

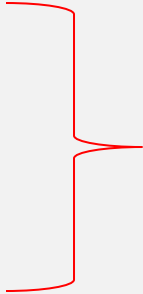
Relational Algebra

- Relational algebra is a **procedural** language ... operations that can be implemented by a database
- SQL is **non-procedural**
- We will consider some operators and the mapping of SQL Select to RA statements and to an RA tree.

Relational Algebra

- Operators:

- SELECT
- PROJECT
- JOIN



There are more ACS-4902

We need to consider a **SQL Select** in terms of **RA operators** select, project, and join

SELECT

- The SELECT Operation
 - Yields a relation, a subset of the tuples from the operand that satisfies a selection condition:

$$\sigma_{\langle \text{selection condition} \rangle}(R)$$

R is a relation or a variable representing a relation

- Boolean expression contains clauses of the form
<attribute name> <comparison op> <constant value>
or
- <attribute name> <comparison op> <attribute name>

AND, OR, and NOT

SELECT

- Example:

$\sigma_{(Dno=4 \text{ AND } Salary > 25000) \text{ OR } (Dno=5 \text{ AND } Salary > 30000)}(EMPLOYEE)$

- <selection condition> applied independently to each individual tuple t in R
 - If condition evaluates to TRUE, tuple selected
- **unary**

PROJECT

- Yields a relation with specific columns (or arithmetic expressions involving columns) from the operand
 - discards the other columns:

$$\pi_{\langle \text{attribute list} \rangle}(R)$$

R is a relation or a variable representing a relation

- **unary**
- **degree**
 - Number of attributes in <attribute list>
- **duplicate elimination**
 - Result of PROJECT operation is a set of distinct tuples

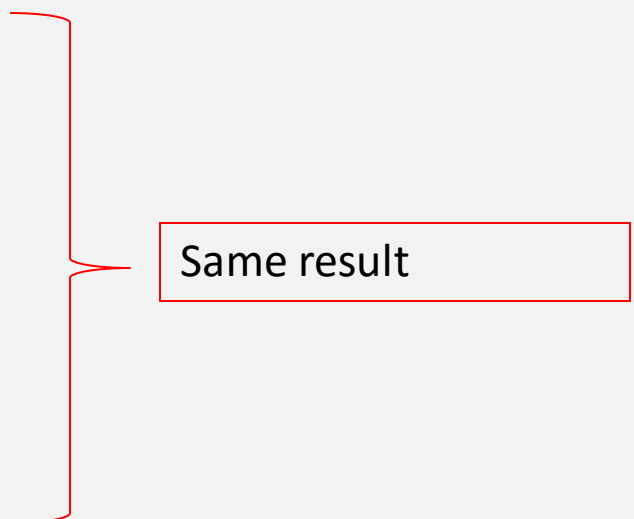
Sequences of Operations

- **In-line expression:**

$\pi_{\text{Fname, Lname, Salary}}(\sigma_{\text{Dno}=5}(\text{EMPLOYEE}))$

- **Sequence of operations:**

$\text{DEP5_EMPS} \leftarrow \sigma_{\text{Dno}=5}(\text{EMPLOYEE})$
 $\text{RESULT} \leftarrow \pi_{\text{Fname, Lname, Salary}}(\text{DEP5_EMPS})$



Same result


The CARTESIAN PRODUCT

- **CARTESIAN PRODUCT**
 - **CROSS PRODUCT** or **CROSS JOIN**
 - Denoted by \times
 - Binary set operation
 - Useful when followed by a selection that matches values of attributes

A Product can yield attributes of the same name ... the dot notation to disambiguate

JOIN

- The **JOIN** Operation

- Denoted by 
- Combine related tuples from two relations
- Inner join
- binary

Join condition

– $\text{DEPT_MGR} \leftarrow \text{DEPARTMENT} \bowtie_{\text{Mgr_ssn}=\text{Ssn}} \text{EMPLOYEE}$
 $\text{RESULT} \leftarrow \pi_{\text{Dname, Lname, Fname}}(\text{DEPT_MGR})$

Theta JOIN

- **Natural join** *
- **THETA JOIN**
 - Each <condition> of the form $A_i \theta B_j$
 - A_i is an attribute of R
 - B_j is an attribute of S
 - A_i and B_j have the same domain
 - θ (theta) is one of the comparison operators:
 - $\{=, <, \leq, >, \geq, \neq\}$

OUTER JOIN SYMBOLS

LEFT



RIGHT



FULL






Notation for Query Trees

- **Query tree for RA operations**
 - Represents the input relations of a query as leaf nodes of the tree
 - Represents the relational algebra operations as internal nodes
 - Execution is bottom-up; results move up the tree as an input/operand

Query 2. For every project located in 'Stafford', list the project number, the controlling department number, and the department manager's last name, address, and birth date.

