

# JOIN AND AGGREGATION ALGORITHMS

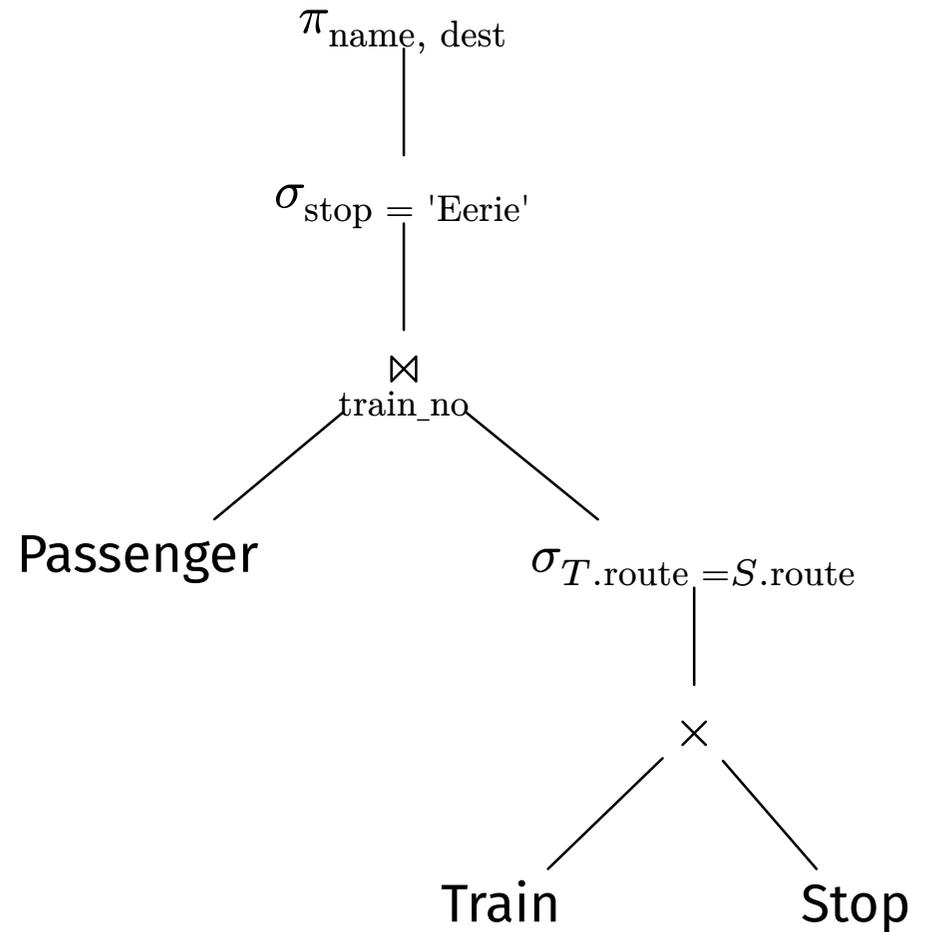
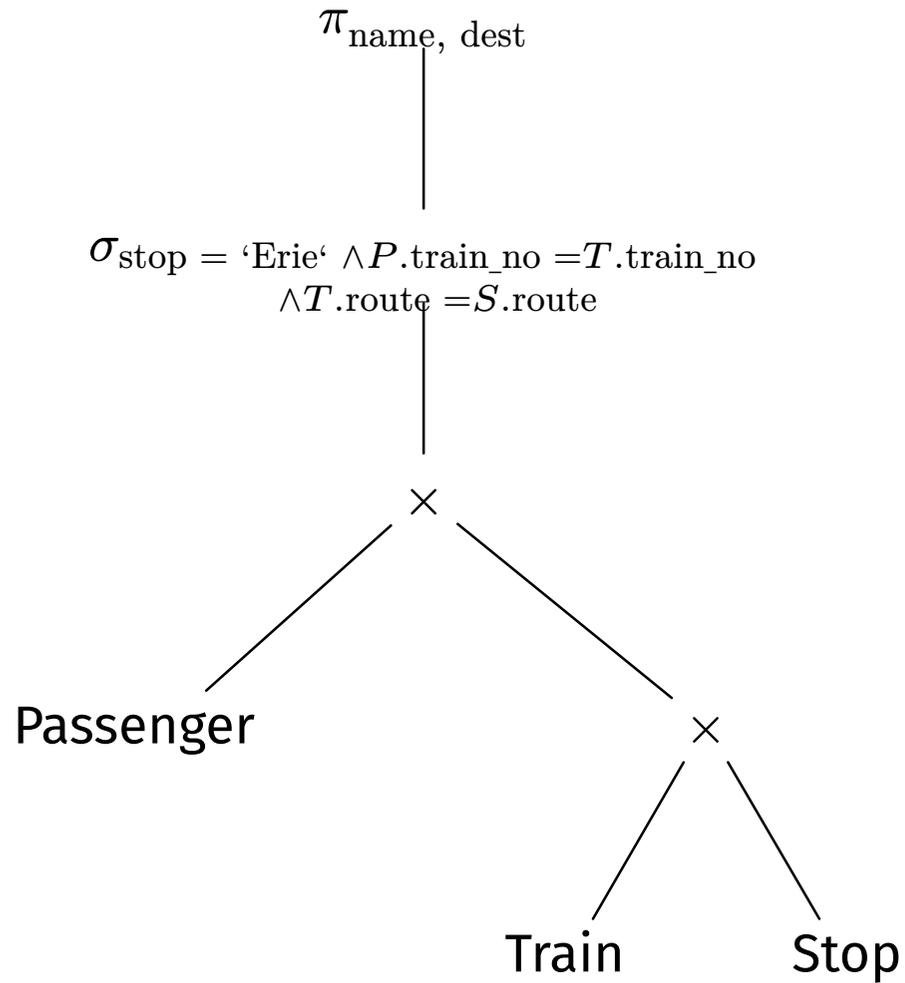
CSE 4/562: Database Systems | Lecture 5

