



CMU SCS

Carnegie Mellon Univ. Dept. of Computer Science 15-415/615 - DB Applications

C. Faloutsos & A. Pavlo
Lecture#7 : *Rel. model - SQL part2*
(*R&G, chapters 5-6*)



CMU SCS

General Overview - rel. model

- Formal query languages
 - rel algebra and calculi
- Commercial query languages
 - SQL
 - QBE, (QUEL)

Faloutsos, Pavlo
CMU SCS 15-415/615
#2



CMU SCS

Overview - detailed - SQL

- DML
 - select, from, where, renaming, ordering,
 - aggregate functions, nested subqueries
 - insertion, deletion, update
- other parts: DDL, authorization, triggers
- embedded SQL

Faloutsos, Pavlo
CMU SCS 15-415/615
#3



CMU SCS

Reminder: our Mini-U db

STUDENT		
Ssn	Name	Address
123	smith	main str
234	jones	forbes ave

CLASS		
c-id	c-name	units
15-413	s.e.	2
15-412	o.s.	2

TAKES		
SSN	c-id	grade
123	15-413	A
234	15-413	B

Faloutsos, Pavlo
CMU SCS 15-415/615
#4



CMU SCS

DML - insertions etc

insert into student
values ("123", "smith", "main")

insert into student(ssn, name, address)
values ("123", "smith", "main")

Faloutsos, Pavlo CMU SCS 15-415/615 #5



CMU SCS

DML - insertions etc

bulk insertion: how to insert, say, a table of
'foreign-student's, in bulk?

Faloutsos, Pavlo CMU SCS 15-415/615 #6



CMU SCS

DML - insertions etc

bulk insertion:

insert into student
 select ssn, name, address
 from foreign-student

Faloutsos, Pavlo CMU SCS 15-415/615 #7



CMU SCS

DML - deletion etc

delete the record of 'smith'

Faloutsos, Pavlo CMU SCS 15-415/615 #8



CMU SCS

DML - deletion etc

delete the record of 'smith':

```
delete from student  
where name='smith'
```

(careful - it deletes ALL the 'smith's!)

Faloutsos, Pavlo CMU SCS 15-415/615 #9



CMU SCS

DML - update etc

record the grade 'A' for ssn=123 and course 15-415

```
update takes  
set grade="A"  
where ssn="123" and c-id="15-415"
```

(will set to "A" ALL such records)

Faloutsos, Pavlo CMU SCS 15-415/615 #10



CMU SCS

DML - joins

so far: 'INNER' joins, eg:

```
select ssn, c-name  
from takes, class  
where takes.c-id = class.c-id
```

Faloutsos, Pavlo CMU SCS 15-415/615 #12



CMU SCS

DML - joins

Equivalently:

```
select ssn, c-name  
from takes join class on takes.c-id = class.c-id
```

Faloutsos, Pavlo CMU SCS 15-415/615 #13

CMU SCS

Joins

```

select [column list]
from table_name
  [inner | {left | right | full} outer ] join
  table_name
  on qualification_list
where...
    
```

Faloutsos, Pavlo CMU SCS 15-415/615 #14

CMU SCS

Reminder: our Mini-U db

STUDENT		
Ssn	Name	Address
123	smith	main str
234	jones	forbes ave

CLASS		
c-id	c-name	units
15-413	s.e.	2
15-412	o.s.	2

TAKES		
SSN	c-id	grade
123	15-413	A
234	15-413	B

Faloutsos, Pavlo CMU SCS 15-415/615 #15

CMU SCS

Inner join

TAKES		
SSN	c-id	grade
123	15-413	A
234	15-413	B

CLASS		
c-id	c-name	units
15-413	s.e.	2
15-412	o.s.	2

SSN	c-name
123	s.e
234	s.e

o.s.: gone!

Faloutsos, Pavlo CMU SCS 15-415/615 #16

CMU SCS

Outer join

TAKES		
SSN	c-id	grade
123	15-413	A
234	15-413	B

CLASS		
c-id	c-name	units
15-413	s.e.	2
15-412	o.s.	2

SSN	c-name
123	s.e
234	s.e.
null	o.s.

