NAME:____

CS 425/ECE 428 Distributed Systems
Mid-Term Exam 1 **Total points: 100**

1. (8 points) This question relates to the "baseball" paper by Doug Terry.

   Suppose that process P1 performs the following write operations in the specified order:

   x := 1
   y := 2
   z : = 3

   Assume that, initially, all variables have value 0.
   Suppose that the above operations have been completed (i.e., P1 has received acknowledgements that these operations are completed).

   At a later time, process P1 performs the following operation:

   a := x + y

   Assume that no other operations performed.

   Under each of the consistency models below, what are the different values possible for variable *a*, after the above update of variable *a* has been completed?

   (a) Eventual consistency (i.e., without read-my-write guarantee)

   (b) Read-my-write

2. (6 points) State true or false:

   (a) In a point-to-point synchronous network of 3 nodes, it is impossible to solve Byzantine Generals problem if one of the nodes may suffer Byzantine failure.

   (b) Suppose that V(e) and L(e) denote the vector and Lamport timestamps of event e, respectively. For any two events e and f, **if** V(e)<V(f) **then** L(e)<L(f).

   (c) Consensus is impossible in an asynchronous system wherein nodes may crash.

3. (15 points) In each part of this question, **<u>if you answer NO</u>, then delete a <u>minimum</u> number of operations to ensure that the modified execution will satisfy the specified property – circle the operations that you want to delete.**

Assume that all variables are initialized to **0**.

(a) Is the execution below linearizable?

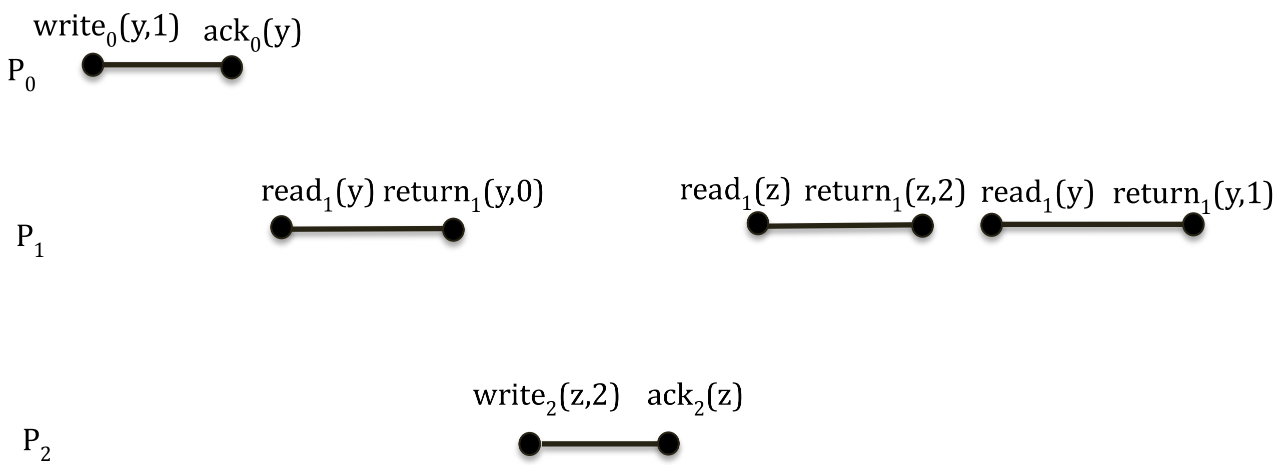

(b) Is the execution below sequentially consistent?