---

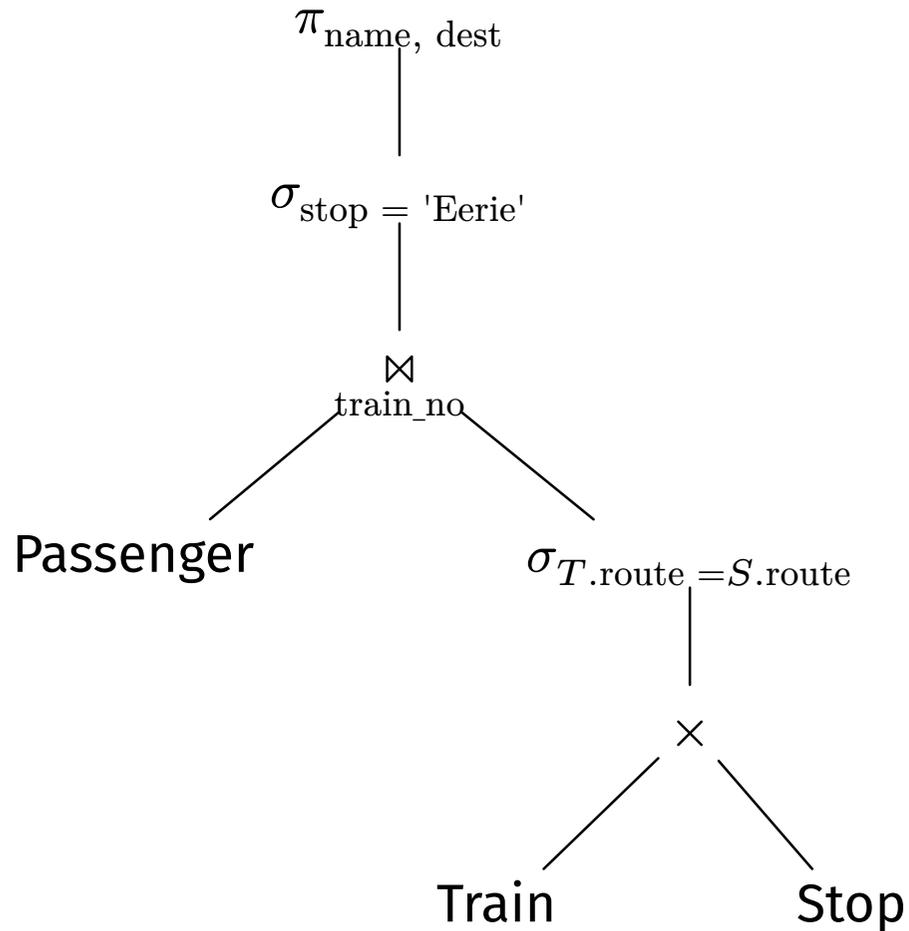
**DB. Sys.: T.C.B.: 16.1-16.3**

**Recap**

```
SELECT name, dest
FROM Passenger P, Train T, Stop S
WHERE stop = 'Erie'
      AND P.train_no = T.train_no
      AND T.route = S.route
```



<b>Operator</b>	<b>Memory</b>	<b>IO</b>
Table	$O(1)$	$O(n)$
Project ( $\pi$ )	$O(1)$	$O(1)$
Filter ( $\sigma$ )	$O(1)$	$O(1)$
Union ( $\cup$ )	$O(1)$	$O(1)$
Nested Loop ( $\times$ )	$O(1)$	$O(n^2)$
Block Nested Loop ( $\times$ )	$O(\text{Buffer})$	$O\left(\frac{n^2}{\text{Buffer}}\right)$
2-Pass Hash Join ( $\bowtie$ )	$O\left(\frac{n}{\text{Buckets}}\right)$	$O(n)$
Sort Merge Join ( $\bowtie$ )	$O(1)+\text{sort}$	$O(1)+\text{sort}$
2-Pass Hash Aggregate ( $\Sigma$ )	$O\left(\frac{n}{\text{Buckets}}\right)$	$O(n)$
Sort Aggregate	$O(1)+\text{sort}$	$O(1)+\text{sort}$



Table(Train)

Table(Stop)

BlockNLoop


Filter

Table(Pa...)

2PHashJoin

Filter

Project



$O(|T|)$  Table(Train)

$+O(|S|)$  Table(Stop)

