

Lecture 12:

HW3 Overview & DB Components

Guest lecturer: Moe Kayali

Announcements

- Review 5 due on Wednesday, 11/5
- HW3 due next Friday, 11/14

Agenda

- High Level Homework Structure
 - What is SimpleDB?
 - Code environment
- Implementation Pieces
 - Database Catalog
 - Buffer Pool Manager
 - Heap File/Page
 - The Tuple Iterator Model
 - Scan
 - Filter
 - Join


SimpleDB

SimpleDB is a simple database system:


- Originally intended for MIT's intro database course
- Written in Java
- Disk-based, row-oriented, single-threaded

If you're not familiar with Java, you should to do a short online tutorial to get (re)acquainted! (One is linked in the HW)

SimpleDB System Overview



SimpleDB System Overview



Code Environment

The code is written as a series of java files in the “starter-code” folder of HW3.

cse544-2024w1 ▾

cse544-master / hw / hw3 / starter-code / src / java / simpledb / + ▾

History

Find file

Edit ▾

↓ ▾

Clone ▾

Name	Last commit	Last update
--		
Aggregate.java	add starter code for hw3	6 years ago
Aggregator.java	add starter code for hw3	6 years ago
BufferPool.java	add starter code for hw3	6 years ago
Catalog.java	add starter code for hw3	6 years ago
CostCard.java	add starter code for hw3	6 years ago
Database.java	add starter code for hw3	6 years ago
DbException.java	add starter code for hw3	6 years ago
DbFile.java	add starter code for hw3	6 years ago
DbFileIterator.java	add starter code for hw3	6 years ago
DbIterator.java	add starter code for hw3	6 years ago
Debug.java	add starter code for hw3	6 years ago
Delete.java	add starter code for hw3	6 years ago

Code Environment

The code is written as a series of java files in the “starter-code” folder of HW3.

We will compile & run the code using “ant”, a java build tool:

```
(base) kylebd99@kyles-laptop:~/GIT/cse544-master/hw/hw3/starter-code$ ant runsystest -Dtest=ScanTest
Buildfile: /home/kylebd99/GIT/cse544-master/hw/hw3/starter-code/build.xml

compile:

testcompile:

runsystest:
[junit] Running simpledb.systemtest.ScanTest
[junit] Testsuite: simpledb.systemtest.ScanTest
```

In the “starter-code” folder

Invoke “ant”

Build the “runsystest” target

W/option “ScanTest”

How to Edit the Code

In the Readme.md for the HW, you'll be pointed to files in the codebase with sections which are incomplete,

```
public void addTable(DbFile file, String name, String pkeyField) {  
    // some code goes here  
}
```

The assignment consists of filling these sections in with the appropriate logic. To do this you may want to:

- Define new functions
- Add attributes to classes
- Import data structures from the standard java library (e.g. a hash table)

DB Catalog

What is a Catalog?

Catalogs are a general term for data structures which store metadata in a DBMS,

- The list of databases
- Database to table mapping
- Table to column mapping
- Table to index mapping
- Table to statistics mapping
- User security privileges

In most systems, these catalogs are themselves stored as tables!

SimpleDB Catalog

In SimpleDB, the catalog only stores a list of tables,


```
public class Catalog {  
    /**  
     * Constructor.  
     * Creates a new, empty catalog.  
     */  
    public Catalog() {  
        // some code goes here  
    }  
}
```

You'll implement:

- The data structure for storing table information
- Adding/removing a table
- Getting information about a table
- Iterating through the tables in the database

Buffer Pools

Recap: The Problem of Storage




The Buffer Pool

A buffer pool is a space in main memory which caches pages that have been retrieved from disk.

When a disk page is requested, the buffer pool

1. Checks whether the page is already in the pool
 - a. If so, immediately returns the page
2. Checks whether the pool is full
 - a. If so, consults the eviction policy and flushes a page to the disk
3. Loads the page from disk

Buffer Pool Design



Eviction Policies

The eviction policy determines which page needs to be flushed to disk when the buffer pool fills up. Potential algorithms for this:

- LRU: Least Recently Used
- FIFO: First In First Out
- Clock: Circular Eviction
- MRU: Most Recently Used

5 Minute Discussion: How would you design a better eviction policy for a DBMS?

The Iterator Model (i.e. “The Pull Model”)

Set-at-a-time Operators

A	B	C	D
...	...	...	...
...			



