INDIAN INSTITUTE OF TECHNOLOGY KHARAGPUR

Stamp / Signature of the Invigilator

EXAMINATION (Mid Semester)	**SEMESTER (Spring 2023)**

Roll Number										**Section**		**Name**	

Subject Number	C	S	6	0	0	0	2	**Subject Name**	*Distributed Systems*

Department / Center of the Student		**Additional sheets**	

Important Instructions and Guidelines for Students

1. You must occupy your seat as per the Examination Schedule/Sitting Plan.
2. Do not keep mobile phones or any similar electronic gadgets with you even in the switched off mode.
3. Loose papers, class notes, books or any such materials must not be in your possession, even if they are irrelevant to the subject you are taking examination.
4. Data book, codes, graph papers, relevant standard tables/charts or any other materials are allowed only when instructed by the paper-setter.
5. Use of instrument box, pencil box and non-programmable calculator is allowed during the examination. However, exchange of these items or any other papers (including question papers) is not permitted.
6. Write on both sides of the answer script and do not tear off any page. **Use last page(s) of the answer script for rough work.** Report to the invigilator if the answer script has torn or distorted page(s).
7. It is your responsibility to ensure that you have signed the Attendance Sheet. Keep your Admit Card/Identity Card on the desk for checking by the invigilator.
8. You may leave the examination hall for wash room or for drinking water for a very short period. Record your absence from the Examination Hall in the register provided. Smoking and the consumption of any kind of beverages are strictly prohibited inside the Examination Hall.
9. Do not leave the Examination Hall without submitting your answer script to the invigilator. **In any case, you are not allowed to take away the answer script with you.** After the completion of the examination, do not leave the seat until the invigilators collect all the answer scripts.
10. During the examination, either inside or outside the Examination Hall, gathering information from any kind of sources or exchanging information with others or any such attempt will be treated as **'unfair means'**. Do not adopt unfair means and do not indulge in unseemly behavior.

Violation of any of the above instructions may lead to severe punishment.

Signature of the Student

To be filled in by the examiner

Question Number	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**Total**
Marks Obtained											

Marks obtained (in words)	**Signature of the Examiner**	**Signature of the Scrutineer**

Write the answers in the boxes only. You can use the designated spaces for rough works. This question has 10 pages including the space for rough works.

1. Answer the following questions briefly. Marks will be deducted for unnecessary descriptions. No marks will be given if the answers are not explained and only Yes/No answer is given.

[10x3=30 Marks]

(a) Distributed Knowledge $\subseteq$ Common Knowledge – is this true? Explain your answer.

(b) Consider a distributed leader election algorithm over a ring-overlay (the agents are connected over a ring topology) as follows. Assume every agent in the ring has a numeric integer ID. One of the agents initiates leader election by sending a marker to its left agent in the ring, by putting its own ID in the marker. The receiving and sending rules for the marker are as follows.
Marker receiving rule at agent *i*: Check if the ID of the agent *i* is less than the ID inscribed within the marker. If not, replace the Marker ID with its own ID.
Marker sending rule at agent *i*: Send the marker to the agent at its left over the ring-overlay.
Consider that an agent becomes the leader once it receives back a marker with its own ID. Now, say agent *i*, receives a marker with its own ID. At this point, is it a common knowledge that agent *i* is the leader? Consider that the environment is synchronous and all the communications are reliable.

(c) Consider a set of agents over a fully connected overlay. The communication channels are reliable, and the message transmissions have a known upper bound (which is a common knowledge to all the agents). However, at most 50% of the agents can show Byzantine behavior (this upper bound is not a common knowledge) by selectively forwarding a message to a subset of their peers. Can a distributed computation attain a common knowledge in this system only through message passing? Explain your answer.

(d) What is the difference between a fair-loss link and an arbitrary link? How can you convert an arbitrary link to a fair-loss link?

(e) Assume a set of N agents divided into two groups G_1 and G_2. Consider a fact π_1 which is a common knowledge in G_1 (and a distributed knowledge to G_2), and another fact π_2 which is a common knowledge in G_2 (and a distributed knowledge in G_1). Assume that there is at least one agent that belongs to both the groups G_1 and G_2. Is $\pi_1 \cup \pi_2$ a common knowledge for the agents in $G_1 \cup G_2$? Explain your answer.

(f) Consider an overlay with a set of agents, where the message transmission has a known upper bound Δ (a common knowledge), although the communication channel is not reliable. Assume that every agent periodically sends their local states to all other agents through message passing. Can we compute a consistent global state from these messages? Explain your answer.

(g) Consider four events e_1, e_2, e_3 and e_4. Say, $e_1 \rightarrow e_2$ and $e_3 \rightarrow e_4$. e_2 and e_3 are the events executed over the same agent, and say $T(e_2) < T(e_3)$, where $T(e)$ is the local timestamp of an event. Then can we say, $e_1 \rightarrow e_4$? Explain your answer.

(h) Let every event e in a distributed system has been associated with an integer ID $I(e)$. Let e_1 and e_2 be two parallel events as per Lamport's clock. We order e_1 and e_2 as $e_1 \rightarrow e_2$ is $I(e_1) < I(e_2)$. With this notion of ordering the parallel events, can we generate a total order of a distributed system with the help of Lamport's clock? Explain.