BlockNLoop


Filter

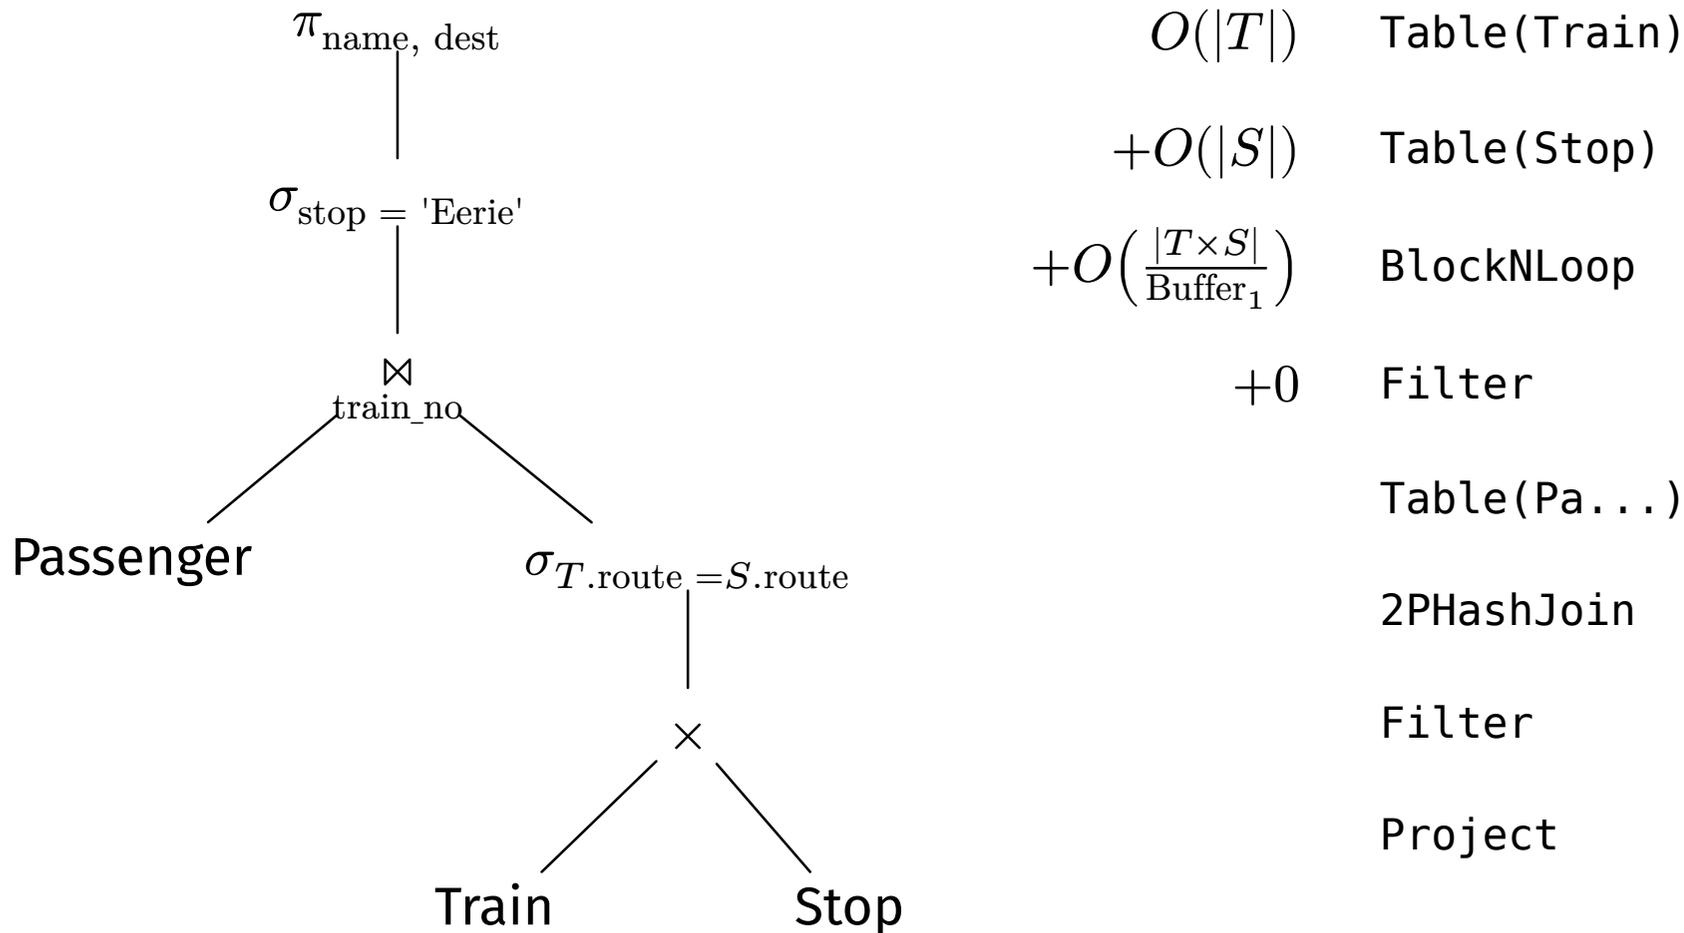
Table(Pa...)

