

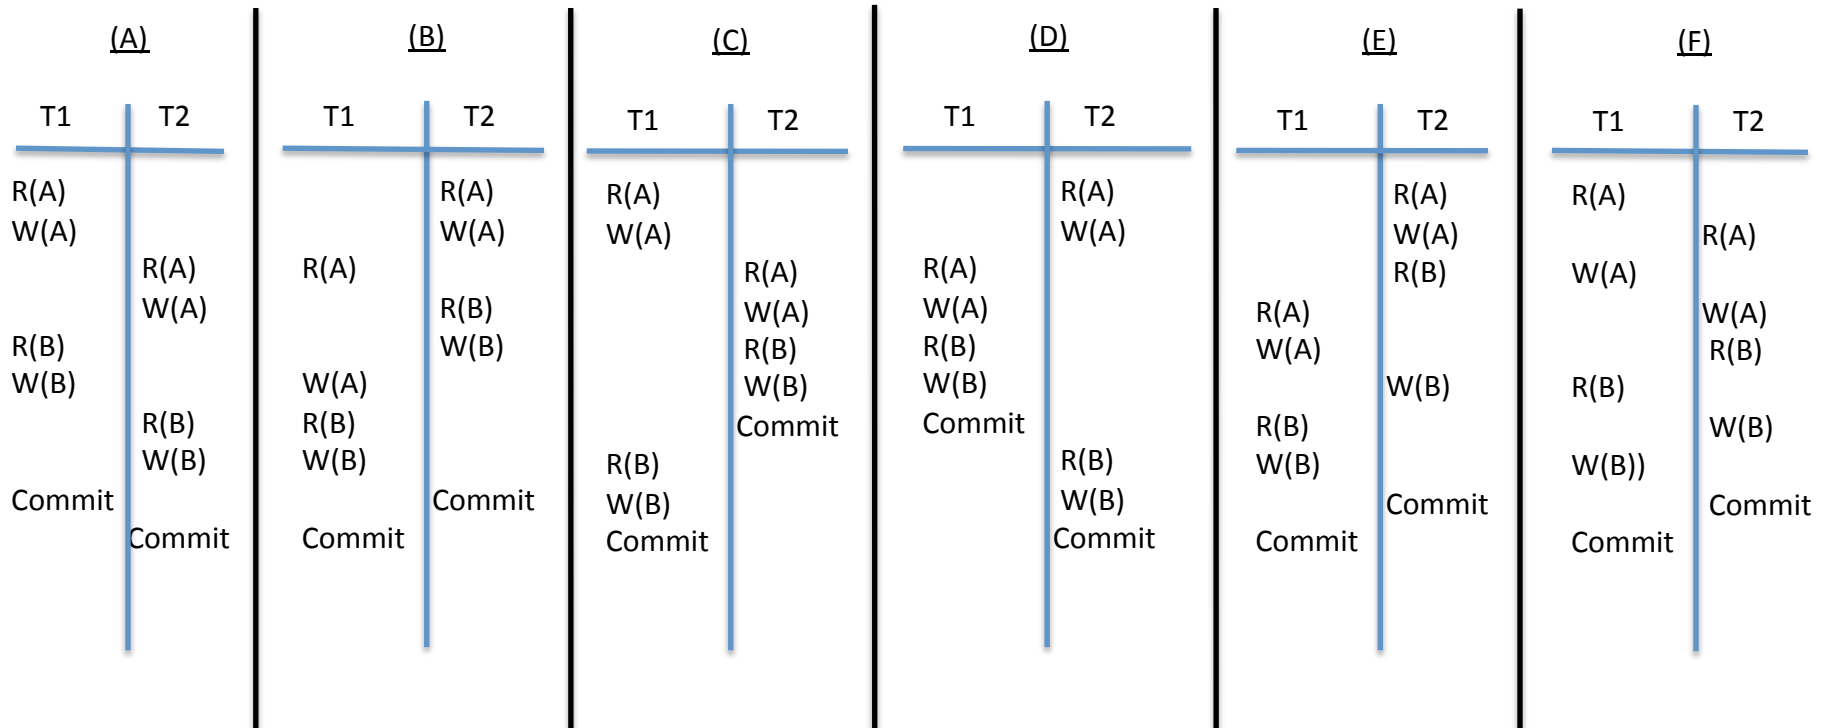
In this question, do not assume that strict 2 phase locking is used. Rather you can assume that locks are released as when they are no longer absolutely needed (after a read for shared, and after a write for exclusive)

9. (4 points) Consider the following two transactions, T1 and T2:

T1: Read(A), Op<sub>11</sub>(A), Write(A), Read(B), Op<sub>12</sub>(B), Write(B), Commit