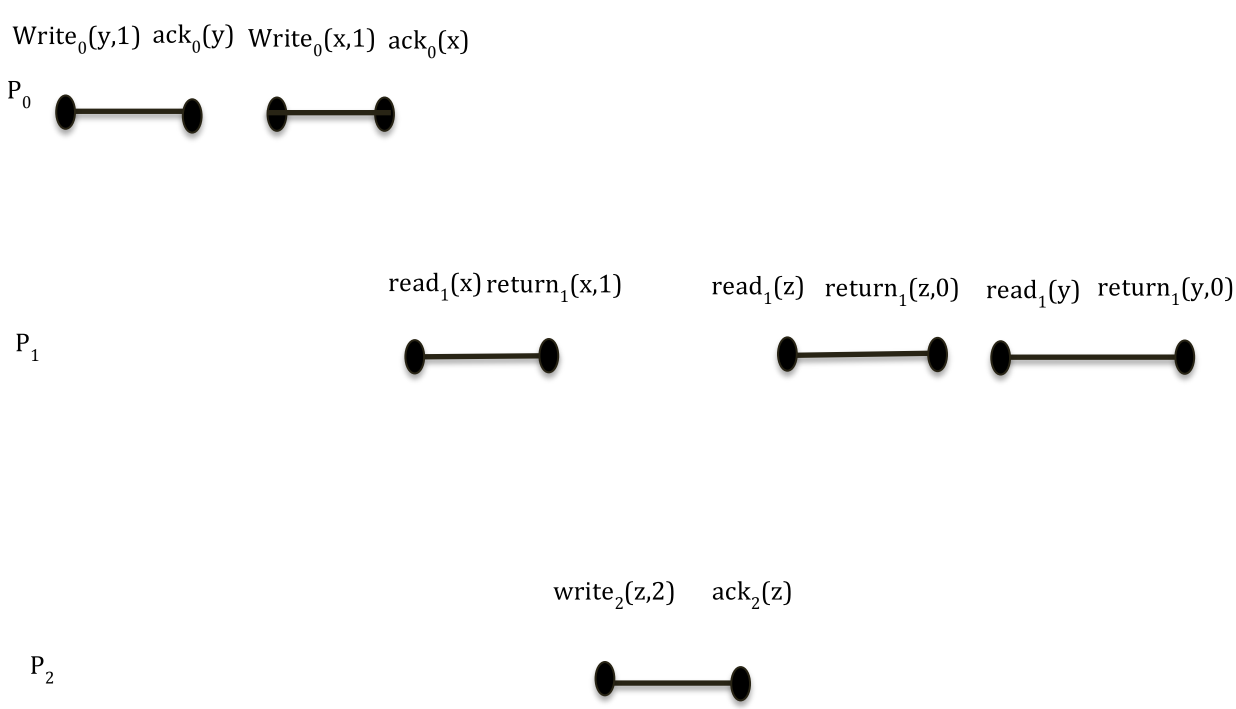

(c) Is the execution below sequentially consistent?

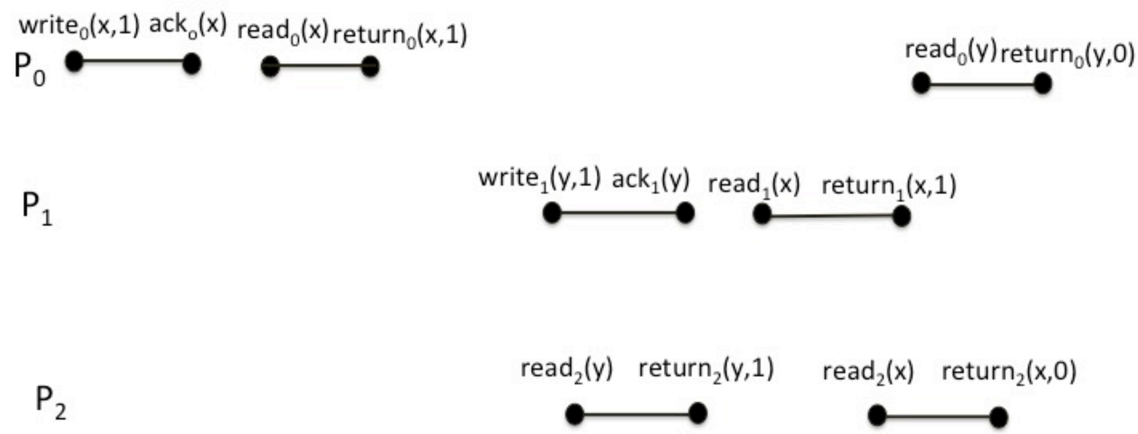

4. (10 points)
(a) In the execution below, draw the consistent cut that contains the largest possible number of events, ***excluding*** **event c**.

