

Scalable and Dynamic Workload Dependent Data Regeneration

Anupam Sanghi Raghav Sood Jayant Haritsa Srikanta Tirthapura¹

Technical Report
TR-2017-01

Database Systems Lab
Dept. of Computational and Data Sciences
Indian Institute of Science
Bangalore 560012, India

<http://dsl.cds.iisc.ac.in>

¹Iowa State University

Abstract

A core requirement of database engine testing is the ability to create synthetic versions of the customer’s data warehouse at the vendor site. A rich body of work exists on synthetic database regeneration, but suffers critical limitations with regard to: (a) scaling to large data volumes, (b) handling complex query workloads, and (c) producing data on demand. In this paper, we present **HYDRA**, a workload-dependent dynamic data regenerator that leverages a declarative approach to data regeneration to assure volumetric similarity, a crucial aspect of statistical fidelity, and materially improves on the prior art by adding scale, dynamism and functionality. Specifically, Hydra uses an optimized linear programming (LP) formulation based on a novel *region-partitioning* approach. This spatial strategy drastically reduces the LP complexity, enabling it to handle query workloads on which contemporary techniques fail. Second, Hydra incorporates deterministic post-LP processing algorithms that provide high efficiency and improved accuracy. Third, Hydra introduces the concept of *dynamic regeneration* by constructing a minuscule *database summary* that can on-the-fly regenerate databases of arbitrary size during query execution, while obeying volumetric specifications derived from the query workload. A detailed experimental evaluation on standard OLAP benchmarks demonstrates that Hydra can efficiently and dynamically regenerate large warehouses that accurately mimic the desired statistical characteristics.

1 Introduction

In industrial practice, a common requirement for database vendors is to adequately test their database engines with representative data and workloads that accurately mimic the data processing environments at customer deployments. This need can arise either in the analysis of problems currently being faced by clients, or in proactively assessing the performance impacts of planned engine upgrades on client applications. While, in principle, clients could transfer their original data and workloads to the vendor for the intended evaluation purposes, this is often infeasible due to privacy and liability concerns. Moreover, even if a client is willing to share the data, transferring and storing the data at the vendor’s site may prove to have impractical space and time overheads, especially in the anticipated Big Data era. For instance, if a customer faces a problem on exabyte (10^{18}) sized relational tables, transferring and storing such data is likely to be infeasible even on the best of systems. Therefore, an important requirement, looking into the future, is to be able to *dynamically* regenerate representative databases, at query execution time that accurately mimic the behavior of the client’s data processing environment.

A rich body of literature exists on data regeneration, beginning with *workload-independent* techniques (e.g [17, 15]), which provide scalable and efficient solutions, but fail to retain complex statistical characteristics such as the sizes of intermediate relations created during execution of a query plan. To address this problem, a particularly potent approach of *workload-dependent* database regeneration was introduced in QAGen [14], and has served as the foundation for many of the practicable systems proposed over the last decade [20, 9]. Workload-dependent techniques aim to generate synthetic data whose behavior is *volumetrically similar* to the client database on the pre-specified query workload. That is, assuming a common choice of query execution plans at the client and vendor sites (ensured through “plan forcing” [7] or “metadata matching” [11]), the output row cardinalities of individual operators in these plans are very similar in the original and synthetic databases. This similarity helps to preserve the multi-dimensional layout and flow of the data, a pre-requisite for achieving similar performance on the client’s workload. As a case

in point, the DataSynth [9, 10] tool from Microsoft expresses such volumetric constraints as a Linear Program (LP) whose solution is used to construct the synthetic database.

A common limitation of contemporary techniques (reviewed in detail in Section 8), is that they run into issues of *scale* and *efficiency* at one stage or the other in the regeneration pipeline. This is partly due to their focus on *materialized* static solutions, making them impractical at large volumes. Further, the ability to scale to large query workloads and data volumes has not been clearly established, and validations have been typically restricted to relatively simple and small benchmarks such as TPC-H [6]. These limitations become especially problematic from a futuristic “Big Data” perspective, where we have to contend with enormous data volumes and complex query workloads.

To materially address this challenge, we present **HYDRA**, a data regeneration tool, which ensures that scale and efficiency are addressed through *the entire regeneration pipeline*. As a concrete example, Hydra was able to accurately regenerate the data processing environment of a 100 GB TPC-DS client database with a workload of 131 distinct representative queries, by generating a *database summary* in less than 2 minutes on a vanilla machine. This summary can be used to statically generate a materialized database, or more potently, to *dynamically* regenerate the desired database during query execution. When the former option is chosen, the static database was successfully created in less than 11 minutes. It is important to note here that the summary construction time is *independent* of the data scale – therefore, even the exabyte-sized data scenario alluded to earlier could be modeled in just a few minutes using Hydra!

The key contributions of Hydra are the following:

Extended Workload Coverage: Hydra incorporates a novel LP formulation technique, *region-partitioning*, that can encode volumetric constraints with an LP of low complexity. When compared with the *grid-partitioning* approach used in DataSynth, region-partitioning reduces the LP complexity by many orders of magnitude. For instance, an LP with more than a *billion* variables in DataSynth is reduced to an LP with a few *thousand* variables in Hydra– in fact, in this case, the LP solver crashes on the DataSynth formulation, but runs to completion in less than a minute on the Hydra formulation. The beneficial outcome of the low LP complexity is that it facilitates the efficient handling of much richer query workloads.

Apart from enhancing the workload scale, Hydra also expands the database scope to include relational schemas that have DAG-structured dependency graphs, and the query scope to include DNF filter predicates as opposed to only simple range predicates.

Database Summary and Dynamic Regeneration: A unique feature of our data regeneration approach is that it delivers a *database summary* as the output, rather than the static data itself. This summary is of negligible size, depends only on the query workload and *not* on the database scale. It can be used either for *dynamically* generating data during query execution, or optionally for materializing static relations when parameters related to I/O processing need to be monitored. This summary-based approach helps eliminate the enormous time and space overheads incurred by prior techniques in generating and storing data before initiating analysis.

Accuracy with Efficiency: Hydra replaces the *sampling-based* approach to data regeneration in DataSynth by a *deterministic alignment* strategy, as discussed in Section 5. The alignment operates directly on the database summary, and is therefore extremely efficient. Fur-

ther, it does not suffer the probabilistic errors that affect the sampling approach, and therefore delivers better fidelity with regard to volumetric similarity.

Enhanced Evaluation: We evaluate Hydra on a diverse workload of 100-plus queries constructed from the complex TPC-DS benchmark, and the results show that it can efficiently regenerate databases for such workloads at various data scales. Further, our evaluation is more comprehensive than prior techniques, which have largely been evaluated on simpler and small-sized query workloads operating on modest databases. For instance, DataSynth has been evaluated on simple TPC-H database environments that resulted, with their formulation, in LPs with only a few thousand variables.

Integration with CODD: CODD [11] is a graphical tool through which database environments with desired meta-data characteristics can be efficiently simulated without persistently generating and/or storing their contents – i.e. a “dataless” approach. We have integrated Hydra with CODD, thus providing an end-to-end system that fully replicates the client data processing environment at the vendor’s site, and is compliant with the CODD’s “dataless” philosophy.

Organization. The remainder of this paper is organized as follows: A brief background on the key underlying concepts is outlined in Section 2. The Hydra architecture is presented in Section 3, and our new region-based LP formulation in Section 4. The database summary generator and the tuple generator are described in Sections 5 and 6, respectively. Our experimental results are analyzed in Section 7. Related work is reviewed in Section 8, and our conclusions are summarized in Section 9.

2 Preliminaries

In this section, we provide background information on the key foundations – Annotated Query Plans [14] and Cardinality Constraints [9] – that lie under this data regeneration framework.

2.1 Annotated Query Plans

Consider a toy scenario (for ease of presentation) where the client has the database schema shown in Figure 1a, where **pk** and **fk** refer to primary-key and foreign-key attributes, respectively.

A sample client query on this schema is shown in Figure 1b, with the corresponding query execution plan in Figure 1c. Note that this execution plan has the output edge of each operator annotated with the associated row cardinality (as evaluated during the client’s execution) – for instance, there are 50000 rows resulting from the join of **R** and (filtered) **S**. Such a plan is referred to as an “Annotated Query Plan” (AQP) in [14]. The goal now is to generate synthetic data at the vendor site such that when the above query is executed on this data, we obtain an identical, or very similar, AQP.

2.2 Cardinality Constraints

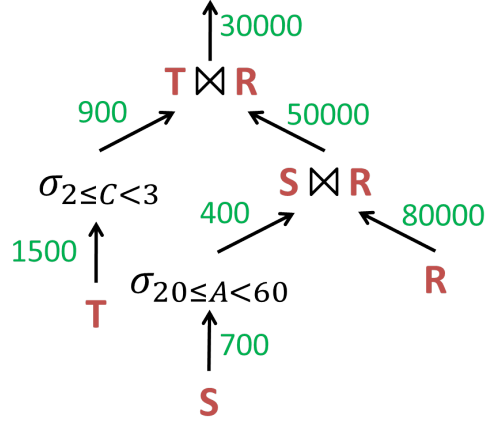
A unified and declarative mechanism for representing AQP data characteristics, called *cardinality constraints* (CCs), was proposed in [9]. For instance, the CCs expressing the AQP of Figure 1c are shown in Figure 1d. The data regeneration technique takes the schematic information and the

$R (\underline{R_pk}, S_fk, T_fk)$	$S (\underline{S_pk}, A, B)$	$T (\underline{T_pk}, C)$
---------------------------------------	-------------------------------	----------------------------

(a) Database Schema

select * from R, S, T where $R.S_fk = S.S_pk$ and $R.T_fk = T.T_pk$ and $S.A \geq 20$ and $S.A < 60$ and $T.C \geq 2$ and $T.C < 3$

(b) Example Query



(c) Annotated Query Plan (AQP)

$ R = 80000$	$ S = 700$	$ T = 1500$
$ \sigma_{S.A \in [20,60)}(S) = 400$	$ \sigma_{T.C \in [2,3)}(T) = 900$	
$ \sigma_{S.A \in [20,60)}(R \bowtie S) = 50000$		
$ \sigma_{S.A \in [20,60) \wedge T.C \in [2,3)}(R \bowtie S \bowtie T) = 30000$		

(d) Cardinality Constraints (CCs)

Figure 1: **Example Database Scenario**

set of CCs from the client site and produces synthetic data that closely meets these CCs. To make the problem tractable, it is assumed that CCs consist of filters on only *non-key* attributes, and that all joins are between primary keys and foreign keys, typically the case in data warehouses.

3 The HYDRA Architecture

In this section, we present an overview of Hydra’s architecture, along with a summary of its various components and their interactions with the database engine. A pictorial view of the architecture is presented in Figure 2 – in this picture, the green boxes represent the new components designed specifically for Hydra. Among these, the primary components are the LP Formulator, the Summary Generator, and the Tuple Generator, all shown with thick borders. The other modules have been sourced from the literature, including the preprocessor (orange) from DataSynth [10], the CODD metadata processor (yellow) [11], and the Z3 solver (blue) [21].

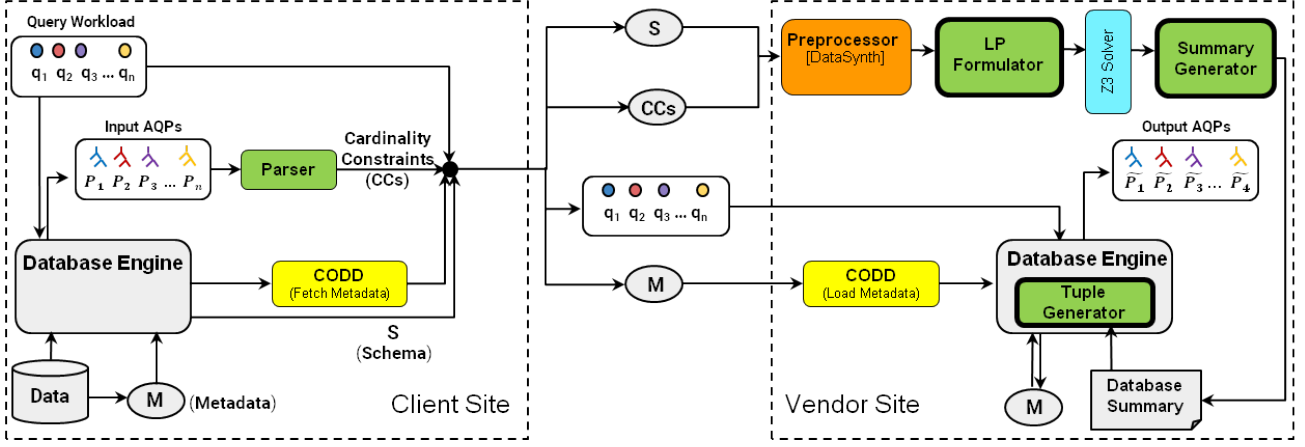


Figure 2: **Hydra Architecture**

3.1 Client Site

The information flow from the client to the vendor is as follows: At the client site, Hydra fetches the schema information (S), and the query workload ($q_1, q_2, q_3, \dots, q_n$) with its corresponding AQPs ($p_1, p_2, p_3, \dots, p_n$) obtained from the database engine. The AQPs are converted to equivalent cardinality constraints (CCs) using a Parser. The metadata (M) from the database catalogs is captured with the help of CODD. All this information is then shipped to the vendor site. If required, privacy concerns can be addressed by passing the information through an appropriate *anonymization* layer at the client.

