2015 SIGMOD PROGRAMMING CONTEST



TASK OVERVIEW

- Given processed transactions on relations
- Concurrent validation requests on the transactions
 - e.g. (Rel. 1: $Q_1 \land Q_2 \land Q_3$) \lor (Rel. 2: $Q_4 \land Q_5$)
 - A Query consists of
 - Column index
 - Operator (==, >, >=, <, <=, !=)
 - Value (64-bit integer)
- Evaluate the validation queries ASAP



서울대학교 컴퓨터공학부 Seoul National University Dept. of Computer Science and Engineering



Eunjin Song, Seoung Gook Sohn, Wonha Ryu sgsohn@dbs.snu.ac.kr eunjin@csap.snu.ac.kr wonha.ryu@gmail.com

TRANSACTION

Transactions are stored as tables for each relation.

- Each relation object holds its own list of transactions
- Store in column-store fashion for cache efficiency
- Group small transactions into one transaction
- to reduce overhead
- Construct bloom filters for each row

SIMPLIFIED BLOOM FILTER





- Each value is viewed as a pair of column index and value
 - $[40, 20, 10] \Rightarrow$ equivalent set: {(1, 40), (2, 20), (3, 10)}
- Only a single hash function is used for performance issues
- Also stored in column-store fashion per transaction



VALIDATION REQUESTS

- Independently process each conjunction in validation requests
- Optimize according to dataset features:
- Reorder queries in a conjunction
- Check if the value of a condition is within minimum and maximum
- Prune candidate rows using bloom filters for equal queries
- Survived candidates are validated naively

REORDERING QUERIES

T _{n+4} Insert rows: 1 Delete rows: 3 	small-sized transactions and exceptional large-sized transactions	C ₂ has lower prob. of conflicting with equal		
Delete rows: 1 T _{n+3} Insert rows: 2380 Delete rows: 1643		3077 $C_1 \stackrel{?}{=} 3076$	5819204839 $C_2 \stackrel{?}{=} 573927302$	
Delete rows: 4 T _{n+1} Insert rows: 1 Delete rows: 2 T _{n+2} Insert rows: 1		3074 3077 3075	8896128387 6811549715 7782881939	

· · · /	- (1 /				
3074	8896128387				
3077	6811549715				
3075	7782881939				
3077	5819204839				
$C_1 \stackrel{?}{=} 3076 \qquad C_2 \stackrel{?}{=} 5739273029$					
C_2 has lower prob. of					
onflicting with equal					

IMPLEMENTATION DESIGN OVERVIEW



Original: $(C_1 < V_1) \land (C_2 == V_2) \land (C_3 == V_3) \land (C_4 >= V_4)$ Reordered: $(C_2 = V_2) \land (C_3 = V_3) \land (C_1 < V_1) \land (C_4 > = V_4)$

Reorder queries in the following order to reject ASAP

- 1. Query with == operator having the widest range of values
- 2. Queries with == operators
- 3. Queries with other operators

PRUNING WITH BLOOM FILTERS



- Only utilize bloom filters for equal queries
- A row is non-conflicting if any corresponding bits of its bloom filter are 0.
 - In above example, it is guaranteed that R_1 doesn't have 234 in the C_1 .
- **Reader** reads input messages and writes to the buffer.
- **Transactors** store transactions and construct bloom filters.
- **Pre-validators** separate and reorder validation requests by conjunctions, and push the result into the queue.
- **Validators** validate the requests from the queue.





- The bloom filters are stored column-wise to easily compute bitwise AND of the columns.
- SIMD bitwise AND can process one column of bloom filters of 256 rows at once.