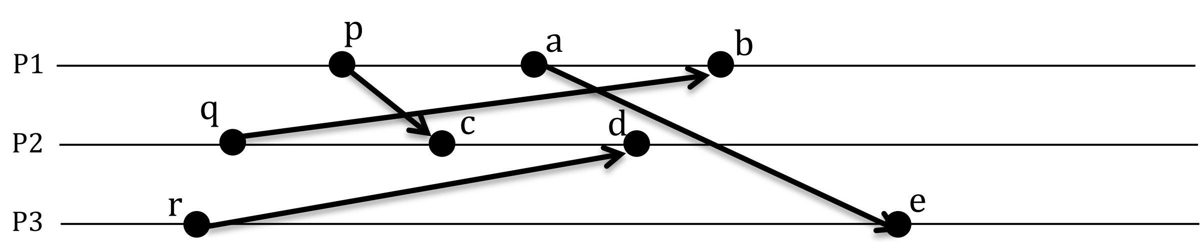

(b) In the execution above, is the cut consisting of events {a.b,e,p,q,r} a consistent cut? *Explain your answer*.

5. (12 points)

(a) In the execution below, assume that logical timestamps of **events p, q and r are 1, 5 and 8**, respectively. Determine the logical timestamps of the remaining events in the figure. Write the logical timestamp of each event next to that event in the figure.

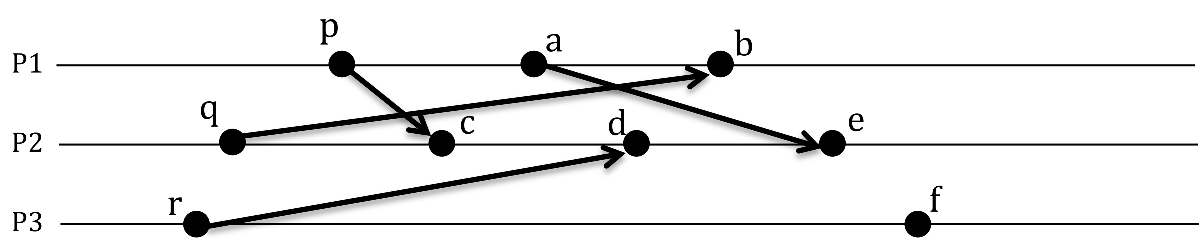

(b) In the execution shown ABOVE, determine all the events that are concurrent with event d.

(c) The figure below shows the vector timestamp of 3 events in the execution. Determine the vector timestamps for the remaining events shown below. Write the vector timestamps of each event next to that event in the figure.

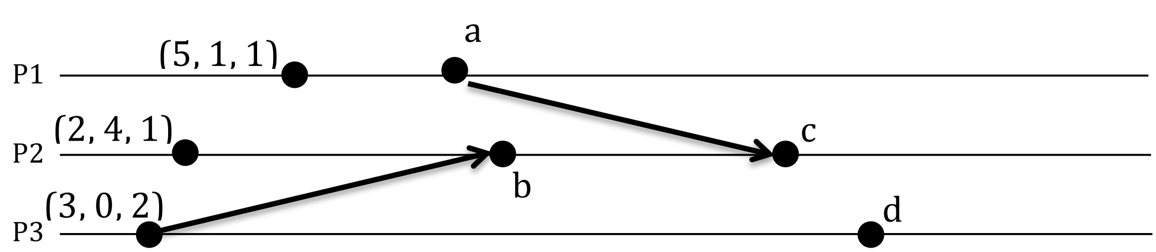

6. (12 points) Consider a synchronous system shown below. The system consists of four nodes (named 0, 1, 2 and 3) and 4 bidirectional links.

