

# CS 186/286 Fall 2017 Final Exam

- Do not turn this page until instructed to start the exam.
- You should receive 1 *answer sheet* and a 15-page *exam packet*.
- All answers should be written on the answer sheet. The exam packet will be collected but not graded.
- You have *3 hours* to complete the final.
- The midterm has *8 questions*, each with multiple parts.
- For each question, place only your *final answer* on the answer sheet; do not show work.
- For multiple choice questions, please fill in the bubble or box completely, **do not mark the box with an X or checkmark**.
- Use the blank spaces in your exam for scratch paper.
- You are allowed **two** 8.5" × 11" double-sided pages of notes.
- No electronic devices are allowed.

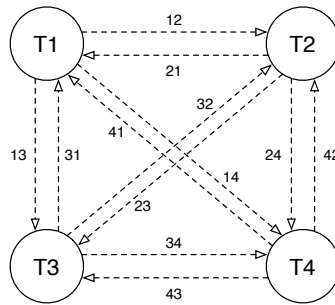
# 1 Concurrency

Consider the following schedule of reads and writes to pages.

	1	2	3	4	5	6	7	8	9	10
$T_1$	R(A)						W(A)	R(C)		
$T_2$		R(A)							R(C)	
$T_3$				R(C)	R(B)					W(B)
$T_4$			R(C)			W(C)				

## Conflicts

- (2 points) What edges are included in the **dependency graph** (*not* waits-for graph) for the above schedule? A reference graph is provided for your convenience. The dotted edges represent all *possible* edges; each edge is labeled with the transaction IDs of the nodes it connects, in order. E.g. the edge from  $T_4 \rightarrow T_1$  is labeled 41. Mark the answer sheet with the labels of the edges that are in the dependency graph of the schedule above.



- (1 point) Which of the following serial schedules are conflict equivalent to the schedule above?
  - $T_4, T_2, T_1, T_3$
  - $T_3, T_2, T_4, T_1$
  - $T_3, T_4, T_2, T_1$
  - $T_1, T_2, T_4, T_3$
  - $T_1, T_4, T_2, T_3$
  - None of the above
- (2.5 points) Mark all statements that are true.
  - In Strict 2PL, we can give up locks after aborting but before rollback is complete.
  - Some conflict serializable schedules cannot be produced when using 2PL.
  - Schedules that are conflict serializable will not produce a cyclic dependency graph.
  - Both Strict 2PL and 2PL enforce conflict serializability.
  - All schedules that are conflict serializable are view serializable.

## Dealing with Deadlock

For the following sequence of lock requests, assume  $T_1$  is older than  $T_2$ , and  $T_2$  is older than  $T_3$ . The oldest transaction has the highest priority. No locks are released in the time frame shown. Note that not all requests may actually be issued—in some cases the transaction is still blocked waiting for a prior request.

Consider the following schedule of lock requests:

	1	2	3	4	5	6	7	8	9
T1	X(D)		S(A)			S(C)			
T2		X(A)		S(B)				S(D)	
T3					S(B)		X(B)		S(C)

4. (1 point) The above schedule could not successfully occur whether one uses deadlock avoidance or not. Suppose we are not using any deadlock avoidance algorithms. List the time steps in the schedule above in which *the lock request could not have been made*: i.e. timesteps when the requesting transaction would have been blocked waiting for a previous lock request.
5. (1 point) Assuming we do not use deadlock avoidance, what is the first time step when deadlock will occur?
6. (1 point) Mark the transactions involved in the first deadlock.
7. (1 point) If we were using the wound-wait deadlock avoidance algorithm, what would happen at time step 3?
  - A. The transaction requesting the lock would wait
  - B. The transaction requesting the lock would abort
  - C. The transaction that holds the lock would abort

## 2 Recovery

1. (0.5 points) Which buffer management policy would allow us to maximize efficiency in terms of speed?
  - A. STEAL, NO FORCE
  - B. STEAL, FORCE
  - C. NO STEAL, FORCE
  - D. NO STEAL, NO FORCE
  
2. (0.5 points) True or False? In ARIES recovery, after the analysis phase, the recLSN of each page in the dirty page table must be larger than the pageLSN of the corresponding page.
  