(i) Consider a distributed overlay where the communication is non-FIFO. Why can't we apply Chandy-Lamport's algorithm to take a consistent global snapshot of this system? Explain with an example and help of consistent cuts.

(j) Consider a distributed computation where you can associate a Lamport's clock value with all the events in the system, based on the *happen-before* relationship. Can there be an unreachable state in this system? Explain.

2. **Common Knowledge** **[Total 10 Marks]**

(a) In the "*Muddy Children Problem*", consider that there are 5 children out of which 3 are muddy. Write down the execution steps of the algorithm (what happens in each round), and explain why the common knowledge $k \geq 1$ is required to make progress in this algorithm. **[4+2 Marks]**

(b) Consider a sender s and a receiver d in a network, where messages are either delivered instantaneously or are delivered at 3 seconds. However, the receiver does not know whether the message will be delivered instantaneously or at 3 seconds. Can we attain common knowledge to this system? Explain your answer. **[4 Marks]**

3. **Distributed Snapshot** **[Total 20 Marks]**

(a) Consider the following sequence of message communications in a distributed system.

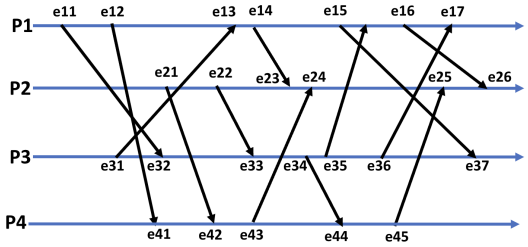

(i) Can we say, $e12 \rightarrow e25$? Explain your answer. **[2 Marks]**

(ii) List down four events, one from each of the four processes, where those four events are parallel events as per Lamport's *happen-before* relationship. **[2 Marks]**

(iii) Write down feasible clock values, for all the events in the above diagram, following the notion of Lamport's logical clock. **[6 Marks]**

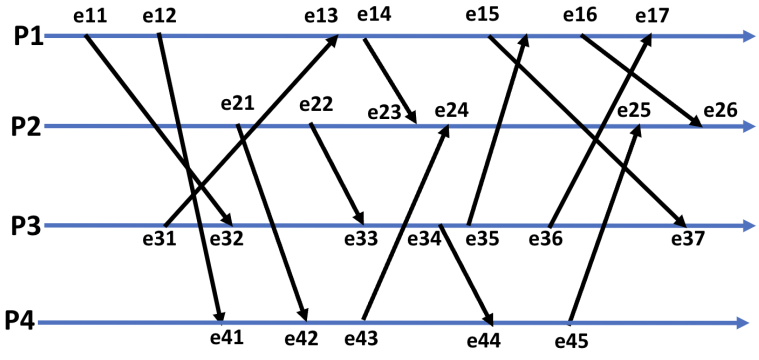

(b) Consider the following distributed system state where the two processes P1 and P4 initiate taking a snapshot (for the same snapshot ID, say S_1) simultaneously. The dotted lines show

the forwarding of the marker messages from P1 and P4 to all other processes. Note that, the processes just record their snapshot on sending and receiving the markers as per the Chandy-Lamport algorithm, and do not forward a marker that it has received from other processes.

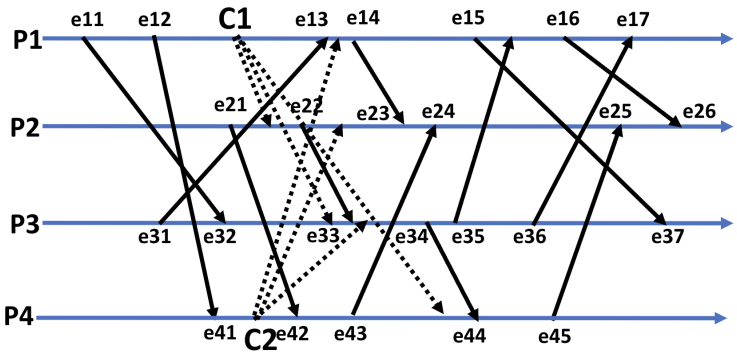

Write down the process and channel states that will be recorded for each of the processes for snapshot S_1. **[6 Marks]**

With the help of consistent cuts, show that Snapshot S_1, as taken above is consistent. Draw the cut, and explain your answer. **[4 Marks]**

Draw the cut on the below diagram.

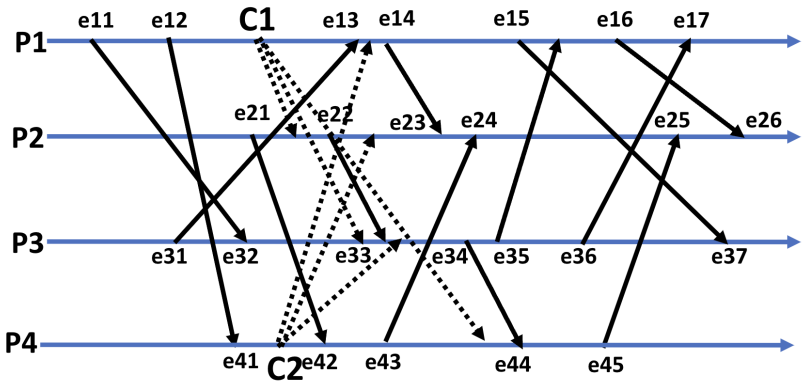

Space for Rough Works

Space for Rough Works