TPC - DS

TPC-DS

Decision Support Benchmark

TPC-Public Relations

2008-05-23

TPC Transaction Processing Performance Council

TPC - DS

Agenda

- Benchmark Objective
- General Assumptions
- Business Model
- Elements of TPC-DS
  - Schema
  - Dataset
  - Queries
  - Data Maintenance
  - Execution
- Positioning of TPC-DS to TPC-H

TPC - DS

Benchmark Objectives

- Measures generally applicable aspects of a Decision Support System
  - Examine large volume of data
  - Give answers to real-world business questions
  - Execute queries of various operational requirements
  - Generate intense activity against the database server component of a system (IO, memory, CPU, Interconnect)
  - Remain closely synchronized with source OLTP database through a periodic database maintenance function

TPC - DS

Overview

- Models the decision support functions of a retail product supplier
- Data contains vital business information such as customer, order and product data
- Models decision support queries and data maintenance

TPC Transaction Processing Performance Council

General Requirements

- 6 years of data
- data is skewed
- fine fact table granularity (lineitem)
- DBMS must demonstrate ACID properties
- update against all tables

Schema

- Reflects the business model
- Structure is representative of today's data warehouse systems
- Stresses optimizer due to size and structure

Schema

- Snow-storm schema
  - TPC-DS consists of multiple snowflake schemas, which are multiple star schema with dimensions linking to dimensions
- Logical schema
  - The actual implementation is flexible as long as views in accordance with the specification are provided
- Large number of tables (26)
  - Large number of columns per table (38)
  - Multiple fact tables to enable joins between large tables (fact to fact joins)