3. (2 points) Select all the correct statements that are true about ARIES.
  - A. CLR's are never undone during ARIES recovery.
  - B. We remove a page from the dirty page table each time we write a CLR to the log.
  - C. The transaction table and dirty page table in the End Checkpoint record are up-to-date as of the time when the End Checkpoint record was logged.
  - D. ARIES never causes a database disk page to be overwritten by a page with the same PageLSN.
  
4. (0.5 points) True or False: Regardless of the number of times we run recovery, each UPDATE in the log will have at most one matching CLR record in the log.
  
5. (1 point) Why would the undoNextLSN field in a CLR be marked null?
  - A. If the transaction committed.
  - B. If this CLR corresponds to the lowest UPDATE LSN of the transaction
  - C. If this CLR corresponds to the highest UPDATE LSN of the transaction
  - D. If this CLR was the highest LSN flushed at the time of crash.
  - E. If this CLR corresponds to a page that was flushed before the checkpoint.

Now consider the following log records which are recovered after a crash.

LSN	Record	PrevLSN
⋮	⋮	⋮
40	UPDATE: $T_3$ writes $P_4$	null
50	UPDATE: $T_2$ writes $P_3$	null
60	Begin Checkpoint	-
70	End Checkpoint	-
80	UPDATE: $T_1$ writes $P_1$	null
90	UPDATE: $T_3$ writes $P_5$	40
100	UPDATE: $T_3$ writes $P_2$	90
110	UPDATE: $T_2$ writes $P_1$	50
120	$T_3$ Commit	100
130	$T_3$ End	120
140	UPDATE: $T_1$ writes $P_5$	80
150	$T_1$ Abort	140
160	CLR: Undo write to $P_5$ ( $T_1$ , LSN 140), undoNextLSN: 80	150
	CRASH!	

The following dirty page table and transaction table are read from the checkpoint:

**Transaction Table**

Transaction	Status	lastLSN
$T_2$	Running	50
$T_3$	Running	40

**Dirty Page Table**

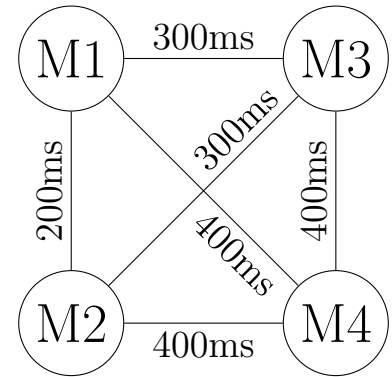
Page	recLSN
$P_4$	40

- (1 point) Assume that at the end of the Analysis phase, we convert all Running status to Aborting in the transaction table. For each transaction, give its status in the transaction table at the end of the analysis phase. List its status as "A" if the transaction is Aborting and "C" if the transaction is Committing. If a transaction is **NOT** in the transaction table, list its status as "0".
- (2.5 points) For each page referenced in the log, give its recLSN in the dirty page table at the end of the analysis phase. If a page is **NOT** in the dirty page table, list its recLSN as "0". Assume the buffer manager does not flush any pages during recovery.
- (1 point) Mark the LSNs of the records that will actually be redone in the redo phase (i.e. that will require us to update a database page).
- (1 point) Mark the LSNs of the UPDATE records that we write CLRs for during the undo phase.

### 3 Two-Phase Commit

Our database runs on 4 machines and uses Two-Phase Commit. Machine 1 is the Coordinator, while Machines 2, 3, and 4 are Participants.

Suppose our machines are connected such that the time it takes to send a message from Machine  $i$  to Machine  $j$  is  $100 \cdot \max(i, j)$  milliseconds (see graph). Assume these communication latencies are symmetric: it takes the same amount of time to send from  $i$  to  $j$  as it takes to send from  $j$  to  $i$ . For example, sending a message between Machine 2 and Machine 4 takes 400 milliseconds in either direction.



Assume that the transaction will commit (i.e. all subordinates vote yes), and that everything is instantaneous except for the time spent sending messages between two machines.

