

Introduction to Data Management CSE 414

Unit 4: RDBMS Internals
Logical and Physical Plans
Query Execution
Query Optimization

(3 lectures)

CSE 414 - Autumn 2018

1

Introduction to Data Management CSE 414

Lecture 15: Introduction to Query
Evaluation

CSE 414 - Autumn 2018

2

Announcements

- WQ5 (datalog) due tomorrow
- HW4 (datalog) due tomorrow
- Midterm review session this evening
– 5:30pm, CSE 2nd Floor Breakout

CSE 414 - Autumn 2018

3

Class Overview

- Unit 1: Intro
- Unit 2: Relational Data Models and Query Languages
- Unit 3: Non-relational data
- Unit 4: RDBMS internals and query optimization
- Unit 5: Parallel query processing
- Unit 6: DBMS usability, conceptual design
- Unit 7: Transactions
- Unit 8: Advanced topics (time permitting)

CSE 414 - Autumn 2018


4

From Logical RA Plans to Physical Plans

CSE 414 - Autumn 2018

5

Query Evaluation Steps Review



6

Logical vs Physical Plans

- Logical plans:
 - Created by the parser from the input SQL text
 - Expressed as a relational algebra tree
 - Each SQL query has many possible logical plans
- Physical plans:
 - Goal is to choose an efficient implementation for each operator in the RA tree
 - Each logical plan has many possible physical plans

CSE 414 - Autumn 2018


7

```
Supplier(sid, sname, scity, sstate)
Supply(sid, pno, quantity)
```

Review: Relational Algebra

```
SELECT sname
FROM Supplier x, Supply y
WHERE x.sid = y.sid
      and y.pno = 2
      and x.scity = 'Seattle'
      and x.sstate = 'WA'
```

Relational algebra expression is also called the "logical query plan"



CSE 414 - Autumn 2018

8

Logical Plan v.s. Physical Plan

- Logical Plan = a Relational Algebra tree
- Physical Plan = a Logical Plan plus annotation of each operator with an algorithm

CSE 414 - Autumn 2018

9

Query Optimization and Execution

- Query optimizer:
 - Choose a good logical plan
 - Refine it to a good physical plan
 - Sometimes these steps are intertwined
- Query execution
 - Execute the physical plan

CSE 414 - Autumn 2018

10

Query Execution

CSE 414 - Autumn 2018

11

Physical Operators

Relational algebra operators:

- Selection, projection, join, union, difference
- Group-by, distinct, sort

Physical operators:

- For each operators above, several possible algorithms
- Main memory algorithms, or disk-based algorithms

CSE 414 - Autumn 2018

12

Supplier(sid, sname, scity, sstate)
Supply(sid, pno, quantity)

Main Memory Algorithms

Logical operator:
Supplier ⋈_{sid=sid} Supply

Propose three physical operators for the join, assuming the tables are in main memory:

- 1.
- 2.
- 3.

CSE 414 - Autumn 2018 13

Supplier(sid, sname, scity, sstate)
Supply(sid, pno, quantity)

Main Memory Algorithms

Logical operator:
Supplier ⋈_{sid=sid} Supply

Propose three physical operators for the join, assuming the tables are in main memory:

1. Nested Loop Join $O(??)$
2. Merge join $O(??)$
3. Hash join $O(??)$

CSE 414 - Autumn 2018 14

Supplier(sid, sname, scity, sstate)
Supply(sid, pno, quantity)

Main Memory Algorithms

Logical operator:
Supplier ⋈_{sid=sid} Supply

Propose three physical operators for the join, assuming the tables are in main memory:

1. Nested Loop Join $O(n^2)$
2. Merge join $O(n \log n)$
3. Hash join $O(n) \dots O(n^2)$

CSE 414 - Autumn 2018 15

BRIEF Review of Hash Tables

Separate chaining:

A (naïve) hash function:
 $h(x) = x \text{ mod } 10$

Operations:
find(103) = ??
insert(488) = ??

