

SC3020 Database Systems Principles Project 2

[GitHub](#) | [Report](#) | [Dataset](#)

Connect to Database

This project needs to connect to PostgreSQL database. If you have already set up the database with TPC-H, you can directly click the 'Connect' button. Otherwise, you need to click the Load Data button, it will download and load the data automatically from Hugging Face. It will take a while as the data is a bit large.

Host	Port	Database	Username	Password
localhost	5432	mydb	sc3020	*****

Connect
Load Data

Status

```

1  {
2  "status": "Connected",
3  "host": "localhost:5432",
4  "dbname": "mydb",
5  "user": "sc3020"
6  }

```

a)

Query

```

1  SELECT
2  n_name,
3  SUM(l_extendedprice * (1 - l_discount)) AS revenue
4  FROM
5  customer,
6  orders,
7  lineitem,
8  supplier,
9  nation,
10 region
11 WHERE
12 c_custkey = o_custkey
13 AND l_orderkey = o_orderkey
14 AND l_suppkey = s_suppkey
15 AND c_nationkey = s_nationkey
16 AND s_nationkey = n_nationkey
17 AND n_regionkey = r_regionkey
18 AND r_name = 'ASIA'
19 AND o_orderdate >= '1994-01-01'
20 AND o_orderdate < '1995-01-01'
21 AND c_acctbal > 10
22 AND s_acctbal > 20
23 GROUP BY
24 n_name
25 ORDER BY
26 revenue DESC;

```

Query Plan

b)

Estimate Startup Cost	Estimate Total Cost
811420.33	1429613.95

What if change scan to ...	What if change join to ...
Default	Default

Explain

Step 1: Perform Parallel Seq Scan on lineitem

Step 2: Perform Seq Scan on supplier and filtering on (s_acctbal > '20'::numeric)

Step 3: Perform Seq Scan on nation

Step 4: Perform Seq Scan on region and filtering on (r_name = 'ASIA'::bpchar)

Step 5: Perform Hash

Step 6: Perform Hash Join and do Hash with condition on (nation.n_regionkey = region.r_regionkey)

Step 7: Perform Hash

Step 8: Perform Hash Join and do Hash with condition on (supplier.s_nationkey = nation.n_nationkey)

c)

Result (Top 100 rows)

n_name	revenue
INDONESIA	45725603.5903
VIETNAM	45448434.1933
CHINA	42484738.5357
INDIA	42390691.0550
JAPAN	37556124.1326

Logs

```

1  {
2  "status": "Success",
3  "message": "Returned 5 rows",
4  "execution_time": "2.985 seconds"
5  }

```

d)