Schema: TPC-DS

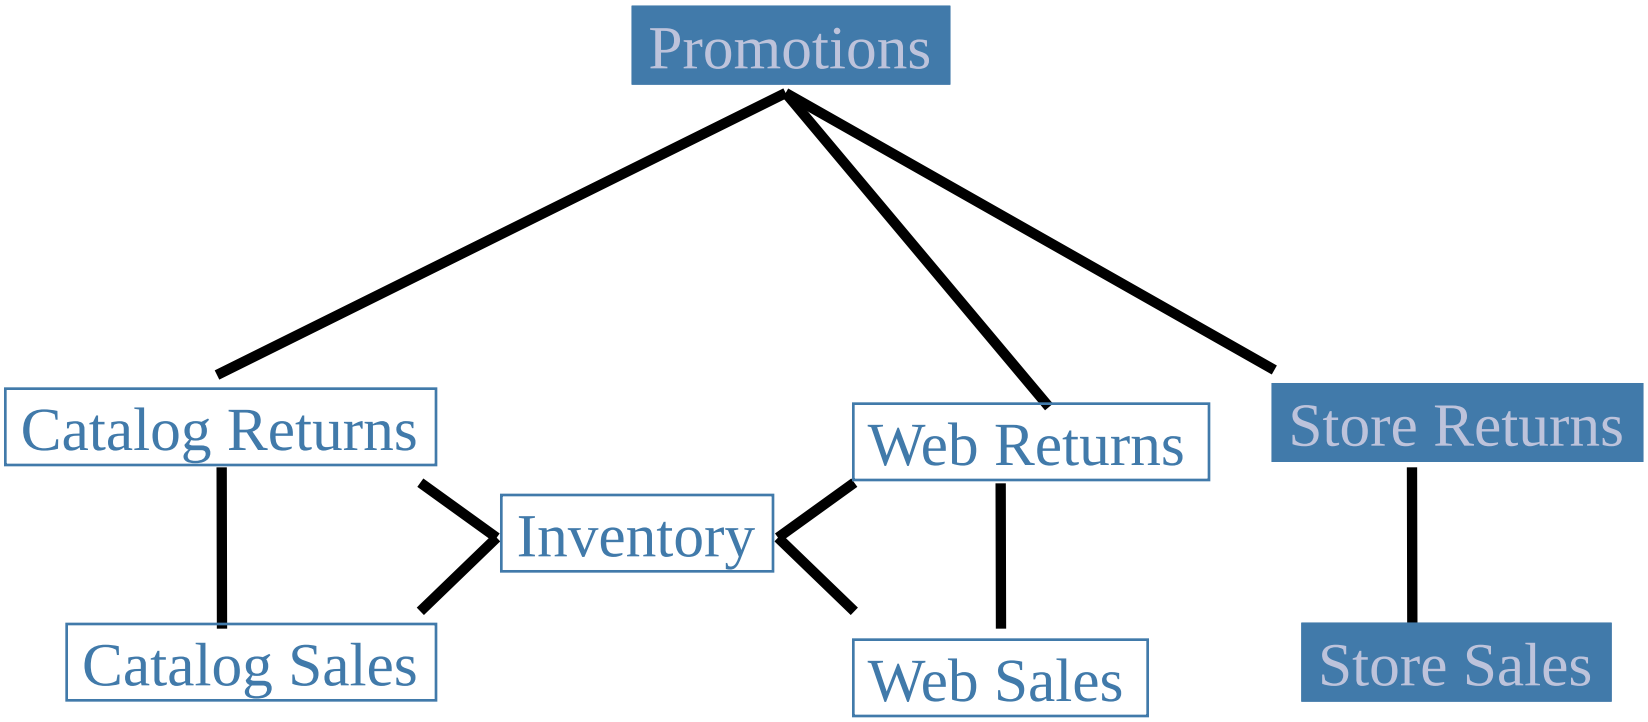

Schema: Store Channel

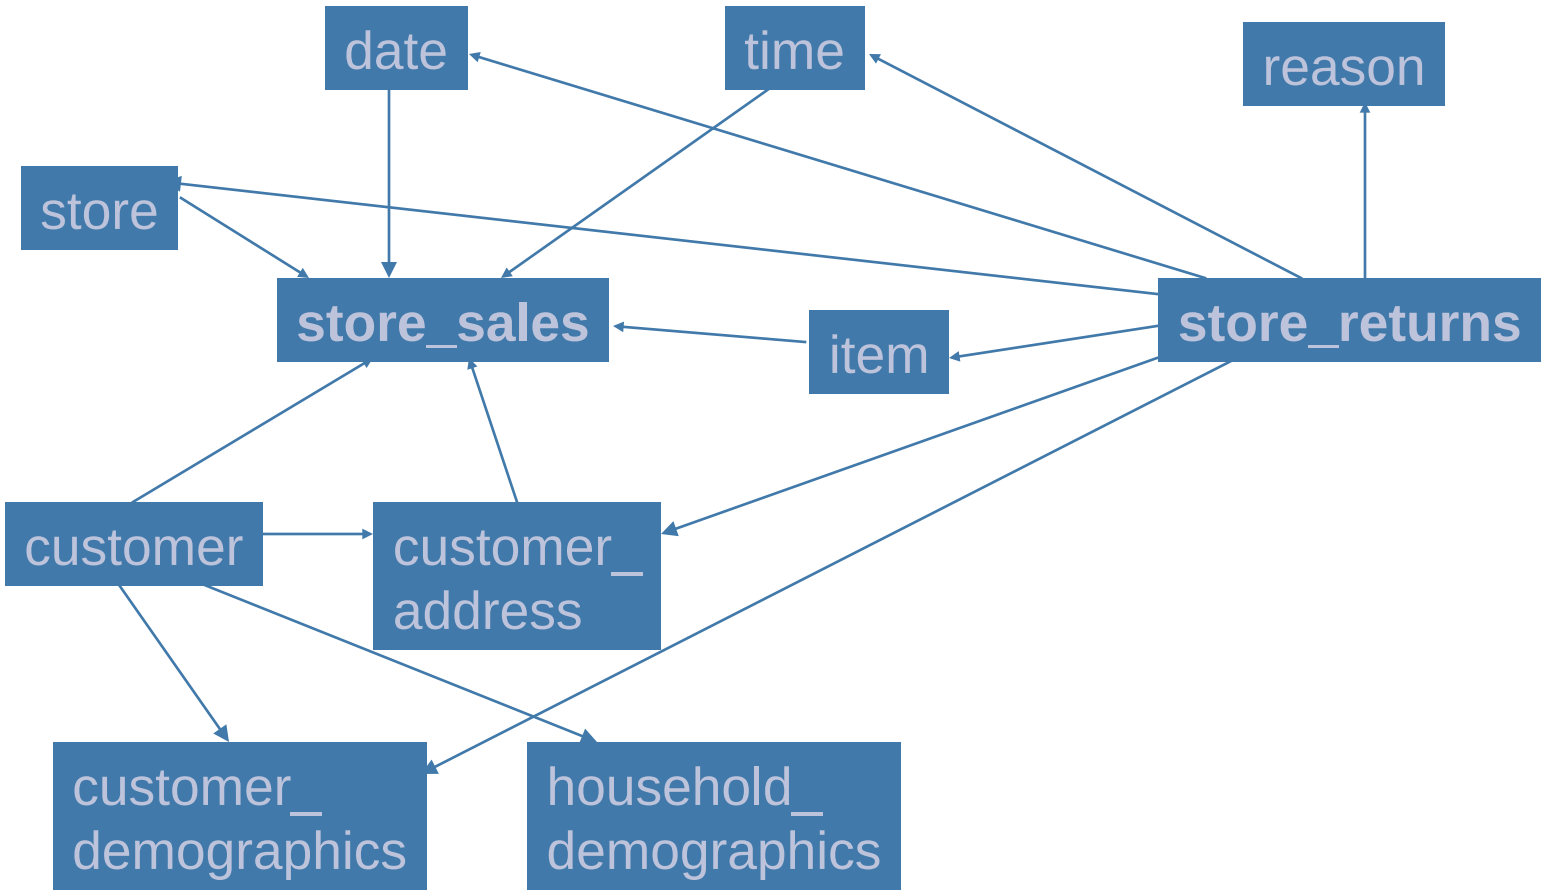

Data Set

- Most data has “real world” content:
  - Names, address, etc.
- Data is skewed:
  - E.g. SS_SOLD_DATE:
    - Low sales: January – July
    - Medium sales: August – October
    - High sales: November/December
  - E.g. I_COLOR:
    - Three groups of low, medium and high likely colors

Realistic Data Generation

- Size is defined in scale-factors
- Scale factor indicates raw data size in GB
  - 100, 300, 1000, 3000, 10000, 30000, 100000
- Row counts for tables scale realistically
  - Fact tables grow linearly
  - Dimension tables grow sub-linearly