Q2:

```
SELECT Pnumber, Dnum, Lname, Address, Bdate
FROM PROJECT, DEPARTMENT, EMPLOYEE
WHERE Dnum=Dnumber
AND Mgr_ssn=Ssn
AND Plocation = 'Stafford';
```

Data Output		Explain	Messages	History	
	pnumber integer	dnum integer	lname character varying	address character varying	bdate date
	30	4	Wallace	291 Berry, Bellaire...	1941-06-20
	10	4	Wallace	291 Berry, Bellaire...	1941-06-20

Q2: SELECT Pnumber, Dnum, Lname, Address, Bdate
FROM PROJECT, DEPARTMENT, EMPLOYEE
WHERE Dnum=Dnumber
AND Mgr_ssn=Ssn
AND
Plocation='Stafford';

*Given the Selects statement,
what is the RA query tree?
Given the RA query tree,
what is the Select
statement?*

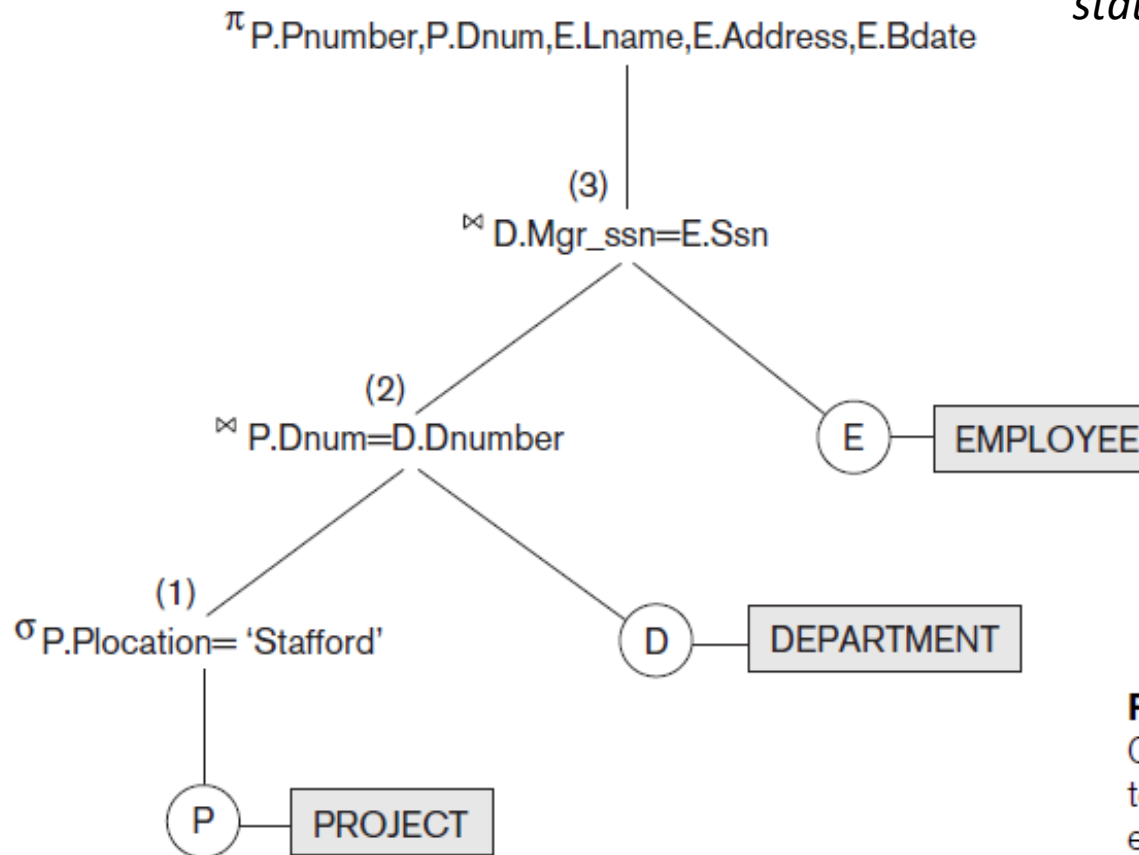


Figure 6.9

Query tree corresponding to the relational algebra expression for Q2.

From **December 2015 exam**:

Consider the XXX database.