A	B
...	...
...	

C	D
...	...
...	

A	B
...	...
...	


σ_P

A	B
...	...
...	

How Do We Combine Them?

Option 1:
materialize intermediate results

Option 2:
treat operators as iterators



Operator Interface

Volcano model:

- `open()`, `next()`, `close()`
- Pull model
- Supported by most DBMS today
- Will discuss next

Newer alternative: push model (won't discuss)

Operator Interface

Open()

- Calls open() on the children
- Creates any local data structures

Next()

- May call next() repeatedly on children
- Returns exactly 1 tuple, or EOF

Close()

- Free any local memory

Iterator Model

A.k.a. Volcano-style execution

```
interface Operator {  
    // initializes operator state  
    // and sets parameters  
    void open (...);  
  
    // calls next() on its inputs  
    // processes an input tuple  
    // produces output tuple(s)  
    // returns null when done  
    Tuple next ();  
  
    // cleans up (if any)  
    void close ();  
}
```

Iterator Model

A.k.a. Volcano-style execution

Example selection operator

```
interface Operator {  
    // initializes operator state  
    // and sets parameters  
    void open (...);  
  
    // calls next() on its inputs  
    // processes an input tuple  
    // produces output tuple(s)  
    // returns null when done  
    Tuple next ();  
  
    // cleans up (if any)  
    void close ();  
}
```

```
class Select implements Operator {...  
    void open (Predicate p,  
               Operator c) {  
        this.p = p; this.c = c; c.open();  
    }  
}
```


Iterator Model

A.k.a. Volcano-style execution

Example selection operator

```
interface Operator {  
    // initializes operator state  
    // and sets parameters  
    void open (...);  
  
    // calls next() on its inputs  
    // processes an input tuple  
    // produces output tuple(s)  
    // returns null when done  
    Tuple next ();  
  
    // cleans up (if any)  
    void close ();  
}
```

```
class Select implements Operator {...  
    void open (Predicate p,  
                Operator c) {  
        this.p = p; this.c = c; c.open();  
    }  
    Tuple next () {  
        boolean found = false;  
        Tuple r = null;  
        while (!found) {  
            r = c.next();  
            if (r == null) break;  
            found = p(r);  
        }  
        return r;  
    }  
}
```

Iterator Model

A.k.a. Volcano-style execution

Example selection operator

```
interface Operator {  
    // initializes operator state  
    // and sets parameters  
    void open (...);  
  
    // calls next() on its inputs  
    // processes an input tuple  
    // produces output tuple(s)  
    // returns null when done  
    Tuple next ();  
  
    // cleans up (if any)  
    void close ();  
}
```

```
class Select implements Operator {...  
    void open (Predicate p,  
                Operator c) {  
        this.p = p; this.c = c; c.open();  
    }  
    Tuple next () {  
        boolean found = false;  
        Tuple r = null;  
        while (!found) {  
            r = c.next();  
            if (r == null) break;  
            found = p(r);  
        }  
        return r;  
    }  
    void close () { c.close(); }  
}
```

Iterator Model

A.k.a. Volcano-style execution

Query plan execution

```
Operator q = parse("SELECT ..."); # sql -> root of an op tree  
q = optimize(q); # op tree -> optimized op tree
```

Iterator Model

A.k.a. Volcano-style execution

Query plan execution

```
Operator q = parse("SELECT ..."); # sql -> root of an op tree
```

```
q = optimize(q); # op tree -> optimized op tree
```

```
q.open();
```

Iterator Model

A.k.a. Volcano-style execution

Query plan execution

```
Operator q = parse("SELECT ..."); # sql -> root of an op tree  
q = optimize(q); # op tree -> optimized op tree
```

```
q.open();  
while (true) {  
    Tuple t = q.next();  
    if (t == null) break; # end of results  
    else printOnScreen(t); # output tuple  
}
```

Iterator Model

A.k.a. Volcano-style execution

Query plan execution

```
Operator q = parse("SELECT ..."); # sql -> root of an op tree  
q = optimize(q); # op tree -> optimized op tree
```

```
q.open();  
while (true) {  
    Tuple t = q.next();  
    if (t == null) break; # end of results  
    else printOnScreen(t); # output tuple  
}  
q.close();
```

