CS145 Autumn 2007

Problem Session 5

Midterm Review

Yusheng Yang
Stanford University
10/29/2007

Outline

- Announcements
- SQL
- Transactions
- DTDs
- Q&A

Announcements

- Project 1 due Wed 10/31
- Gradiance – one due Fri 11/2, two due Wed 11/7
- Midterm Wed 10/31 11am-12:15pm Gates B01
  - Up to and including 10/24 lecture on XPath
  - Open notes/book/laptop. Closed Internet.

SQL

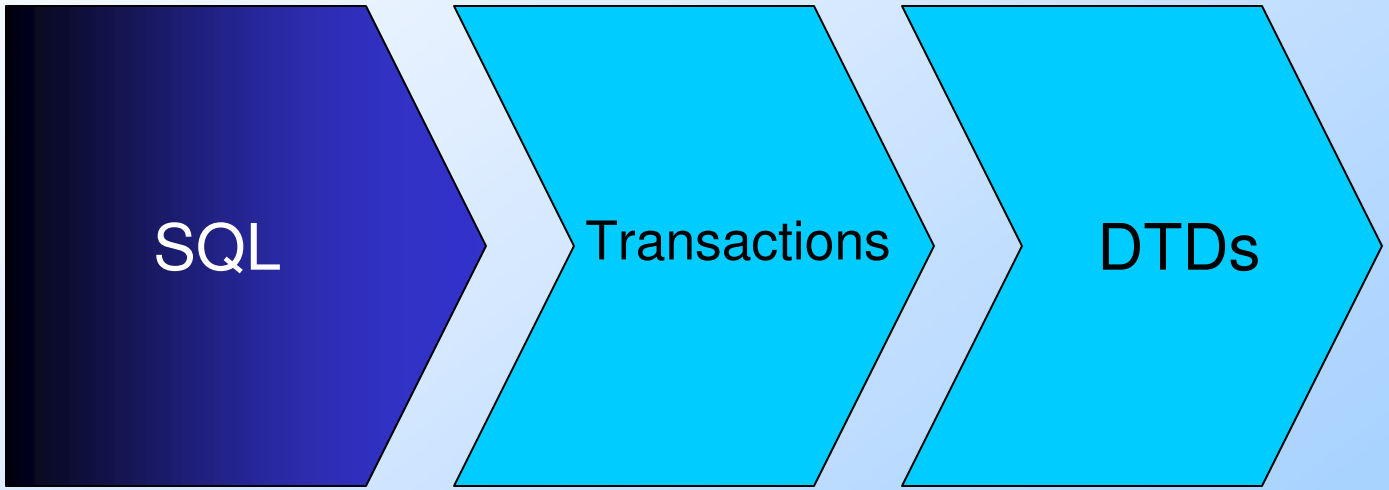

From 2005 Midterm

Consider a table **Exams(student, score)**.

Write a SQL query to find **the student with the highest score differential**, i.e. the student with the largest spread between his or her highest and lowest scores,

among all students with scores in the table.

Assume there is a unique student with the highest spread and return that student only once.

Sample Solution

```
SELECT student
FROM Exams
GROUP BY student
HAVING MAX(score) - MIN(score) >= ALL(
        SELECT MAX(score) - MIN(score)
        FROM Exams
        WHERE score IS NOT NULL
        GROUP BY student
)
```