Consider the SQL statement:

```
select city  
from Offices  
where state = "CA";
```

Translate the SQL statement into a sequence of relational algebra **statements**.

From **December 2015 exam**:

Consider the SQL statement:

```
select city, amount
from customers
inner join orders
on (customers.customerNumber = orders.customerNumber)
inner join payments
on (customers.customerNumber = payments.customerNumber)
where city = "NYC"
and status = "shipped"
;
```

Illustrate a relational algebra **query tree** for the execution of the statement.

Example and next 3 slides

Example : Consider this query that retrieves from the text's Company database the SSN and last name of employees who earn more than 10000 and have a daughter and a son

```
SELECT ssn, lname
```

```
FROM dependent x
```

```
INNER JOIN employee e
```

```
      ON (x.essn = e.ssn)
```

```
INNER JOIN dependent y
```

```
      ON (y.essn = e.ssn)
```

```
where x.relationship = 'Son'
```

```
and y.relationship = 'Daughter'
```

```
and e.salary > 10000
```

Create a query tree for the above

Example

First: how do we apply our study of relational algebra to this?

SELECT ssn, lname **a projection**

FROM dependent x
INNER JOIN employee e **inner join between x and e**
 ON (x.essn = e.ssn)
INNER JOIN dependent y **inner join between y and e**
 ON (y.essn = e.ssn)

where x.relationship = 'Son' **a selection on x**
and y.relationship = 'Daughter' **a selection on y**
and e.salary > 10000 **a selection on e**

Example

First: how do we apply our study of relational algebra to this?

SELECT ssn, lname **a projection**

FROM dependent x
INNER JOIN employee e **inner join between x and e**
 ON (x.essn = e.ssn)
INNER JOIN dependent y **inner join between y and e**
 ON (y.essn = e.ssn)

where x.relationship = 'Son' **a selection on x**
and y.relationship = 'Daughter' **a selection on y**
and e.salary > 10000 **a selection on e**

Our approach:

We need to create *the query tree with the appropriate joins, selections, and projections, and then push operations downwards as far as possible*

Example

a projection

inner join between x and e

inner join between y and e

a selection on x

a selection on y

a selection on e

Query trees

In class:

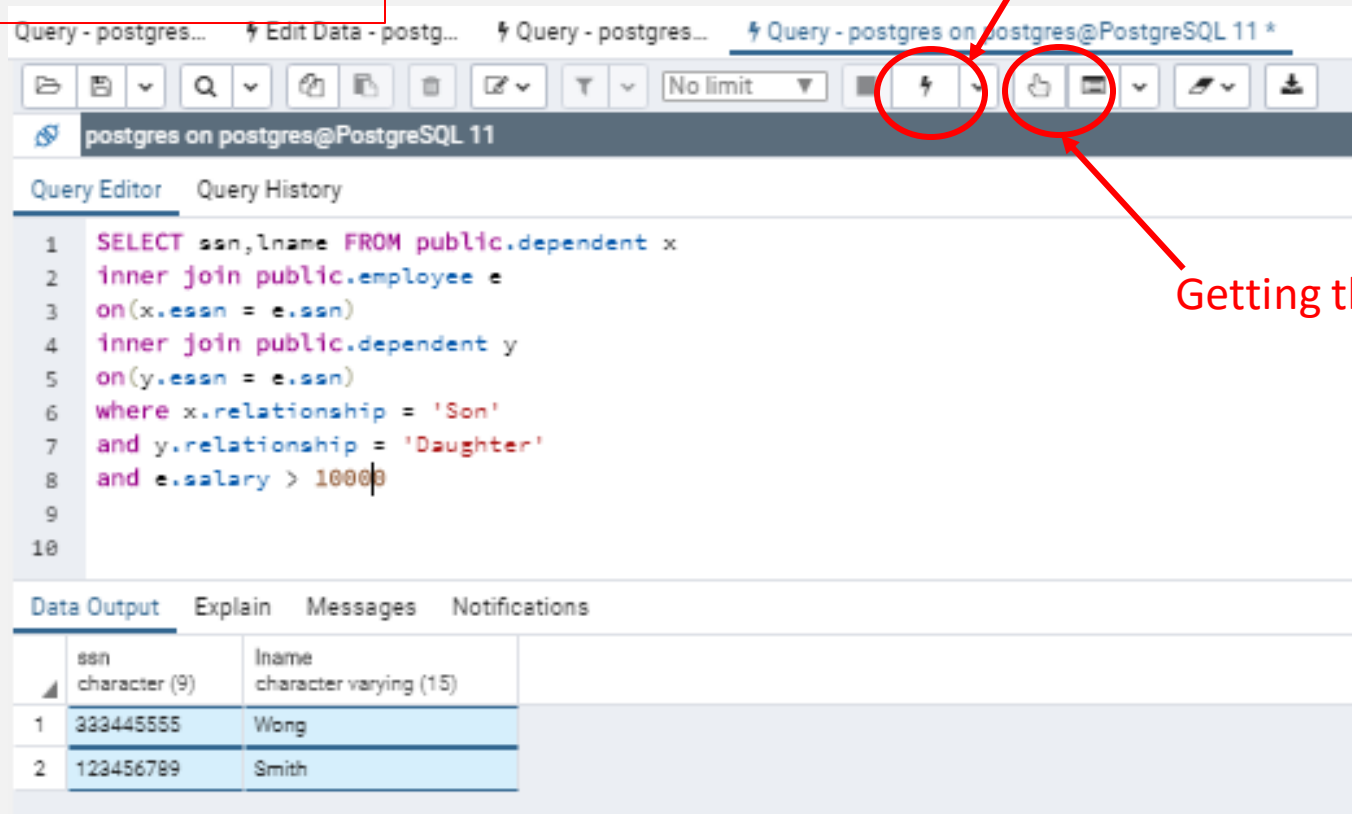
Create the diagram, then check that operations are as low as possible in tree.