3.2 Vendor Site

The main modules at the vendor site are as follows:

Preprocessor [10]: In this module, sourced from DataSynth, the schema information and CCs obtained from the client are processed to create the input for the LP Formulator. Each relation is solved independently, and this process is initiated by first creating a *view* comprised of its own non-key attributes, augmented with the non-key attributes of the relations on which it depends through referential constraints (both directly or transitively). This transformation results in replacing the join-expression present in a CC with a view that covers all the attributes (non-key) featured in the relations participating in the join-expression. As a case in point, following views are generated for the example in Figure 1:

$$R_view(A, B, C) \quad S_view(A, B) \quad T_view(C)$$

Further, the last two constraints in Figure 1d can be rewritten as:

$$\begin{aligned} |\sigma_{A \in [20,60]}(R_view)| &= 50000 \\ |\sigma_{A \in [20,60] \wedge C \in [2,3]}(R_view)| &= 30000 \end{aligned}$$

An LP is independently formulated for each view created by the above process. Since the LP complexity is adversely affected by the number of attributes in the view, the view is first decomposed into a set of *sub-views* to reduce the effective complexity. This is achieved as follows: Construct a “view-graph” by first creating a node for each attribute, and then inserting an edge between a pair of nodes if the corresponding attributes appear together in one or more CCs .

Further, additional edges are added (if required) to make the view-graph to be *chordal*, a property required to ensure acyclicity in the subsequent processing. Now, the sub-views are identified as the *maximal cliques* in the view-graph. For instance, the sub-views of the view-graph in Figure 3a are AEC, BED and CED, as shown in Figure 3b.




Figure 3: View Decomposition

LP Formulator and Solver: For each view, the LP Formulator takes as input the corresponding set of subviews and applicable CCs, and then constructs the LP. The domain corresponding to each sub-view is partitioned into regions using a novel *region-partitioning* algorithm that takes as input the different cardinality constraints. There is one variable for each region, corresponding to the number of tuples chosen from the region. Each cardinality constraint is encoded as an LP constraint on these variables, and the solution of the LP is used in deciding which tuples to include in the sub-view. The complete details of this algorithm are enumerated in Section 4.

Our region-partitioning strategy is in marked contrast to the *grid-partitioning* strategy used in DataSynth. Grid-Partitioning first intervalizes the domain of each attribute based on the constants appearing in the CCs, and divides the domain into a grid aligned with the interval boundaries for each attribute. If a sub-view has n attributes, and each attribute gets divided into ℓ intervals, then the domain of the sub-view is partitioned into a grid of ℓ^n cells. For each cell in the grid, a variable is created that represents the number of data rows present in that cell. In contrast, our region-partitioning strategy divides the domain into only the number of regions required to precisely write out each cardinality constraint, and assigns one variable to each region – this typically leads to far fewer variables than grid-partitioning.

To make the above concrete, consider a single view “Person” with the following three selection CCs:

$$\begin{aligned} |\text{age} < 40 \wedge \text{salary} < 40\text{K} (\text{Person})| &= 1000 \\ |20 \leq \text{age} < 60 \wedge 20\text{K} \leq \text{salary} < 60\text{K} (\text{Person})| &= 2000 \\ |\text{Person}| &= 8000 \end{aligned}$$

Grid-partitioning divides the domain of the view as shown in Figure 4a. With a variable assigned to each grid cell, there is a total of 16 variables. In contrast, the region-partitioning strategy partitions the space into 4 regions as shown in Figure 4b, resulting in a tally of only 4 variables.

The CCs of Person, expressed in terms of LP constraints, are shown below in Figure 5a and 5b for grid-partitioning and region-partitioning, respectively.

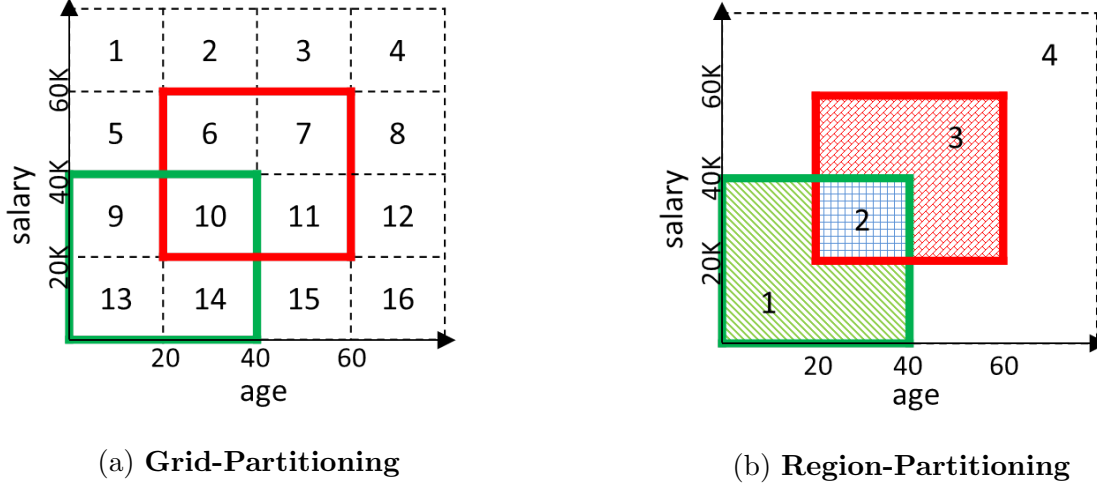


Figure 4: Grid-Partitioning vs Region-Partitioning

$$\begin{aligned}
 x_9 + x_{10} + x_{13} + x_{14} &= 1000 \\
 x_6 + x_7 + x_{10} + x_{11} &= 2000 \\
 x_1 + x_2 + \dots + x_{16} &= 8000
 \end{aligned}$$

(a) **Grid-Partitioning**

$$\begin{aligned}
 y_1 + y_2 &= 1000 \\
 y_2 + y_3 &= 2000 \\
 y_1 + y_2 + y_3 + y_4 &= 8000
 \end{aligned}$$

(b) **Region-Partitioning**

Figure 5: LP Constraints

The LPs are passed on to the solver, which provides one of the feasible solutions as the output – we have used Z3 [21], a popular SMT solver, to implement this functionality. With region-partitioning, the LP is usually much simpler due to the smaller number of variables. Further, as the cardinality constraints get more complex, the differences in complexity of the LPs produced by region-partitioning and grid-partitioning become more pronounced. This effect is quantified in Section 7.

Summary Generator: This module generates the database summary from the LP solutions obtained on the views. Since partitioning is carried out at a sub-view level, the LP solution, which is expressed in terms of sub-view variables, needs to be mapped to equivalents in the original view space. A sampling-based approach was proposed in [9] for this purpose – for example, say a view (A, B, C) is split into a pair of sub-views (A, B) and (B, C) , the algorithm computes the distributions $Prob(A, B)$ and $Prob(C|B)$. Then, each tuple is generated by first sampling a point from the former distribution, and then sampling a point from the latter conditioned on this outcome.

However, we have chosen not to take this approach since the computational overheads incurred are enormous, and the sampling process introduces errors in volumetric fidelity. Instead, we have designed and implemented an alternative data-scale free, deterministic alignment algorithm (details in Section 5), which produces an intermediate database summary in the output. This component is also responsible for ensuring that the generated summary obeys referential integrity. Finally, summarized relations from corresponding view summary are obtained. An example

database summary finally obtained from the AQP shown in Figure 1c, along with additional two AQPs, is shown in Figure 6. Here, entries of the type $a - b$ in the primary key columns (e.g. 101-250 for S_pk in table S), mean that the relation has $b - a + 1$ tuples with values $(a, a + 1, a + 2, \dots, b)$ for that column, keeping the other columns unchanged.

R			S			T	
R_pk	S_fk	T_fk	S_pk	A	B	T_pk	C
1 – 30K	321	1	1–100	0	15	1–600	0
30001 – 50K	621	601	101–250	20	15	601–1500	2
50001 – 60K	71	601	251–500	20	10		
60001 – 70K	121	1	501–700	0	5		
70001 – 80K	1	1					

Figure 6: **Example Database Summary**

Tuple Generator: The Tuple Generator resides in the database engine. It ensures that whenever a query is fired, data is not fetched from the disk but instead gets generated *on-demand*, using the database summary. The details of this component and its implementation in PostgreSQL are presented in Section 6.

We note in closing that in order to ensure the execution plan chosen at the vendor site is the same as that in the client site, metadata matching is implemented in Hydra using CODD’s metadata transfer feature.

4 LP Formulation

An LP for a view V is constructed as follows: For each sub-view s in V , every CC that is within its scope is formulated as an LP constraint. Since sub-views may share common attributes, additional *consistency constraints* are added to the LP to ensure that the marginal distributions along the common set of attributes are identical in the solutions for the sub-views.

In this section, we first present the mathematical basis underlying our formulation of LP constraints for a set of CCs applicable on a sub-view. We then present an algorithm that partitions the domain into the minimum number of regions required to capture each CC precisely, resulting in an LP with the optimal number of variables. Finally, we discuss the formulation of additional consistency constraints to ensure consistency across multiple sub-views belonging to V .

4.1 Mathematical Basis for LP Formulation

Let n denote the number of attributes in the given sub-view s , \mathcal{D}_i the domain of the i th attribute, and \mathcal{D} the data universe $\mathcal{D}_1 \times \mathcal{D}_2 \times \dots \times \mathcal{D}_n$.

We are given a set of m CCs that are applicable on s . For $1 \leq j \leq m$, each constraint C_j is a pair $\langle \sigma_j, k_j \rangle$ where σ_j is a selection predicate and k_j is a non-negative integer equal to the number of rows satisfying predicate σ_j . We assume that each predicate is in disjunctive normal form (DNF).

Simple LP Formulation Let us first consider a simple way of formulating an LP that encodes all CCs. For each tuple $t \in \mathcal{D}$, assign a variable x_t that denotes the number of copies of t in the sub-view s . Then, the LP formulation shown in Figure 7 ensures that a feasible solution satisfies all CCs, including a constraint on the total size of s .

The problem with this formulation is that the number of variables in the resulting LP is as large as the size of the universe \mathcal{D} . Hence, it is infeasible to work directly with this formulation.

$$\begin{aligned}
(1) & \text{ For each } t \in \mathcal{D}, x_t \geq 0 \\
(2) & \left[\sum_{t \in \mathcal{D}} x_t \right] = k \\
(3) & \text{ For each } j, 1 \leq j \leq m, \left[\sum_{t: \sigma_j(t)=\text{true}} x_t \right] = k_j
\end{aligned}$$

Figure 7: Simple LP formulation

Reduced LP Formulation We can derive an LP with far fewer variables as follows: We first note that in the simple formulation, variables corresponding to a pair of points $t_1, t_2 \in \mathcal{D}$ that behave identically with respect to a constraint C_j (i.e. $\sigma_j(t_1) = \sigma_j(t_2)$) can be combined together as $(x_{t_1} + x_{t_2})$, for the purposes of satisfying constraint C_j . If this is true that with respect to every constraint C_j for $j = 1 \dots m$, $\sigma_j(t_1) = \sigma_j(t_2)$, then there is no need to treat t_1 and t_2 separately – instead, they can be combined into a single region, and the variables x_{t_1} and x_{t_2} can be merged into a single variable $(x_{t_1} + x_{t_2})$ in every equation, leading to fewer variables in the LP. By repeating this variable merging process recursively until it is no further possible, we arrive at a vastly reduced LP.

We hasten to add that the above LP construction process based on merging variables is only for illustrating the concept – the actual algorithm employed in our system *directly* derives the regions, as described in Section 4.2.

For constraint C and $t \in \mathcal{D}$, let $C(t)$ be an indicator variable:

$$C(t) = \begin{cases} \text{true} & \text{if } t \text{ satisfies } C \\ \text{false} & \text{otherwise} \end{cases}$$

Definition For a pair of points $p, q \in \mathcal{D}$ and a set of constraints \mathbb{C} , we say $pR^{\mathbb{C}}q$ if for each $C \in \mathbb{C}$, $C(p) = C(q)$.