CSE 414 - Autumn 2018 16

BRIEF Review of Hash Tables

- insert(k, v) = inserts a key k with value v
- Many values for one key
 - Hence, duplicate k's are OK
- find(k) = returns the list of all values v associated to the key k

CSE 414 - Autumn 2018 17

Query Execution


- Join $R \bowtie S$: e.g. using hash-join:
 - Nested-loop: forall x in R forall y in S do ...
 - Hash-join: build a hash table on S, probe R
- Selection: $\sigma(R)$: e.g. "on-the-fly"
- But what about a larger plan?
 - Each operator implements the Iterator Interface

CSE 414 - Autumn 2018 18

Implementing Query Operators with the Iterator Interface

Each operator implements three methods:

- open()
- next()
- close()



CSE 414 - Autumn 2018

19

Implementing Query Operators with the Iterator Interface

Example "on the fly" selection operator

```
interface Operator {
```

```
}
```

CSE 414 - Autumn 2018

20

Implementing Query Operators with the Iterator Interface

Example "on the fly" selection operator

```
interface Operator {
    // initializes operator state
    // and sets parameters
    void open (...);
}
```

CSE 414 - Autumn 2018

21

Implementing Query Operators with the Iterator Interface

Example "on the fly" 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 ();
}
```

CSE 414 - Autumn 2018

22

Implementing Query Operators with the Iterator Interface

Example "on the fly" 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 ();
```

CSE 414 - Autumn 2018

23

Implementing Query Operators with the Iterator Interface

Example "on the fly" 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 ();
```

CSE 414 - Autumn 2018

24

Implementing Query Operators with the Iterator Interface

Example "on the fly" 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 child) {
        this.p = p; this.child = child;
    }
    Tuple next () {
        // ...
    }
}
CSE 414 - Autumn 2018 25
```

Implementing Query Operators with the Iterator Interface

Example "on the fly" 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 child) {
        this.p = p; this.child = child;
    }
    Tuple next () {
        boolean found = false;
        Tuple r = null;
        while (!found) {
            r = child.next();
            if (r == null) break;
            found = p(in);
        }
        return r;
    }
}
CSE 414 - Autumn 2018 26
```

Implementing Query Operators with the Iterator Interface

Example "on the fly" 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 child) {
        this.p = p; this.child = child;
    }
    Tuple next () {
        boolean found = false;
        Tuple r = null;
        while (!found) {
            r = child.next();
            if (r == null) break;
            found = p(in);
        }
        return r;
    }
}
CSE 414 - Autumn 2018 27
```

Implementing Query Operators with the Iterator Interface

Example "on the fly" 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 child) {
        this.p = p; this.child = child;
    }
    Tuple next () {
        boolean found = false;
        Tuple r = null;
        while (!found) {
            r = child.next();
            if (r == null) break;
            found = p(in);
        }
        return r;
    }
    void close () { child.close(); }
}
CSE 414 - Autumn 2018 28
```

Implementing Query Operators with the Iterator Interface

Query plan 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 ();
}


Operator q = parse("SELECT ...");
q = optimize(q);
q.open();
while (true) {
    Tuple t = q.next();
    if (t == null) break;
    else printOnScreen(t);
}
q.close();
CSE 414 - Autumn 2018 29
```

Supplier(sid, sname, scity, sstate) Supply(sid, sname, quantity) Pipelining

(On the fly)

