

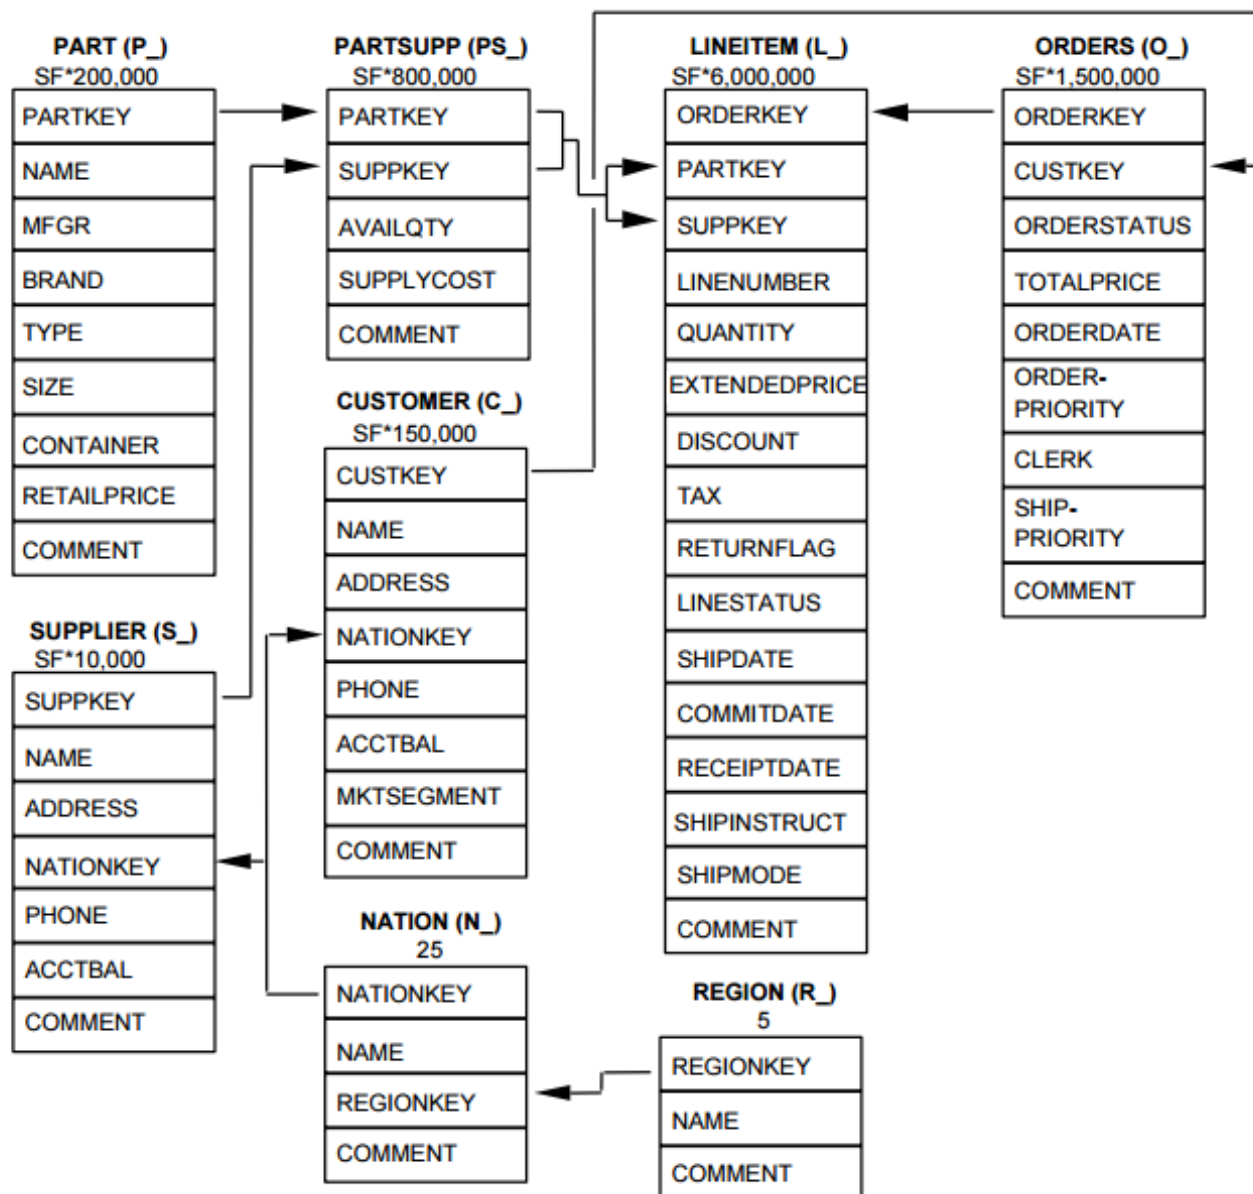


CS425 – Project Assignment

- n Tune the TPC-H queries in a database of your choice...
- n What is TPC-H?
 - | Industry standard benchmark for testing database performance for complex SQL
 - ▶ Decision support
 - ▶ Analytics
 - ▶ Ad-hoc
- n TPC-H specification is here:
 - | http://www.tpc.org/tpc_documents_current_versions/pdf/tpc-h_v2.17.1.pdf
 - ▶ We're not actually going to run the full benchmark according to specs
 - ▶ Just focus on making the queries run faster