Remaining slides are for interest only

For interest only

Aside: Seeing the query plan in PostgreSQL

Two employees have both
a daughter and son and
earn more than 10000



The screenshot shows the PostgreSQL Query Editor interface. The query is as follows:

```
1 SELECT ssn, lname FROM public.dependent x
2 inner join public.employee e
3 on(x.ssn = e.ssn)
4 inner join public.dependent y
5 on(y.ssn = e.ssn)
6 where x.relationship = 'Son'
7 and y.relationship = 'Daughter'
8 and e.salary > 10000
9
10
```

The results are displayed in the Data Output tab:

	ssn character (9)	lname character varying (15)
1	333445555	Wong
2	123456789	Smith

Annotations in the image:

- A red arrow points to the lightning bolt icon in the toolbar, labeled "Running the query".
- A red arrow points to the thumbs-up icon in the toolbar, labeled "Getting the plan".

For interest only

ASIDE: query plans in PostgreSQL

To get the query plan

The screenshot shows the PostgreSQL Query Editor interface. The top toolbar contains various icons, with the 'Explain' icon (a document with a magnifying glass) circled in red. Below the toolbar, the 'Query Editor' tab is active, displaying a SQL query:

```
1 SELECT ssn, lname FROM public.dependent x
2 inner join public.employee e
3 on(x.ssn = e.ssn)
4 inner join public.dependent y
5 on(y.ssn = e.ssn)
6 where x.relationship = 'Son'
7 and y.relationship = 'Daughter'
8 and e.salary > 10000
9
10
```

Below the query editor, the 'Data Output' tab is active, showing the 'Explain' view of the query plan. The plan consists of several steps:


- public.dependent**: A table scan operation.
- Hash**: A hash operation on the **public.dependent** table.
- public.employee**: A table scan operation.
- Hash Inner Join**: A hash inner join operation combining the **public.employee** table and the **Hash** operation.
- public.dependent_pkey**: A primary key scan operation on the **public.dependent** table.
- Nested Loop Inner Join**: A nested loop inner join operation combining the **Hash Inner Join** and the **public.dependent_pkey** operation.

On the right, if one lets the mouse hover over a node you get more information... for example, the output from each step, algorithm, filtering

... to get the query plan in text → next slide

For interest only

ASIDE: query plans in PostgreSQL



Execute the query but prefix the Select with “explain”