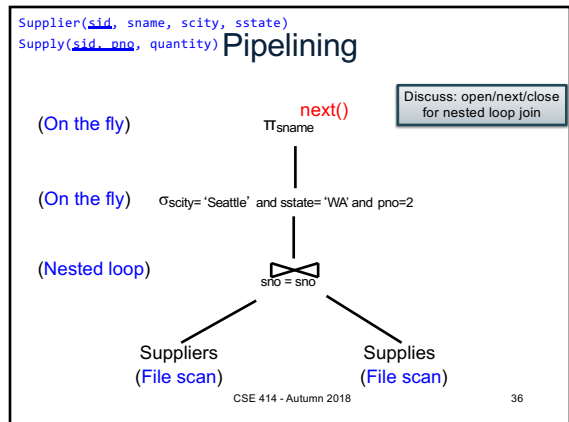
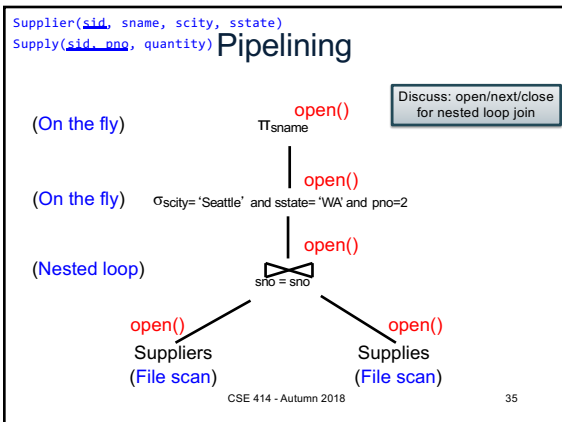
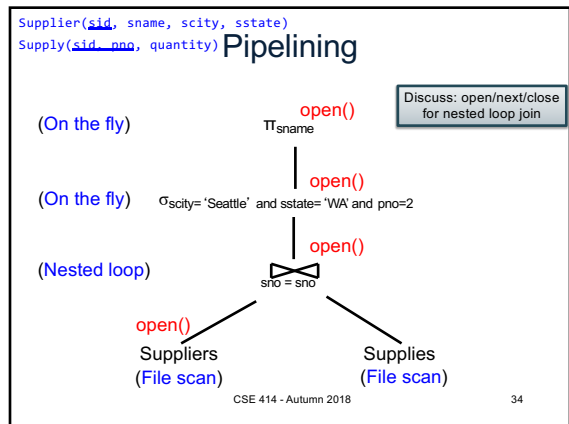
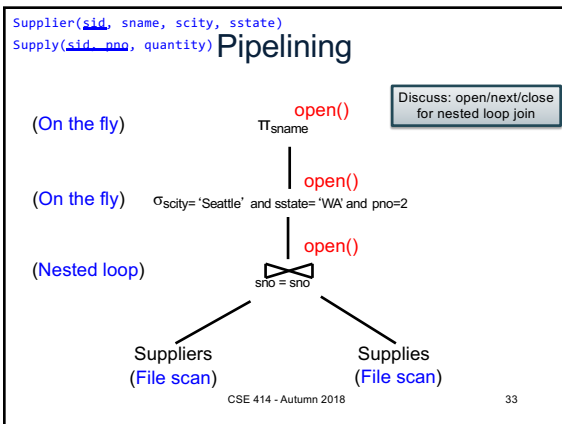
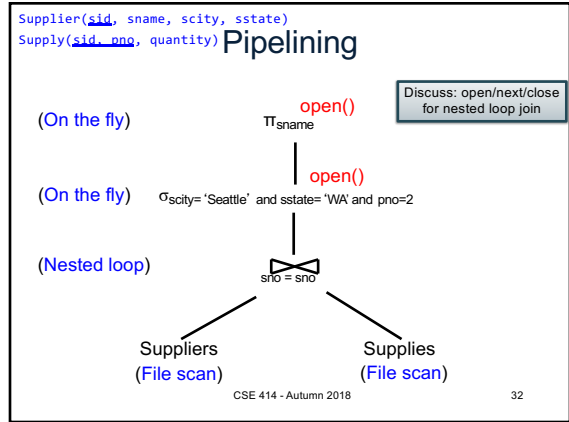
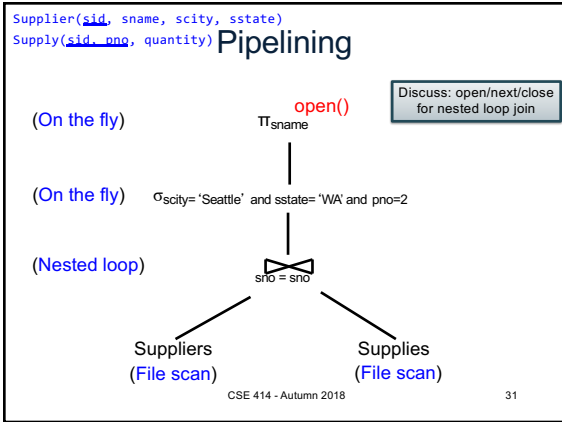
(On the fly)

