CS315: Principles of Database Systems, IIT Kanpur				**Midsem** (17 Sep 2024)
Name				**40 marks** Page **1** of **4**
Roll No		Dept.		

Instructions:

1. This question paper contains 2 pages (4 sides of paper). Please verify.
2. Write your name, roll number, department above in **block letters neatly with ink**.
3. Write your final answers neatly **with a blue/black pen**. Pencil marks may get smudged.
4. Don't overwrite/scratch answers especially in MCQ – ambiguous cases may get 0 marks.
5. **Hardcoding attempts will not get any credit**.
6. **Be extremely precise in your answers and be careful not to make spelling or punctuation mistakes**. We may type your answers as SQLite queries to actual DB and give marks based on how correct the retrieved results are.

(DBs can do Math!) Deebo has an SQLite table `mth` with 4096 rows. The first column `num` contains integers between 1 and 4096 (both included). Each number occurs exactly once – no duplicates or missing numbers – but the numbers are not in sorted order. The second column `fav` has integers that are 0 or 1 indicating if that number is Deebo's favourite or not (1 = favourite, 0 = not). **Note: SQLite supports modular arithmetic** – if a, b are integers, then the expression $a \% b$ (or $\mathrm{mod}(a, b)$) will give the remainder of a when divided by b.

mth

num	fav
1	0
1729	1
42	1
…	
2607	0

Q1. Write an SQLite query to retrieve all even numbers from `num` sorted in ascending order. Your result should have a single column. **Using the mod operator will incur a 1 mark penalty**.(**3 marks**)

Q2. Write a query to retrieve all primes from `num` sorted in descending order (2 is a prime 1 is not). Your result should have a single column. **Using mod operator will incur 1 mark penalty**.(**4 marks**)

Q3. For each value n in the `num` column, count how many numbers $\leq n$ are Deebo's favourite using an SQLite query. Your result should have two columns – the first containing values from the `num` column sorted in descending order and the second containing the favourite counts.(**5 marks**)

For Q4,5,6,7, assume that the results of Q1, Q2 are available in views named `even` ***and*** `prime`***. Both views contain a single column containing all even numbers and primes respectively, sorted in ascending and descending order respectively. You may use these views to shorten your queries.***

Q4. Let's verify Goldbach's conjecture – *Every even number greater than 2 is the sum of two primes*. Write a query to retrieve 3 columns n, p, q. $n > 2$ should take even values from `num`, p, q must be primes with $p \leq q, n = p + q$. If n is a prime sum in multiple ways e.g. $14 = 3 + 11 = 7 + 7$, then there should be those many rows for n. If $p = q$, don't create cloned rows e.g. for 14, there should be only 2 rows $(14,3,11), (14,7,7)$, not 3 rows $(14,3,11), (14,7,7), (14,7,7)$. Sort results by n asc. If n has many rows then sort those by p asc e.g. $(14,3,11)$ comes just before $(14,7,7)$. (**5 marks**)

Q5. Create a view `succ` with 3 columns – the first with values n from `num` in ascending order, the second with the successor of n if it exists in `num` and `null` otherwise and the third containing the successor of n if it exists in `num` and is also Deebo's favourite and `null` otherwise. **(6 marks)**

Q6. Fill one box, give brief justification. Assume `succ` is a table, not a view. Deebo dislikes at least the numbers 1 and 2607, maybe others too. PK ≡ PRIMARY KEY, U ≡ UNIQUE**(3 x (1+1) = 6 marks)**

Can the first column of the `succ` table become a PRIMARY KEY or satisfy UNIQUE constraint?

- ☐ Only PK (not U)
- ☐ Only U (not PK)
- ☐ Both PK and U
- ☐ Neither PK nor U

Give justification here

Can the second column of the `succ` table become a PRIMARY KEY or satisfy UNIQUE constraint?

- ☐ Only PK (not U)
- ☐ Only U (not PK)
- ☐ Both PK and U
- ☐ Neither PK nor U

Give justification here

Can the third column of the `succ` table become a PRIMARY KEY or satisfy UNIQUE constraint?

- ☐ Only PK (not U)
- ☐ Only U (not PK)
- ☐ Both PK and U
- ☐ Neither PK nor U

Give justification here

Q7. Write an SQLite query to retrieve a bitmap index for primes. Your result should have 2 columns, the first having values from `num` sorted in descending order and the second containing only 0 or 1 depending on whether the number is prime or not (prime ⇒ 1, not prime ⇒ 0). **(5 marks)**

Q8. Dooba has written a relational expression to find Deebo's favourite perfect squares from `num` i.e. n s.t. $n = m^2$ for some m and n is a favourite. ⋈ without a θ expression does a natural join.

$$\pi_{\text{M1.num}}\left(\sigma_{(\text{M1.num}=\text{M2.num}*\text{M2.num})\vee(\text{M2.fav}=1)}\left(\rho_{\text{M1}}(\text{mth}) \bowtie \rho_{\text{M2}}(\text{mth})\right)\right)$$

Deebo suspects that Dooba's expression will not give the output as intended. Help Deebo make all corrections to the expression by filling the dashed boxes. Using your corrected expression, write an SQLite query to retrieve all favourite perfect squares sorted in ascending order. (**4+2=6 marks**)

$\pi_{\text{M1.num}}(\sigma_{(\quad)}$

$\rho_{\text{M1}}(\text{mth})\quad\quad\rho_{\text{M2}}(\text{mth}))$

SQLite query