Observation 1 $R^{\mathbb{C}}$ is an equivalence relation on \mathcal{D} .

Proof It can be easily seen that $R^{\mathbb{C}}$ is reflexive and symmetric. For transitivity, suppose that for $p, q, r \in \mathcal{D}$, $pR^{\mathbb{C}}q$ and $qR^{\mathbb{C}}r$. Note that for each $C \in \mathbb{C}$, it must be true that $C(p) = C(q)$ and $C(q) = C(r)$. Therefore, it must be true that $C(p) = C(r)$ for each $C \in \mathbb{C}$, showing that the relation is transitive.

A partition of \mathcal{D} is a set of subsets of \mathcal{D} such that every element $x \in \mathcal{D}$ is in exactly one of these subsets. The individual sets in a partition are called *blocks*.

Definition A set of points b is said to be valid with respect to a set of constraints \mathbb{C} if for any two points $p, q \in b$, $pR^{\mathbb{C}}q$. Given a set of constraints \mathbb{C} , a partition \mathbb{P} of \mathcal{D} is said to be a *valid partition* if for each block $b \in \mathbb{P}$, b is valid with respect to \mathbb{C} .

In a valid partition of \mathcal{D} with respect to \mathbb{C} , any pair of points within the same block satisfy the same set of CCs. Once we obtain a valid partition \mathbb{P} , the LP can be re-formulated as shown in Figure 8. Instead of a variable for each point $t \in \mathcal{D}$, there is now a single variable x_b for each block $b \in \mathbb{P}$ representing the number of tuples of the sub-view that are contained in this block. Note that the tuples in a sub-view need not be unique, therefore x_b may include duplicates in its count.

$$\begin{aligned}
(1) & \text{ For each } b \in \mathbb{P}, x_b \geq 0 \\
(2) & \left[\sum_{b \in \mathbb{P}} x_b \right] = k \\
(3) & \text{ For each } j, 1 \leq j \leq m, \left[\sum_{b: \sigma_j(b)=\text{true}} x_b \right] = k_j
\end{aligned}$$

Figure 8: Reduced LP formulation

The total number of variables in the reduced LP shown in Figure 8 is equal to the number of blocks in the partition \mathbb{P} and is potentially much smaller than the number of variables in the original LP, shown in Figure 7. Since we desire an LP with the smallest number of variables, we look for a valid partition of \mathcal{D} with the minimum number of blocks. A valid partition with respect to \mathbb{C} is an *optimal partition* if it has the smallest number of blocks from among all valid partitions of \mathcal{D} with respect to \mathbb{C} .

Lemma 4.1 *The quotient set of \mathcal{D} by $R^{\mathbb{C}}$ is the (unique) optimal partition of \mathcal{D} with respect to \mathbb{C} .*

Proof Let \mathbb{P}_1 denote the quotient set² of \mathcal{D} by $R^{\mathbb{C}}$. By the definition of an equivalence relation, for any block $b \in \mathbb{P}_1$, all points in b are related to each other by $R^{\mathbb{C}}$, and hence \mathbb{P}_1 is a valid partition.

Suppose that \mathbb{P}_1 is not the unique optimal partition. Then, there must exist another valid partition \mathbb{P}_2 such that $\mathbb{P}_2 \neq \mathbb{P}_1$ and $|\mathbb{P}_2| \leq |\mathbb{P}_1|$. This implies that there exist two points $p, q \in \mathcal{D}$ such that p and q are in different blocks in \mathbb{P}_1 , but in the same block in \mathbb{P}_2 . Since p and q belong to different blocks in \mathbb{P}_1 , it must be true that p and q are not related by $R^{\mathbb{C}}$. But, in \mathbb{P}_2 points p and q belong to the same block, which implies that \mathbb{P}_2 cannot be a valid partition, a contradiction.

4.2 Deriving the Optimal Partition

We now present an algorithm to derive the optimal partition of \mathcal{D} with respect to \mathbb{C} . Each constraint $C \in \mathbb{C}$ is in DNF, and is expressed as the union of many smaller “sub-constraints”.

²The quotient set is the set of equivalence classes resulting from $R^{\mathbb{C}}$ on \mathcal{D} .

Each sub-constraint is the conjunction of many per-attribute constraints, and each per-attribute constraint is a constraint on the values that the attribute is permitted to take. For example, the following constraint on attributes A_1 and A_2 :

$$((A_1 \leq 20) \wedge (A_2 > 30)) \vee (A_1 > 50)$$

is divided into the basic sub-constraints:

$$(A_1 \leq 20) \wedge (A_2 > 30) \text{ and } (A_1 > 50)$$

Algorithm 1 (Optimal Partition) takes a set of DNF constraints as input, and returns a partition with the smallest number of regions with respect to this set. Internally, it invokes Algorithm 2 (Valid Partition) that takes a set of sub-constraints as input and returns a valid partition of the domain with respect to this set.

Algorithm 1: Optimal Partition(\mathcal{D}, \mathbb{C})

Input: Universe \mathcal{D} , set of DNF constraints \mathbb{C}

Output: An optimal partition \mathbb{P}^* of \mathcal{D} subject to \mathbb{C}

- 1 Generate the set of sub-constraints \mathbb{C}' resulting from the constraints in \mathbb{C} ;
 - 2 Construct a valid partition \mathbb{P}' of \mathcal{D} subject to \mathbb{C}' using Valid-Partition(\mathcal{D}, \mathbb{C}') (Algorithm 2);
 - 3 For each block $b \in \mathbb{P}'$, compute the label $\ell(b)$, equal to the set of all constraints in \mathbb{C} that b satisfies. Let L denote the set of all distinct labels from $\{\ell(b) | b \in \mathbb{P}'\}$;
 - 4 Coarsen partition \mathbb{P}' into \mathbb{P}^* as follows: For each label $l \in L$, merge all blocks in \mathbb{P}' whose labels equal l into a single block;
 - 5 Return \mathbb{P}^* ;
-

Lemma 4.2 *Given a set of DNF constraints \mathbb{C} , Algorithm 1 returns an optimal partition of \mathcal{D} with respect to \mathbb{C} .*

Proof As in the algorithm, let \mathbb{C}' denote the set of sub-constraints resulting from constraints in \mathbb{C} . From Lemma 4.3, we know that \mathbb{P}' is a valid partition with respect to \mathbb{C}' . Consider any block $b \in \mathbb{P}'$. Since b is valid with respect to \mathbb{C}' , and each constraint in \mathbb{C}' is stricter than a corresponding constraint in \mathbb{C} , b is valid with respect to \mathbb{C} . Hence, \mathbb{P}' is a valid partition with respect to \mathbb{C} .

Next, consider that each block b^* in \mathbb{P}^* was obtained by merging blocks in \mathbb{P}' that have the same label. For any pair of points p, q in b^* , it is true they satisfy the same set of constraints in \mathbb{C} , showing that \mathbb{P}^* is a valid partition wrt \mathbb{C} . Also, any two blocks in \mathbb{P}^* have distinct labels (if they had the same label, they would have been merged). Therefore, we conclude using arguments similar to Lemma 4.1 that \mathbb{P}^* is an optimal partition of \mathcal{D} with respect to \mathbb{C} .

Deriving a Valid Partition for a Set of Sub-Constraints: We now present an algorithm for deriving a valid partition with a small number of blocks, for a set of sub-constraints \mathbb{C} .

Definition For a sub-constraint C and dimension i , let C^i denote the restriction (projection) of C to dimension i . Further, let $C_1^i = \bigwedge_{k=1 \dots i} C^k$ denote the restriction of C to dimensions $1, 2, \dots, i$. For instance, if $C = (A_1 \geq 1) \wedge (A_2 \geq 4) \wedge (A_2 \leq 5) \wedge (A_3 > 6)$, then $C^2 = (A_2 \geq 4) \wedge (A_2 \leq 5)$, and $C_1^2 = (A_1 \geq 1) \wedge (A_2 \geq 4) \wedge (A_2 \leq 5)$. For convenience, if C does not have a constraint along dimension i , then C^i is defined to be “true”.

Our algorithm, described in Algorithm 2, proceeds iteratively, one dimension at a time. Before processing dimension i , it has a partition of \mathcal{D} that is a valid partition subject to constraints along dimensions 1 till $(i - 1)$. In processing dimension i , it refines the current partition as follows: For each block b in the current partition, it appropriately divides the block along dimension i if there is a constraint $C \in \mathbb{C}$ such that there are some points in b that satisfy constraint C^i , and some that do not.

Definition A constraint C is said to split a block $b \subseteq \mathcal{D}$ if there exist a pair of points $p_1, p_2 \in b$ such that $C(p_1) = \text{true}$ and $C(p_2) = \text{false}$. If C splits b , then refining b by C partitions b into two subsets $b^+(C) = \{x \in b | C(x) = \text{true}\}$ and $b^-(C) = \{x \in b | C(x) = \text{false}\}$.

Algorithm 2: Valid-Partition(\mathcal{D}, \mathbb{C})

Input: Universe \mathcal{D} , set of sub-constraints \mathbb{C}

Output: A valid partition \mathbb{P} of \mathcal{D} subject to set of sub-constraints \mathbb{C}

```

1  $\mathbb{P}^0 = \{\mathcal{D}\}$  // A partition with one set,  $\mathcal{D}$ .
2 for  $i$  from 1 to  $n$  do
3    $M \leftarrow \mathbb{P}^{i-1}$ ;
4   foreach  $C \in \mathbb{C}$  do
5      $M' \leftarrow \emptyset$ ;
6     foreach block  $b \in M$  do
7       if  $C^i$  splits  $b$  then
8         Let  $b^+$  and  $b^-$  result from refining  $b$  with  $C^i$ ;
9         Add  $b^+$  and  $b^-$  to  $M'$ ;
10      else
11        Add  $b$  to  $M'$ ;
12     $M \leftarrow M'$ ;
13  $\mathbb{P}^i \leftarrow M$ ;
14 Return  $\mathbb{P}^n$ ;

```

Lemma 4.3 *Given a set of sub-constraints \mathbb{C} , Algorithm 2 returns a valid partition of \mathcal{D} with respect to \mathbb{C} .*

Proof For $1 \leq i \leq n$, let $\mathbb{C}_1^i = \{C_1^i | C \in \mathbb{C}\}$. We show by induction on i that after the i th iteration of the outermost for loop in the algorithm, \mathbb{P}^i contains a valid partition of \mathcal{D} with respect to \mathbb{C}_1^i . Since $\mathbb{C}_1^n = \mathbb{C}$, it follows that after n iterations, \mathbb{P}^n contains a valid partition of \mathcal{D} with respect to \mathbb{C} . We consider $i = 0$ as the base case, and the set \mathbb{C}_1^0 as a set of “always true” constraints. Hence, \mathbb{P}^0 , which consists of only one element, \mathcal{D} , is a valid partition with respect to \mathbb{C}_1^0 .

For the inductive step, suppose that for $i > 0$, \mathbb{P}^{i-1} is a valid partition of \mathcal{D} with respect to \mathbb{C}_1^{i-1} . For each block $b \in \mathbb{P}^{i-1}$, two cases are possible: (1) b is not split by C^i , for any $C \in \mathbb{C}$. Then b is valid with respect to \mathbb{C}_1^i , and will be retained in \mathbb{P}^i . (2) b is split by one more constraints C^i . The algorithm iterates through all such constraints that split b , and partitions block b such that every resulting block is valid with respect to each C^i , $C \in \mathbb{C}$.

We next note that \mathbb{P}^i is indeed a partition of \mathcal{D} (i.e. the union of all blocks equals \mathcal{D}). To see this observe that each block $b \in \mathbb{P}^{i-1}$ is either present in \mathbb{P}^i or has been refined and all its constituent blocks (whose union equals b) are in \mathbb{P}^i . Thus, \mathbb{P}^i is a valid partition with respect to \mathbb{C}_1^i . This proves the inductive step.