SS_SOLD_DATE Distribution

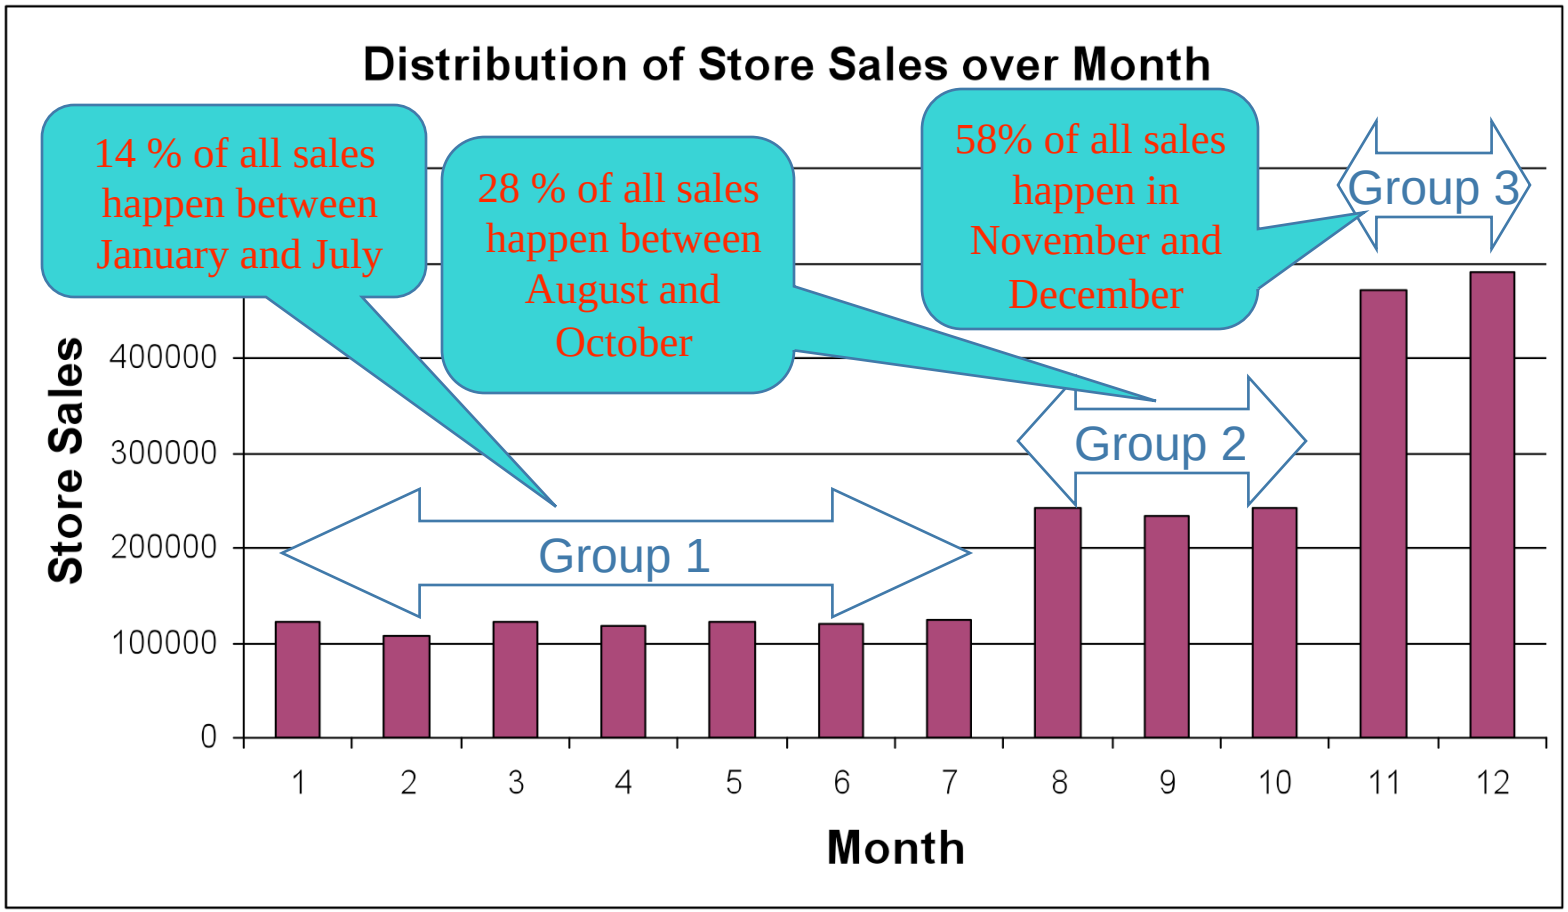

I_COLOR

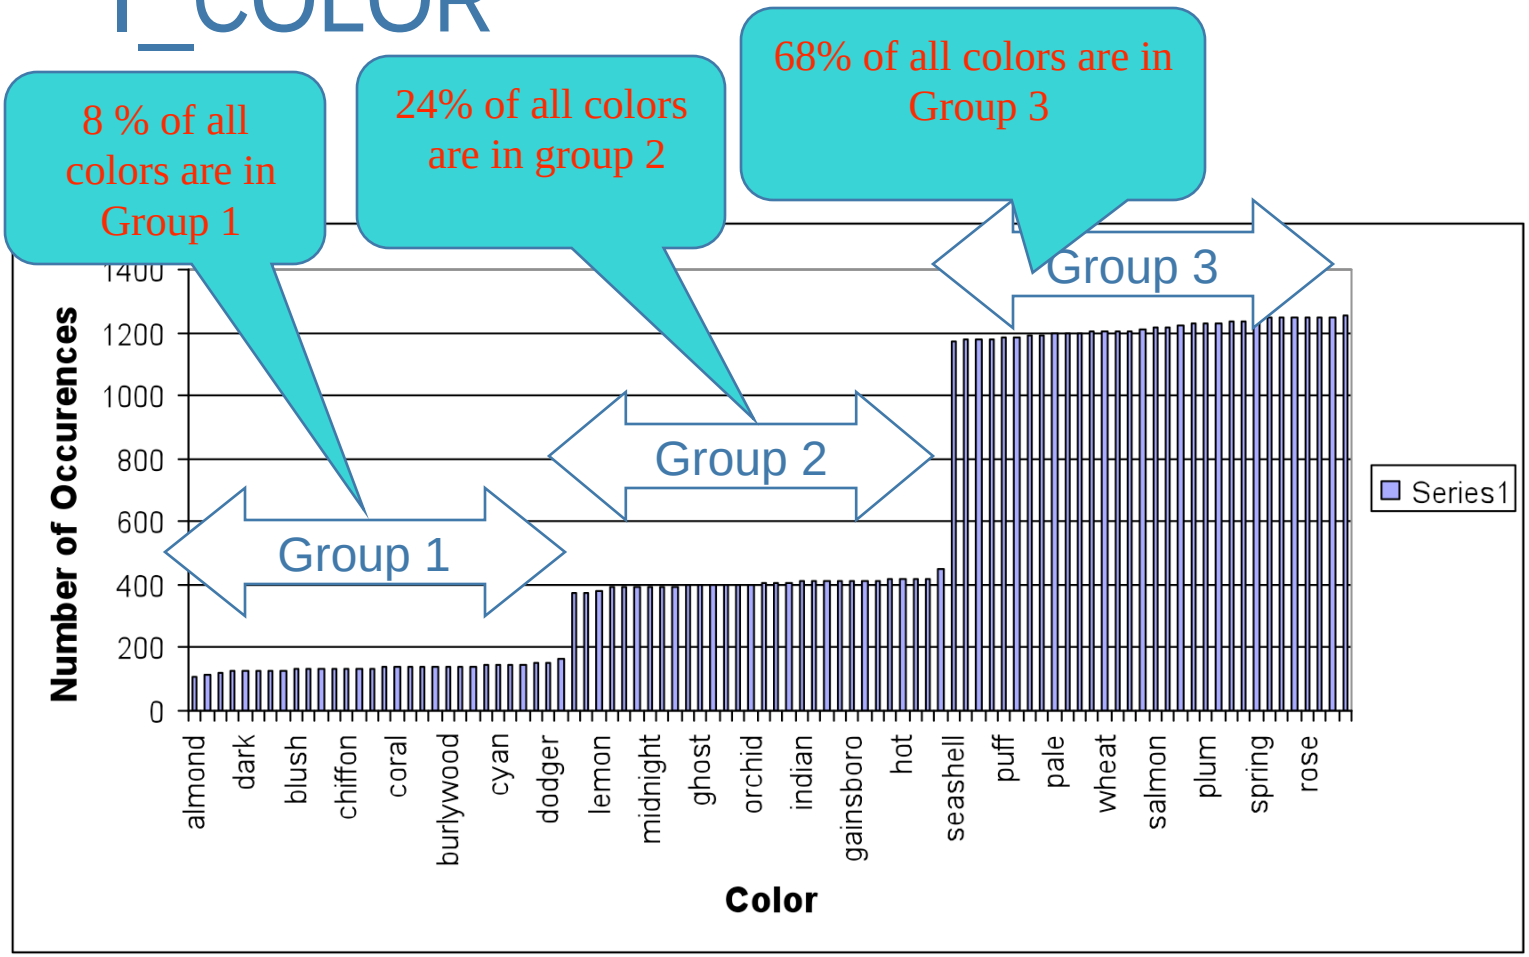
8 % of all colors are in Group 1
24% of all colors are in group 2
68% of all colors are in Group 3
Group 1
Group 2
Group 3
Number of Occurences
0
200
400
600
800
1000
1200
1400
almond
dark
blush
chiffon
coral
burlywood
cyan
dodger
lemon
midnight
ghost
orchid
indian
gainsboro
hot
seashell
puff
pale
wheat
salmon
plum
spring
rose
Color
Series1

TPC - DS

Database Scaling (fact tables)

- Scale Factor correspond to the nominal database size in GB
- I.e. Scale Factor 100 is approx. 100GB, not including indexes, materialized views and temporary tables
- Indexes, materialized views and temporary tables significantly increase the total disk capacity (3x to 5x)

