- Multiple read operations can be performed simultaneously.
- Multiple write operations can be performed simultaneously.

2. Which of the following statements is correct about Broadcast in MPI?

Select one alternative

- The root process sends the same message to all processes in a communicator.
- The root process sends different message to a group of processes in a communicator.
- It is a point - to - point communication.
- The root process combines message from other processes in a communicator into a single message.

3. Which of the following statements is correct about the client-server model?

Select one alternative

- The client offers the service.
- Either TCP or UDP protocol is used in the transport layer to deliver data packets.
- The server runs several services on the same port.
- The server sends a socket which includes the server's IP address and port number to the client.

4. Which of the following is NOT correct about atomic variables?

Select one alternative

- When a thread is doing an operation with an atomic variable, other threads can modify that variable.
- Using atomic variables in Java is a method to avoid the data race problem in a multithread program.
- In some cases, atomic variables offer a better performance than other synchronization mechanisms such as locks.
- Atomic variables in Java provide provide atomic actions on single variables.

5. Which of the following is NOT correct about the Critical Section?

Select one alternative

- Critical section is a space in the main memory that can be accessed by multiple threads in a multithread program.
- The number of threads which can execute codes in the critical section is limited.
- The access to the share-memory area has to be in the critical section to avoid the data race problem.
- Making a critical section is a solution to solve data race problem in a multithread program.

Maximum marks: 10

8 Programming - Part 2

1. Is the deadlock in the following program or not? Explain why or why not.

```
public class TestThread {
    public static Object resource1 = new Object();
    public static Object resource2 = new Object();

    public static void main(String args[]) {
        ThreadDemo1 T1 = new ThreadDemo1();
        ThreadDemo2 T2 = new ThreadDemo2();
        T1.start();
        T2.start();
    }

    private static class ThreadDemo1 extends Thread {
        public void run() {
            synchronized (resource1) {
                System.out.println("Thread 1: Holds resource 1...");

                try {
                    Thread.sleep(10);
                } catch (InterruptedException e) {}
                System.out.println("Thread 1: Waits for resource 2...");

                synchronized (resource2) {
                    System.out.println("Thread 1: Holds resource 1 and 2...");
                }
            }
        }
    }

    private static class ThreadDemo2 extends Thread {
        public void run() {
            synchronized (resource1) {
                System.out.println("Thread 2: Holds resource 1...");

                try {
                    Thread.sleep(10);
                } catch (InterruptedException e) {}
                System.out.println("Thread 2: Waits for resource 2...");













                synchronized (resource2) {
                    System.out.println("Thread 2: Holds resource 1 and 2...");
                }
            }
        }
    }
}
```


2. In the following snippet of code, the synchronized keyword is used to implement mutual exclusion in a multithread program. Supposed that the program runs correctly. How many critical section does the program create? How many methods can be executed at the time?

```
public class MonitorSynchronization {
```

```
public synchronized void set(){  
    ...  
}  
  
public synchronized void wait(){  
    ...  
}  
  
public synchronized void get(){  
    ...  
}  
}
```

Fill in your answer here

Format ▾ | **B** | *I* | U | x_2 | x^2 | I_x |  |  |  |  |  |  |  |  |  |  |  |  |



Words: 0










Maximum marks: 5

9 Logical Reasoning

Supposed that you are designing a multithread program in a shared-memory system to reduce the program execution time and that you have to use a synchronization mechanism to synchronize the threads. In doing so, you will need to divide the problem being solved into a number of tasks.

Should the problem be divided into a smaller number of larger tasks with low intercommunication (among the tasks) or a larger number of smaller tasks with high intercommunication? What factors should be considered when making this decision?

Fill in your answer here

Format | **B** | *I* | U | x_2 | x^2 | I_x |  |  |  |  |  |  | Ω |  |  | Σ | 










Words: 0


Maximum marks: 5

10 Your comments

If you have any comments regarding your answers to the exam questions, please write down in the text box below.

Your comments:

Format | **B** | *I* | U | x_2 | x^2 | I_x |  |  |  |  |  |  |  |  |  |

Σ | 

Words: 0

Maximum marks: 0