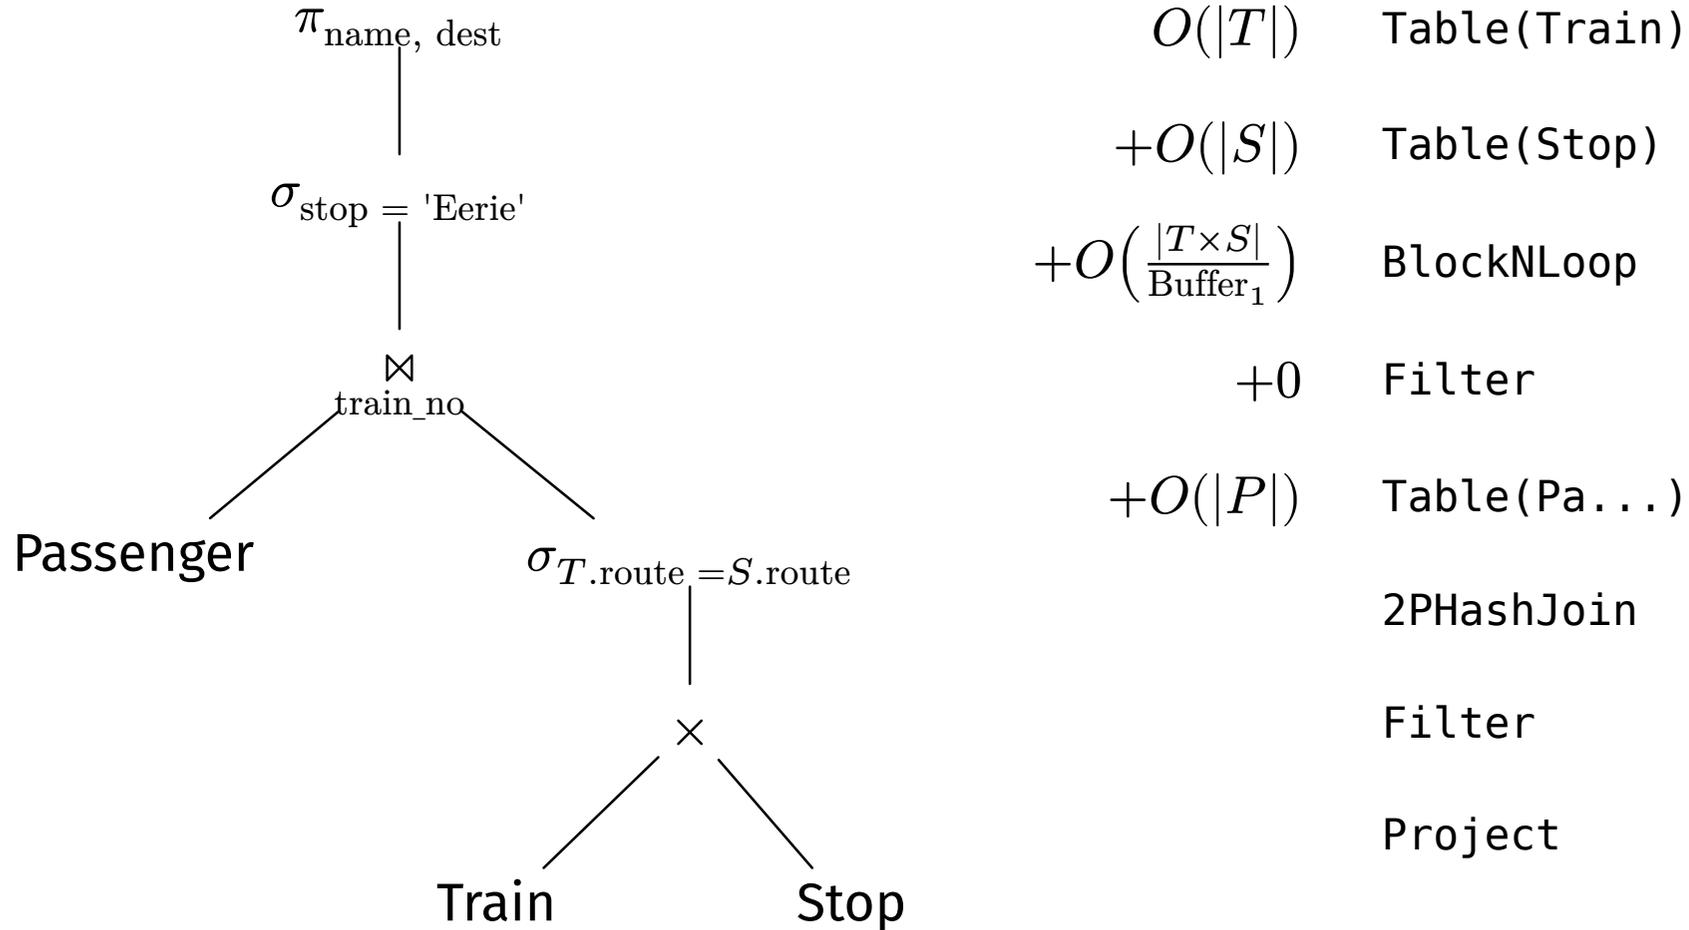
2PHashJoin


Filter

Project

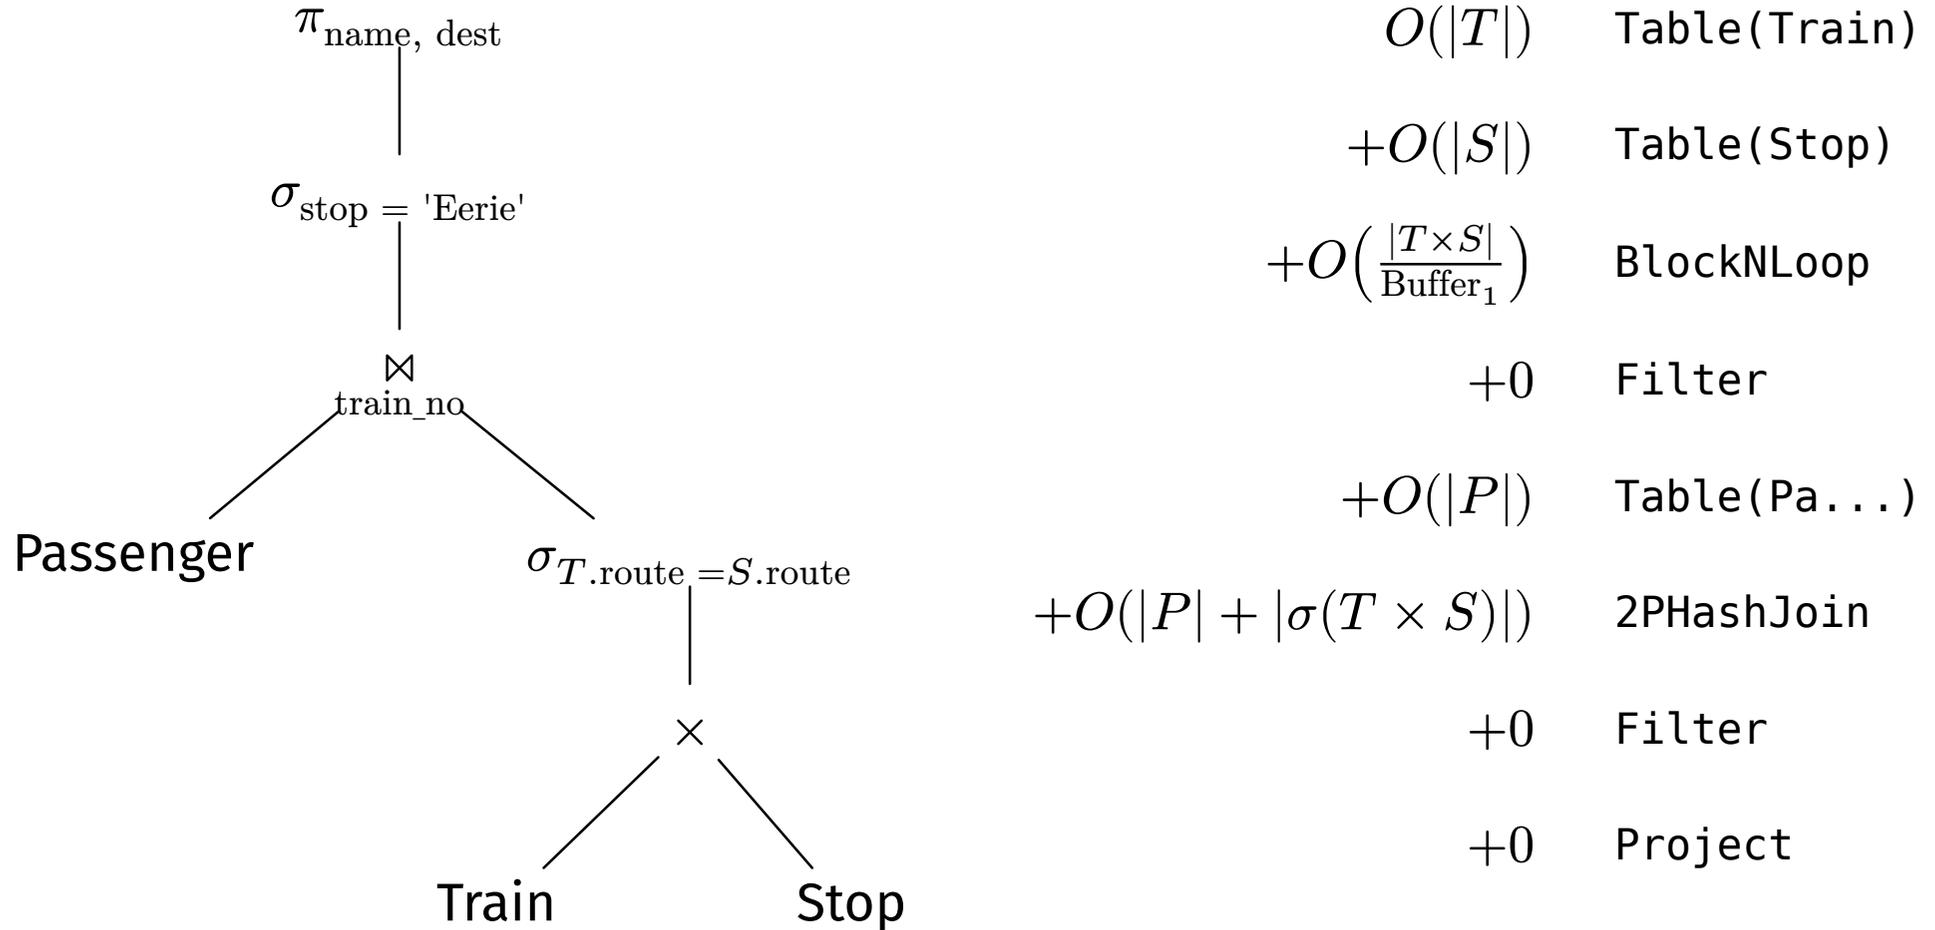


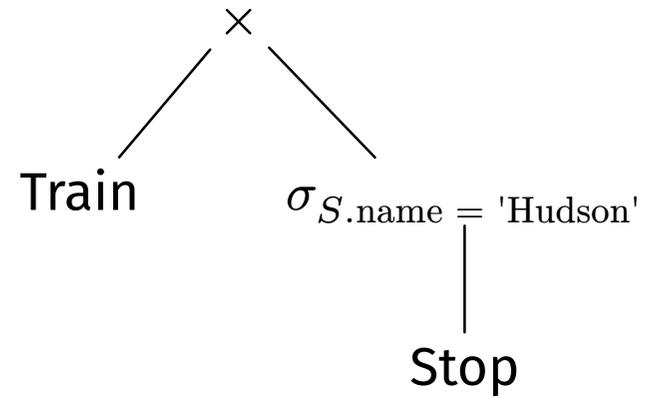
