Lecture 8: SQL (Advanced)

CS348 Spring 2025: Introduction to Database Management

> Instructor: Xiao Hu Sections: 001, 002, 003

Announcement

• Assignment 1 Due on Jun 1

SQL features covered so far

- Basic topics:
 - Data-definition language (DDL)
 - Data-manipulation language (DML)
 - Constraints
- Advanced topics:
 - View

Post-Lecture: Views and Table Subquery

- By Default, a view is simply a stored query and the query result is NOT physically stored separately
 - But, systems can choose to materialize view (i.e., physically store the query results)
- A table subquery is a temporary table that physically stores data and it only exists in the duration of the whole query that uses it

POST-Lecture: View

Given any DML that attempts to modify view but violate the view's filter:

- If WITH CHECK OPTION: reject outright
- If WITH CHECK OPTION is not specified: it is possible to "sneak" rows into the base table through the view -- these rows simply won't appear in the view
 - However, such rows can only be inserted into the base table if they still satisfy the table's constraints

POST-Lecture: View

User (<u>uid</u> int, name string, age int, pop float) Group (<u>gid</u> string, name string) Member (<u>uid</u> int, <u>gid</u> string)

• Assume there is a CHECK constraint on User table such that age is above 0 and below 140

CREATE VIEW youngUsers AS (SELECT * FROM User WHERE age < 25) WITH CHECK OPTION;

• What happens to the following statements?

INSERT INTO youngUsers VALUES (835, 'Alex', 30, 0.2);

INSERT INTO youngUsers VALUES (923, 'James', 150, 0.3);

- (835, 'Alex', 30, 0.2) will not be inserted into User, and it also will not appear in youngUser
- Same for (923, 'James', 150, 0.3)

POST-Lecture: View

User (<u>uid</u> int, name string, age int, pop float) Group (<u>gid</u> string, name string) Member (<u>uid</u> int, <u>gid</u> string)

 Assume there is a CHECK constraint on User table such that age is above 0 and below 140

CREATE VIEW youngUsers AS (SELECT * FROM User WHERE age < 25);

• What happens to the following statements?

INSERT INTO youngUsers VALUES (835, 'Alex', 30, 0.2);

INSERT INTO youngUsers VALUES (923, 'James', 150, 0.3);

- (835, 'Alex', 30, 0.2) can be inserted into User table but it will not appear in youngUser
- (923, 'James', 150, 0.3) will not be inserted into User (since there is a CHECK constraint on the User table itself) and it also will not appear in youngUser

What is next?

- Basic topics:
 - Data-definition language (DDL)
 - Data-manipulation language (DML)
 - Constraints
- Advanced topics:
 - View
 - Trigger

(Recap) Referential Integrity

Example: Member.uid references User.uid

- Delete or update a User row whose uid is referenced by some Member row
 - Multiple Options (in SQL)



Option 3: Set NULL

CREATE TABLE Member (uid INTNOT NULL REFERENCES User(uid) ON DELETE CASCADE,);

Option 2: Cascade (ripple changes to all referring rows)

Can we generalize it?

Referential constraints

Delete/update a User row

Whether its uid is referenced by some Member row

If yes: Reject/ Delete cascade/Null



General constraints

Some user's popularity is updated

Whether the user is a member of "popgroup" and *pop* drops below 0.5

If yes: kick that user out of popgroup!

Triggers

- A trigger is an event-condition-action (ECA) rule
 - When event occurs, test condition; if condition is TRUE, execute action



Trigger option 1 – possible events

- Possible events include:
 - INSERT ON table; DELETE ON table; UPDATE [OF column] ON table



Trigger option 2 – timing

- Timing -- action can be executed:
 - AFTER or BEFORE the triggering event
 - INSTEAD OF the triggering event on views (more later)



Trigger option 3 – granularity

- Granularity the trigger can be activated:
 - FOR EACH ROW: execute its action for each individual affected rows in the event



Trigger option 3 – granularity

- Granularity -- the trigger can be activated:
 - FOR EACH STATEMENT: execute its action once for the entire SQL statement (rather than for each individual row affected by that statement)



Can only be used with **AFTER** triggers

INSTEAD OF triggers for views



UPDATE AveragePop SET pop = 0.5;