Supplier(sno,sname,scity,sstate)
Supply(sno,pno,price)


Example

"Normal" hash-join

```
for x in Supplier do
  insert(x.sno, x)

for y in Supply do
  x = find(y.sno);
  output(x,y);
```

Pipelining changes
the order significantly



Supplier(sno,sname,scity,sstate)
Supply(sno,pno,price)


Example

"Normal" hash-join

```
for x in Supplier do
  insert(x.sno, x)

for y in Supply do
  x = find(y.sno);
  output(x,y);
```

Pipelining changes
the order significantly



Supplier(sno,sname,scity,sstate)
Supply(sno,pno,price)


Example

"Normal" hash-join

```
for x in Supplier do
  insert(x.sno, x)

for y in Supply do
  x = find(y.sno);
  output(x,y);
```

Pipelining changes
the order significantly



Supplier(sno,sname,scity,sstate)
Supply(sno,pno,price)


Example

"Normal" hash-join

```
for x in Supplier do
  insert(x.sno, x)

for y in Supply do
  x = find(y.sno);
  output(x,y);
```

Pipelining changes
the order significantly



Supplier(sno,sname,scity,sstate)
Supply(sno,pno,price)

Example


"Normal" hash-join

```
for x in Supplier do
  insert(x.sno, x)

for y in Supply do
  x = find(y.sno);
  output(x,y);
```

Pipelining changes
the order significantly

Details



Supplier(sno,sname,scity,sstate)
Supply(sno,pno,price)


Example

"Normal" hash-join

```
for x in Supplier do
  insert(x.sno, x)

for y in Supply do
  x = find(y.sno);
  output(x,y);
```

Pipelining changes
the order significantly



Supplier(sno,sname,scity,sstate)
Supply(sno,pno,price)


Example

"Normal" hash-join

```
for x in Supplier do
  insert(x.sno, x)

for y in Supply do
  x = find(y.sno);
  output(x,y);
```

Pipelining changes
the order significantly



Supplier(sno,sname,scity,sstate)
Supply(sno,pno,price)


Example

"Normal" hash-join

```
for x in Supplier do
  insert(x.sno, x)

for y in Supply do
  x = find(y.sno);
  output(x,y);
```

Pipelining changes
the order significantly



Supplier(sno,sname,scity,sstate)
Supply(sno,pno,price)


Example

"Normal" hash-join

```
for x in Supplier do
  insert(x.sno, x)

for y in Supply do
  x = find(y.sno);
  output(x,y);
```

Pipelining changes
the order significantly



Supplier(sno,sname,scity,sstate)
Supply(sno,pno,price)


Example

”Normal” hash-join

```
for x in Supplier do
  insert(x.sno, x)

for y in Supply do
  x = find(y.sno);
  output(x,y);
```

Pipelining changes
the order significantly



Supplier(sno,sname,scity,sstate)
Supply(sno,pno,price)


Example

"Normal" hash-join

```
for x in Supplier do
  insert(x.sno, x)

for y in Supply do
  x = find(y.sno);
  output(x,y);
```

Pipelining changes
the order significantly



Supplier(sno,sname,scity,sstate)
Supply(sno,pno,price)


Example

"Normal" hash-join

```
for x in Supplier do
  insert(x.sno, x)

for y in Supply do
  x = find(y.sno);
  output(x,y);
```

Pipelining changes
the order significantly



Supplier(sno,sname,scity,sstate)
Supply(sno,pno,price)


Example

"Normal" hash-join

```
for x in Supplier do
  insert(x.sno, x)

for y in Supply do
  x = find(y.sno);
  output(x,y);
```

Pipelining changes
the order significantly



Supplier(sno,sname,scity,sstate)
Supply(sno,pno,price)


Example

”Normal” hash-join

```
for x in Supplier do
  insert(x.sno, x)

for y in Supply do
  x = find(y.sno);
  output(x,y);
```

Pipelining changes
the order significantly



Supplier(sno,sname,scity,sstate)
Supply(sno,pno,price)


Example

”Normal” hash-join

```
for x in Supplier do
  insert(x.sno, x)

for y in Supply do
  x = find(y.sno);
  output(x,y);
```

Pipelining changes
the order significantly



Supplier(sno,sname,scity,sstate)
Supply(sno,pno,price)