Transactions

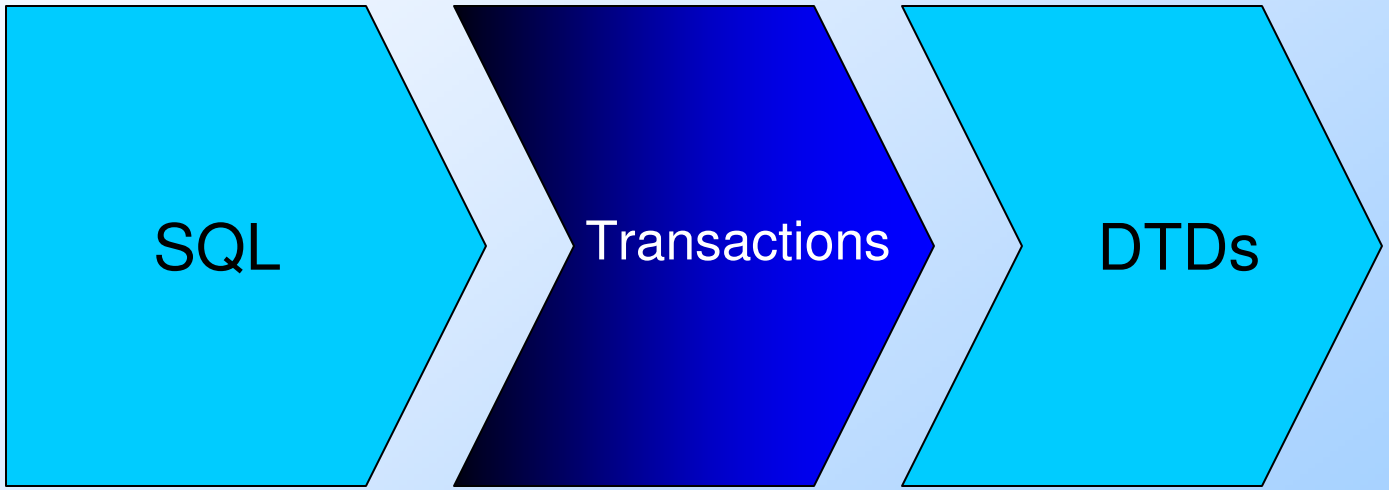

From Lecture

Joe_Sells(beer, price). Initially: (Bud, 2.50) and (Miller, 3).

Sally: BEGIN TRANSACTION
S1: SELECT MAX(price) FROM Joe_Sells
S2: SELECT MIN(price) FROM Joe_Sells
COMMIT

Joe: BEGIN TRANSACTION
S3: DELETE FROM Joe_Sells
S4: INSERT INTO Joe_Sells VALUES('Heineken', 3.50);
COMMIT

Suppose S1,S3,S4,Joe commits,S2,Sally commits.

Solution

a) Sally: SERIALIZABLE: MAX = $3.00, MIN = $2.50.
b) Sally: REPEATABLE READ: MAX = $3.00, MIN = $2.50. Sally saw a **phantom** tuple: ('Heineken', $3.50).
c) Sally: READ COMMITTED: MAX = $3.00, MIN = $3.50. Sally saw Joe's committed deletion.
d) Sally: READ UNCOMMITTED: MAX = $3.00, MIN = $3.50. Sally saw Joe's uncommitted deletion.

Question: What isolation level do you think Oracle supports as a default?
Answer: REPEATABLE READ. Guarantees no loss of data.

