

CSE 444: Database Internals

Lecture 3 DBMS Architecture

CSE 444 – Spring 2018

1

Announcements

- Lab 1 part 1 due Monday
- Turn-in script has a bug:
 - “failed to push some refs ...”
 - Probably due to our repo setup and not your personal copy
 - The script still tags your last commit correctly, we'll figure out the bug and update you

CSE 444 – Spring 2018

4

What we already know...

- **Database** = collection of related files
- **DBMS** = program that manages the database

CSE 444 – Spring 2018

6

What we already know...

- **Data models**: relational, semi-structured (XML), graph (RDF), key-value pairs
- **Relational model**: defines only the logical model, and does not define a physical storage of the data

CSE 444 – Spring 2018

7

What we already know...

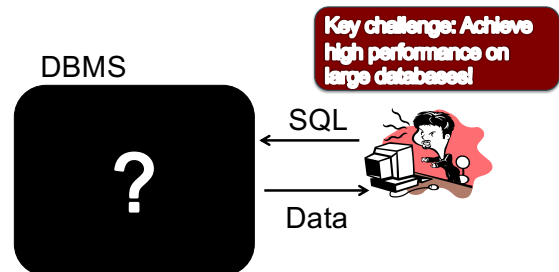
Relational Query Language:

- **Set-at-a-time**: instead of tuple-at-a-time
- **Declarative**: user says what they want and not how to get it
- **Query optimizer**: from *what* to *how*

CSE 444 – Spring 2018

8

How to Implement a Relational DBMS?



CSE 444 – Spring 2018


9

DBMS Architecture

CSE 444 – Spring 2018

10

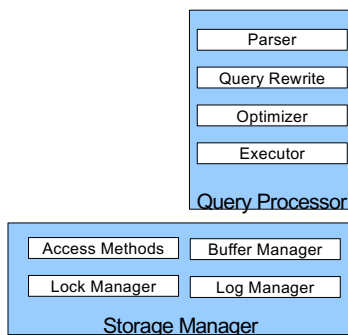
DBMS Architecture



CSE 444 – Spring 2018

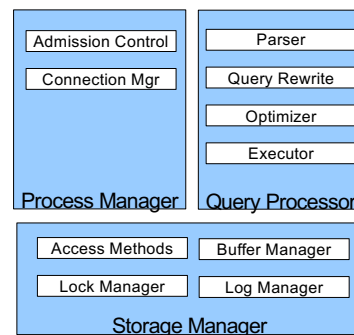
11

DBMS Architecture



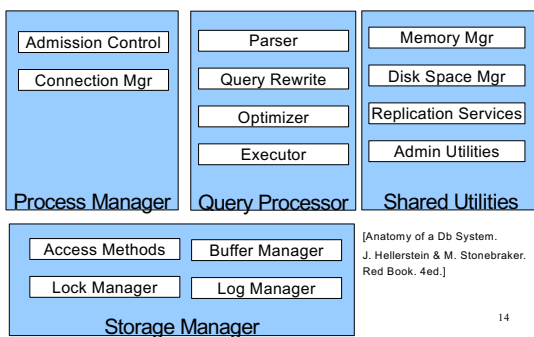
12

DBMS Architecture



13

DBMS Architecture



[Anatomy of a Db System.
J. Hellerstein & M. Stonebraker.
Red Book. 4ed.]

14

Goal for Today

Overview of query execution

Overview of storage manager

CSE 444 – Spring 2018

15

Query Processor

CSE 444 – Spring 2018

16

Example Database Schema

Supplier (sno, sname, scity, sstate)
Part (pno, pname, psize, pcolor)
Supplies (sno, pno, price)

View: Suppliers in Seattle

```
CREATE VIEW NearbySupp AS
SELECT sno, sname
FROM Supplier
WHERE scity='Seattle' AND sstate='WA'
```

17

Example Query

- Find the names of all suppliers in Seattle who supply part number 2

```
SELECT sname
FROM NearbySupp
WHERE sno IN ( SELECT sno
                FROM Supplies
                WHERE pno = 2 )
```

CSE 444 – Spring 2018

18

Query Processor

- Step 1: Parser**
 - Parses query into an internal format
 - Performs various checks using **catalog**
- Step 2: Query rewrite**
 - View rewriting, flattening, etc.