Example

"Normal" hash-join

```
for x in Supplier do
  insert(x.sno, x)

for y in Supply do
  x = find(y.sno);
  output(x,y);
```

Pipelining changes
the order significantly



Supplier(sno,sname,scity,sstate)
Supply(sno,pno,price)


Example

”Normal” hash-join

```
for x in Supplier do
  insert(x.sno, x)

for y in Supply do
  x = find(y.sno);
  output(x,y);
```

Pipelining changes
the order significantly



Supplier(sno,sname,scity,sstate)
Supply(sno,pno,price)


Example

"Normal" hash-join

```
for x in Supplier do
  insert(x.sno, x)

for y in Supply do
  x = find(y.sno);
  output(x,y);
```

Pipelining changes
the order significantly



Supplier(sno,sname,scity,sstate)
Supply(sno,pno,price)


Example

"Normal" hash-join

```
for x in Supplier do
  insert(x.sno, x)

for y in Supply do
  x = find(y.sno);
  output(x,y);
```

Pipelining changes
the order significantly



Supplier(sno,sname,scity,sstate)
Supply(sno,pno,price)


Example

"Normal" hash-join

```
for x in Supplier do
  insert(x.sno, x)

for y in Supply do
  x = find(y.sno);
  output(x,y);
```

Pipelining changes
the order significantly



Supplier(sno,sname,scity,sstate)
Supply(sno,pno,price)


Example

"Normal" hash-join

```
for x in Supplier do
  insert(x.sno, x)

for y in Supply do
  x = find(y.sno);
  output(x,y);
```

Pipelining changes
the order significantly



Supplier(sno,sname,scity,sstate)
Supply(sno,pno,price)


Example

”Normal” hash-join

```
for x in Supplier do
  insert(x.sno, x)

for y in Supply do
  x = find(y.sno);
  output(x,y);
```

Pipelining changes
the order significantly



Supplier(sno,sname,scity,sstate)
Supply(sno,pno,price)


Example

"Normal" hash-join

```
for x in Supplier do
  insert(x.sno, x)

for y in Supply do
  x = find(y.sno);
  output(x,y);
```

Pipelining changes
the order significantly



Supplier(sno,sname,scity,sstate)
Supply(sno,pno,price)


Example

"Normal" hash-join

```
for x in Supplier do
  insert(x.sno, x)

for y in Supply do
  x = find(y.sno);
  output(x,y);
```

Pipelining changes
the order significantly



Supplier(sno,sname,scity,sstate)
Supply(sno,pno,price)


Example

"Normal" hash-join

```
for x in Supplier do
  insert(x.sno, x)

for y in Supply do
  x = find(y.sno);
  output(x,y);
```

Pipelining changes
the order significantly



Supplier(sno,sname,scity,sstate)
Supply(sno,pno,price)


Example

“Normal” hash-join

```
for x in Supplier do
  insert(x.sno, x)

for y in Supply do
  x = find(y.sno);
  output(x,y);
```

Pipelining changes
the order significantly



Supplier(sno,sname,scity,sstate)
Supply(sno,pno,price)


Example

"Normal" hash-join

```
for x in Supplier do
  insert(x.sno, x)

for y in Supply do
  x = find(y.sno);
  output(x,y);
```

Pipelining changes
the order significantly



Supplier(sno,sname,scity,sstate)
Supply(sno,pno,price)


Example

"Normal" hash-join

```
for x in Supplier do
  insert(x.sno, x)

for y in Supply do
  x = find(y.sno);
  output(x,y);
```

Pipelining changes
the order significantly



Supplier(sno,sname,scity,sstate)
Supply(sno,pno,price)


Example

"Normal" hash-join

```
for x in Supplier do
  insert(x.sno, x)

for y in Supply do
  x = find(y.sno);
  output(x,y);
```

Pipelining changes
the order significantly



Iterate v.s. Materialize

- Iterate

- Disk IOs: $O(\text{Input})$
- Memory Footprint: $O(\sim \text{Input})$ <- keeps hash tables in memory

- Materialize

- Disk IOs: $O(\text{Intermediates})$
- Memory Footprint: $O(1)$

Summary

- Start SimpleDB Early! (Due Feb 23rd)
- The catalog stores metadata
- The buffer pool caches disk pages
- The iterator model reduces disk I/O by using memory
 - Flexible model for implementation