

Distributed Systems (Spring 2020)

Task 2: processes and clock synchronization

Practical information:

Due date: Monday, April 13, noon

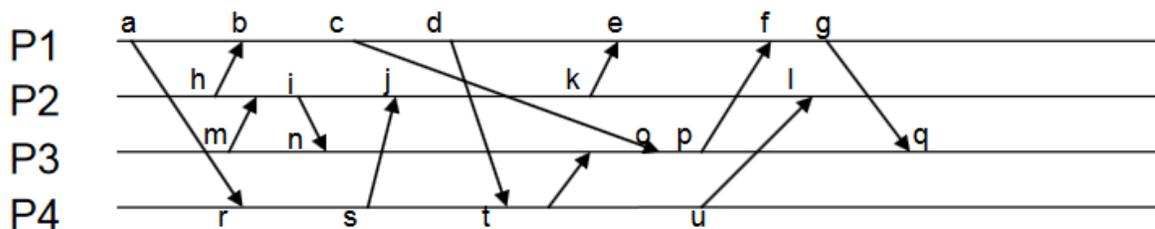
Each hour late after midnight is 10% less points

- After 6 hours, value of the task is zero points
- This task can be done in a team of max 3 persons ONLY.

Submit your solution to mohan.liyanage@ut.ee, (CC) huber.flores@ut.ee

Instructions: Please respond the following questions. Be clear and concise in your answers. Provide enough explanation to support your arguments. Ambiguous answers to fill up space are considered wrong and no points are granted.

1. Consider *Figure 1* that shows four processes (P1, P2, P3, P4) with events a, b, c,... and messages communicating between them. Assume that initial logical clock values are all initialized to 0 (20 points)
 - (5 Points) List the Lamport timestamps for each event shown in Figure 1. Assume that each process maintains a logical clock as a single integer value as a Lamport clock. Provide timestamps for each labeled event.
 - (10 Points) List the Vector Clock timestamps for each event shown in Figure 1. Provide timestamps for each labeled event.
 - (3 Points) Is there the potential for a causal violation? Explain why.
 - (2 Points) List the processes that are concurrent and explain detail why?



- Figure 1: 4 processes P1, P2, P3, P4 run events a,b,c,d,... to send and receive messages
2. Measurements taken during one hour from a Web server indicate that the utilization of the CPU and the two disks are: $U_{CPU} = 0.25$, $U_{disk1} = 0.35$, and $U_{disk2} = 0.30$. The Web server log shows that 41,600 requests were processed during the measurement interval. What are the service demands at the CPU and both disks, what is the maximum throughput, and what was the response time of the Web server during the measurement interval? (5 points)
 3. A computer system is measured for 30 minutes. During this time, 7,500 transactions are completed and 20,900 I/O operations are executed on a certain disk that is 35% utilized. What is

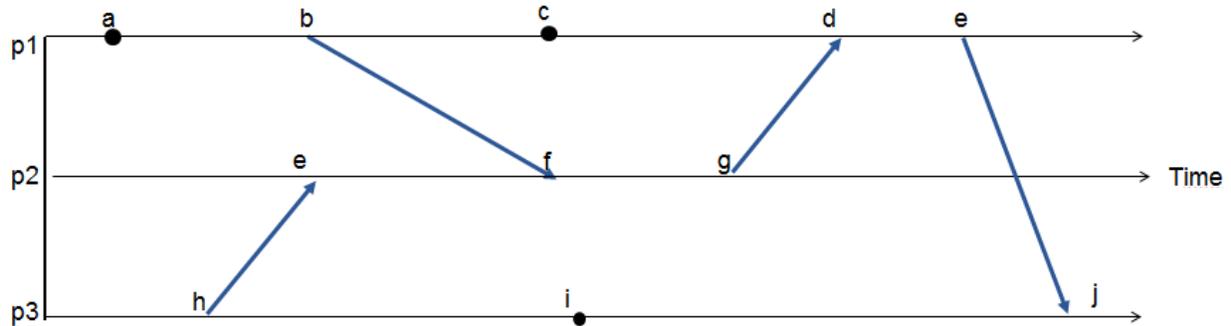
the average number of I/O operations per transaction on this disk? What is the average service time per transaction on this disk? (5 points)