- (0.25 points) What is the first message Machine 1 sends?  
A. VOTE YES      B. PREPARE      C. COMMIT      D. None of these
  - (0.25 points) What is the second message Machine 1 sends?  
A. VOTE YES      B. PREPARE      C. COMMIT      D. None of these
  - (1 point) How much time passes from when Machine 1 sends its first message to when Machine 1 sends its second message?
  - (0.25 points) What is the first message Machine 2 sends?  
A. VOTE YES      B. PREPARE      C. COMMIT      D. None of these
  - (0.25 points) What is the second message Machine 2 sends?  
A. VOTE YES      B. PREPARE      C. COMMIT      D. None of these
  - (1 point) How much time passes from when Machine 2 sends its first message to when Machine 2 sends its second message?
  - (0.5 points) True or False. A transaction is considered committed even if over half of the participants do not acknowledge the commit.
- Now suppose that our implementation of 2-Phase Commit has an off-by-one bug where the Coordinator receives, but does not use, Machine 4's vote. That is, Machine 4's vote does not affect whether or not the transaction commits or aborts.
- (0.5 points) True or False. A transaction that should normally commit may be aborted instead.
  - (0.5 points) True or False. A transaction that should normally abort may be committed instead.
  - (0.5 points) True or False. A transaction that should normally commit may be committed properly.
  - (0.5 points) True or False. A transaction that should normally abort may be aborted properly.

## 4 Joins

### 4.1 Young Justice... again

For the following questions in this section, assume that we are streaming our query output to a terminal. (Do not consider the cost of writing the final output)

```
CREATE TABLE JusticeLeague (  
    member_id INTEGER PRIMARY KEY,  
    code_name CHAR(33),  
    birth_planet CHAR(20),  
    power_level INTEGER  
);  
  
CREATE TABLE Teaches (  
    member_id INTEGER PRIMARY KEY REFERENCES JusticeLeague,  
    teacher_id INTEGER REFERENCES JusticeLeague,  
    since DATE  
);
```

We assume that

- JusticeLeague has  $[J] = 100$  pages
  - JusticeLeague has  $|J| = 1000$  members
  - Teaches has  $[T] = 200$  pages
  - Teaches has  $|T| = 2000$  members
  - Buffer size = 12 pages
  - An Alternative-1 B+ tree for Teaches indexed on `T.member_id` has the height 2. *In this problem, we interpret height as the number of I/Os it takes to travel from the root to the leaf page.*
1. (2 points) What is the I/O cost of using Index Nested Loop Join to perform a natural join of Teaches and JusticeLeague? Consider table JusticeLeague as the outer relation. Note that we want the actual cost in I/O requests, not a Selinger-style estimate. Treat each I/O request the same, regardless of whether you think the request will hit in the buffer pool. *Your answer must be a single integer.*
  2. (2 points) Assume we introduce an Alternative 2, clustered B+ tree for JusticeLeague indexed on `member_id` with height 2. What is the I/O cost of using Index Nested Loop Join to perform a natural join of Teaches and JusticeLeague? This time consider table Teaches as the outer relation. *Your answer must be a single integer.*
  3. (1 point) If the previous question's B+ tree for JusticeLeague had used an unclustered index, would the I/O cost have been different? (Answer only in Yes or No)

## 4.2 General Join Algorithm Questions

4. (2 points) In this question, we explore the strengths and weaknesses of various join algorithms. Assume throughout that we are joining two tables  $R$  and  $S$ , and we have  $B$  buffers of memory for our join. Do not assume any additional information about the tables that is not explicitly provided.
- (a) (1 point) We are performing an equijoin.  $[R] < B$ ,  $[S] > B^3$ . Which join algorithm will provide the correct answer with the fewest I/Os:
- A. Block Nested Loops Join
  - B. Grace Hash Join
  - C. Sort-Merge Join
- (b) (1 point) We are performing an equijoin.  $B < [R] < (B - 1)^2$ , and  $B < [S] < (B - 1)^2$ . Both  $R$  and  $S$  have only 2 distinct values in the join key column. Which join algorithm will provide the correct answer with the fewest I/Os:
- A. Block Nested Loops Join
  - B. Grace Hash Join
  - C. Sort-Merge Join
5. (2 points) Consider the following “Index On Demand” join algorithm. Given an equijoin on tables  $R$  and  $S$  (with  $[R] < [S]$ ), but no indexes on the join columns initially, the algorithm will (a) bulk-load an Alternative-1 B+-tree on  $R$ , (b) perform an Index Nested Loops Join on  $S \bowtie R$  using the index, and (c) delete the index on  $R$ . Mark all of the following that you believe to be true for joining  $R$  and  $S$ , assuming  $B < [R] < [S] < (B - 1)^2$ :
- A. This algorithm generates more I/Os than Grace Hash Join (without recursive partitioning).
  - B. This algorithm generates more I/Os than a 2-pass Sort-Merge join.
  - C. This algorithm generates more *random* I/Os than Block Nested Loops Join.
  - D. If we omit step (c), this algorithm might be of interest for its benefits on subsequent queries, even though it could be sub-optimal for the current query.