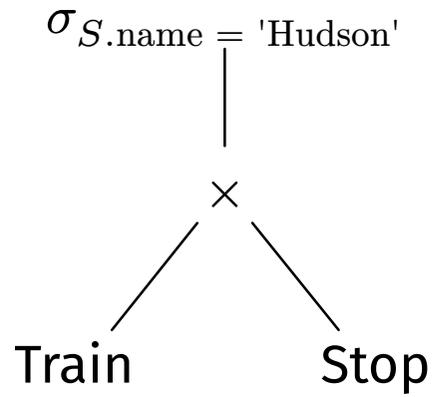


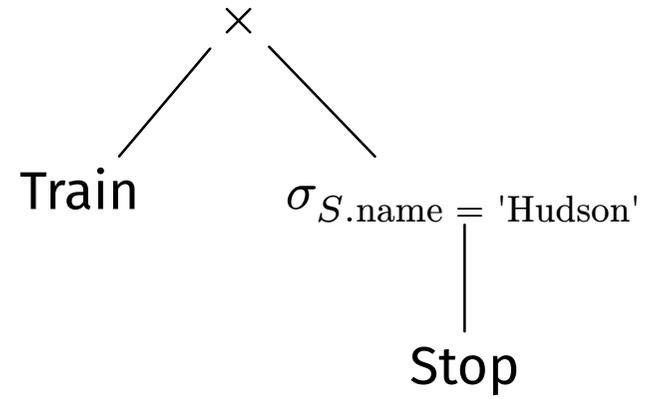
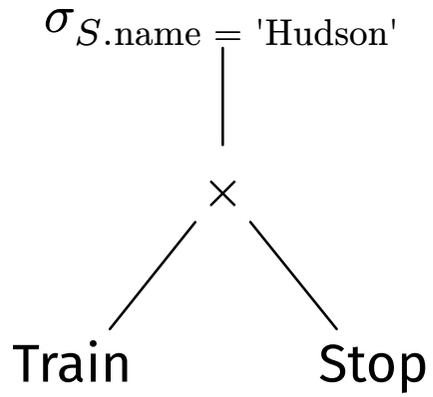