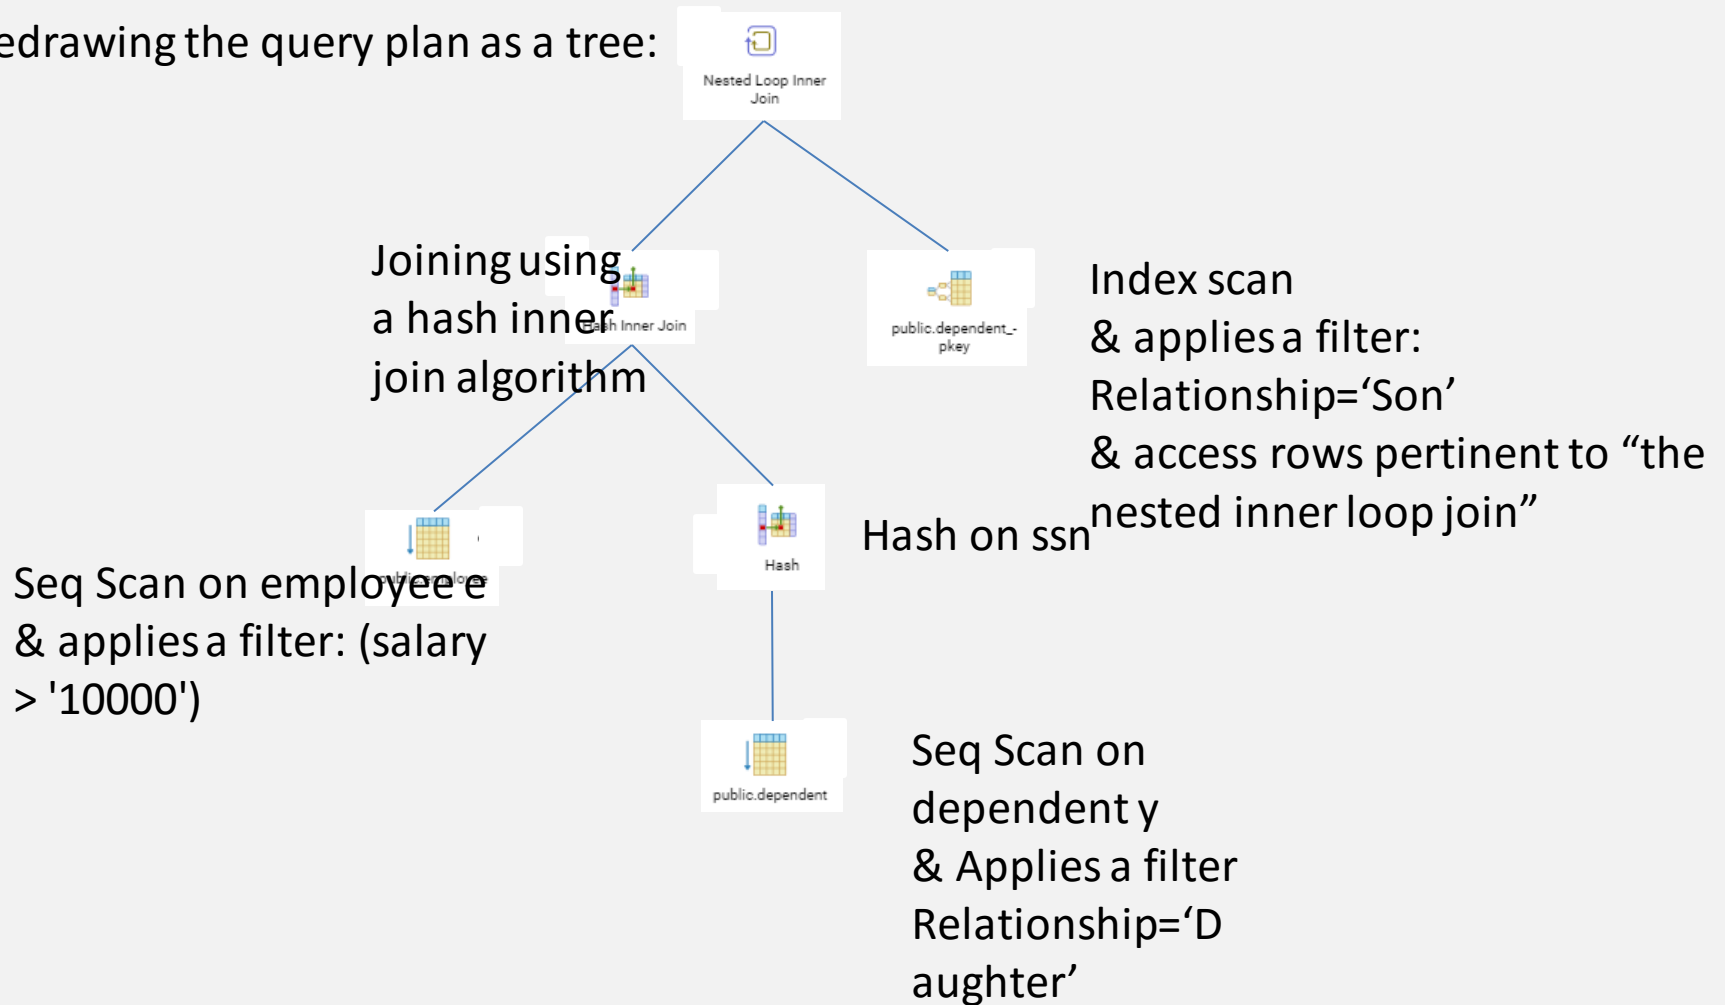
Query plan in **text** form – note that it shows join algorithms and filters

```
Nested Loop (cost=16.42..30.50 rows=1 width=88)
-> Hash Join (cost=16.27..29.72 rows=1 width=128)
    Hash Cond: (e.ssn = y.ssn)
    -> Seq Scan on employee e (cost=0.00..13.13 rows=83 width=88)
        Filter: (salary > '10000'::numeric)
    -> Hash (cost=16.25..16.25 rows=2 width=40)
        -> Seq Scan on dependent y (cost=0.00..16.25 rows=2 width=40)
            Filter: ((relationship)::text = 'Daughter'::text)
    -> Index Scan using dependent_pkey on dependent x (cost=0.15..0.77 rows=1
width=40)
        Index Cond: (ssn = e.ssn)
        Filter: ((relationship)::text = 'Son'::text)
```