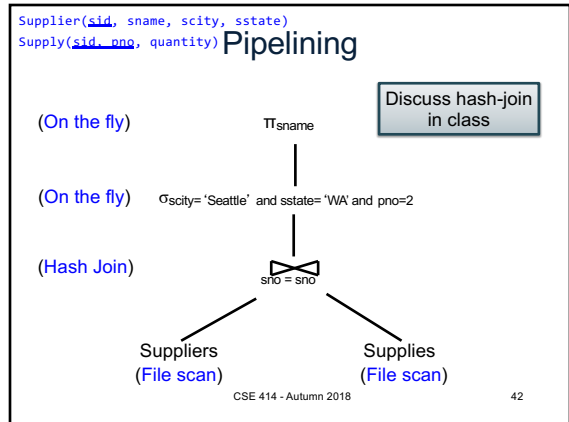
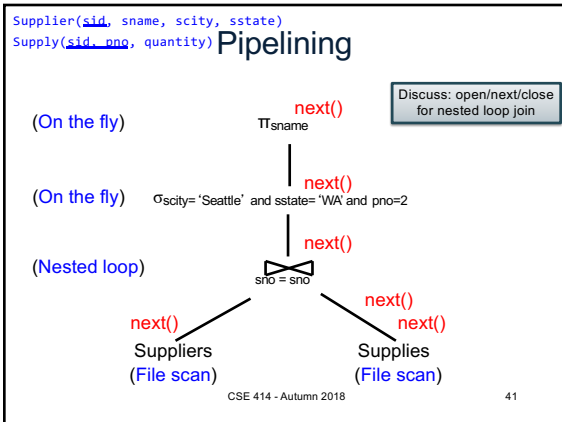
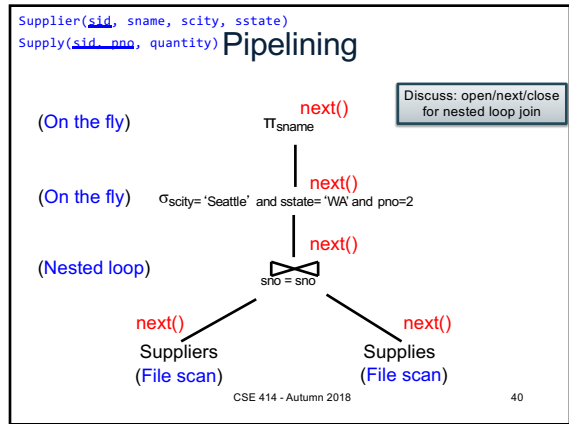
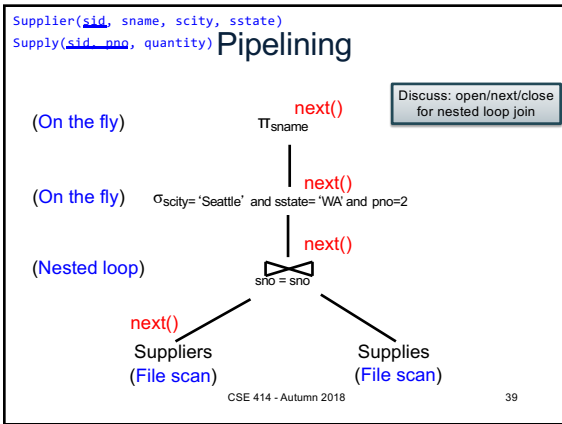
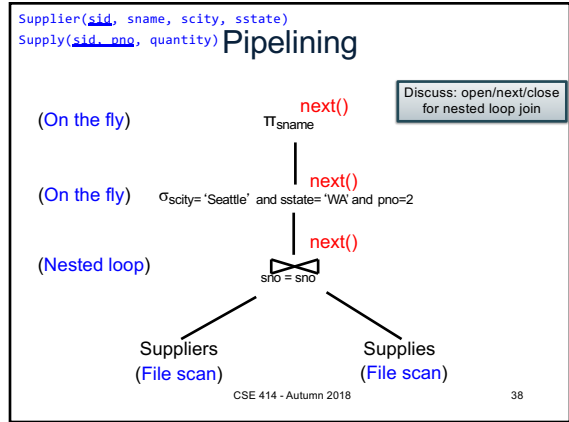
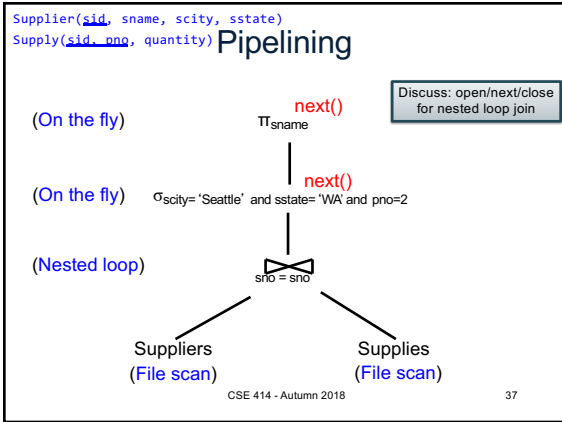
(Nested loop)