## 5 Query Optimization

For the following questions, assume the following:

- The System R assumptions about uniformity and independence from lecture hold
- We use System R defaults when selectivity estimation is not possible
- Primary key IDs are sequential in each table, starting from 1, and there are no gaps in the ID sequence in any table
- Our system implements only implements Grace Hash Join. It allocates 5 pages of memory for the join operator.
- Assume all indices are alternative 3 indices.
- Assume the optimizer statistics are fully accurate w.r.t. the current content of the database.

Table Schema	Records	Pages	Indices
CREATE TABLE User ( uid INTEGER PRIMARY KEY, fullname VARCHAR(32), country VARCHAR(4)) )	1,000	100	None
CREATE TABLE Item ( sku INTEGER PRIMARY KEY, itemname VARCHAR(32), price NUMERIC )	10,000	1,000	<ul style="list-style-type: none"> <li>• Index 0: Unclustered Index on sku.</li> <li>• Index 1: Clustered Index price Low 0, High 100</li> </ul>
CREATE TABLE Order ( uid INTEGER REFERENCES User, sku INTEGER REFERENCES Item, quantity INTEGER) )	100,000	5,000	<ul style="list-style-type: none"> <li>• Index 2: Clustered Index on uid</li> </ul>
CREATE TABLE Tax ( sku INTEGER PRIMARY KEY REFERENCES Item, taxrate NUMERIC )	10,000	1,000	None

1. (1 point) After applying the System R optimizer, which query has a *lower* estimated final cost for the optimal query plan.

A. SELECT \*  
FROM Item  
WHERE Item.itemname = 'WidgetA'

B. SELECT \*  
FROM Item  
WHERE Item.price < 3

2. (1 point) After applying the System R optimizer, which query has a *lower* estimated final cost for the optimal query plan.

A. SELECT \*  
FROM Item  
ORDER BY Item.price

B. SELECT \*  
FROM Tax  
ORDER BY Tax.taxrate

3. (0.5 points) (True or False) The System R optimizer will apply a projection to the Tax table that preserves only the `taxrate` attribute before performing a join for the following query:

```
SELECT taxrate
FROM Item, Tax
WHERE Tax.sku = Item.sku
```

4. (0.5 points) (True or False) For the following query, one can determine the cardinality of the join result (i.e the number of tuples in the result) exactly:

```
SELECT taxrate
FROM Item, Tax
WHERE Tax.sku = Item.sku
```

5. (0.5 points) (True or False) In general, even if an eligible index exists, a sequential scan can be selected rather than an index scan.

6. (1.5 points) Consider the following query:

```
SELECT *
FROM User, Order
WHERE User.uid = Order.uid AND
      User.country = 'USA'
```

During Pass 2 of the System R Optimizer, what is the estimated cost of executing the query using a Grace hash join? Be sure to include the cost of the heap scans and any other upstream operators.

7. (1 point) Consider the following query:

```
SELECT *
FROM User, Order, Tax, Item
WHERE User.uid = Order.uid AND
      Item.sku = Order.sku AND
      Tax.sku = Item.sku AND
      Tax.sku = Order.sku
```

Without making assumptions about costs and selectivities, which of the following join orders could **NEVER** be considered by the System R optimizer to answer this query? *Mark all that apply*

- A. (((User ⋈ Order) ⋈ Item) ⋈ Tax)
  - B. (User ⋈ Order) ⋈ (Item ⋈ Tax)
  - C. (((User ⋈ Item) ⋈ Order) ⋈ Tax)
  - D. (((User ⋈ Order) ⋈ Tax) ⋈ Item)
  - E. (((Tax ⋈ Item) ⋈ Order) ⋈ User)
  - F. None of the above
8. (0.5 points) (True or False) Ignoring order by statements, group by statements, and nested queries, the System R optimizer is guaranteed to return the left-deep query plan with the lowest IO cost with respect to the estimated costs.