For interest only

ASIDE: query plans in PostgreSQL → similar to a query tree

Redrawing the query plan as a tree:



For interest only

Aside: Query plans in real database systems are trees where nodes are operations the database system implements

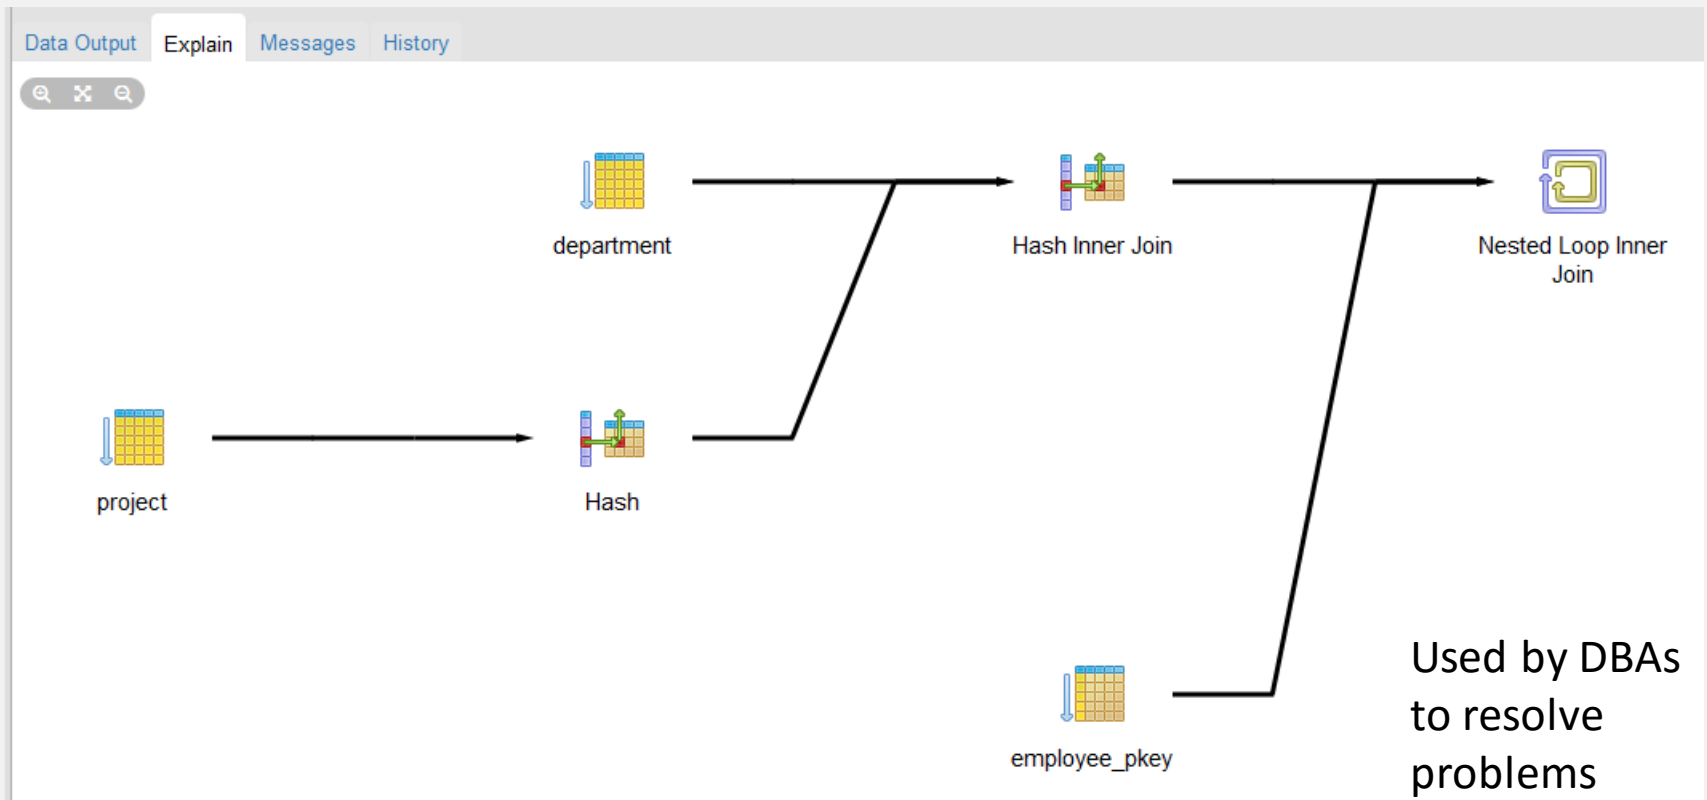
Some examples→

<https://www.postgresql.org/docs/current/static/using-explain.html>

For interest only

Query 2. For every project located in 'Stafford', list the project number, the controlling department number, and the department manager's last name, address, and birth date.

An actual explain plan from PostgreSQL



For interest only

Query 2. (again...rerun today)

4902 on postgres@PostgreSQL 11

Query Editor Query History

```
1 SELECT Pnumber, Dnum, Lname, Address, Bdate