←

Faloutsos, Pavlo CMU SCS 15-415/615 #17

CMU SCS

Outer join

select ssn, c-name
from takes **right outer join** class **on** takes.c-id=class.c-id

SSN	c-name
123	s.e
234	s.e.
null	o.s.

Faloutsos, Pavlo CMU SCS 15-415/615 #18

CMU SCS

Outer join

- left outer join
- right outer join
- full outer join
- natural join

Faloutsos, Pavlo CMU SCS 15-415/615 #19

CMU SCS

Null Values

- **null** -> unknown, or inapplicable, (or ...)
- Complications:
 - 3-valued logic (true, false and *unknown*).
 - **null = null** : false!!

Faloutsos, Pavlo CMU SCS 15-415/615 #20

CMU SCS

Overview - detailed - SQL

- DML
 - select, from, where, renaming, ordering,
 - aggregate functions, nested subqueries
 - insertion, deletion, update
- other parts: **DDL**, authorization, triggers
- embedded SQL

Faloutsos, Pavlo CMU SCS 15-415/615 #21



CMU SCS

Data Definition Language

```
create table student
(ssn char(9) not null,
 name char(30),
 address char(50),
 primary key (ssn) )
```

Faloutsos, Pavlo CMU SCS 15-415/615 #22



CMU SCS

Data Definition Language

```
create table r( A1 D1, ..., An Dn,
 integrity-constraint1,
 ...
 integrity-constraint-n)
```

Faloutsos, Pavlo CMU SCS 15-415/615 #23



CMU SCS

Data Definition Language

Domains:

- **char(n), varchar(n)**
- **int, numeric(p,d), real, double precision**
- **float, smallint**
- **date, time**

Faloutsos, Pavlo CMU SCS 15-415/615 #24



CMU SCS

Data Definition Language

```
delete a table: difference between
drop table student

delete from student
```

Faloutsos, Pavlo CMU SCS 15-415/615 #25



CMU SCS

Data Definition Language

modify a table:

alter table student **drop** address

alter table student **add** major char(10)

Faloutsos, Pavlo CMU SCS 15-415/615 #26



CMU SCS

Data Definition Language

integrity constraints:

- **primary key**
- **foreign key**
- **check(P)**

Faloutsos, Pavlo CMU SCS 15-415/615 #27



CMU SCS

Data Definition Language

create table takes

(ssn **char**(9) **not null**,
 c-id **char**(5) **not null**,
 grade **char**(1),
primary key (ssn, c-id),
check grade in (“A”, “B”, “C”, “D”, “F”))

Faloutsos, Pavlo CMU SCS 15-415/615 #28



CMU SCS

Referential Integrity constraints

‘foreign keys’ - eg:

create table takes(
 ssn **char**(9) **not null**,
 c-id **char**(5) **not null**,
 grade **integer**,
primary key(ssn, c-id),
foreign key ssn **references** student,
foreign key c-id **references** class)

Faloutsos, Pavlo CMU SCS 15-415/615 #29



CMU SCS

Referential Integrity constraints

...

foreign key ssn references student,
foreign key c-id references class)

Effect:

- expects that ssn to exist in 'student' table
- blocks ops that violate that - how??
 - insertion?
 - deletion/update?

Faloutsos, Pavlo CMU SCS 15-415/615 #30



CMU SCS

Referential Integrity constraints

...

foreign key ssn references student
on delete cascade
on update cascade,

...

- -> eliminate all student enrollments
- other options (set to null, to default etc)

Faloutsos, Pavlo CMU SCS 15-415/615 #31



CMU SCS

Overview - detailed - SQL

- DML
 - select, from, where, renaming, ordering,
 - aggregate functions, nested subqueries
 - insertion, deletion, update
- other parts: DDL, authorization, **triggers**
- embedded SQL

Faloutsos, Pavlo CMU SCS 15-415/615 #32



CMU SCS

Weapons for IC:

- assertions
 - **create assertion** <assertion-name> **check** <predicate>
- triggers (~ assertions with 'teeth')
 - **on** operation, **if** condition, **then** action

Faloutsos, Pavlo CMU SCS 15-415/615 #33



CMU SCS

Triggers - example

define trigger zerograde **on update** takes
(**if new** takes.grade < 0
then takes.grade = 0)

Faloutsos, Pavlo CMU SCS 15-415/615 #34



CMU SCS

Triggers - discussion

- more complicated: “managers have higher salaries than their subordinates” - a trigger can automatically boost mgrs salaries
- triggers: tricky (infinite loops...)