Consistency Constraints Since different sub-views can have common attribute(s), additional constraints need to be added to ensure that their distributions for the common attribute(s) are the same. In order to do so, we may need to further refine the partition generated from the above procedure. Specifically, consider a pair of sub-views s_1 and s_2 with attribute sets \mathbb{A}_1 and \mathbb{A}_2 respectively, such that $\mathbb{A}_1 \cap \mathbb{A}_2 \neq \emptyset$. Let $\mathcal{D}^1 = \prod_{i \in \mathbb{A}_1} \mathcal{D}_i$, and $\mathcal{D}^2 = \prod_{j \in \mathbb{A}_2} \mathcal{D}_j$ be the corresponding domains for s_1 and s_2 respectively, and $\mathcal{D}^{1,2} = \prod_{i \in \mathbb{A}_1 \cap \mathbb{A}_2} \mathcal{D}_i$. Let the partitions obtained on \mathcal{D}^1 and \mathcal{D}^2 be \mathbb{P}_1 and \mathbb{P}_2 , respectively. In order to keep \mathbb{P}_1 and \mathbb{P}_2 consistent with each other, we need to ensure that their region boundaries are aligned with each other, and this is achieved by refining \mathbb{P}_1 and \mathbb{P}_2 so that they have common boundaries along dimensions $\mathbb{A}_1 \cap \mathbb{A}_2$. We consider the union of the “split points” of \mathbb{P}_1 and \mathbb{P}_2 along dimensions $\mathbb{A}_1 \cap \mathbb{A}_2$ and further for each block in \mathbb{P}_1 (and \mathbb{P}_2), we refine this block until it no longer crosses such a split point. Finally, we add LP constraints that equate distributions of the common attributes in \mathbb{P}_1 and \mathbb{P}_2 .

4.3 Implementation Details

Data Representation To represent the different split points within a specific dimension, we use an array of integers, called as *Bucket*. A list of these Buckets, where a Bucket is assigned to each attribute, is called as *BucketStructure*. Further, a set of these BucketStructures constitutes a *Region*. The set of points in \mathcal{D} that satisfy a CC is represented by a Region. Also, each block that is present in the outcome of the region-partitioning algorithm is represented by a Region. Finally, a *Partition* is a set of such Region objects.

Note that a CC or block in the partition need not be a contiguous hyper-rectangular region. It is simply a collection of points in the \mathcal{D} , which can be stored in a Region object. This modeling choice helps to handle DNF constraints easily.

5 Database Summary Generator

This component takes the LP solution for each view as the input and generates the database summary, which as mentioned previously, can be used for dynamically generating data for query execution, or can optionally be used to generate the materialized database.

Recall that a variable in the LP (for a view) represents an underlying block in a sub-view’s partition, and its assigned value is the number of rows present in that block – this value is hereafter referred to generically as NUMTUPLES. The collection of NUMTUPLES values represent the sub-view solutions, and these solutions are integrated to obtain the solution for the complete view. However, since each view is solved independently, the referential constraints that exist between the corresponding relations may be lost in these view solutions. Therefore, they may have to be modified to ensure global consistency. Finally, it is necessary to extract relations from the views in order to populate the database. Accordingly, the summary generator component in Hydra is responsible for the following sequence of tasks:

1. Constructing a solution for complete views

2. Instantiating view summaries
3. Making view summaries consistent wrt each other
4. Extracting relation summaries from view summaries

5.1 Constructing Solution for the View

For integrating the sub-view solutions to obtain the collective solution for the complete view, we first *order* the sub-views. Then, we iteratively build the view-solution by *aligning* and *merging* the next sub-view solution in the given order. Let \mathbb{S} denote the input list of sub-view solutions, and *viewSol* be the final view solution that we wish to compute. Algorithm 3 describes the high-level process for constructing *viewSol* from \mathbb{S} , and its ordering, aligning and merging procedures are described in the remainder of this sub-section.

Algorithm 3: View Solution Construction

```

1  $\mathbb{S} \leftarrow \text{ORDERSUBVIEWS}(\mathbb{S});$ 
2  $\text{viewSol} \leftarrow \emptyset;$ 
3 foreach  $s \in \mathbb{S}$  do
4    $\text{viewSol}, s \leftarrow \text{ALIGN}(\text{viewSol}, s);$ 
5    $\text{viewSol} \leftarrow \text{MERGE}(\text{viewSol}, s);$ 
```

5.1.1 Sub-View Ordering

Ordering is implemented through a greedy iterative algorithm where we can start with any sub-view as the first choice. Subsequently, at iteration i , let the set of visited sub-views until now be \mathbb{S} . A sub-view s from outside this set can be chosen to be the next in the ordering only if it satisfies the following condition: On removing the common vertices between s and \mathbb{S} in the (chordal) view-graph, there should not exist any path between the remaining vertices of s and the remaining vertices of \mathbb{S} .

For instance, in the example of Figure 3, there are three sub-views: AEC, ECD and BED. Here, if we start with sub-view AEC as the first choice in the ordering sequence, then BED cannot be the next selection – this is because even after removing the common attribute E between BED and AEC, the edge DC continues to connect the two components. The sub-view ECD, on the other hand, satisfies the required condition, and is chosen to be the second in the order. After ECD is chosen, then BED follows as the last sub-view in the sequence.

5.1.2 Aligning

After obtaining the sub-view merge order as per above, in every iteration we merge the next sub-view solution (s) in the sequence to the current view-solution (*viewSol*), after a process of alignment. The alignment algorithm is a two step exercise, as shown in the example of Figure 9:

Solution Sorting: First, the *viewSol* and s solutions are each sorted on their common set of attributes to facilitate direct comparison of their matching ranges. For instance, the solutions A, B and A, C in Figure 9a are each sorted on the intervals enumerated in the common attribute A .

Row Splitting: Our addition of consistency constraints during the LP formulation ensured that the distribution of tuples along the common set of attributes is the same in the various sub-views. Therefore it easy to see that the sum of NUMTUPLES values in any interval of the common attributes is the same for the sub-view solutions under alignment. For example, in Figure 9a, the total number of tuples with $A = [40, 60)$ is 30K in both the A, B and A, C solutions. Likewise, the other entries in column A also have matching total number of tuples across the solutions.

The align step *splits* the rows in these solutions such that the corresponding rows in both solutions have the same number of tuples. The sub-view solutions of Figure 9a are shown in Figure 9b after undergoing the alignment process, with both solutions now having identical NUMTUPLES in the corresponding rows.

A	B	NUMTUPLES	A	C	NUMTUPLES
[60, inf)	[0, inf)	30K	[60, inf)	[0, inf)	30K
[40, 60)	[0, 5) U [15, inf)	20K	[40, 60)	[2, 3)	30K
[40, 60)	[5, 15)	10K	[20, 40)	[0, 2) U [3, inf)	20K
[20, 40)	[5, 10)	10K			
[20, 40)	[15, inf)	10K			

(a) **Sub-view Solution**

A	B	NUMTUPLES	A	C	NUMTUPLES
[60, inf)	[0, inf)	30K	[60, inf)	[0, inf)	30K
[40, 60)	[0, 5) U [15, inf)	20K	[40, 60)	[2, 3)	20K
[40, 60)	[5, 15)	10K	[40, 60)	[2, 3)	10K
[20, 40)	[5, 10)	10K	[20, 40)	[0, 2) U [3, inf)	10K
[20, 40)	[15, inf)	10K	[20, 40)	[0, 2) U [3, inf)	10K

(b) **View Alignment**

A	B	C	NUMTUPLES
[60, inf)	[0, inf)	[0, inf)	30K
[40, 60)	[0, 5) U [15, inf)	[2, 3)	20K
[40, 60)	[5, 15)	[2, 3)	10K
[20, 40)	[5, 10)	[0, 2) U [3, inf)	10K
[20, 40)	[15, inf)	[0, 2) U [3, inf)	10K

(c) **Merged View Solution**

Figure 9: Align and Merge Example

5.1.3 Merging

This is the last step in the construction of the view solution. Here we simply merge the two solutions obtained after alignment through a “position” based join, where the physically corresponding rows in each solution are combined, with the common attributes being represented once. For example, the aligned solutions of Figure 9b are merge-joined using the positions (or row identifiers) to deliver the final view solution of Figure 9c.

As discussed earlier, DataSynth adopted a sampling algorithm for constructing the view solutions post LP solving. In marked contrast, Hydra *deterministically* generates the view solutions, facilitating us to operate purely in the summary space. There are two tangible benefits of this deterministic strategy: (a) elimination of the time and space overheads due to sampling, and (b) elimination of sampling-based errors in satisfying CCs.

5.2 Instantiating View Summaries

As shown in Figure 9c, each row in the view solution is comprised of a series of intervals (across various attributes) and the number of tuples in the region represented by these intervals. We now need to decide as to how these tuples are distributed *within* the attribute intervals. Our current solution is very simple: Assign the *entire* cardinality to the *left boundaries* of the intervals. For example, the third row in Figure 9c would result in generation of 10000 tuples all having $A = 40, B = 5, C = 2$ values.

Note that, in principle, we could have used a more sophisticated cardinality distribution within the intervals. However, our simple deterministic choice helps to reduce the subsequent additive errors that are incurred while ensuring referential integrity across views (described in next subsection). This is so because choosing values deterministically within a bucket minimizes the likelihood of encountering an **fk** value that is not present in the corresponding **pk** column.

5.3 Making View Summaries Consistent

Since the solution for each view is obtained independently, there could be inconsistencies across them. For example, referring back to the view schema shown in Section 3.2, R_view has attributes borrowed from S_view and T_view , and its solution may feature values that are not present in the corresponding attributes of these two views. To address this problem, we first carry out a *topological sort* on the “referential dependency graph”³ and then iteratively make the current view consistent with its predecessors. Since a topological sort is employed, Hydra can handle dependency graphs that are DAGs unlike DataSynth which is restricted to tree traversals.

To make a pair of views V_i and V_j consistent with each other, where V_i is dependent on V_j , we iterate over the rows in the view solution of V_i and look for the value combination that each row has for the attributes borrowed from V_j . If that value combination is not present in the solution of V_j , we add a new row in its solution with the corresponding NUMTUPLES attribute set to 1. This results in an additive error in the total number of tuples in the view as compared to the original AQP at the client. But we hasten to add that the error is a fixed number of rows, determined by the nature of the constraints and the LP solution, and *not* by the data scale. Therefore, at Big Data volumes, the discrepancy can be expected to be minuscule, and our experiments empirically confirm this expectation.

³A graph where each relation is represented by a node and an edge (u, v) is added if relation u is dependent on relation v through a referential constraint.

The inter view consistency component is present in DataSynth as well, but since its view solutions are comprised of complete database instantiations, and not just summaries, the time and space overheads incurred for making the views consistent can be large. Moreover, the additive error in DataSynth is amplified due to its inherent sampling errors. Our experiments also capture this distinction between the errors incurred due to referential constraints in Hydra and DataSynth.

5.4 Constructing Relation Summaries

After constructing consistent solutions across all the views, we next need to obtain the corresponding relation summaries. For this, we create a summarized relation schema \widetilde{R}_i for each relation R_i . This schema consists of all attributes in R_i except the primary key attribute, and additionally, the NUMTUPLES value for each entry in \widetilde{R}_i , as sourced from the view solutions.

For the common attributes between the summarized relation and the corresponding view solution, the value combinations and corresponding NUMTUPLES value are directly borrowed from the solution. What remains are the foreign key attributes. For filling a foreign key attribute fk , we need to first consult the view corresponding to fk 's target relation, say V_j . To fill the fk value in row r of \widetilde{R}_i , we extract the value combination in row r of view solution of V_i , and then project the attributes corresponding to V_j – let this be denoted by v . Now, we iterate over the solution set of V_j and compute the cumulative sum of the cardinality entries till v is reached. This sum provides the fk value corresponding to the r th row of \widetilde{R}_i , and we thus obtain \widetilde{R}_i for each relation R_i .

The set of relation summaries, computed as described above, provides the entire database summary – a sample such summary was previously shown in Figure 6 (for simplicity, the figure shows the PK columns instead of the number of tuples).

Like before, DataSynth again iterates over the complete instantiated (consistent) views to construct the corresponding materialized relations. Obviously, this leads to enormous time and space overheads in contrast to our data-scale independent summary based approach.

6 Tuple Generator

The Tuple Generator component resides inside the database engine, and needs to be explicitly incorporated in the engine codebase by the vendor. As a proof of concept, we have implemented it for the PostgreSQL v9.3 engine by adding a new feature called *datagen*, which is included as a property for each relation in the database. Whenever this feature is enabled for a relation, the scan operator for that relation is replaced with the dynamic generation operator. As a result, during query execution, the executor does not fetch the data from the disk but is instead supplied by the Tuple Generator in an *on-demand* manner, using the available relation summary.

Each row in the relation summary has a value combination and an associated NUMTUPLES entry. We consider the **pk** values to be the row numbers of the relation. Therefore, to get the r th tuple of a relation R , the **pk** is chosen as r and the rest of the attributes come from the relation summary. We iterate over the rows of \widetilde{R} and take the cumulative sum of the NUMTUPLES entries until the sum exceeds r . Say the summation crosses the value r in j th row of \widetilde{R} . Then the rest of the values of the r th tuple are assigned to be precisely the same as those present in the j th row of \widetilde{R} . For example, the 120th row of relation S in Figure 6, would be $\langle 120, 20, 15 \rangle$.

Note that this form of tuple generation is expected to be efficient since the attribute value

assignments are deterministic and independent, and these expectations are confirmed in the experiments shown in the following section.