CSE 444 – Spring 2018

19

Rewritten Version of Our Query

Original query:

```
SELECT sname
FROM NearbySupp
WHERE sno IN ( SELECT sno
                FROM Supplies
                WHERE pno = 2 )
```

Rewritten query (expanding NearbySupp view):

```
SELECT S.sname
FROM Supplier S, Supplies U
WHERE S.scity='Seattle' AND S.sstate='WA'
AND S.sno = U.sno
AND U.pno = 2;
```

CSE 444 – Spring 2018

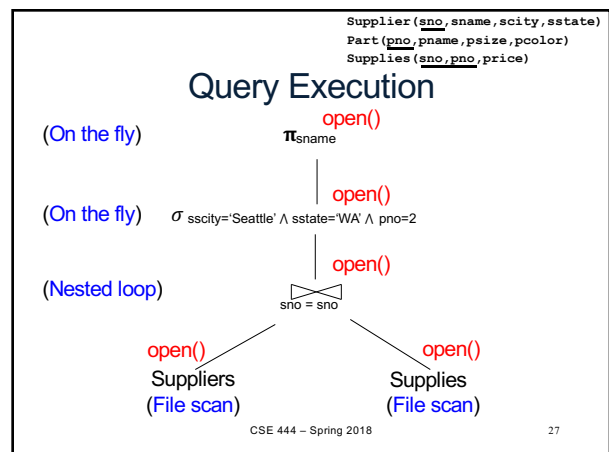
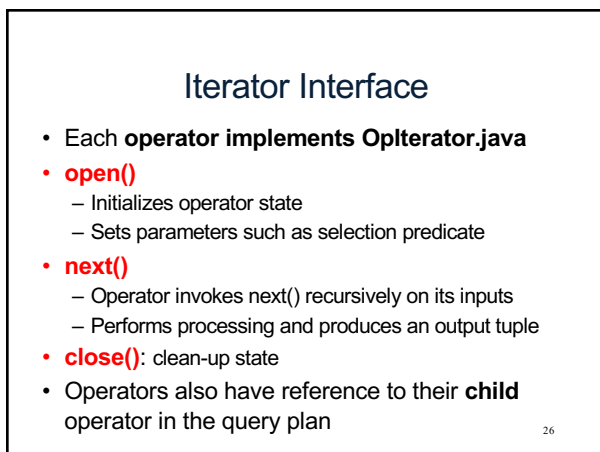

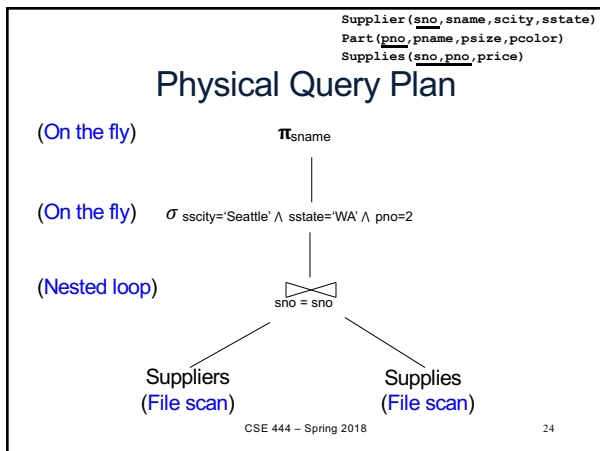
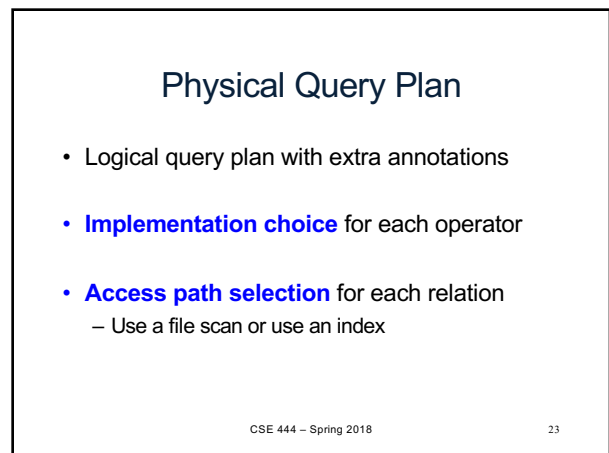
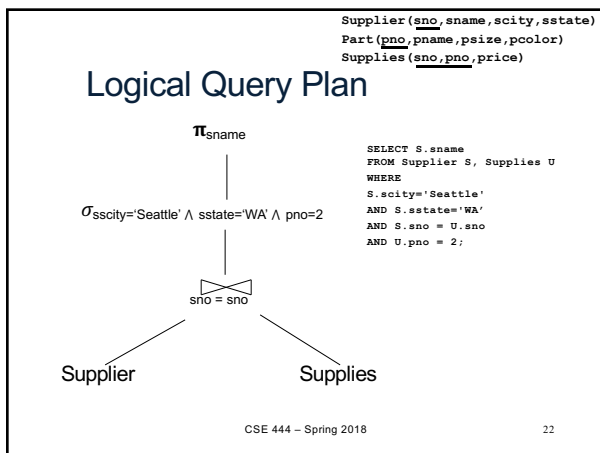
20

