

2006 Paper 9 Question 4

Distributed Systems

Members of an open process group manage distributed replicas of data values stored in persistent memory. To allow the system to operate in the presence of transient failures of some replica managers, a quorum assembly scheme is used. Replica managers are assumed to be non-malicious and fail-stop.

To update a managed data item, the operations provided by the managing process include:

```
lock(item)
update(item, value, timestamp)
read(item, timestamp)
unlock(item)
```

- (a) Suppose the data item is an initially empty list of values and the update operation appends a value. Illustrate the quorum assembly scheme for five replicas, showing a number of update and read operations. [8 marks]
- (b) How is a total order of updates achieved by quorum assembly in the presence of concurrent update requests by clients to the open group? Discuss how any problems that might arise can be solved. [4 marks]
- (c) When can `unlock(item)` be executed safely by the initiating replica manager? Describe any additional protocol that is needed. [5 marks]
- (d) Suppose that the process group is managing non-overlapping partitions of a distributed database instead of replicas. Can quorum assembly play any part in making the related updates required for distributed transactions? Justify your answer. [3 marks]