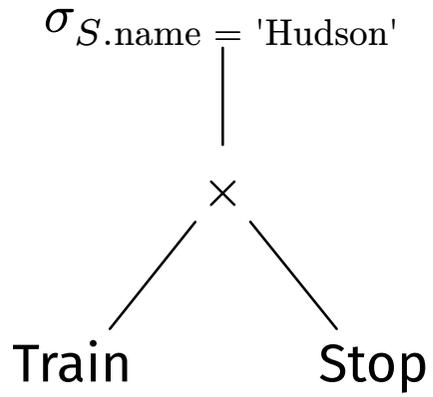
$O( T )$	Table(Train)
$+O( S )$	Table(Stop)
$+O\left(\frac{ T \times S }{\text{Buffer}_1}\right)$	BlockNLoop
$+0$	Filter
$+O( P )$	Table(Pa...)
$+O( P  +  \sigma(T \times S) )$	2PHashJoin
	Filter
	Project



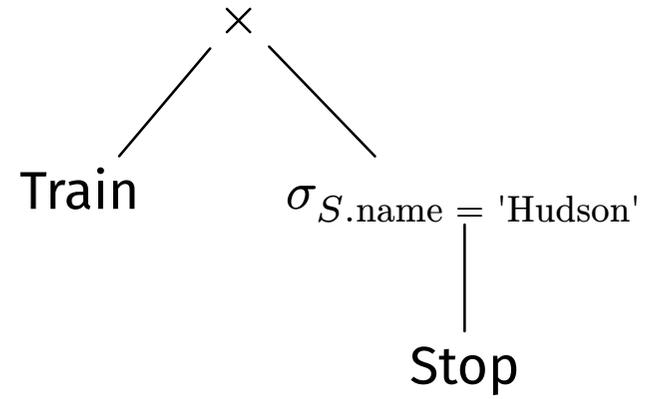




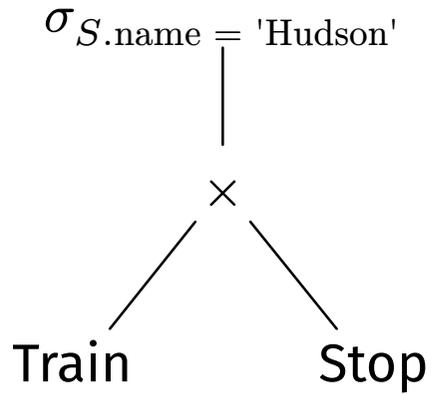
$$|T| + |S| + \frac{|T| \times |S|}{\text{Buffer}} + 0$$




$$|T| + |S| + \frac{|T| \times |S|}{\text{Buffer}} + 0$$



$$|T| + |S| + 0 + \frac{|T| \times |\sigma S|}{\text{Buffer}}$$



$$|T| + |S| + \frac{|T| \times |S|}{\text{Buffer}} + 0$$



$$|T| + |S| + 0 + \frac{|T| \times |\sigma S|}{\text{Buffer}}$$

$$|S| > |\sigma S|$$

**Which query plans  
are the same?**

$$a \times (b + (c \times (e + f)))$$

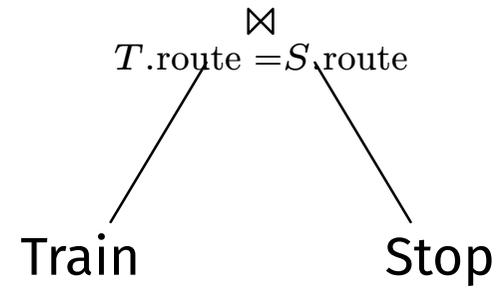
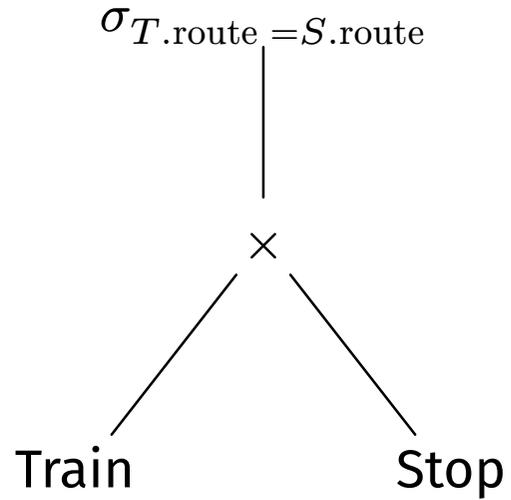
$$ace + ba + fac$$

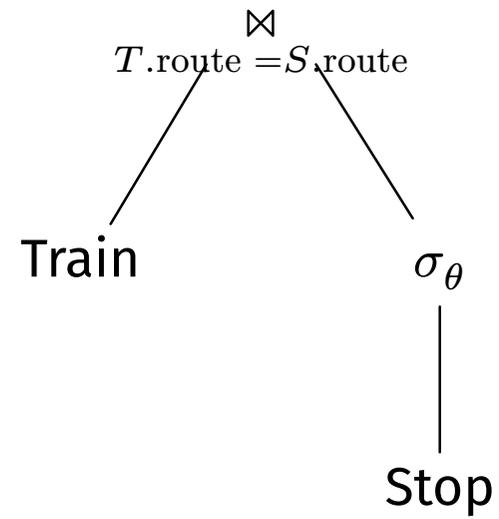
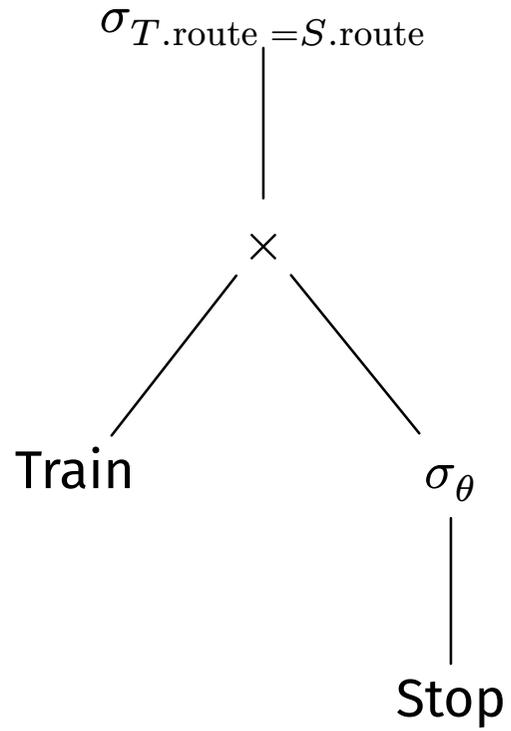
Are these the same?