Faloutsos, Pavlo CMU SCS 15-415/615 #35



CMU SCS

Overview - detailed - SQL

- DML
 - select, from, where, renaming, ordering,
 - aggregate functions, nested subqueries
 - insertion, deletion, update
- other parts: DDL, **authorization**, triggers
- embedded SQL

Faloutsos, Pavlo CMU SCS 15-415/615 #36



CMU SCS

Authorization

- **grant** <priv.-list> **on** <table-name> **to** <user-list>
- privileges for tuples: ??
- privileges for tables: ??

Faloutsos, Pavlo CMU SCS 15-415/615 #37



CMU SCS

Authorization

- **grant** <priv.-list> **on** <table-name> **to** <user-list>
- privileges for tuples: read / insert / delete / update
- privileges for tables: create, drop, index

Faloutsos, Pavlo CMU SCS 15-415/615 #38



CMU SCS

Authorization – cont'd

- variations:
 - **with grant option**
 - **revoke** <priv.-list> **on** <t-name> **from** <user_ids>
- Eg:
- **grant read, update on takes to csSecy with grant option**

Faloutsos, Pavlo CMU SCS 15-415/615 #39



CMU SCS

Overview - detailed - SQL

- DML
 - select, from, where, renaming, ordering,
 - aggregate functions, nested subqueries
 - insertion, deletion, update
- ➔ other parts: **DDL/views**, authorization, triggers
- embedded SQL

Faloutsos, Pavlo CMU SCS 15-415/615 #40



CMU SCS

Views

find the ssn with the highest GPA -
we can create a permanent, virtual table:

create view helpfulTable(ssn, gpa) **as**
select ssn, **avg**(grade)
from takes
group by ssn

Faloutsos, Pavlo CMU SCS 15-415/615 41



CMU SCS

Views

- views are recorded in the schema, for ever (ie., until ‘**drop view...**’)
- typically, they take little disk space, because they are computed on the fly
- (but: materialized views...)

Faloutsos, Pavlo CMU SCS 15-415/615 42



CMU SCS

Subtle issue

DML - view update

consider the student_gpa view:
create view student_gpa **as**
 (select ssn, avg(grade) **as** gpa **from** takes)
 view updates are tricky - typically, we can only
 update views that have no joins, nor aggregates
 Example?

ssn	c-id	grade
123	15-415	4
123	15-412	3

ssn	gpa
123	3.5

Faloutsos, Pavlo CMU SCS 15-415/615 #43



CMU SCS

Overview - detailed - SQL

- DML
 - select, from, where, renaming, ordering,
 - aggregate functions, nested subqueries
 - insertion, deletion, update
- other parts: DDL, authorization, triggers
- embedded SQL**; application development

Faloutsos, Pavlo CMU SCS 15-415/615 #44



CMU SCS

Embedded SQL

from within a ‘host’ language (eg., ‘C’, ‘VB’)
 EXEC SQL <emb. SQL stmt> END-EXEC

Q: why do we need embedded SQL??

Faloutsos, Pavlo CMU SCS 15-415/615 #45



CMU SCS

Embedded SQL

from within a 'host' language (eg., 'C', 'VB')

```
EXEC SQL <emb. SQL stmt> END-EXEC
```

Q: why do we need embedded SQL??

A1: pretty formatting (eg., receipts, paychecks)

A2: plots

A3: forecasting; fancy math/stat computations

Faloutsos, Pavlo CMU SCS 15-415/615 #46



CMU SCS

Overview - detailed - SQL

- DML
 - select, from, where, renaming, ordering,
 - aggregate functions, nested subqueries
 - insertion, deletion, update
- other parts: DDL, authorization, triggers
- embedded SQL; **application development**

Faloutsos, Pavlo CMU SCS 15-415/615 #52



CMU SCS

Overview

- concepts of SQL programs
- walkthrough of embedded SQL example

Faloutsos, Pavlo CMU SCS 15-415/615 #53



CMU SCS

Outline of an SQL application

- establish connection with db server
- authenticate (user/password)
- execute SQL statement(s)
- process results
- close connection