7 Experiments

We have implemented the Hydra design, described in the previous sections, in a Java tool running to over 15K lines of code. The popular Z3 [21] solver is leveraged to compute solutions for the LP formulations. In this section, we evaluate Hydra’s empirical performance, using our implementation of DataSynth as the comparative yardstick in the analysis.

Database Environment The TPC-DS [5] decision-support benchmark database, with a default size of 100GB, is used as the baseline in our experiments. The database is hosted on a PostgreSQL v9.3 engine [4] with the hardware platform being a vanilla HP workstation (3.2 GHz 16 core processor, 32 GB memory, 500 GB SSD hard drive) running Ubuntu Linux 16.04.3.

A complex query workload, WL_c , featuring 131 distinct queries (enumerated in Appendix A), was created by customizing the 99 queries of the benchmark such that only non-key filter predicates and PK-FK joins were retained, and all nested queries were separated into independent sub-queries⁴. The AQPs for these queries were generated on the PostgreSQL query processor, resulting in 351 cardinality constraints. The distribution of the cardinalities for these CCs are shown in Figure 10, with the cardinalities measured on a log-scale. The figure clearly indicates that a wide range of cardinalities are present in the constraints, going from a few tuples to almost a billion.

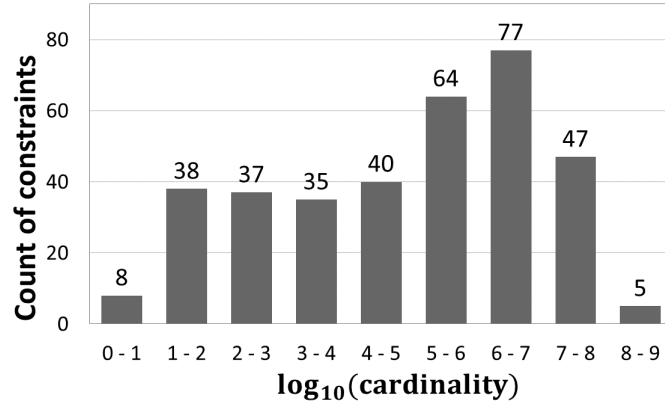


Figure 10: Distribution of Cardinality in CCs (WL_c)

The above constraints result in a large number of geometrically overlapping regions. Hydra, due to its region-partitioning approach, comfortably handles this scenario. In marked contrast, DataSynth, due to its grid-partitioning construction, generates a very large number of LP variables (in the several billion) from the constraints, overwhelming the solver’s capabilities. We therefore also created an alternative simplified query workload, called WL_s , with 311 CCs, wherein the variables created by DataSynth were less than a million, and therefore well within the solver’s processing power.

⁴Similar to DataSynth, the restriction to non-key-based filters is because the conversion from relations to views lose the key attributes. Likewise, only PK-FK joins are supported since they are inherently present in the design of views.

7.1 Quality of Volumetric Similarity

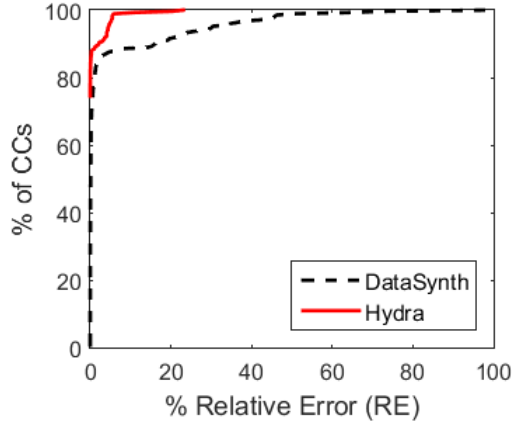


Figure 11: Quality of Volumetric Similarity (WL_c)

We begin by investigating how closely the volumetric similarity, with regard to operator output cardinalities, is achieved between the client and vendor sites for the WL_s workload by the Hydra and DataSynth regenerators. This behavior is captured in Figure 11, which plots the percentage of CCs that are within a given relative error of volumetric similarity. From the plot it is evident that Hydra satisfies around 90 percent of the CCs with virtually no error, and the remaining CCs are also satisfied within a relative error of less than 10%. This is in contrast to DataSynth, which accurately satisfies around 80 percent of the CCs, but then incurs as much as 60% relative error to achieve complete coverage of the remaining CCs.

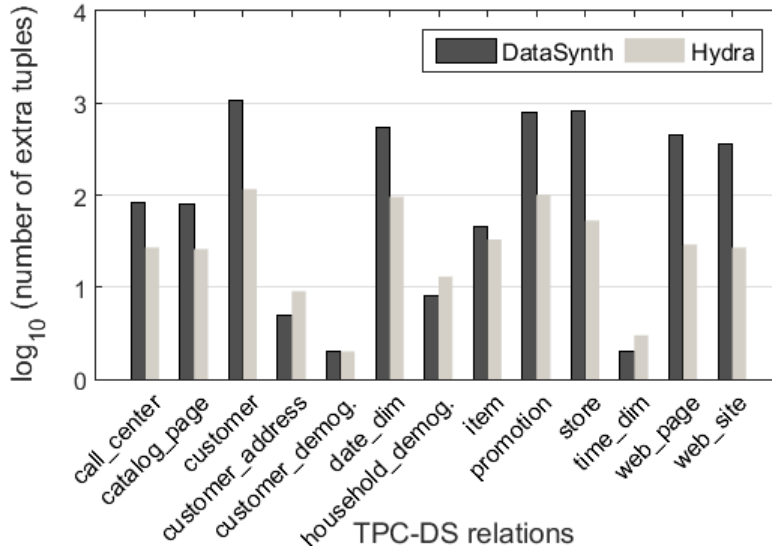


Figure 12: Extra tuples for Referential Integrity (WL_c)

There are two reasons for the error-prone behavior of DataSynth: (1) the probabilistic sampling technique, and (2) the maintenance of referential integrity. While Hydra also is forced to insert additional tuples to maintain referential integrity, the number is substantially smaller than those injected by DataSynth. This is because the integrity errors are *amplified* by the impact of the sampling errors. This effect is quantified in Figure 12, where the number of extra tuples inserted is plotted on a log-scale for representative TPC-DS tables. We see here that Hydra is often

an *order-of-magnitude* smaller with regard to the addition of these extra tuples as compared to DataSynth.

Recall that integrity errors in Hydra are independent of the data scale. Figure 13 shows the distribution of absolute volumetric errors for the constraints set at different database scales. We can see that this distribution is very similar for the three different scales. Hence, in the relative sense, the errors become minuscule as the data is scaled to large volumes.

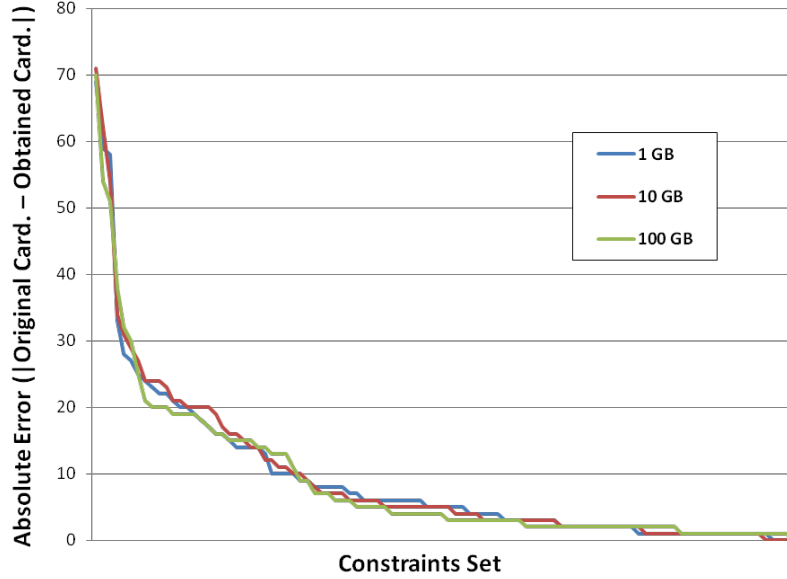


Figure 13: Distribution of Absolute Errors

As a final observation, it is interesting to note that DataSynth has to contend with both *negative* (volumes less than desired) and *positive* (volume greater than desired) relative errors, due to its sampling strategy – in fact, about one-third of the CCs suffered negative relative errors. In contrast, Hydra only generates positive errors due to the inclusion of extra tuples for satisfying referential integrity. From a practical standpoint, it is perhaps preferable to have positive errors since they induce greater stress on the data processing elements in the engine.

7.2 Scalability with Workload Complexity

We now turn our attention to evaluating the complexity of the underlying LP that is formulated by Hydra and DataSynth. Since LP complexity is essentially proportional to the number of variables in the problem, we compare this number for the two techniques. Further, since LP complexity is, to the first degree of approximation, independent of the database size, we present the comparison only for the 100 GB instance.⁵

The number of LP variables for a representative set of TPC-DS relations, including the major fact and dimension tables (`catalog_sales`, `store_sales`, `item`) is captured, on a log-scale, in Figure 14 for the WL_c complex workload. We observe here that the LPs formulated using the region-partitioning strategy in Hydra generate *several orders of magnitude* fewer variables than the corresponding LPs derived from the grid-partitioning in DataSynth. As a case in point, consider the `catalog_sales` table – the number of variables created by DataSynth was almost 5.5 million,

⁵Of course, the database engine’s choice of query plans may change to some extent with database size, leading to a slightly different set of CCs.

which is reduced to as low as 1620 by Hydra. Even more dramatic is the change for `item` table, where the number of variables is reduced from an enormous 10^{11} to around 3700.

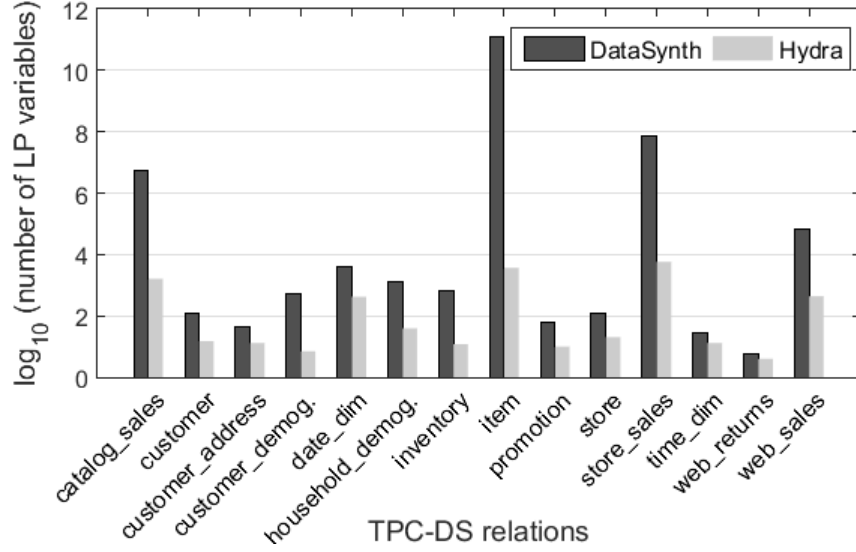


Figure 14: Number of variables in the LP (WL_c)

From an absolute perspective also, the large number of variables created by DataSynth is a critical problem since, as mentioned previously, the LP solver crashed in handling these cases. In marked contrast, the few thousands of LP variables generated by Hydra were easily solvable in less than a minute. Moreover, even when we switched to the simple workload, WL_s , the LP solution time for DataSynth was almost an hour, whereas Hydra completed in a few seconds as shown in Figure 15.

Complex Workload (WL_c)		Simple Workload (WL_s)	
DataSynth	Hydra	DataSynth	Hydra
<i>crash</i>	58 sec	50 min	13 sec

Figure 15: LP Processing Time

7.3 Scalability with Materialized Data Size

This experiment compares the data instantiation times, post LP solution, of DataSynth and Hydra on the WL_s workload. While Hydra, in principle, due to its summary-based approach, does *not* have to instantiate the data immediately, we assume in this experiment that the vendor requires complete materialization.

The experimental results are shown in Figure 16, where we also present, for comparative purposes, the performance with 10 GB and 1000 GB databases, apart from the default 100 GB database. We see here that there is a huge reduction in the materialization time of Hydra at all scales. Further, even in absolute terms, Hydra is able to output a 100 GB database in around 11 minutes, whereas DataSynth takes 42 hours to complete the same task.

The marked difference in the efficiency of the two techniques is attributed to the fact that DataSynth instantiates complete views through sampling, subsequently performs *several passes* on these instantiations to ensure referential integrity, and to derive relations from them. Hydra

on the other hand, after LP-solving, constructs the database summary in just a few seconds, and then instantiates the materialized database directly from it.

