

Problem A: 28 points—4 points each question. Whenever applicable show the appropriate graph to justify your answer.

Consider the following schedule, where T1 starts before T2 which starts before T3

<p>T1</p> <p>start</p> <p>write(A)</p> <p>read (C)</p> <p>write(B)</p>	<p>$L_x(A)$</p>	<p>T2</p> <p>start</p> <p>$L_S(A)$ blocked</p> <p>read(A)</p> <p>read(C)</p>	<p>start</p> <p>$L_x(C)$</p> <p>write(C)</p> <p>read(A)</p>	<p>T3</p>
--	----------------------------	--	---	-----------




Please answer the following questions:

- Q1. Is this schedule conflict-serializable (explain your answer with a graph)?
- Q2. Show what will happen if they try to execute this schedule under strict 2PL and no deadlock prevention (assume that a transaction locks a resource just before it needs it). If there is no deadlock show the complete schedule otherwise prove deadlock by the appropriate graph
- Q3. What will happen if we instead execute these transactions using strict 2PL and a wound-wait deadlock prevention scheme (assume that a transaction locks a resource just before it needs it)?
- Q4. In general, does the protocol used in Q3 guarantee cascadeless schedules?
- Q5. What will happen if we instead execute these transactions using a time-stamp based protocol to guarantee the serializability of the transactions? Show the resulting schedule.
- Q6. Does the protocol used in Q5 guarantee freedom from deadlocks (yes/no)?
- Q7. In the wait-die deadlock prevention protocol, transactions that are rolled back keep their old timestamp. What is the reason for that?

Problem B: 24 points: questions 1–6 two points; questions 7–10 three points.

Consider the following ER-diagram:



Mark which of the following six statements are TRUE or FALSE, according to the ER-diagram.

- (1) Each City has at most one mayor. ✓
- (2) Two different states cannot have cities that have the same name. ✗
- (3) A city with a given **cname** can be In at most one state. ✗
- (4) No two states can have the same name (i.e., identical **sname**). ✓
- (5) No two cities can share the same mayor. ✗
- (6) No two cities In the same state can have the same name (i.e., identical **cname**). ✓

Now answer the following four questions about our running example:

Say that we decide to implement the relational database storing the information described by our ER-diagram using the following schema, where the attributes belonging to the key are underlined, `citiesStates(cname, Mayor, sname, popl)` and answer the following questions:

- (7) is this relation 3NF (justify your answer)? ✗
- (8) Say that “popl” denotes the state population, as per the last census. Then, what update anomaly will we experience when we have to update with a new census results the **popl** of a state where there are 629 cities?
- (9) Design a BCNF schema, i.e. a set of BCNF relations, that eliminates those anomalies.
- (10) Show all the keys and foreign keys (with their referenced attributes) that must be declared in your schema to assure that all the constraints defined by the ER-diagram are enforced. (no SQL required)

Problem C: Transaction Recovery: 20 Points—4 points per question

The log used by the transaction manager contains information about commits, aborts, updates, and checkpoints. Please answer the following questions about checkpoints as TRUE or FALSE and also give justification for your answer:

- 1) An important function of checkpoints is to manage the cascading of rollbacks for transactions that had read dirty data from other transactions.
- 2) A key function of checkpoints is to give the database administrator the ability to check on the current state of transactions and decide whether they must abort or commit.
- 3) The main function of checkpoints is to expedite the recovery process.
- 4) The transaction manager performs (fuzzy) checkpoint operations by forcing each running transaction to commit and making sure that their updates are actually recorded on disk.
- 5) Assume that for unforeseen reasons a few checkpoint records of a lengthy log were damaged in the last crash, and therefore they cannot be read by the recovery manager. If only checkpoint records were damaged, then the recovery manager will be able to complete the recovery process anyway and bring back the database to a consistent state, in spite of such losses.

Problem D: FDs, and Schema Design: 28 Points—7 points each question:

Given the following FDs for $R(A, B, C, D, E)$

1. $AC \rightarrow B$
2. $B \rightarrow A$
3. $D \rightarrow E$
4. $C \rightarrow D E$

Answer the following four questions by considering these FDs in ascending order:

- Q1. Is this schema BCNF (justify your answer)?
- Q2. Check if these given FDs are in canonical form (only one attribute on the right side) required by the BCNF design algorithm. if they are not replace them with an equivalent set of FDs that are in canonical form.
- Q3. Compute a lossless decompositions of $R(A, B, C, D, E)$ into BCNF relations.
- Q4. Is your decomposition FD-preserving (if applicable list the lost FDs)?

Extra Credit Problem : 6 points—each question 2 points.

D is a cohort in a 2PC where A is the coordinator and B and C are the other two cohorts. Each cohort can communicate with the coordinator and the other cohorts, and all the sites and the communication lines operate with great reliability EXCEPT for the communication line between D and A which undergoes failures that require a long time to repair. The distributed DB (DDB) administrator is concerned about what happens when a failure of the line between A and D occurs and ask the following questions. Please, select the correct answer to each question.

EC1. At the end of the 2PC protocol, is it possible for site D to have committed even though A,B,C aborted?

(A) No. Because the coordinator or one of the cohorts aborted, Coordinator A must issue an ABORT request in phase 2 that will undo (ABORT) any commit site D may have made during phase 1.

(B) Yes. Site D can commit once it has received a COMMIT request and release its resources as soon as it has completed the transaction issued by a COMMIT-REQUEST. Aborts by the coordinator or other cohorts will not affect this commit.

(C) Yes. Site D can prepare the transaction after receiving a COMMIT-REQUEST in phase 1. After coordinator A receives an AGREE message from site D, it will issue a COMMIT message to site D to commit the transaction.

EC2. If the Coordinator does not get the ready/abort message from a Cohort (say, D) after timeout in phase 1, then:

(A) It can decide to commit independently if it gets "ready" from all other cohorts, and send the decision to D.

(B) It needs to wait indefinitely until it gets decision from all cohorts.

(C) It aborts the transaction only if it gets at least one abort from another cohort.

(D) It can abort the transaction.

EC3. If the Cohort site does not get the commit/abort message in phase 2, which of following is invalid?

(A) It can learn the decision from other cohorts.

(B) It can decide whether to commit or abort independently, and send the decision to coordinator.

(C) It must wait until its communication with coordinator recovers.

(D) The coordinator can still consider the transaction is committed.