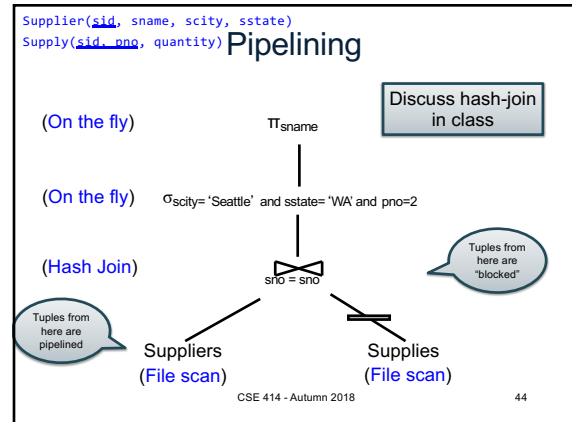
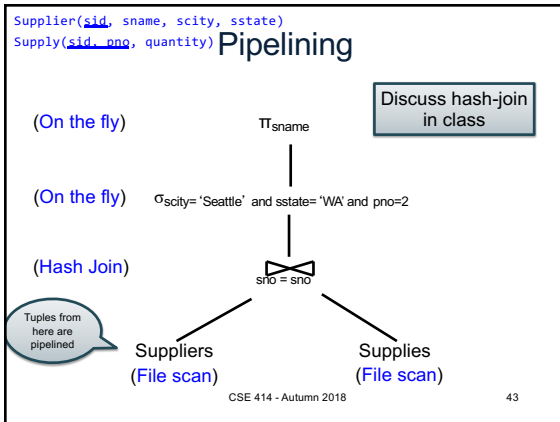


CSE 414 - Autumn 2018

30







- ## Pipeline v.s. Blocking
- Pipeline
 - A tuple moves all the way through up the query plan
 - Advantages: speed
 - Disadvantage: need all hash at the same time in memory
 - Blocking
 - The entire result of the subplan is computed (and stored to disk) before the first tuple is sent up the plan
 - Advantage: saves memory
 - Disadvantage: slower
- CSE 414 - Autumn 2018 47

- ## Discussion on Physical Plan
- More components of a physical plan:
- **Access path selection** for each relation
 - Scan the relation or use an index (next lecture)
 - **Implementation choice** for each operator
 - Nested loop join, hash join, etc.
 - **Scheduling decisions** for operators
 - Pipelined execution or intermediate materialization
- CSE 414 - Autumn 2018 48