## 6 SQL

Consider a table `Friend` representing a social network. Each record of the `Friend` table consists of two integers corresponding to a friendship relationship between two users (identified by the `id` number). Assume that no person is friends with themselves.

```
CREATE TABLE Friend(  
    id1 INTEGER,  
    id2 INTEGER,  
    PRIMARY KEY (id1,id2));
```

1. (1.5 points) Some of the friendship relationships in this table are not reciprocated, i.e.,  $(id1=a, id2=b)$  exists in the table but  $(id1=b, id2=a)$  does not exist. Which of the following SQL queries identifies all tuples where the reciprocal relationship does not exist in the table. *Mark all that apply.*

- A. 

```
SELECT *  
FROM Friend f1  
WHERE NOT EXISTS (SELECT *  
                  FROM Friend f2  
                  WHERE f1.id1 = f2.id2 AND f1.id2 = f2.id1)
```
- B. 

```
SELECT t.id1, t.id2  
FROM (  
    (SELECT id1, id2  
     FROM Friend f1)  
    UNION ALL  
    (SELECT f2.id2 as id1, f2.id1 as id2  
     FROM Friend f2)  
    ) as t  
GROUP BY t.id1, t.id2  
HAVING count(*) < 2
```
- C. 

```
SELECT t.id1, t.id2  
FROM (  
    SELECT f2.id2 as id1, f2.id1 as id2  
    FROM Friend f2  
    ) as t FULL OUTER JOIN Friend  
ON t.id1 = Friend.id1 AND t.id2 = Friend.id2 AND Friend.id1 <> NULL
```
- D. None of the above

2. (1 point) Consider the same table, but without the primary key constraint. Continue to assume that no person is friends with themselves.

```
CREATE TABLE Friend(  
    id1 INTEGER,  
    id2 INTEGER);
```

This table can contain multiple copies of the same friendship relationship, e.g., two tuples with  $(id1=a, id2=b)$ . Therefore, each distinct  $(id1=a, id2=b)$  pair will have a count of the number of times it occurs in the relation. We want to write a query that finds all distinct pairs where the count of  $(id1=a, id2=b)$  is not equal to the count of  $(id1=b, id2=a)$ . *Mark all that apply.*

- A. `SELECT id1, id2, count(*)  
FROM Friend  
GROUP BY id1, id2  
  
EXCEPT  
  
SELECT id2, id1, count(*)  
FROM Friend  
GROUP BY id1, id2`
- B. `SELECT id1, id2, count(*)  
FROM Friend  
GROUP BY id1, id2  
  
EXCEPT ALL  
  
SELECT id2, id1, count(*)  
FROM Friend  
GROUP BY id1, id2`
- C. `(SELECT id1, id2, count(*)  
FROM Friend  
GROUP BY id1, id2  
  
EXCEPT  
  
SELECT id2, id1, count(*)  
FROM Friend  
GROUP BY id1, id2)  
  
UNION ALL  
  
(SELECT id2, id1, count(*)  
FROM Friend  
GROUP BY id1, id2  
  
EXCEPT  
  
SELECT id1, id2, count(*)  
FROM Friend  
GROUP BY id1, id2)`
- D. None of the above

3. (1.5 points) Assume the following `Friend` table with the primary key constraint.

```
CREATE TABLE Friend(  
  id1 INTEGER,  
  id2 INTEGER,  
  PRIMARY KEY (id1,id2));
```

A “triangle” of friends is defined as Person  $a$  is friends with Person  $b$ , Person  $b$  is friends with Person  $c$ , and Person  $c$  is friends with Person  $a$ . Which of the following queries lists of all of the distinct triangles output in alphabetical order—i.e., if  $a, b, c$  is in the result then  $c, a, b$  should not be in the result. *Mark all that apply.*

- A. `SELECT f1.id1, f2.id1, f3.id1  
FROM Friend f1, Friend f2, Friend  
f3  
WHERE f1.id2 = f2.id1 AND  
f2.id2 = f3.id1 AND  
f3.id2 = f1.id1 AND  
f1.id1 < f2.id1 AND  
f2.id1 > f3.id1`
- B. `SELECT f1.id1, f2.id1, f3.id1  
FROM Friend f1, Friend f2, Friend  
f3  
WHERE f1.id2 = f2.id1 AND  
f2.id2 = f3.id1 AND  
f3.id2 = f1.id1`
- C. `SELECT f1.id1, f2.id1, f3.id1  
FROM Friend f1, Friend f2, Friend  
f3  
WHERE f1.id2 = f2.id1 AND  
f2.id2 = f3.id1 AND  
f3.id2 = f1.id1 AND  
f1.id1 < f2.id1 AND  
f2.id1 < f3.id1`
- D. None of the above