TPC-H Data Model





TPC-H Scale Factor

- n Scale Factor (SF) 1 is approximately 1GB of data
- n Defined scale factors are:
 - | 1, 10, 30, 100, 300, 1000, 3000, 10000, 30000, 100000

Table Name	Cardinality (in rows)	Length (in bytes) of Typical ² Row	Typical ² Tab Size (in MB)
SUPPLIER	10,000	159	2
PART	200,000	155	30
PARTSUPP	800,000	144	110
CUSTOMER	150,000	179	26
ORDERS	1,500,000	104	149
LINEITEM ³	6,001,215	112	641
NATION ¹	25	128	< 1
REGION ¹	5	124	< 1
Total	8,661,245		956

Scale Factor (SF)	Cardinality of LINEITEM Table
1	6001215
10	59986052
30	179998372
100	600037902
300	1799989091
1000	5999989709
3000	18000048306
10000	59999994267
30000	179999978268
100000	599999969200



Defining the TPC-H Tables

- n Section 1.4 in the specification gives all the info you need for figuring out what your CREATE TABLE statements should look like.
- n Whether or not to implement primary keys and foreign keys is your choice
- n ORDERS Table Layout
- n Column Name Datatype Requirements Comment
- n O_ORDERKEY Identifier SF*1,500,000 are sparsely populated
- n O_CUSTKEY Identifier Foreign Key to C_CUSTKEY
- n O_ORDERSTATUS fixed text, size 1
- n O_TOTALPRICE Decimal
- n O_ORDERDATE Date
- n O_ORDERPRIORITY fixed text, size 15
- n O_CLERK fixed text, size 15
- n O_SHIPPRIORITY Integer
- n O_COMMENT variable text, size 79
- n Primary Key: O_ORDERKEY



Loading Your Tables

- n First you have to generate the data to be loaded
- n This is done with the dbgen tool
 - | http://www.tpc.org/tpc_documents_current_versions/download_programs/tools-download-request.asp?BM=TPC-H&mode=CURRENT-ONLY
- n DBGEN is delivered as source code and has to be compiled
- n On linux copy makefile.suite to Makefile, and edit
 - | DATABASE=DB2
 - | MACHINE =LINUX
- n The run, make.
- n This assumes you have gcc installed (yum install gcc.x86_64)
- n `./dbgen -s 1` (would generate data for SF1)
- n Output files end in `.tbl`
- n They are pipe delimited text files



Loading Your Tables 2

- n Now that you have the data generated, you need to load the data into your tables
- n This process is different for each database
- n For Postgres, see the COPY command
- n For MySQL, see the LOAD DATA command
- n For DB2, see the LOAD command
- n Etc...



TPC-H Queries

- n There are 22 TPC-H queries
- n The SQL is given in section 2.4 of the specification
- n The queries have substitution parameters
- n Use the parameters specified in the query validation info for each query
- n Change the SQL as little as possible
- n The biggest challenge is that the syntax for the date stuff is different from database to database
- n Some other calls like SUBSTRING (SUBSTR) might need to be modified slightly
- n You also may want to qualify all the table names with a schema
- n Before you do any changes more severe than this type of thing, check with me



How to Benchmark

- n Save your SQL in a text file
- n All database have a way that you can ask them to execute whatever SQL is in a text file
 - | Postgres – `psql -f`
 - | The mysql command plus input redirection
 - | Etc...
- n But, we also need to know how long each query takes
 - | Postgres - put “\timing on” in your script
 - ▶ Should also add “\pset pager off”
 - | MySQL – gives you elapsed time by default
- n Probably want to redirect your output as well, so that it gets saved
- n One run is usually not good enough for us to have much confidence in the numbers
- n A better approach is to do 3 runs and take the mode



How to Tune

- n You have many options available to make the SQL statements faster
 - | Database specific extensions to the CREATE TABLE statement
 - | Look at options on CREATE DATABASE / CREATE TABLESPACE statements
 - | Buffer manager buffer sizes
 - | Indexes
 - | System parameters
 - ▶ postgresql.conf
 - ▶ my.cnf
 - ▶ Etc...



How to Tune 2

- n Some system parameters can usually be changed online via some command
- n Others cannot...
- n Regardless, to make the change permanent, you normally need to edit the configuration file
- n Good benchmarking practice says to make as few changes as possible and then retest
 - | Sometimes, changing 1 parameter requires changes to something else for it all to work effectively and that's OK
 - | But, make your change set as small as possible
- n If re-testing shows improvement, keep the change
- n If not, revert it
- n Keep track of all the changes you keep
 - | Also track the performance #s before and after



What to Turn In

- n Start with SF1 and tune it the best you can
- n Then go to SF10 and see if there's any additional tuning that can be done (there may or may not be)

- n Here's what you will need to submit on Blackboard
 - | HW specs – physical and virtual (cores, mem)
 - | Database type and version
 - | The exact SQL you ended up using for the queries
 - | For the best results you could get at SF1...
 - ▶ All your DDL
 - ▶ Any system parms that were changed from defaults
 - ▶ Elapsed time for each query + the total for all queries
 - ▶ A log of how you got there (changed x to y and #s went from a to b)



What to Turn In 2

- n Also include results for SF10
 - | If any DDL was different, state this
 - | If any system parms were different, state this
 - | Elapsed time for each query + the total for all queries
 - | If there were changes from SF1, a log of how you got there starting from SF1 settings

- n Lessons learned
 - | Successes and failures 😊

- n For classroom section students, please prepare a short slide deck of no more than 10 minutes to share your results and what you learned with the class
 - | This part will not be graded, it's really so that you can learn from each other