2 FROM PROJECT, DEPARTMENT, EMPLOYEE
3 WHERE Dnum=Dnumber
4 AND Mgr_ssn=Ssn
5 AND
6 Plocation='Stafford'
```

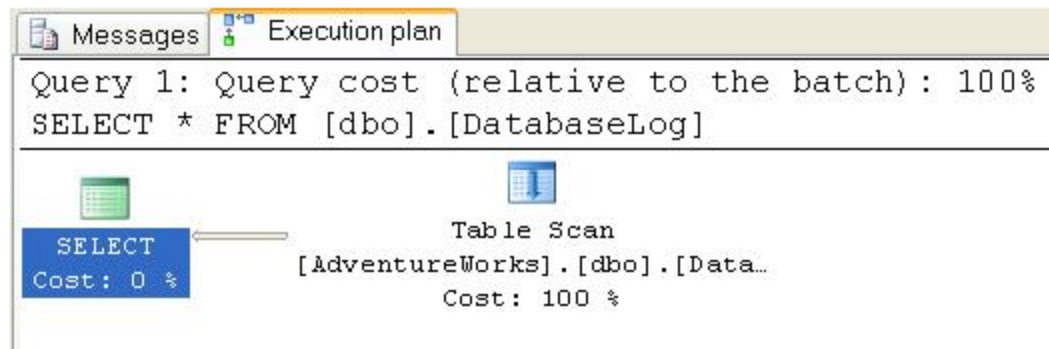
Data Output Explain Messages Notifications

The diagram illustrates the execution plan for the query. It shows three tables: 'project', 'department_pkey', and 'employee_pkey'. The 'project' table is joined with 'department_pkey' using a 'Nested Loop Inner Join'. The result of this join is then joined with 'employee_pkey' using another 'Nested Loop Inner Join'. The tables are represented by icons: a grid for 'project', a grid with a key for 'department_pkey', and a grid with a key for 'employee_pkey'. Arrows indicate the flow of data from the tables to the join operations.