0.3 + 0.1

Transition variables/tables

- OLD ROW: the modified row before the triggering event
- NEW ROW: the modified row after the triggering event
- OLD TABLE: a (hypothetica) read-only table containing all old rows modified by the triggering event
- NEW TABLE: a (hypothetica) table containing all modified rows after the triggering event

Event	Row	Statement	
Delete	old r	old t	
Insert	new r	new t	
Update	old/new r	old/new t	
AFTER Trigger			

Event	Row	Statement		
Delete	old r	-		
Insert	new r	-		
Update	old/new r	-		
BEFORE Trigger				

In class exercises

• If a user with pop>0.5 is added to the User table, they must automatically belong to the 'popgroup'. Create a trigger using FOR EACH ROW to achieve this behavior.



In class exercises

• If a user with pop>0.5 is added to the User table, they must automatically belong to the 'popgroup'. Create a trigger using FOR EACH STATEMENT to achieve this behavior.



Statement- vs. Row-level triggers

- Simple row-level triggers are easier to implement
 - Statement-level triggers: require a significant amount of state to be maintained in OLD TABLE and NEW TABLE
- However, in some cases, a row-level trigger may be less efficient
 - E.g., 4B rows and a trigger may affect 10% of the rows. Recording an action for 400M rows, one at a time, is not feasible due to resource constraints.
- Certain triggers only possible at the statement level

Certain triggers are only possible at statement level



SQL features covered so far

- Basic topics:
 - Data-definition language (DDL)
 - Data-manipulation language (DML)
 - Integrity constraint
- Advanced topics:
 - View
 - Triggers (System issues)
 - Recursive firing of triggers
 - Interaction with constraints (tricky to get right!)
 - Best to avoid when alternatives exist
 - Recursion

A motivating example



- Example: find Bart's ancestors
- "Ancestor" has a recursive definition
 - X is Y's ancestor if
 - X is Y's parent, or
 - X is Z's ancestor and Z is Y's ancestor

Recursion in SQL

- SQL2 had no recursion
 - You can find Bart's parents, grandparents, great grandparents, etc.

```
SELECT p1.parent AS grandparent
FROM Parent p1, Parent p2
WHERE p1.child = p2.parent AND p2.child = 'Bart';
```

- But you cannot find all his ancestors with a single query
- SQL3 introduced recursion
 - WITH RECURSIVE clause
 - Many systems support recursion but limited functionality

Example of Ancestor Query



Finding ancestors

						parent	child
WITH RECURSIVE Ancestor(anc, desc) AS (SELECT parapt, child EPOM Parapt)					Homer	Bart	
					Homer	Lisa	
(SELECT parent, child FROM Parent)					Marge	Bart	
(SELECT al.anc. a2.desc					Marge	Lisa	
FROM Ancestor a1, Ancestor a2					Abe	Homer	
WHERE a1.desc = a2.anc))					Orville	Abe	
recursive step					anc	desc	
						Homer	Bart
			anc	desc		Homer	Lisa
Anc	estor		Homer	Bart		Marge	Bart
anc	desc	\rightarrow	Homer	Lisa	\rightarrow	Marge	Lisa
			Marge	Bart		Abe	Homer
			Marge	Lisa		Orville	Abe
			Abe	Homer		Abe	Bart
			Orville	Abe		Abe	Lisa
						Orville	Homer

		$\Big)$
	anc	desc
	Homer	Bart
	Homer	Lisa
	Marge	Bart
	Marge	Lisa
	Abe	Homer
•	Orville	Abe
	Abe	Bart
	Abe	Lisa
	Orville	Homer
	Orville	Bart
	Orville	Lisa

Parent

Another example of Ancestor Query



Fixed point of a function

- If $f: D \rightarrow D$ is a function from a type D to itself, a fixed point of f is a value x such that f(x) = x
 - Example: what is the fixed point of f(x) = x/2?
 - Answer: 0 since f(0)=0
- To compute a fixed point of *f* :
 - Start with a "seed": $x \leftarrow x_0$
 - Compute f(x)
 - If f(x) = x, stop; x is fixed point of f
 - (Similar to base case in recursive programming)
 - Otherwise, $x \leftarrow f(x)$; repeat

Fixed point of a query

- A query q is a function that maps an input table to an output table
- A fixed point of q is a table T such that q(T) = T
- To compute the fixed point of *q*:
 - Start with executing q on a base table $T_0: T \leftarrow q(T_0)$
 - Evaluate q over T
 - If q(T) = T, stop; T is a fixed point of q
 - Otherwise, let q(T) be the new input table; repeat