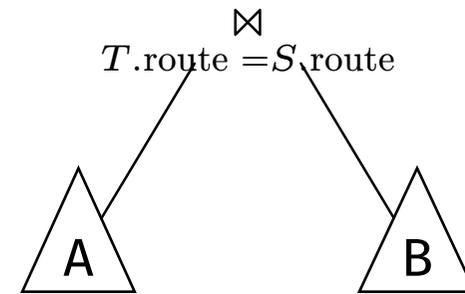

$$a \times (b + (c \times (e + f)))$$



$$ace + ba + fac$$

Is one easier to solve than the other?









- $\sigma_{\theta}(R \times S) \rightarrow R \bowtie_{\theta} S$ 
  - 🏃 Usually faster if  $\theta$  is an equality test
  - Strategy 1 (Fixed Point; Spark)
- $R \times S \rightarrow S \times R$ 
  - 🙋 Sometimes a good idea?
  - Strategy 2 (Cost-Based; Postgres)

- Search through the plan for an occurrence of  $\sigma_\theta(Q_1 \times Q_2)$ .
- Replace it with  $Q_1 \bowtie_\theta Q_2$ .
- Repeat until done.

- For (**Pattern**  $\rightarrow$  **Rewrite**) in **Rules**
  - Search through the plan for a subplan  $Q$  that matches **Pattern**<sup>1</sup>.
  - Replace  $Q$  with **Rewrite**( $Q$ ).
- Repeat until done.

---

<sup>1</sup>We'll go into more detail when we talk about Checkpoint 3, and after we talk about IVM.

- **Plans** ← Naive Plan
- Loop until done (or threshold)
  - **Plans** ← Apply a new rewrite to a plan in **Plans**
- Estimate the cost of all **Plans** and use the cheapest.

**What are valid rewrites?**

If, for any  $Q_1, Q_2$  where:

- $Q_1$  matches **Pattern**
- $Q_2 = \mathbf{Rewrite} (Q_1)$

...it is also true that, for any database  $D$ :

- $Q_1(D) = Q_2(D)$

...then we say that the rewrite rule (**Pattern, Rewrite**) is valid.

$$Q_1(D) = Q_2(D)$$

$$Q_1(D) = Q_2(D)$$

The relation obtained by evaluating  $Q_1$  on **any**  $D$  is equivalent to that obtained by evaluating  $Q_2$

- **Bag-Equivalence:** The same number of each record. Order does not matter.
- **Set-Equivalence:** The same records, ignoring duplicates. Order does not matter.
- **List-Equivalence:** The same records in the same order.

$$Q_1(D) = Q_2(D)$$

The relation obtained by evaluating  $Q_1$  on **any**  $D$  is equivalent to that obtained by evaluating  $Q_2$

- **Bag-Equivalence:** The same number of each record. Order does not matter.
- **Set-Equivalence:** The same records, ignoring duplicates. Order does not matter.
- **List-Equivalence:** The same records in the same order.

For now, we use **Bag** equivalence

$$Q_1(D) = Q_2(D)$$

The relation obtained by evaluating  $Q_1$  on **any**  $D$  is equivalent to that obtained by evaluating  $Q_2$

- **Bag-Equivalence:** The same number of each record. Order does not matter.
- **Set-Equivalence:** The same records, ignoring duplicates. Order does not matter.
- **List-Equivalence:** The same records in the same order.

For now, we use **Bag** equivalence

We also ignore attribute order ( $R(A, B, C)$  is the same as  $R(C, B, A)$ )

## Query Equivalence...

$Q_1(D) \equiv Q_2(D)$  if and only if, for any database  $D$ :

- $Q_1(D)$  is a valid query if and only if  $Q_2(D)$  is a valid query.
- The set of attributes of  $Q_1(D)$  is the same as the set of attributes of  $Q_2(D)$ .
- The bag of records of  $Q_1(D)$  is the same as the bag of records of  $Q_2(D)$ .

## Rewrite Validity

**(Pattern, Rewrite)** is valid if for any database  $D$  and any valid query  $Q(D)$  that matches **Pattern, Rewrite**:

- $(\text{Rewrite } (Q))(D) \equiv Q(D)$

## Selection

$$\sigma_{\theta_1 \wedge \theta_2}(R) \equiv \sigma_{\theta_1}(\sigma_{\theta_2}(R)) \quad \text{Decomposability}$$

## Projection

$$\pi_A(R) \equiv \pi_A(\pi_{A \cup B}(R)) \quad \text{Idempotence}$$

## Cartesian Product

$$R \times (S \times T) \equiv (R \times S) \times T \quad \text{Associativity}$$

$$R \times S \equiv S \times R \quad \text{Commutativity}$$

$$\sigma_{\theta}(R \times S) \equiv R \bowtie_{\theta} S \quad \text{Join Conversion}$$

## Union

$$R \cup (S \cup T) \equiv (R \cup S) \cup T \quad \text{Associativity}$$

$$R \cup S \equiv S \cup R \quad \text{Commutativity}$$

**Try it: Show that...**

$$R \times (S \times T) \equiv T \times (S \times R)$$

**Try it: Show that...**

$$\sigma_{\theta_1}(\sigma_{\theta_2}(R)) \equiv \sigma_{\theta_2}(\sigma_{\theta_1}(R))$$

**Try it: Show that...**

$$R \otimes_{\theta} S \equiv S \otimes_{\theta} R$$

**Try it: Show that...**

$$\sigma_{R.B=S.B \wedge R.A>3}(R \times S) \equiv \sigma_{R.A>3} \left( R \underset{B}{\bowtie} S \right)$$