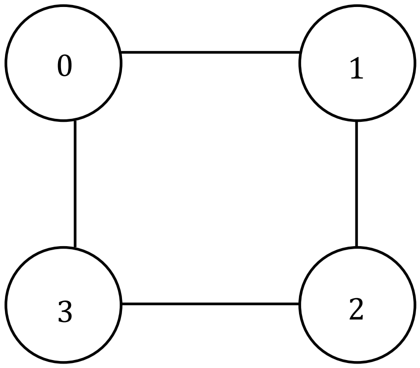

(a) Assume the following: (i) Nodes 0 and 1 never fail, and (ii) nodes 2 and 3 may both crash, either simultaneously or individually.
Is it possible to solve consensus in this system? If you answer no, explain why.
If you answer yes, provide a sketch of an algorithm to solve consensus on this system.

(b) Assume the following: (i) Nodes 0 and 2 never fail, and (ii) nodes 1 and 3 may both crash, either simultaneously or individually.
Is it possible to solve consensus in this system? If you answer no, explain why.
If you answer yes, provide a sketch of an algorithm to solve consensus on this system.

7. (a) (6 points) The execution below shows the time at which four multicast messages are delivered to processes P0, P1, P2 and P3. Does the message deliveries satisfy the following properties? In each case, answer Yes or No -- if you answer No, explain why.