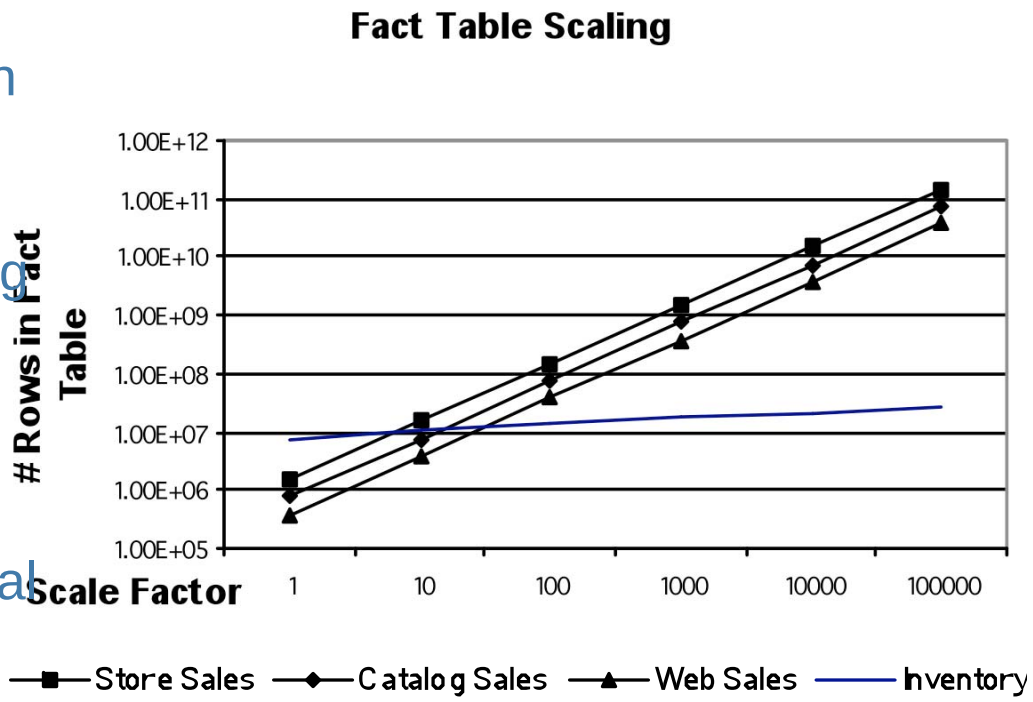

- fact table amount for the majority of the data (99%)

TPC Transaction Processing Performance Council

Database Scaling (Dimensions)

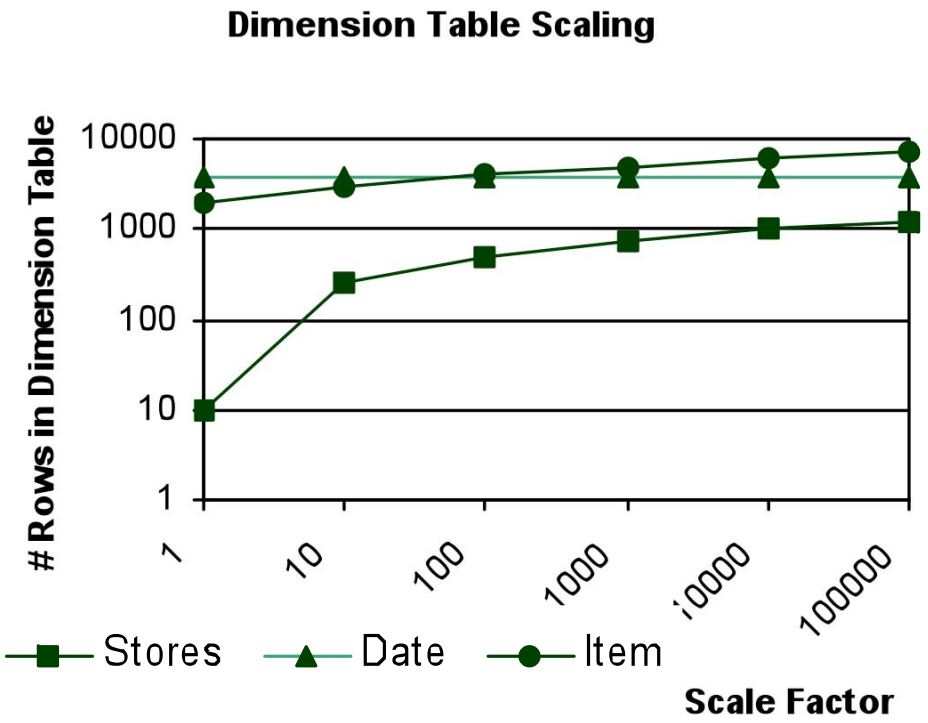

- scale sub-linearly
- amount for a fraction of the fact tables

Query Model

- Query Language: SQL99 + OLAP extensions
- Variants are allowed, but need to be approved in advance
- Four Classes of queries
  - Reporting
  - Ad-doc (scenario) OLAP
  - Ad-hoc (individual) DSS
  - Data extraction
- 99 different queries per user