## 7 Replication

1. (3 points) Which of the following are reasons to replicate data:
  - A. To increase the odds that some node will be able to respond to any request at any given time
  - B. To allow multiple nodes to split up the work for a large volume of requests.
  - C. To reduce latency by allowing clients to fetch data from a nearby server
  - D. To handle ever-increasing database sizes.
  - E. To hide transient performance problems.
  - F. To recover from failures.
2. (2 points) Consider a single-master system that performs 2PC upon transaction commit. Which of the following problems could occur in a single-master system, but are prevented by 2PC?
  - A. The standby node may have an older value for some key than the master.
  - B. Transactions could commit while the standby node has failed.
  - C. The standby node may have processed an update that conflicts with the master node.
  - D. The master node may fail and block progress.
3. (2 points) Which of the following drawbacks are true of 2PC-controlled replication?
  - A. If it is not strict 2PC, you could get cascading aborts.
  - B. A failed participant node can cause live nodes to abort.
  - C. A failed coordinator node can freeze up the entire system indefinitely.
  - D. Message latency for 2PC across wide-area networks can be high, and lead to slow transaction performance.
4. (0.5 points) True or False: In multi-master replication, it is possible for the system to get “split brain”: two nodes may disagree about the value associated with some key.
5. (0.5 points) True or False: In the absence of 2PC, multi-master replication is a high-bandwidth approach because many nodes can service writes to the same key in parallel, without sending each other messages.

## 8 B+ Trees

### 8.1 True/False

- (3 points) Mark all statements that are true.
  - A B+ tree is always balanced-height
  - The keys on an internal node are always in sorted order
  - A B+ tree built using bulk loading is always more bushy than a B+ trees built from standard insert operations
  - Alternative 3 is always better than alternative 2
  - The number of records in the table must be known before performing bulk loading
  - A leaf node can be a root node

### 8.2 Best and Worst Case IOs

Consider an alternative 2 B+ tree index of height  $h = 3$  and order  $d = 4$ . In this problem, the height of a tree corresponds to the number of I/Os to traverse to a leaf. Assume that the index key is the primary key of the table. Note that index pages include both leaf and non-leaf pages. The maximum number of leaf pages that can exist in this index is  $(2d + 1)^3 = 729$ . The minimum number of leaf pages that can exist in this index is  $(d + 1)^3 = 125$ .

- (0.5 points) In the **worst case**, what is the **largest** number of index pages that will be **read** in an equality search?

A. 0	C. 3	E. 729	G. 732
B. 1	D. 4	F. 731	H. none of the above
- (0.5 points) In the **best case**, what is **fewest** number of index pages that will be **read** in an equality search?

A. 0	C. 3	E. 125	G. 128
B. 1	D. 4	F. 127	H. none of the above
- (0.5 points) In the **worst case**, what is the **largest** number of index pages that you will **write to** during an insert? Include both existing index pages that are written to and index pages that are created as part of the write in your answer.

A. 2	C. 4	E. 6	G. 8
B. 3	D. 5	F. 7	H. none of the above
- (0.5 points) In the **best case**, what is the **fewest** number of index pages that you will **write to** during an insert? Include both existing index pages that are written to and index pages that are created as part of the write in your answer.

A. 0	C. 2	E. 4	G. 6
B. 1	D. 3	F. 5	H. none of the above

6. (0.5 points) In the **worst case**, what is the **largest** number of index pages that will be **read** during a range search that returns exactly 7 records?

- |      |      |      |                      |
|------|------|------|----------------------|
| A. 2 | C. 4 | E. 6 | G. 8                 |
| B. 3 | D. 5 | F. 7 | H. none of the above |

7. (0.5 points) In the **best case**, what is the **fewest** number of index pages that will be **read** during a range search that returns exactly 7 records?

- |      |      |      |                      |
|------|------|------|----------------------|
| A. 2 | C. 4 | E. 6 | G. 8                 |
| B. 3 | D. 5 | F. 7 | H. none of the above |