Size (in GB)	DataSynth	Hydra
10	4 hours	2 min
100	42 hours	11 min
1000	> 1 week	1.6 hours

Figure 16: Data Materialization Time

7.4 Scalability to Big Data Volumes

In our next experiment, we validated the ability of Hydra, thanks to its summary-based technique, to scale to Big Data volumes. To demonstrate this feature, we modeled an exabyte-sized (10^{18} bytes) data scenario as follows: We used CODD, which is capable of modeling arbitrary metadata scenarios, to obtain the optimizer-chosen plans at the exabyte database scale for all the workload queries. To get AQPs for this database, we executed the obtained plans on the 100 GB instance and scaled the intermediate row counts with the appropriate scale factor. Hydra was able to formulate and solve the LPs (one per relation), and generate the database summary in less than **2 minutes**. Once the summary is generated, the database can begin to submit the workload queries since the data required for the execution can be produced on-the-fly by the Tuple Generator.

7.5 Dynamism in Data Generation

Our next experiment evaluates Hydra’s ability, due to the Tuple Generator and Database Summary architecture, to produce tuples *on-the-fly* instead of first materializing them, and then reading from the disk. To verify whether dynamic generation can indeed produce data at rates that are practical for supporting query execution, we compared the total time that Hydra’s tuple generator took to construct and supply tuples to the executor, while running simple aggregate queries, as compared to the standard sequential scan from the disk.

Rel. Name	Size	Row count (in millions)	Scan time (secs)	
	(in GB)		Disk	Dynamic
store_returns	3	29	16	8
web_sales	10	72	43	25
inventory	19	399	107	74
catalog_sales	20	144	46	48
store_sales	34	288	168	87

Figure 17: Data Supply Times

The results of this experiment are shown in Figure 17 for the five biggest relations in the 100 GB database instance. We see here that the tuple generator is not only competitive with a materialized solution, but is in fact typically *faster*. Therefore, using dynamic generation can prove to be a good option since it can help to eliminate the large time and space overheads incurred in: (1) dumping generated data on the disk, and (2) loading the data on the engine under test.

7.6 Performance on JOB Benchmark

A legitimate concern with regard to the above encouraging results for Hydra is that they may be an artifact of the TPC-DS database, and perhaps might under-perform on other datasets. To address this concern, we consider in our final experiment, a schematically highly different database, namely the JOB benchmark [19], which is based on the IMDB real-world dataset. Here, we created a workload of 260 queries, resulting in 523 CCs, whose cardinality distribution is again highly varied as seen in Figure 18.

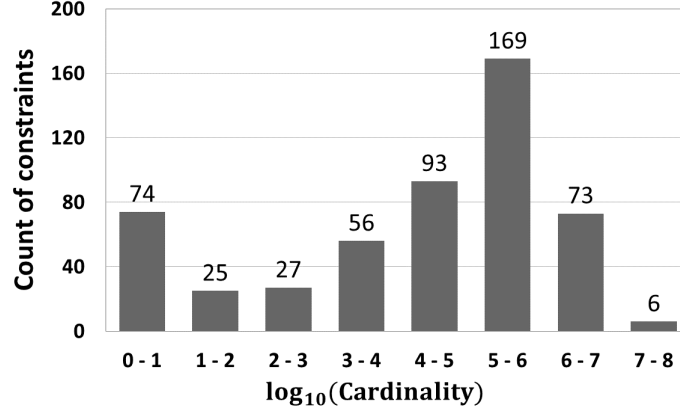


Figure 18: Cardinality distribution of CCs in JOB

We found that Hydra efficiently solved this workload as well, with the number of variables in each view being typically in the few thousands, and never exceeding a hundred thousand, as shown quantitatively in Figure 19. The overall database summary was quickly generated in around 20 seconds, and produced a database of high fidelity that satisfied all the constraints with no more than 2 percent relative error.




Figure 19: Number of Variables for JOB

8 Related Work

Over the past few decades, a rich corpus of literature has developed on synthetic database construction. There are two broad streams of research on the topic, one dealing with the *ab initio* generation of new databases using standard mathematical distributions (e.g. [17, 15]), and the other with *regeneration* of an arbitrary existing database. In the latter category, there are two approaches, one of which uses only schematic and statistical information from the original database (e.g. [22, 24]). The other uses both the original database and the query workload to achieve statistical fidelity during evaluation (e.g [14, 9]) – our work on Hydra falls into this class. In this section, we briefly review recent literature on this spectrum of research categories.

Ab Initio Generation. Descriptive languages for the definitions of data dependencies and column distributions were proposed in [15, 18, 23]. For example, [15] proposed a special purpose language called Data Generation Language (DGL) that is used by the tool to generate synthetic data distributions by utilizing the concept of iterators. It supports a broad range of dependencies between relations but the construction of dependent tables always requires access to the referenced table, creating a bottleneck on the data generation speed.

In contrast to the above, MUDD [26] and PSDG [18] generate all related data at the same time. However, this approach can also be rendered inefficient if the referenced tables are large in size. MUDD proposes algorithms to parallelize the data generation process, and to efficiently generate dense-unique-pseudo-random sequences and derive nonuniform distributions. Both MUDD and PSDG decouple data generation details from data description, facilitating customization of the tool to suit user needs.

In the distributed setting, a faster way of generating references is through recomputing since it eliminates the I/O costs incurred to satisfy referential constraints across relations that are present across different nodes. PDGF [23] was designed with this goal of achieving scalability and decoupling. In PDGF, the user specifies two XML configuration files, one for the data model and one for the formatting instructions. The generation strategy is based on the exploitation of determinism in pseudo-random number generators (PRNG), which enables regeneration of the same sequences, hence eliminating the scan overheads. PDGF supports the generation of data with cyclic dependencies as well, but incurs high computation costs for generating the associated keys. Finally, PDGF comes with a set of fixed generators for different datatypes and basic distribution functions.

A similar generator is Myriad [8], which implements an efficient parallel execution strategy leveraged by extensive use of PRNGs with random access support. With these PRNGs, Myriad distributes the generation process across the compute nodes and ensures that they can run independently from each other, without imposing any restrictions on the data modeling language.

Finally, a rule-based probabilistic approach, based on an extension of Datalog, has been recently proposed in [12], which is capable of generating data characterized by parametrized classical discrete distributions – however, it is not always feasible to assign such distributions to real-world data, especially over multivariate spaces.

Database-dependent Regeneration. DBSynth[22] is an extension to PDGF, which builds data models from an existing database by extracting schema information, and using sampling to construct histograms and dictionaries of text-valued data. Further, if the textual data contains multiple words, Markov chain generators are used to analyze the word combination frequencies

and probabilities. Finally, after the model construction is complete, PDGF is invoked to generate the corresponding data.

Like DBSynth, RSGen [24] takes a metadata dump, including 1-D histograms, as the input, and generates database tables along with a loading script as the output. It uses a bucket based model at its core, which is able to generate trillions of records with minimum memory footage. However, the proposed technique works well only for queries with only a single range predicate. Further, due to the inaccurate statistical models in the query optimizer, the volumetric similarity is poor for queries involving predicates on correlated attributes.

UpSizeR [25] is a graph-based tool that uses attribute correlations extracted from an existing database to generate an equivalent synthetic database. A derivative work, Rex [16] produces an extrapolated database given an integer scaling factor and the original database, while maintaining referential constraints and the distributions between the consecutive linked tables. Dscaler [28] addresses the problem of generating a non-uniformly⁶ scaled version of a database using fine-grained, per-tuple correlations for key attributes, but such information is typically hard to come by. Moreover, all these techniques only generate the *key* attributes, whereas the non-key values are sampled from the original database using these key values. Hence, the approach becomes impractical in Big Data and security-conscious environments. Finally, Dscaler fails to retain accuracy for some common query classes.

Query-dependent Regeneration. Apart from the above techniques, another line of work [13, 14, 20, 9] is based on workload dependence (as in the case of Hydra). Here the aim is to generate a database given a workload of queries such that volumetric similarity is achieved on these queries. In particular, RQP [13] gets a query and a result as input, and returns a possible database instance that could have produced the result for that query. The idea of using cardinalities from a query plan tree was first introduced in QAGen [14]. They start by constructing a *symbolic database*⁷, and then translate the input AQPs to constraints over the symbols in the database. Subsequently, a constraint satisfaction program (CSP) is invoked to identify values for symbols that satisfy all the constraints.

On the positive side, these generators are capable of handling complex operators as they use a general CSP, but the performance cost is huge since the number of CSP calls also increases with the database size. Further, it requires operating on a symbolic database of matching size to the original database, and processing of the entire database during the algorithm execution. This makes it impractical for Big Data environments. Finally, QAGen supports only one query plan in the input. This limitation was addressed in a follow-up tool called MyBenchmark [20], which creates a symbolic database on a per query basis and at the end tries to heuristically merge the various databases into a small number of databases. Clearly, generating a database on a per query basis has enormous time and space overheads, and further, a single database is not guaranteed in the output.

DataSynth [9] identified the declarative property of cardinality constraints and its ability to specify data characteristics. Given a large number of cardinality constraints as input, the paper proposed algorithms based on the LP solver and graphical models to instantiate tables that satisfy those constraints. However, it suffers from high LP complexity, data scale dependencies, and inaccuracies with regard to volumetric similarity, as we have discussed. Hydra materially extends the DataSynth approach by adding dynamism, scale and functionality.

⁶In non-uniform scaling, individual tables are scaled by different factors.

⁷A symbolic database is similar to a regular database, but its attribute values are symbols (variables), not constants.

9 Conclusions

The ability to synthetically regenerate data that accurately conforms to the volumetric behavior on queries at client sites is of crucial importance to database vendors, and will become even more so with the advent of Big Data applications. In this paper, we have proposed Hydra, a data regeneration tool that takes a substantial step forward towards achieving this goal. Specifically, by reworking the basic LP problem formulation into a region-based variable assignment, Hydra improves on the state-of-the-art DataSynth’s performance by orders of magnitude with regard to problem complexity, data materialization time, and scalability to large volumes. Secondly, by using a deterministic alignment technique for database consistency, it provides far better accuracy in meeting volumetric constraints as compared to the probabilistic approach employed in DataSynth. Finally, its summary-based framework organically supports the dynamic regeneration of streaming data sources, an essential pre-requisite for efficiently testing contemporary deployments.

References

- [1] Big data. en.wikipedia.org/wiki/Big_data
- [2] Lernaean Hydra — Wikipedia, The Free Encyclopedia. en.wikipedia.org/wiki/Lernaean_Hydra
- [3] Hydra Database Regenerator. dsl.cds.iisc.ac.in/projects/HYDRA
- [4] PostgreSQL. postgresql.org/docs/9.3
- [5] TPC-DS. tpc.org/tpcds
- [6] TPC-H. tpc.org/tpch
- [7] USE PLAN SQL Server. [technet.microsoft.com/en-us/library/ms186954\(v=sql.105\).aspx](http://technet.microsoft.com/en-us/library/ms186954(v=sql.105).aspx)
- [8] A. Alexandrov, K. Tzoumas and V. Markl. Myriad: Scalable and Expressive Data Generation. *PVLDB*, 5(12):1890-1893, 2012.
- [9] A. Arasu, R. Kaushik, and J. Li. Data Generation using Declarative Constraints. In *Proc. of ACM SIGMOD Conf.*, 2011, pgs. 685-696.
- [10] A. Arasu, R. Kaushik, and J. Li. Datasynth: Generating synthetic data using declarative constraints. *PVLDB*, 4(12):1418-1421, 2011.
- [11] S. Ashoke and J. R. Haritsa. CODD: A Dataless Approach to Big Data Testing. *PVLDB*, 8(12):2008-2011, 2015.
- [12] V. Barany, B. Cate, B. Kimelfeld, D. Olteanu and Z. Vagena. Declarative Probabilistic Programming with Datalog. In *Proc. of 19th ICDT Conf.*, 2016, pgs. 7:1-7:19.
- [13] C. Binnig, D. Kossmann, E. Lo. Reverse Query Processing. In *Proc. of 23rd ICDE Conf.*, 2007, pgs. 506-515.
- [14] C. Binnig, D. Kossmann, E. Lo, and M. T. Özsu. QAGen: Generating Query-Aware Test Databases. In *Proc. of ACM SIGMOD Conf.*, 2007, pgs. 341-352.
- [15] N. Bruno and S. Chaudhuri. Flexible Database Generators. In *Proc. of 31st VLDB Conf.*, 2005, pgs. 1097-1107.
- [16] T. S. Buda, T. Cerqueus, J. Murphy and M. Kristiansen. ReX: Extrapolating Relational Data in a Representative Way. In *Proc. of BICOD Conf.*, 2015, pgs. 95-107.
- [17] J. Gray, P. Sundaresan, S. Englert, K. Baclawski, and P. J. Weinberger. Quickly Generating Billion-Record Synthetic Databases. In *Proc. of ACM SIGMOD Conf.*, 1994, pgs. 243-252.
- [18] J. E. Hoag and C. W. Thompson. A Parallel General-Purpose Synthetic Data Generator. In *Proc. of ACM SIGMOD Conf.*, 2007, pgs. 19-24.
- [19] V. Leis, A. Gubichev, A. Mirchev, P. Boncz, A. Kemper and T. Neumann. How Good Are Query Optimizers, Really? *PVLDB*, 9(3):204-215, 2015.