Data Maintenance Model

(Extraction) - Transformation - Load

Normalized Data

Cleansing
Transformation
Slowly Changing
Dimensions

Snow Flake Schema

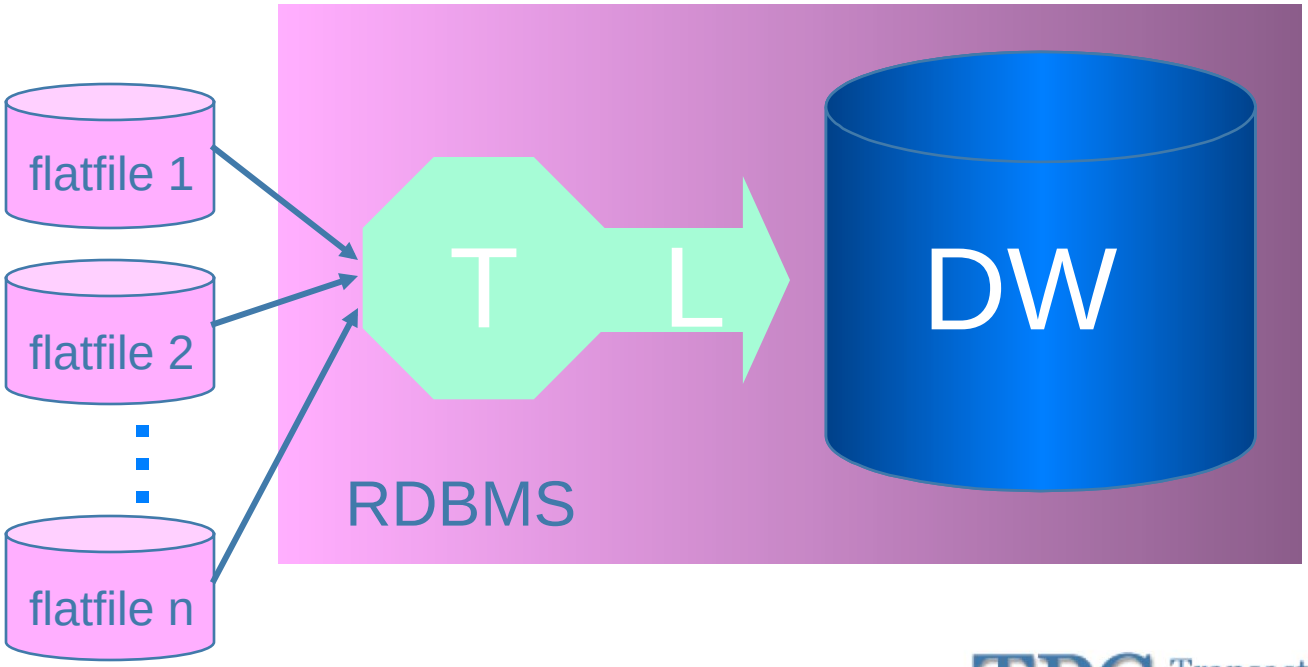

Data Maintenance Model

- Server-centric refresh model (E)TL
- No measurement of an extraction process
- Transformations and loads from a normalized dataset (flat-files) into a star schema
- Transformations can be done in any SQL or procedural language
- Data Maintenance Model includes slowly changing dimensions
- Updates all non-static tables

Benchmark Execution

- 5 Phases (timed and un-timed)
- Sequence is:

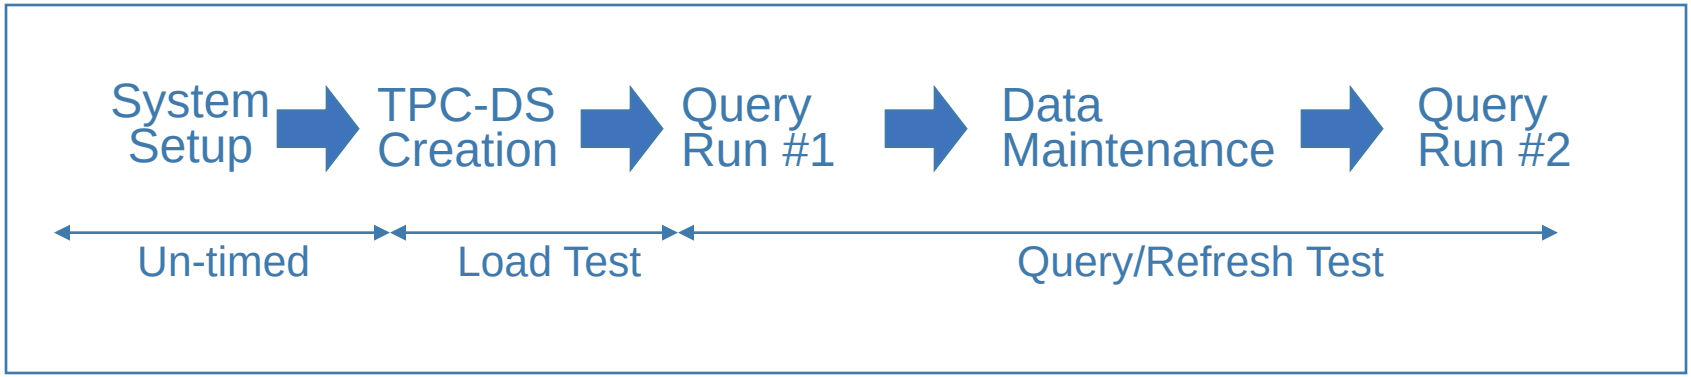

- Data Maintenance can run concurrently with Queries (Trickle)