4. Get acquaintance with the Java reflection utility that comes with any JVM. Implement code migration at method level of any sorting algorithm, e.g., bubble sort, quick sort, etc. The migration needs to happen via sockets between a client and a server. In this process, the client captures the details of code execution during runtime and sends it to the server for processing; the server then takes those details and executes the code. The result is send back to the client, such that it can continue with the execution of the algorithm transparently. (30 points)
 - a. The migration process can be implemented in a single machine, but specify two separate instances, one for a client, and one for the server
 - b. Here you can find some insights about the process (<https://gist.github.com/huberflores/9829019> – This is just to provide you an idea, not to enforce you to follow the same implementation pattern).
 - c. Please provide enough information to execute your program (code, binaries, and instructions). Provide also a short video that shows the execution of your program and an overall description of code.

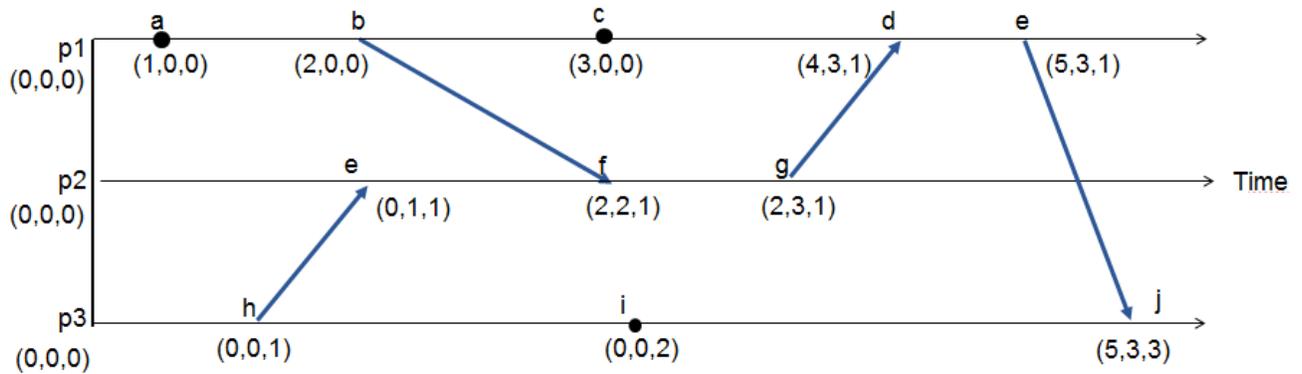
5. A Web server, composed of a single CPU and single disk, was monitored for one hour. The main workload of the server can be divided into HTML files and requests for image files. During the measurement interval 14,040 requests for HTML files and 1,034 requests for image files are processed. An analysis of the Web server log shows that HTML files are 3,000-bytes long and image files are 15,000-bytes long on average. The average disk service time is 12 msec for 1,000-byte blocks. The CPU demand, in seconds, per HTTP request, is given by the expression $CPU_{Demand} = 0.008 + 0.002 \times RequestSize$, where RequestSize is given in the number of 1000-byte blocks processed. This expression for the CPU demand indicates that there is a constant time associated to processing a request (i.e., 0.008 seconds) regardless of the size of the file being requested. This constant time involves opening a TCP connection, analyzing the HTTP request, and opening the requested file. The second component of the CPU demand is proportional to the file size since the CPU is involved in each I/O operation. What is the response time for HTML and image file requests for the current load and for a load five times larger? (10 points)

Tip: Since the workload is characterized as being composed of two types of requests, a two-class queuing network model is required. Should and open or closed model be used?

6. Implement the vector clock algorithm. The program must provide a graphical interface to draw any number of processes P_n , and events (local (point) or message (arrow) events [a...z; A...Z]) (30 points). For instance



Once the intercommunication of processes is drawn, then the program should calculate automatically the vector timestamps associated to each event. The output should look like this



- Provide an executable binary of your program along with the code (You can use the programming language of your preference)
- Provide enough instructions to execute your program
- Provide a short video showing your working program (Using the provided example)