- [20] E. Lo, N. Cheng, W. W. Lin, W.-K. Hon, and B. Choi. MyBenchmark: generating databases for query workloads. *The VLDB Journal*, 23(6):895-913, 2014.
- [21] L. De Moura and N. Bjørner. Z3: An efficient SMT solver. In *Proc. of TACAS Conf.*, 2008, pgs. 337-340.
- [22] T. Rabl, M. Danisch, M. Frank, S. Schindler and H. Jacobsen. Just can't get enough - Synthesizing Big Data. In *Proc. of ACM SIGMOD Conf.*, 2015, pgs. 1457-1462.
- [23] T. Rabl, M. Frank, H. M. Sergieh and H. Kosch. Data Generator for Cloud-Scale Benchmarking. In *Proc. of 2nd TPCTC Conf.*, 2010, pgs. 41-56.
- [24] E. Shen and L. Antova. Reversing statistics for scalable test databases generation. In *Proc. of 6th DBTest*, 2013, pgs. 7:1-7:6.
- [25] J. M. Stephens and M. Poess. MUDD: A Multi-dimensional Data Generator. In *Proc. of 4th WOSP*, 2004, pgs. 104-109.
- [26] Y. C. Tay, B. T. Dai, D. T. Wang, E. Y. Sun, Yong Lin and Yuting Lin. UpSizeR: Synthetically scaling an empirical relational database. *Inf. Syst.*, 38(8):1168-1183, 2013.
- [27] R. S. Trivedi, I. Nilavalagan, and J. R. Haritsa. CODD: COnstructing Dataless Databases. In *Proc. of 5th DBTest*, 2012, pgs. 4:1-4:6.
- [28] J. W. Zhang and Y. C. Tay. Dscaler: Synthetically Scaling A Given Relational Database. *PVLDB*, 9(14):1671-1682, 2016.

Appendix

A Query Workload

Following are the queries we have used in conducting the experiments. These queries have been adapted from the TPC-DS benchmark.

Query No.	SQL Text
1	<pre>select * from store_sales, household_demographics, time_dim, store where ss_sold_time_sk = time_dim.t_time_sk and ss_hdemo_sk = household_demographics.hd_demo_sk and ss_store_sk = s_store_sk and time_dim.t_hour = 8 and time_dim.t_minute >= 30 and household_demographics.hd_dep_count = 5 and store.s_store_name = 'ese'</pre>
2	<pre>select * from store_sales, customer_demographics, date_dim, item, promotion where ss_sold_date_sk = d_date_sk and ss_item_sk = i_item_sk and ss_cdemo_sk = cd_demo_sk and ss_promo_sk = p_promo_sk and cd_gender = 'M' and cd_marital_status = 'M' and cd_education_status = '4 yr Degree' and (p_channel_email = 'N' or p_channel_event = 'N') and d_year = 2001</pre>
3	<pre>select * from inventory, item, warehouse, date_dim where inv_item_sk = i_item_sk and inv_warehouse_sk = w_warehouse_sk and inv_date_sk = d_date_sk and d_year =1999</pre>
4	<pre>select * from catalog_sales, item, date_dim where i_manufact_id = 291 and i_item_sk = cs_item_sk and d_date between '2000-03-22' and '2000-06-22' and d_date_sk = cs_sold_date_sk</pre>
5	<pre>select * from date_dim, store_sales, item where d_date_sk = ss_sold_date_sk and ss_item_sk = i_item_sk and i_manager_id=91 and d_moy=12 and d_year=2002</pre>
6	<pre>select * from web_sales, date_dim d1, item where d1.d_month_seq between 1211 and 1211+11 and d1.d_date_sk = ws_sold_date_sk and i_item_sk = ws_item_sk</pre>
7	<pre>select * from store_returns, date_dim where sr_returned_date_sk = d_date_sk and d_year =1998</pre>

8	select * from catalog_returns, date_dim where cr_returned_date_sk = d_date_sk and d_year = 2000 and d_moy = 12
9	select * from customer, customer_address, customer_demographics where cd_demo_sk = c_current_cdemo_sk and ca_address_sk = c_current_addr_sk and ((cd_marital_status = 'M' and cd_education_status = 'Unknown') or (cd_marital_status = 'W' and cd_education_status = 'Advanced Degree')) and ca_gmt_offset = -6
10	select * from inventory, warehouse, item, date_dim where i_current_price between 0.99 and 1.49 and i_item_sk = inv_item_sk and inv_warehouse_sk = w_warehouse_sk and inv_date_sk = d_date_sk and d_date between '2001-02-14' and '2001-05-14'
11	select * from date_dim, store_sales, store where d_date_sk = ss_sold_date_sk and s_store_sk = ss_store_sk and s_gmt_offset = -5 and d_year = 2000
12	select * from store_sales, customer_demographics, date_dim, store, item where ss_sold_date_sk = d_date_sk and ss_item_sk = i_item_sk and ss_store_sk = s_store_sk and ss_cdemo_sk = cd_demo_sk and cd_gender = 'F' and cd_marital_status = 'U' and cd_education_status = 'Secondary' and d_year = 1999 and s_state in ('TN')
13	select * from web_sales ws1, date_dim, customer_address, web_site where d_date between '2002-4-01' and '2002-6-01' and ws1.ws_ship_date_sk = d_date_sk and ws1.ws_ship_addr_sk = ca_address_sk and ca_state = 'MN' and ws1.ws_web_site_sk = web_site_sk and web_company_name = 'pri'
14	select * from web_sales, date_dim where ws_sold_date_sk = d_date_sk and d_qoy = 2 and d_year = 1999
15	select * from date_dim where d_date = '2002-03-09'
16	select * from store_sales, date_dim d1, item, store where d1.d_year = 2001 and d1.d_date_sk = ss_sold_date_sk and i_item_sk = ss_item_sk

	and s_store_sk = ss_store_sk and s_state in ('TN')
17	select * from store_sales, date_dim, customer_address, item where i_category in ('Home') and ss_item_sk = i_item_sk and ss_sold_date_sk = d_date_sk and d_year = 2000 and d_moy = 2 and ss_addr_sk = ca_address_sk and ca_gmt_offset = -5
18	select * from catalog_sales, date_dim, customer_address, item where i_category in ('Home') and cs_item_sk = i_item_sk and cs_sold_date_sk = d_date_sk and d_year = 2000 and d_moy = 2 and cs_bill_addr_sk = ca_address_sk and ca_gmt_offset = -5
19	select * from web_sales, date_dim, customer_address, item where i_category in ('Home') and ws_item_sk = i_item_sk and ws_sold_date_sk = d_date_sk and d_year = 2000 and d_moy = 2 and ws_bill_addr_sk = ca_address_sk and ca_gmt_offset = -5
20	select * from store_sales, date_dim, store, household_demographics where store_sales.ss_sold_date_sk = date_dim.d_date_sk and store_sales.ss_store_sk = store.s_store_sk and store_sales.ss_hdemo_sk = household_demographics.hd_demo_sk and (household_demographics.hd_dep_count = 8 or household_demographics.hd_vehicle_count = -1) and date_dim.d_dow in (6, 0) and date_dim.d_year in (2000, 2000+1, 2000+2) and store.s_city in ('Midway', 'Fairview')
21	select * from web_sales, date_dim where d_month_seq between 1212 and 1212 + 11 and ws_ship_date_sk = d_date_sk
22	select * from catalog_sales cs1, date_dim, customer_address, call_center where d_date between '2000-3-01' and '2000-5-01' and cs1.cs_ship_date_sk = d_date_sk and cs1.cs_ship_addr_sk = ca_address_sk and ca_state = 'MN' and cs1.cs_call_center_sk = cc_call_center_sk and cc_county in ('Williamson County')
23	select * from customer c, customer_address ca where c.c_current_addr_sk = ca.ca_address_sk and ca_county in ('Yellowstone County', 'Montgomery County', 'Divide County', 'Cedar County', 'Manassas Park city')
24	select * from store_sales, date_dim

	where ss_sold_date_sk = d_date_sk and d_year = 2000 and d_moy between 2 and 2+3
25	select * from web_sales, date_dim where ws_sold_date_sk = d_date_sk and d_year = 2000 and d_moy between 2 and 2+3
26	select * from catalog_sales, date_dim where cs_sold_date_sk = d_date_sk and d_year = 2000 and d_moy between 2 and 2+3
27	select * from item, store_sales, date_dim where ss_item_sk = i_item_sk and ss_sold_date_sk = d_date_sk and d_month_seq in (1206, 1206+1, 1206+2, 1206+3, 1206+4, 1206+5, 1206+6, 1206+7, 1206+8, 1206+9, 1206+10, 1206+11) and ((i_category in ('Books', 'Children', 'Electronics') and i_class in ('personal', 'portable', 'refernece', 'self-help') and i_brand in ('scholaramalgamalg #14', 'scholaramalgamalg #7', 'exportiunivamalg #9', 'scholaramalgamalg #9')) or (i_category in ('Women', 'Music', 'Men') and i_class in ('accessories', 'classical', 'fragrances', 'pants') and i_brand in ('amalgimporto #1', 'edu packscholar #1', 'exportiimporto #1', 'importoa-malg #1'))))
28	select * from customer c, customer_address ca where c.c_current_addr_sk = ca.ca_address_sk and ca_state in ('SD', 'KY', 'MO')
29	select * from store_sales, date_dim where ss_sold_date_sk = d_date_sk and d_year = 2004 and d_moy between 3 and 3+2
30	select * from web_sales, date_dim where ws_sold_date_sk = d_date_sk and d_year = 2004 and d_moy between 3 and 3+2
31	select * from catalog_sales, date_dim where cs_sold_date_sk = d_date_sk and d_year = 2004 and d_moy between 3 and 3+2
32	select * from store_sales, date_dim, customer_address, item where i_category in ('Children') and ss_item_sk = i_item_sk and ss_sold_date_sk = d_date_sk and d_year = 1998 and d_moy = 8 and ss_addr_sk = ca_address_sk and ca_gmt_offset = -5
33	select * from catalog_sales, date_dim, customer_address, item where i_category in ('Children') and cs_item_sk = i_item_sk

	and cs_sold_date_sk = d_date_sk and d_year = 1998 and d_moy = 8 and cs_bill_addr_sk = ca_address_sk and ca_gmt_offset = -5
34	select * from web_sales, date_dim, customer_address, item where i_category in ('Children') and ws_item_sk = i_item_sk and ws_sold_date_sk = d_date_sk and d_year = 1998 and d_moy = 8 and ws_bill_addr_sk = ca_address_sk and ca_gmt_offset = -5
35	select * from item, inventory, date_dim, catalog_sales where i_current_price between 4 and 10 and inv_item_sk = i_item_sk and d_date_sk=inv_date_sk and d_date between '2002-01-18' and '2002-03-18' and i_manufact_id in (744, 691, 853, 946) and inv_quantity_on_hand between 100 and 500 and cs_item_sk = i_item_sk
36	select * from store_sales , item , date_dim where ss_item_sk = i_item_sk and i_category in ('Music', 'Jewelry', 'Women') and ss_sold_date_sk = d_date_sk and d_date between '1999-04-26' and '1999-05-26'
37	select * from store_sales, date_dim d1 where d1.d_month_seq between 1181 and 1181+11 and d1.d_date_sk = ss_sold_date_sk
38	select * from store_sales, date_dim, item where ss_sold_date_sk=d_date_sk and ss_item_sk=i_item_sk and d_month_seq between 1214 and 1214+11
39	select * from store_sales where ss_quantity between 0 and 5 and (ss_list_price between 51 and 51+10 or ss_coupon_amt between 12565 and 12565+1000 or ss_wholesale_cost between 52 and 52+20)
40	select * from store_sales where ss_quantity between 6 and 10 and (ss_list_price between 135 and 135+10 or ss_coupon_amt between 3897 and 3897+1000 or ss_wholesale_cost between 79 and 79+20)
41	select * from store_sales where ss_quantity between 11 and 15 and (ss_list_price between 106 and 106+10 or ss_coupon_amt between 10740 and 10740+1000 or ss_wholesale_cost between 16 and 16+20)
42	select * from store_sales where ss_quantity between 16 and 20