DTDs

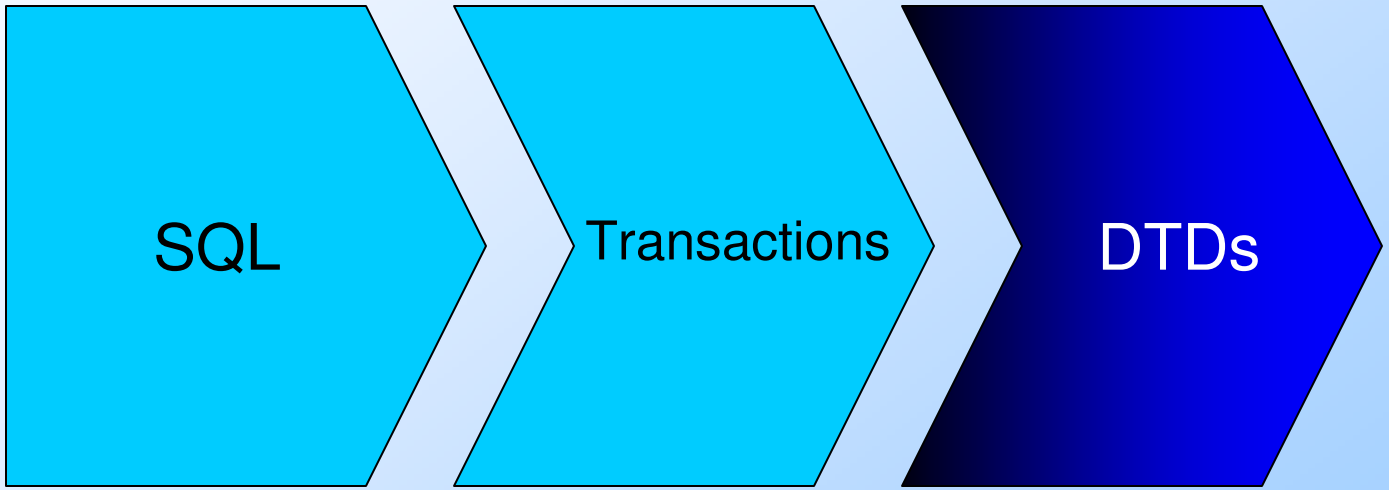

From 2006 Midterm

DTD1: <!DOCTYPE SP [
<!ELEMENT SP (Project*)>
<!ELEMENT Project (Title, Student+)>
<!ATTLIST Project ProjNum ID>
<!ELEMENT Title (#PCDATA)>
<!ELEMENT Student>
<!ATTLIST Student StudID ID Name CDATA>]>

1) For each project, there is
   a) exactly one student; b) at least one student
2) For each student, there is
   a) exactly one project; b) at least one project

Answer: 1b, 2a

From 2006 Midterm

DTD2: <!DOCTYPE SP [
<!ELEMENT SP (Student*)>
<!ELEMENT Student (Project)>
<!ATTLIST Student StudID ID Name CDATA>
<!ELEMENT Project (Title)>
<!ATTLIST Project ProjNum ID>
<!ELEMENT Title (#PCDATA)>]>

1) For each project, there is
   a) exactly one student; b) at least one student
2) For each student, there is
   a) exactly one project; b) at least one project

Answer: 1a, 2a

From 2006 Midterm

DTD3: <!DOCTYPE SP [
<!ELEMENT SP (Project*, Student*)>
<!ELEMENT Project (Title)>
<!ATTLIST Project ProjNum ID stud IDREF>
<!ELEMENT Title (#PCDATA)>
<!ELEMENT Student>
<!ATTLIST Student StudID ID Name CDATA>]>

1) For each project, there is
   a) exactly one student; b) at least one student
2) For each student, there is
   a) exactly one project; b) at least zero projects

Answer: 1a, 2b

From 2006 Midterm

DTD4: <!DOCTYPE SP [
<!ELEMENT SP (Student*, Project*)>
<!ELEMENT Student>
<!ATTLIST Student StudID ID Name CDATA proj IDREFS>
<!ELEMENT Project (Title)>
<!ATTLIST Project ProjNum ID>
<!ELEMENT Title (#PCDATA)>]>

1) For each project, there is
   a) exactly one student; b) at least zero students
2) For each student, there is
   a) exactly one project; b) at least one project

Answer: 1b, 2b

Midterm Topics

* Relational Algebra: union/intersect/difference/select/project/product/join/rename
* SQL: select/from/where
* SQL: multirelational queries
* SQL: subqueries
* SQL: outerjoins
* SQL: group by/having
* SQL: insert/delete/update
* SQL: constraints
* SQL: triggers
* SQL: transactions
* SQL: views
* SQL: indexes
* XML: DTDs
* XML: XML Schema
* XML: XPath

Q & A

Questions?

1. lecture notes
2. Coursework Discussion
3. Office Hours (Mon 1-2 Gates 433, Tue 1-4, 8-11pm Gates B24A)
4. cs145-aut0708-staff@lists.stanford.edu