Faloutsos, Pavlo CMU SCS 15-415/615 #54

CMU SCS

```

csv2sql.py
-----
#!/usr/bin/python
#####
# Author: christos faloutsos
# Date: Jan. 2012
# Purpose: Mainly wants to illustrate cursors
# Specifically,
# * expects a csv file, and
# * loads it into a sqlite db file
# And answers a few queries, for fun
#####

import sqlite3
import csv

fname='tst.csv'
dbname='tst.db'

# conn = sqlite3.connect('memory:')
conn = sqlite3.connect(dbname)

conn.execute('create table if not exists tst (name text, address text, state text, s
alary integer)')

csvreader = csv.reader(open(fname), delimiter=',', quotechar='"')

print " --- csv2sql started on ", fname
for row in csvreader:
    conn.execute('insert into tst values' +
        ' (%s, %s, %s, %s)' % row)

print " --- csv2sql inserted ", fname
print " --- printing all tuples --- "
cur = conn.cursor()
cur.execute('select * from tst')
for row in cur:
    for elem in row:
        print elem, '\t',
    print

print " --- printing rich tuples (as%20set), sorted ---- "
cur.execute('select * from tst where salary > 2000 order by salary desc ')
for row in cur:
    for elem in row:
        print elem, '\t',
    print

print " --- printing sum-salary per state ---- "
cur.execute('select state, sum(salary) from tst group by state')
for row in cur:
    for elem in row:
        print elem, '\t',
    print

conn.commit()
cur.close()
conn.close()
    
```

Faloutsos, Pavlo #57

CMU SCS

Check python code

- <http://www.cs.cmu.edu/~christos/courses/dbms.S15/PYTHON-examples/csv2sql.py>
- Or follow instructions at
- <http://www.cs.cmu.edu/~christos/courses/dbms.S15/PYTHON-examples/>

Faloutsos, Pavlo CMU SCS 15-415/615 #58

CMU SCS

```

#####
# Author: christos faloutsos
# Date: Jan. 2012
# Purpose: Mainly wants to illustrate cursors
# Specifically,
# * expects a csv file, and
# * loads it into a sqlite db file
# And answers a few queries, for fun
#####

import sqlite3
import csv

fname='tst.csv'
dbname='tst.db'

# conn = sqlite3.connect('memory:')
conn = sqlite3.connect(dbname)

conn.execute('create table if not exists tst (name text, address text, s
alary integer)')
    
```

Faloutsos, Pavlo CMU SCS 15-415/615 #59

CMU SCS

```

print " --- csv2sql inserted ", fname
print " --- printing all tuples --- "
cur = conn.cursor()
cur.execute('select * from tst')
for row in cur:
    for elem in row:
        print elem, "\t",
    print " "
    
```

Faloutsos, Pavlo CMU SCS 15-415/615 #60

CMU SCS

```

print " --- csv2sql inserted ", fname
print " "
print " --- printing all tuples --- "
cur = conn.cursor()
cur.execute('select * from tst')
for row in cur:
    for elem in row:
        print elem, "\t",
    print " "

```

Faloutsos, Pavlo CMU SCS 15-415/615 #61

CMU SCS

```

print " --- csv2sql inserted ", fname
print " "
print " --- printing all tuples --- "
cur = conn.cursor()
cur.execute('select * from tst')
for row in cur:
    for elem in row:
        print elem, "\t",
    print " "

```

Faloutsos, Pavlo CMU SCS 15-415/615 #62

CMU SCS

```

print " "
print " --- printing sum-salary per state ----"
cur.execute('select state, sum(salary) from tst group by state')
for row in cur:
    for elem in row:
        print elem, "\t",
    print " "

```

```

conn.commit()
cur.close()
conn.close()

```

Faloutsos, Pavlo CMU SCS 15-415/615 #63

CMU SCS

Conclusions

Outline of an SQL application:

- establish connection with db server
- authenticate (user/password)
- execute SQL statement(s) (using **cursors**)
- process results
- close connection

Faloutsos, Pavlo CMU SCS 15-415/615 #64