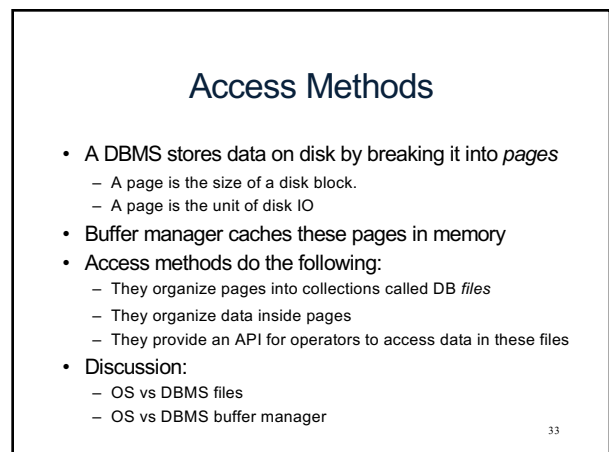
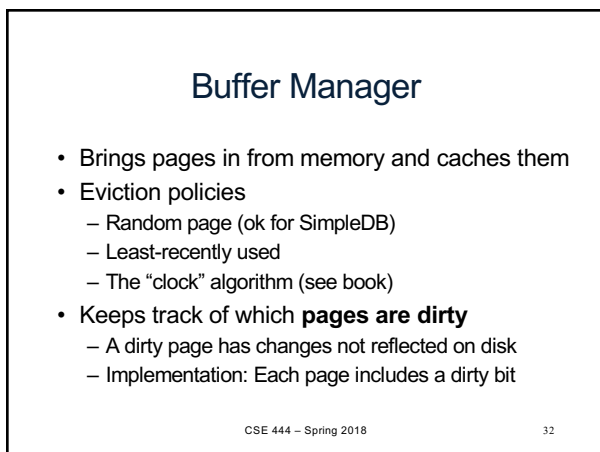
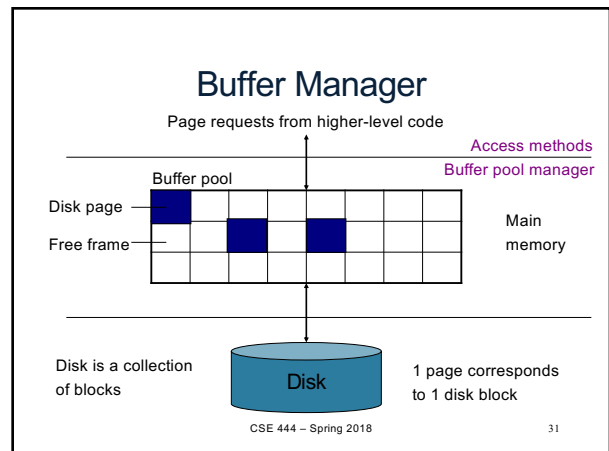
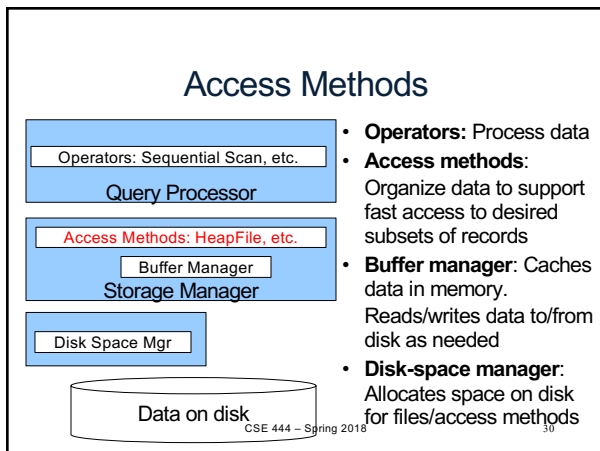
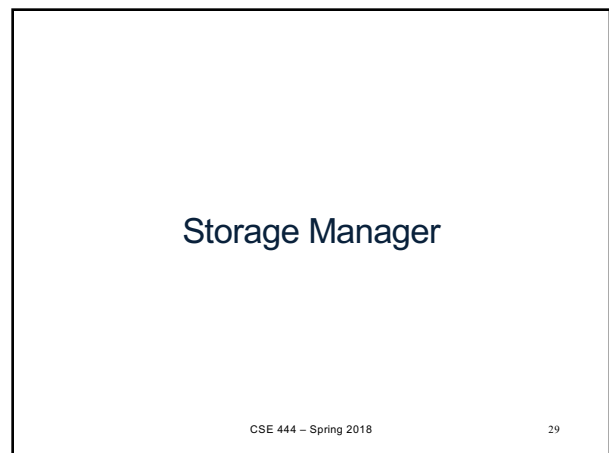
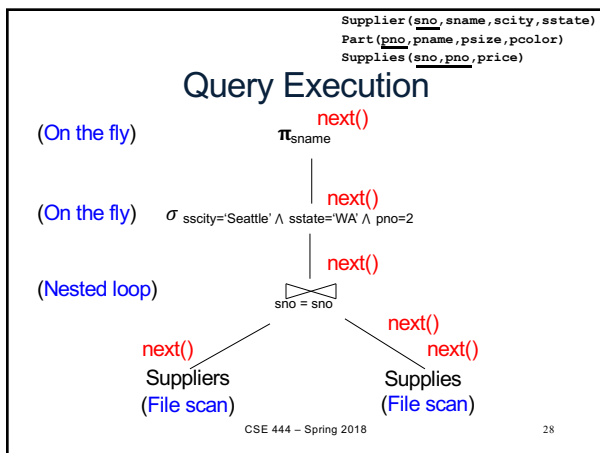
Query Processor