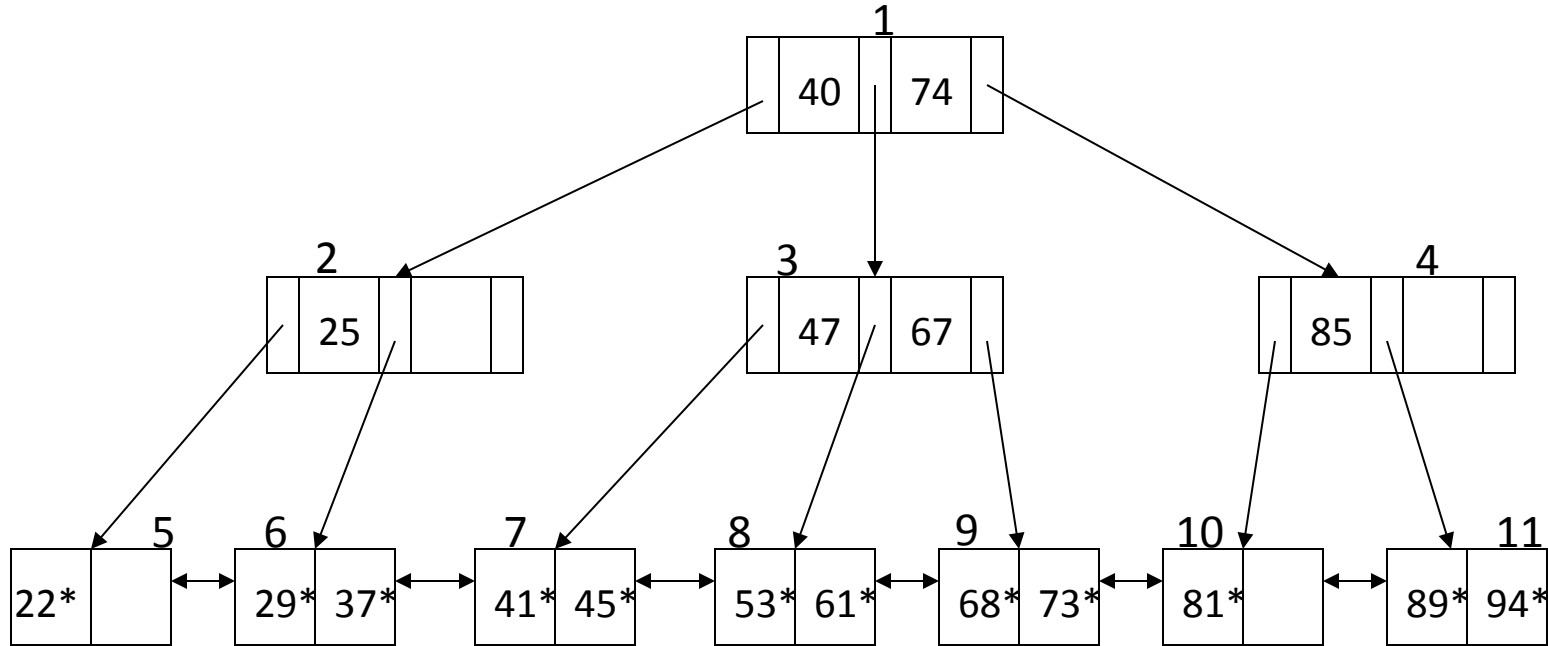
T2: Read(A), Op<sub>21</sub>(A), Write(A), Read(B), Op<sub>22</sub>(B), Write(B), Commit

Circle all schedules (just showing disk reads and writes) that clearly result in *serializable* behavior even without knowing when the Ops are performed. **A and B are distinct (but are shared across transactions).**



(G) NONE OF THE ABOVE

**21. (4 pts)** Consider the B+ tree index for attribute A of table T. Above each node is a numeric label for the node (1 through 11), which you will use in answering this question.



For each of the following operations, list the *nodes* (by label), in proper order, that would be locked (shared or exclusive) in strict two-phase locking (2PL) when performing the respective operation. If no nodes need be locked, then write *None*. Ignore data nodes, which are not shown, and do not list new nodes that might be introduced. Assume that this index for attribute A is used in evaluating each operation below. Write S(label) for a shared lock, and X(label) for an exclusive lock. Note that the same node, A, may be listed twice, first as S(A) then as X(A), for the same operation. Do not show order of lock release (we'll assume all locks released at end of operation, upon commit). Treat each operation as independent, and not as a sequence of actions. Assume that redistribution is never used.

(a) SELECT T.A FROM T WHERE T.A > 75 :

(b) UPDATE T SET T.B = T.B + 100 WHERE T.A = 37 :

(c) UPDATE T SET T.A = T.A + 5 WHERE T.A = 29 :

(d) INSERT INTO T (A, B, C) VALUES (70, 20, 10) :