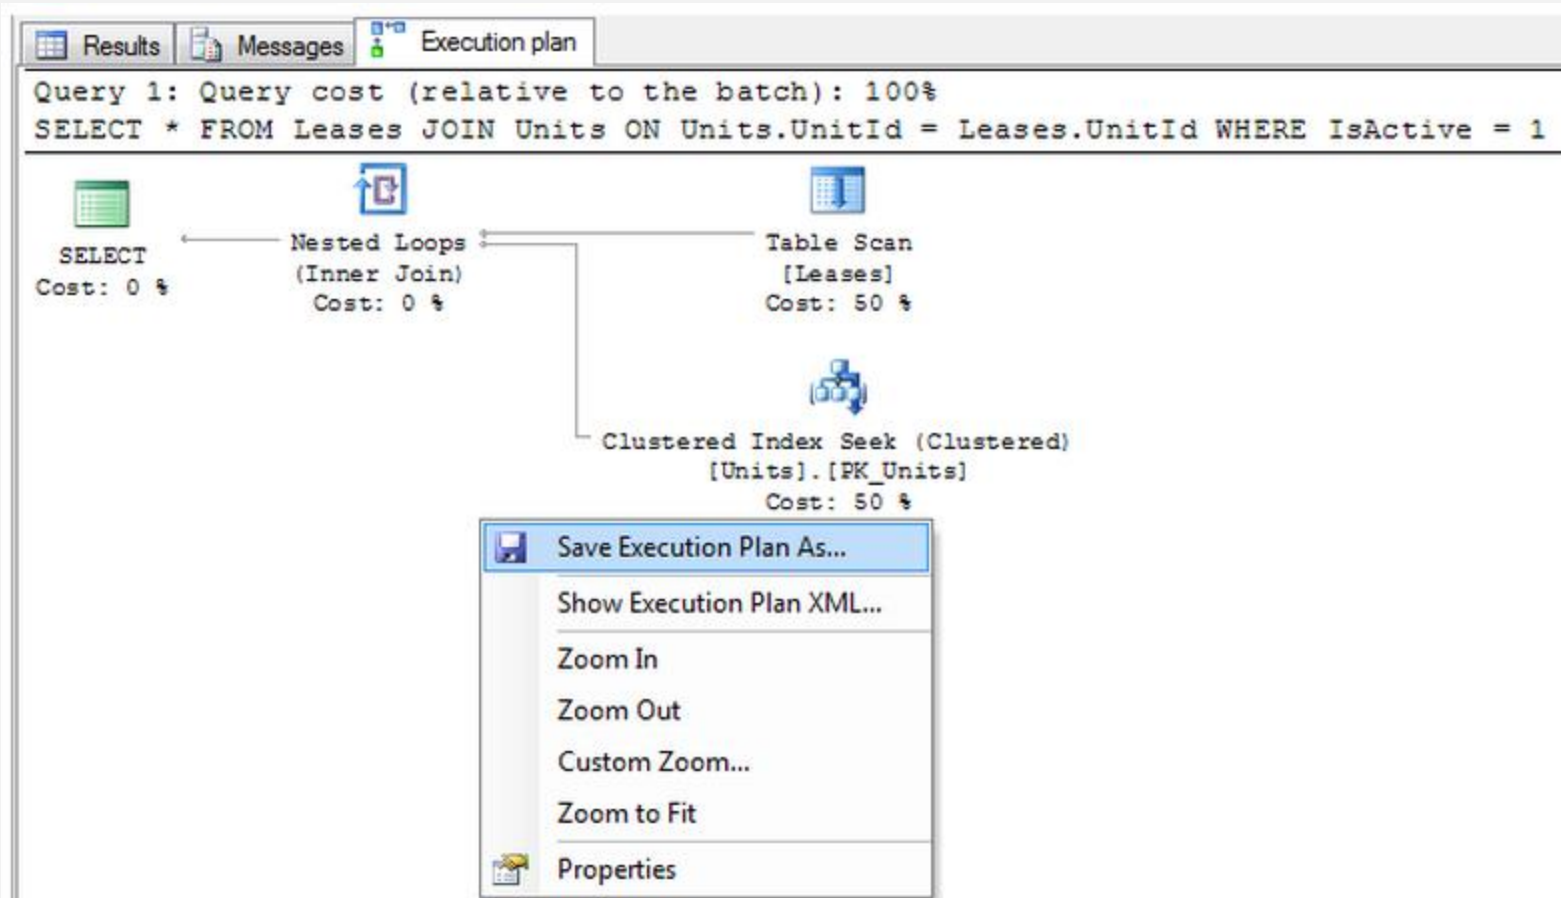
For interest only

```
SELECT *  
FROM [dbo].[DatabaseLog];
```

I tend to click the icon more often than not but, either way, we see our very first **Estimated execution plan**, as in Figure 1.



For interest only



For interest only

```
SELECT CONCAT(customer.last_name, ', ', customer.first_name) AS customer, address.phone, film.title
FROM rental
INNER JOIN customer ON rental.customer_id = customer.customer_id
INNER JOIN address ON customer.address_id = address.address_id
INNER JOIN inventory ON rental.inventory_id = inventory.inventory_id
INNER JOIN film ON inventory.film_id = film.film_id
WHERE rental.return_date IS NULL
AND rental_date + INTERVAL film.rental_duration DAY < CURRENT_DATE()
LIMIT 5;
```