## Selection + Projection

$$\pi_A(\sigma_\theta(R)) \equiv \sigma_\theta(\pi_A(R)) \quad \text{Commutativity}$$

... but only if  $A$  and  $\theta$  are compatible

$A$  must include all columns referenced by  $\theta$  ( $\text{cols}(\theta)$ )

**Try it: Show that...**

$$\pi_A(\sigma_\theta(R)) \equiv \pi_A\left(\sigma_\theta\left(\pi_{A \cup \text{cols}(\theta)}(R)\right)\right)$$

## Selection+Cross Product

$$\sigma_{\theta}(R \times S) \equiv (\sigma_{\theta}(R)) \times S \quad \text{Selection "Push Down"}$$

... but only  $\theta$  references only columns of  $R$

$$\text{cols}(\theta) \subseteq \text{cols}(R)$$

**Try it: Show that...**

$$\sigma_{R.B=S.B \wedge R.A>3}(R \times S) \equiv (\sigma_{R.A>3}(R)) \bowtie_B S$$

## Projection+Cross Product

$$\pi_A(R \times S) \equiv \left( \pi_{A_R}(R) \right) \times \left( \pi_{A_S}(S) \right) \quad \text{Projection "Push Down"}$$

... where  $A_R = A \cap \text{cols}(R)$  and  $A_S = A \cap \text{cols}(S)$

**Try it: Show that...**

$$\pi_A(R \bowtie S) \equiv \left( \pi_{A_R}(R) \right) \bowtie \left( \pi_{A_S}(S) \right)$$

(Is this always true?)

## Union

$$\sigma_{\theta}(R \cup S) \equiv (\sigma_{\theta}(R)) \cup (\sigma_{\theta}(S))$$

Selection “Push Down”

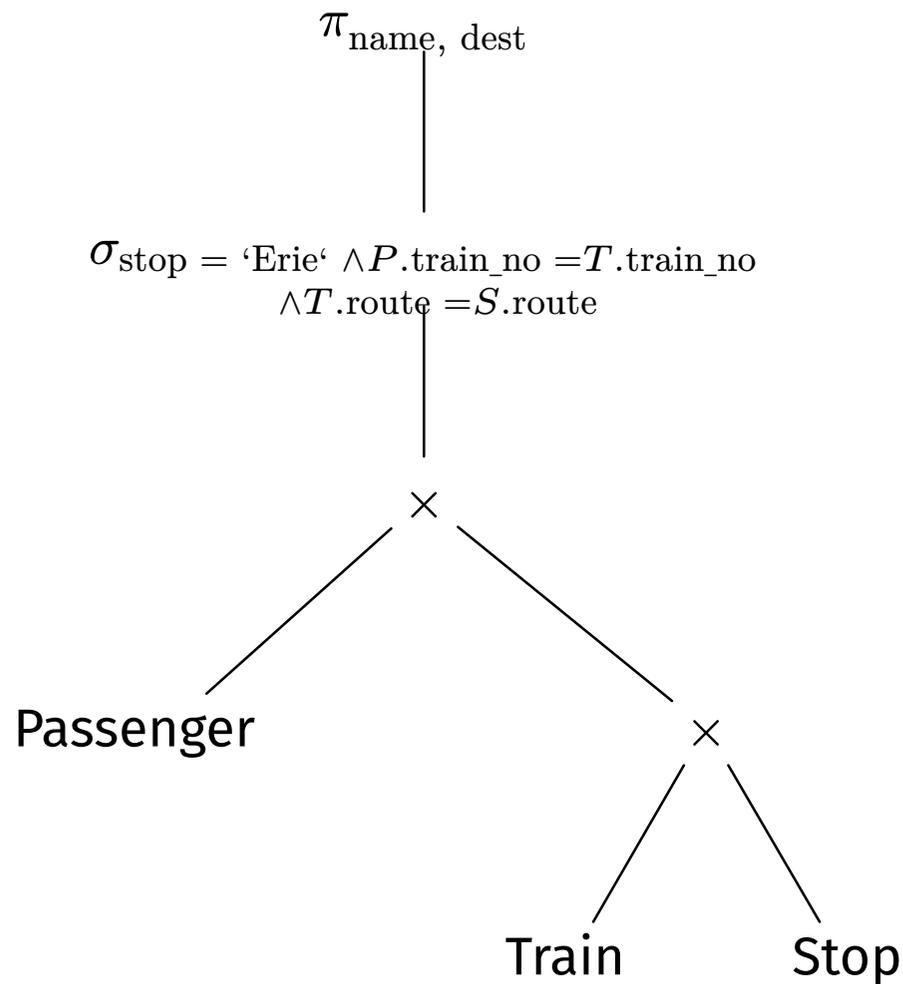
$$\pi_A(R \cup S) \equiv (\pi_A(R)) \cup (\pi_A(S))$$

Projection “Push Down”

$$R \times (S \cup T) \equiv (R \times S) \cup (R \times T)$$

Distributivity

```
SELECT p.name, t.id
FROM passenger p, train t, stop s
WHERE p.train_no = t.train_no
      AND t.route = s.route
      AND stop = 'Erie'
```



## Selection Pushdown & Join Construction

Almost always a good idea

## Projection Pushdown

May remove redundant columns (data copies), and may avoid loading data

## Join Algorithm Selection

Remember that  $\times$  and  $\bowtie$  are actually several different algorithms

## Join/Union Ordering

$R \times S$  and  $S \times R$  may have different memory requirements.

$R \bowtie (S \bowtie T)$  and  $(R \bowtie S) \bowtie T$  have *different* memory/IO requirements.

## Access Paths

$\sigma_{\theta}(R)$  and  $Q(\dots) \bowtie_{\theta} R$  are special cases that we can sometimes optimize

## **Deliverables**

- AI Quiz and Checkpoint 0 Past Due!
  - Reach out to me if you haven't finished it yet!
- Checkpoint 1 due Monday!
  - Don't forget to schedule a code review (times available tomorrow and next Friday).