System Setup

- This phase prepares the system for a performance measurement run (un-timed)
  - Flat-file creation
  - Hardware setup
    - Server, storage, network
  - DBMS setup – no data load
  - Disk setup

TPC-DS creation

- This phase prepares the database for a performance run
  - Load raw data into base tables
  - Create auxiliary data structures
  - Analyze data

Query Run #1

This test measures the system's performance when executing queries in a multi-user fashion

Stream 1: Q_{31},Q_{21},Q_1,Q_3,Q_8, … , Q_{47},Q_{123}

Stream 2: Q_1,Q_{25},Q_4,Q_4,Q_9, … , Q_{12},Q_3

…

Stream n (Query permutation for stream n)

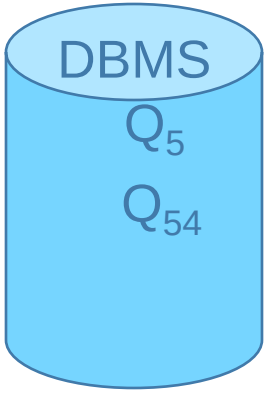

Query Run #1

- Each stream runs all 99 queries
  - In different order
  - With different substitution parameters
- Only one query in each stream is executing at any given time

  → n queries run concurrently

Data Maintenance Run (ETL)

- The data maintenance run represents the integration and consolidation of data from source systems and the removal of obsolete data from a data warehouse
- It performs the following tasks:
  - Maintaining dimensional data
    - Loading changed and new dimensional data
    - Transforming dimensional data
    - Updating/inserting of dimension data
  - Maintaining new fact data
    - Loading new fact data
    - Transforming fact data
    - Inserting fact data
    - Purging fact data

Data Maintenance Run

- Serves to assure that auxiliary data structures are properly maintained
- Reads data from flat files
- Must be implemented in the database (no external tools)

TPC - DS

Primary Metrics

1. Performance Metric: Queries per Hour

$$QphDS@SF = \frac{99 * 2 * S * 3600 * SF}{\left(T_{TT1} + T_{TT2} + 0.2 * T_{Load}\right)}$$

   – Where:
     - SF=scale factor
     - T_{TT1} = Elapsed time of the first throughput run
     - T_{TT2} = Elapsed time of the second throughput run
     - T_{Load} = Elapsed time of the load test

3. Price Performance: Price per QphDS

$$\$/QphDS@SF = \frac{P}{QphDS@SF}$$

   – P is the price of the Priced Configuration

5. System Availability Date

TPC Transaction Processing Performance Council

Comparison to TPC-H: Schema

Characteristic	TPC-H	TPC-DS
Business model	Retail	Retail
Normalization	3rd normal	Star schema
Number of tables	8	26
Number of columns (largest table)	15	38

Comparison to TPC-H: Data Set

Characteristic	TPC-H	TPC-DS
Data distribution	Uniform	Skewed
Data generation	Synthetic	Pseudo realistic
Row count scaling	Linear	Linear/sub-linear

Comparison to TPC-H: Queries

Characteristic	TPC-H	TPC-DS
SQL Dialect	SQL92	SQL99+OLAP
# of queries	22	150+
Query types	Ad-hoc*	Mix*
Substitution parameters	Yes	Yes
Query repeatability	Yes	Yes

* Achieved by limiting database technology

** Reporting, Ad-Hoc, iterative OLAP and extraction

Comparison to TPC-H: Data Maintenance

Characteristic	TPC-H	TPC-DS
Number tables updated	2 out of 8	20 out of 26
Insert	Yes	Yes
Delete	Yes	Yes
Update	No	Yes
Random inserts/deletes?	Yes	No

Comparison to TPC-H: Execution

Characteristic	TPC-H	TPC-DS
Multi user test	Yes	Yes
Single user test	Yes	Not decided
Measurement	Queries/hour	Queries/hour

Comparison to TPC-H: Miscellaneous

Characteristic	TPC-H	TPC-DS
Restrictions on Auxiliary Structures	Yes	No
ACID	Yes	Yes