Restrictions on recursive queries

Lecture 3

- A recursive query *q* must be monotonic
 - If more tuples are added to the recursive relation, *q* must return at least the same set of tuples as before, and possibly return additional tuples
- The following is not allowed in q:
 - Aggregation on the recursive relation
 - NOT EXISTS/NOT IN in generating the recursive relation
 - Set difference (EXCEPT) whose right-hand side uses the recursive relation

SQL features covered so far

- Basic topics:
 - Data-definition language (DDL)
 - Data-manipulation language (DML)
 - Integrity constraint
- Advanced topics:
 - View
 - Triggers
 - Recursion
 - Index

Motivation of using indexes

SELECT * FROM User WHERE name = 'Bart';

 Can we go "directly" to User rows with name='Bart' instead of scanning the entire table? Index on User.name

SELECT * FROM User, Member WHERE User.uid = Member.uid AND Member.gid = 'popgroup';

- For each Member row, can we go "directly" to Member rows with gid = 'popgroup'? Index on Member.gid
- For each Member row, can we "directly" look up User rows with matching Member.uid? Index on User.uid

Indexes

- An index is an auxiliary persistent data structure that helps with efficient searches
 - Search tree (e.g., B⁺-tree), lookup table (e.g., hash table), etc.

More on indexes later in this course!

CREATE INDEX unique_name ON USER(name);

DROP INDEX unique name;

• Typically, the DBMS will automatically create indexes for PRIMARY KEY and UNIQUE constraint declarations

Indexes

- An index on R(A) can speed up accesses of the form
 - A = x
 - A > x (depends on the index)
 - Can be extended to multi-attribute index $R(A_1, A_2, ..., A_n)$
- Questions (Indexing Lecture):
 Crdering of index columns is important -- is an index on R(A, B) equivalent to one on R(B, A)?
 How about an index on R(A) plus another on R(B)?
 More indexes = better performance?

Summary of SQL

- Basic topics
 - Data-definition language (DDL): define/modify schemas, drop relations
 - Data-manipulation language (DML): query data
 - SELECT-FROM-WHERE
 - DISTINCT, UNION/EXCEPT/INTERSECT (ALL)
 - Table, Scalar, IN, EXISTS, ALL, ANY)
 - GROUP BY, HAVING
 - ORDER
 - NULL and JOIN and modify data (INSERT/DELETE/UPDATE)
 - Constraints (NOT NULL, UNIQUE, PRIMARY/FOREIGN KEY, CHECK, ASSERTION)
- Advanced topics
 - View, Triggers, Recursion, Index, Programming (optional)

SQL Programming (optional)

- Pros and cons of SQL
 - Very high-level, possible to optimize
 - Not intended for general-purpose computation
- Can SQL and general-purpose programming languages (PL) interact with each other?

YES!!

Dynamic SQL Build SQL statements at runtime using APIs provided by DBMS

Embedded SQL

SQL statements embedded in general-purpose PL; identified at compile time

Mismatch b/w SQL and PLs (optional)

- SQL operates on a set of records at a time
- Typical low-level general-purpose programming languages operate on one record at a time

Solution: cursor

- Open (a result table), Get next, Close
- Found in virtually every database language/API
 - With slightly different syntaxes

Dynamic SQL

Working with SQL through an API

- Example: Python psycopg2, JDBC, ODBC (C/C++)
 - All based on the SQL/CLI (Call-Level Interface) standard
- The application program sends SQL commands to the DBMS at runtime
- Responses/results are converted to objects in the application program

Example API: Python psycopg2



Example API: Python psycopg2



Example API: Python psycopg2

```
# "commit" each change immediately—need to set this option just once at
the start of the session
                                                       Perform parsing,
conn.set_session(autocommit=True)
                                                       semantic analysis,
# ...
                                                         optimization,
uid = input('Enter the user id to update: ').strip()
                                                       compilation, and
name = input('Enter the name to update: ').strip()
                                                        final execution
pop = float(input('Enter new pop: '))
cur.execute("
        UPDATE User
        SET pop = %s
        WHERE uid = %s AND name = %s", (pop, uid, name))
```

What's next?

• Lecture 9: Database Design