   (i) FIFO ordering

   (ii) Causal ordering

   (iii) Total ordering

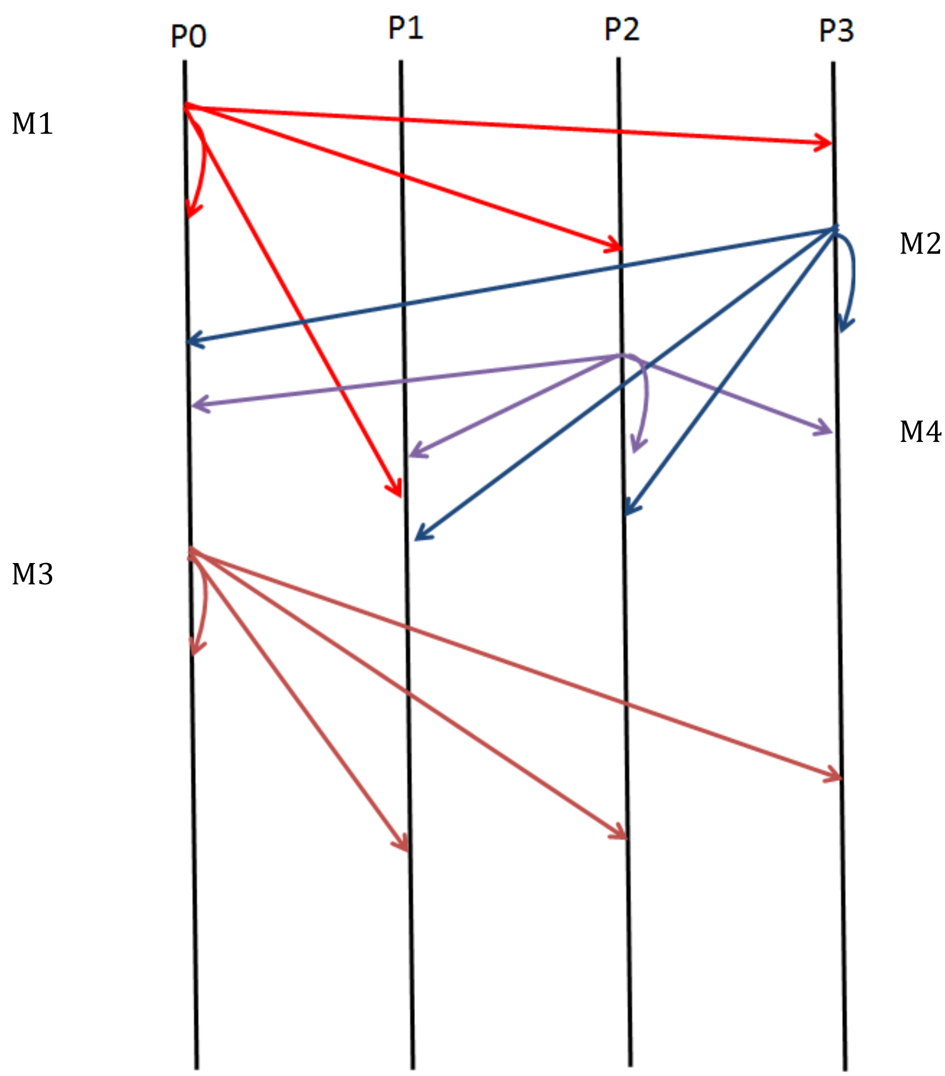

(b) (7 points) The algorithm below is a modified version of the algorithm in the *last slide* in the handout provided to you for this exam. The modification is to delete a part of the code in the "wait until" statement, as shown below. The original (unmodified) algorithm achieves causally ordered multicast.

Does the modified algorithm achieve any of the following properties for the multicast? For each property, answer Yes or No.

(i) FIFO ordering

(ii) Causal ordering

(iii) Total ordering

Algorithm for group member p_i $(i = 1, 2\ldots, N)$

On initialization

$V_i^g[j] := 0$ $(j = 1, 2\ldots, N)$; The number of group-g messages from process j that have been seen at process i so far

To CO-multicast message m to group g

$V_i^g[i] := V_i^g[i] + 1$;

B-multicast$(g, \langle V_i^g, m\rangle)$;

On B-deliver$(\langle V_j^g, m\rangle)$ *from* p_j, *with* $g = group(m)$

place $\langle V_j^g, m\rangle$ in hold-back queue;

wait until $V_j^g[j] = V_i^g[j] + 1$ $\cancel{\text{and } V_j^g[k] \le V_i^g[k]\ (k \ne j)}$;

CO-deliver m; // after removing it from the hold-back queue

$V_i^g[j] := V_i^g[j] + 1$;

8. (12 points) Consider the systems shown in the figure below. Assume the following: (i) the message delay on link between nodes 0 and 1 is unbounded (maybe arbitrarily large), and (ii) on each of the remaining links, the message delay is always less than 50 ms. You may assume that processing delays are 0.

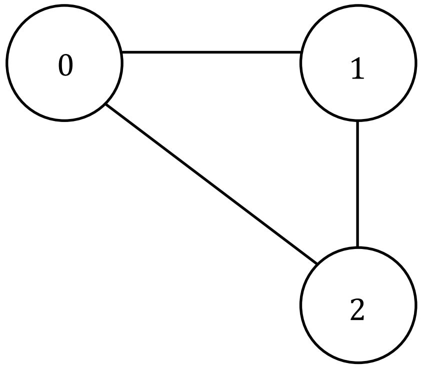

(a) Is it possible for the three nodes to synchronize their clocks to within 100 ms of each other? (i.e., skew <= 100 ms). Answer YES or NO.
If you answer NO, explain why.
If you answer YES, sketch an algorithm for achieving the above objective.

(b) If we also know that message delays are always at least 10 ms, is it possible for the three nodes to synchronize their clocks to within 90 ms of each other? (i.e., skew <= 90 ms). Answer YES or NO.

9. (6 points) Consider the Chandy-Lamport algorithm on page 2 of the handout provided to you for this exam. Will this algorithm record a correct snapshot if the communication channels are not FIFO? Justify your answer.

10. (6 points) Does there exist any algorithm to achieve approximate consensus in an asynchronous system with up to f crash failures, if the total number of nodes is f+1. Justify your answer.

This page intentionally left blank.