- Step 3: Optimizer**
 - Find an efficient query plan for executing the query
 - A query plan is**
 - Logical:** An extended relational algebra tree
 - Physical:** With additional annotations at each node
 - Access method to use for each relation
 - Implementation to use for each relational operator
- Step 4: Executor**
 - Actually executes the physical plan


CSE 444 – Spring 2018

21

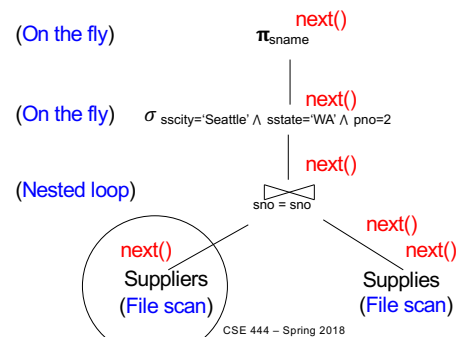




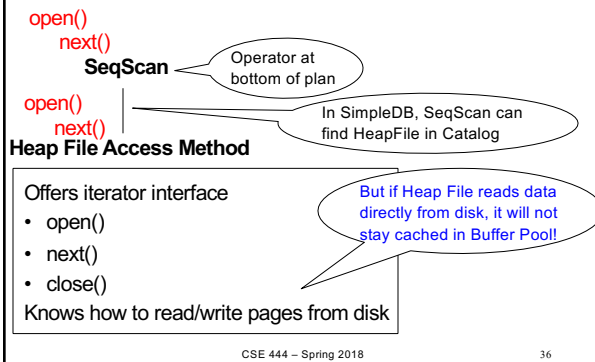
Query Execution How it all Fits Together



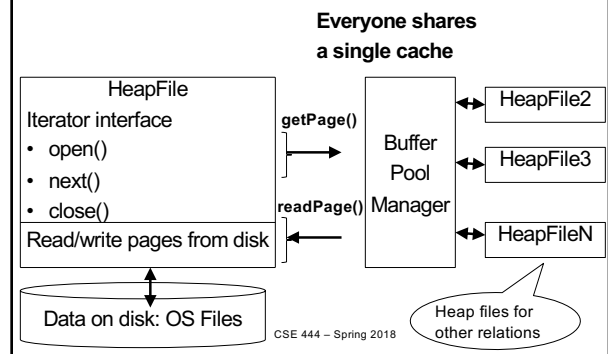
Query Execution How it all Fits Together



Query Execution In SimpleDB



Query Execution In SimpleDB



HeapFile In SimpleDB

- Data is stored on disk in an OS file. HeapFile class knows how to "decode" its content
 - Control flow:
 - SeqScan calls methods such as "iterate" on the HeapFile Access Method
 - During the iteration, the HeapFile object needs to call the BufferManager.getPage() method to ensure that necessary pages get loaded into memory.
 - The BufferManager will then call HeapFile.readPage()/writePage() page to actually read/write the page.
- CSE 444 – Spring 2018 38