	and (ss.list_price between 16 and 16+10 or ss.coupon_amt between 13313 and 13313+1000 or ss.wholesale_cost between 8 and 8+20)
43	select * from store_sales where ss.quantity between 21 and 25 and (ss.list_price between 153 and 153+10 or ss.coupon_amt between 4490 and 4490+1000 or ss.wholesale_cost between 14 and 14+20)
44	select * from store_sales where ss.quantity between 26 and 30 and (ss.list_price between 111 and 111+10 or ss.coupon_amt between 8115 and 8115+1000 or ss.wholesale_cost between 35 and 35+20)
45	select * from store_sales, date_dim where ss.sold_date_sk = d.date_sk and d.month_seq between 1202 and 1202 + 11
46	select * from web_sales, date_dim, time_dim , ship_mode where ws.sold_date_sk = d.date_sk and ws.sold_time_sk = t.time_sk and ws.ship_mode_sk = sm.ship_mode_sk and d.year = 1999 and t.time between 46185 and 46185+28800 and sm.carrier in ('DIAMOND', 'ZOUROS')
47	select * from catalog_sales, date_dim, time_dim , ship_mode where cs.sold_date_sk = d.date_sk and cs.sold_time_sk = t.time_sk and cs.ship_mode_sk = sm.ship_mode_sk and d.year = 1999 and t.time between 46185 and 46185+28800 and sm.carrier in ('DIAMOND', 'ZOUROS')
48	select * from web_sales, household_demographics, time_dim, web_page where ws.sold_time_sk = time_dim.t.time_sk and ws.ship_hdemo_sk = household_demographics.hd_demo_sk and ws.web_page_sk = web_page.wp_web_page_sk and time_dim.t.hour between 11 and 11+1 and household_demographics.hd_dep_count = 8 and web_page.wp_char_count between 5000 and 5200
49	select * from web_sales, household_demographics, time_dim, web_page where ws.sold_time_sk = time_dim.t.time_sk and ws.ship_hdemo_sk = household_demographics.hd_demo_sk and ws.web_page_sk = web_page.wp_web_page_sk and time_dim.t.hour between 13 and 13+1 and household_demographics.hd_dep_count = 8 and web_page.wp_char_count between 5000 and 5200
50	select * from item, store_sales, date_dim where ss.item_sk = i.item_sk and ss.sold_date_sk = d.date_sk and (d.year = 2000 or (d.year = 2000-1 and d.moy =12))

	or (d_year = 2000+1 and d_moy =1))
51	select * from web_sales ws1, date_dim, customer_address, web_site where d_date between '2000-2-01' and '2000-4-01' and ws1.ws_ship_date_sk = d_date_sk and ws1.ws_ship_addr_sk = ca_address_sk and ca_state = 'NE' and ws1.ws_web_site_sk = web_site_sk and web_company_name = 'pri'
52	select * from web_sales, item, date_dim where i_manufact_id = 248 and i_item_sk = ws_item_sk and d_date between '2000-02-02' and '2000-05-02' and d_date_sk = ws_sold_date_sk
53	select * from date_dim dt, store_sales, item where dt.d_date_sk = store_sales.ss_sold_date_sk and store_sales.ss_item_sk = item.i_item_sk and item.i_manufact_id = 783 and dt.d_moy=11
54	select * from store_sales, date_dim where ss_sold_date_sk = d_date_sk and d_year = 2001 and d_qoy < 4
55	select * from web_sales, date_dim where ws_sold_date_sk = d_date_sk and d_year = 2001 and d_qoy < 4
56	select * from catalog_sales, date_dim where cs_sold_date_sk = d_date_sk and d_year = 2001 and d_qoy < 4
57	select * from store_sales, date_dim where store_sales.ss_sold_date_sk = date_dim.d_date_sk and d_month_seq between 1191 and 1191 + 11
58	select * from catalog_sales, date_dim where catalog_sales.cs_sold_date_sk = date_dim.d_date_sk and d_month_seq between 1191 and 1191 + 11
59	select * from web_sales, date_dim where web_sales.ws_sold_date_sk = date_dim.d_date_sk and d_month_seq between 1191 and 1191 + 11
60	select * from inventory, date_dim, item, warehouse where inv_date_sk=d_date_sk and inv_item_sk=i_item_sk and inv_warehouse_sk = w_warehouse_sk and d_month_seq between 1199 and 1199 + 11
61	select * from item, store_sales, date_dim where ss_item_sk = i_item_sk and ss_sold_date_sk = d_date_sk and d_year in (2002) and ((i_category in ('Jewelry', 'Women', 'Shoes'))

	and i_class in ('mens watch', 'dresses', 'mens')) or (i_category in ('Men', 'Sports', 'Music') and i_class in ('sports-apparel', 'sailing', 'pop'))))
62	select * from date_dim where d_year = 1998 and d_moy = 5
63	select * from date_dim dt, store_sales, item where dt.d_date_sk = store_sales.ss_sold_date_sk and store_sales.ss_item_sk = item.i_item_sk and item.i_manager_id = 1 and dt.d_moy=12 and dt.d_year=2000
64	select * from date_dim dt , store_sales , item where dt.d_date_sk = store_sales.ss_sold_date_sk and store_sales.ss_item_sk = item.i_item_sk and item.i_manager_id = 1 and dt.d_moy=11 and dt.d_year=2001
65	select * from item where i_manufact_id between 770 and 770+40
66	select * from item where (i_category = 'Women' and (i_color = 'olive' or i_color = 'grey') and (i_units = 'Bundle' or i_units = 'Cup') and (i_size = 'petite' or i_size = 'small')) or (i_category = 'Women' and (i_color = 'hot' or i_color = 'thistle') and (i_units = 'Box' or i_units = 'Each') and (i_size = 'large' or i_size = 'medium')) or (i_category = 'Men' and (i_color = 'chiffon' or i_color = 'yellow') and (i_units = 'Carton' or i_units = 'Dozen') and (i_size = 'NA' or i_size = 'extra_large')) or (i_category = 'Men' and (i_color = 'bisque' or i_color = 'turquoise') and (i_units = 'Case' or i_units = 'Tsp') and (i_size = 'petite' or i_size = 'small'))
67	select * from item

	where (i_category = 'Women' and (i_color = 'chocolate' or i_color = 'lemon') and (i_units = 'Unknown' or i_units = 'Oz') and (i_size = 'petite' or i_size = 'small')) or (i_category = 'Women' and (i_color = 'light' or i_color = 'ivory') and (i_units = 'Ounce' or i_units = 'Ton') and (i_size = 'large' or i_size = 'medium')) or (i_category = 'Men' and (i_color = 'rose' or i_color = 'sandy') and (i_units = 'Pound' or i_units = 'Lb') and (i_size = 'NA' or i_size = 'extra_large')) or (i_category = 'Men' and (i_color = 'wheat' or i_color = 'burnished') and (i_units = 'Dram' or i_units = 'Pallet') and (i_size = 'petite' or i_size = 'small'))
68	select * from store_sales, date_dim where ss_sold_date_sk = d_date_sk and d_qoy = 1 and d_year = 2001
69	select * FROM customer_address, customer where ca_address_sk = c_current_addr_sk and c_preferred_cust_flag='Y'
70	select * from web_sales , item , date_dim where ws_item_sk = i_item_sk and i_category in ('Children', 'Sports', 'Music') and ws_sold_date_sk = d_date_sk and d_date between '2002-04-01' and '2002-05-01'
71	select * from catalog_sales, item, date_dim where cs_item_sk = i_item_sk and i_category in ('Shoes', 'Women', 'Music') and cs_sold_date_sk = d_date_sk and d_date between '1999-06-03' and '1999-07-03'
72	select * from store_sales, household_demographics, time_dim, store where ss_sold_time_sk = time_dim.t_time_sk and ss_hdemo_sk = household_demographics.hd_demo_sk and ss_store_sk = s_store_sk and time_dim.t_hour = 8

	and time_dim.t_minute >= 30 and ((household_demographics.hd_dep_count = 4 and household_demographics.hd_vehicle_count<=4+2) or (household_demographics.hd_dep_count = -1 and household_demographics.hd_vehicle_count<=-1+2) or (household_demographics.hd_dep_count = 0 and household_demographics.hd_vehicle_count<=0+2)) and store.s_store_name = 'ese'
73	select * from store_sales, household_demographics, time_dim, store where ss_sold_time_sk = time_dim.t_time_sk and ss_hdemo_sk = household_demographics.hd_demo_sk and ss_store_sk = s_store_sk and time_dim.t_hour = 9 and time_dim.t_minute < 30 and ((household_demographics.hd_dep_count = 4 and household_demographics.hd_vehicle_count<=4+2) or (household_demographics.hd_dep_count = -1 and household_demographics.hd_vehicle_count<=-1+2) or (household_demographics.hd_dep_count = 0 and household_demographics.hd_vehicle_count<=0+2)) and store.s_store_name = 'ese'
74	select * from store_sales, household_demographics, time_dim, store where ss_sold_time_sk = time_dim.t_time_sk and ss_hdemo_sk = household_demographics.hd_demo_sk and ss_store_sk = s_store_sk and time_dim.t_hour = 9 and time_dim.t_minute >= 30 and ((household_demographics.hd_dep_count = 4 and household_demographics.hd_vehicle_count<=4+2) or (household_demographics.hd_dep_count = -1 and household_demographics.hd_vehicle_count<=-1+2) or (household_demographics.hd_dep_count = 0 and household_demographics.hd_vehicle_count<=0+2)) and store.s_store_name = 'ese'
75	select * from store_sales, household_demographics, time_dim, store where ss_sold_time_sk = time_dim.t_time_sk and ss_hdemo_sk = household_demographics.hd_demo_sk and ss_store_sk = s_store_sk and time_dim.t_hour = 10 and time_dim.t_minute < 30 and ((household_demographics.hd_dep_count = 4 and household_demographics.hd_vehicle_count<=4+2) or (household_demographics.hd_dep_count = -1 and household_demographics.hd_vehicle_count<=-1+2) or (household_demographics.hd_dep_count = 0 and household_demographics.hd_vehicle_count<=0+2)) and store.s_store_name = 'ese'
76	select * from store_sales, household_demographics, time_dim, store where ss_sold_time_sk = time_dim.t_time_sk and ss_hdemo_sk = household_demographics.hd_demo_sk

	and ss_store_sk = s_store_sk and time_dim.t_hour = 10 and time_dim.t_minute >= 30 and ((household_demographics.hd_dep_count = 4 and household_demographics.hd_vehicle_count<=4+2) or (household_demographics.hd_dep_count = -1 and household_demographics.hd_vehicle_count<=-1+2) or (household_demographics.hd_dep_count = 0 and household_demographics.hd_vehicle_count<=0+2)) and store.s_store_name = 'ese'
77	select * from store_sales, household_demographics, time_dim, store where ss_sold_time_sk = time_dim.t_time_sk and ss_hdemo_sk = household_demographics.hd_demo_sk and ss_store_sk = s_store_sk and time_dim.t_hour = 11 and time_dim.t_minute < 30 and ((household_demographics.hd_dep_count = 4 and household_demographics.hd_vehicle_count<=4+2) or (household_demographics.hd_dep_count = -1 and household_demographics.hd_vehicle_count<=-1+2) or (household_demographics.hd_dep_count = 0 and household_demographics.hd_vehicle_count<=0+2)) and store.s_store_name = 'ese'
78	select * from store_sales, household_demographics, time_dim, store where ss_sold_time_sk = time_dim.t_time_sk and ss_hdemo_sk = household_demographics.hd_demo_sk and ss_store_sk = s_store_sk and time_dim.t_hour = 11 and time_dim.t_minute >= 30 and ((household_demographics.hd_dep_count = 4 and household_demographics.hd_vehicle_count<=4+2) or (household_demographics.hd_dep_count = -1 and household_demographics.hd_vehicle_count<=-1+2) or (household_demographics.hd_dep_count = 0 and household_demographics.hd_vehicle_count<=0+2)) and store.s_store_name = 'ese'
79	select * from store_sales, household_demographics, time_dim, store where ss_sold_time_sk = time_dim.t_time_sk and ss_hdemo_sk = household_demographics.hd_demo_sk and ss_store_sk = s_store_sk and time_dim.t_hour = 12 and time_dim.t_minute < 30 and ((household_demographics.hd_dep_count = 4 and household_demographics.hd_vehicle_count<=4+2) or (household_demographics.hd_dep_count = -1 and household_demographics.hd_vehicle_count<=-1+2) or (household_demographics.hd_dep_count = 0 and household_demographics.hd_vehicle_count<=0+2)) and store.s_store_name = 'ese'
80	select * from store_sales, store, promotion, date_dim, item

	where ss_sold_date_sk = d_date_sk and ss_store_sk = s_store_sk and ss_promo_sk = p_promo_sk and ss_item_sk = i_item_sk and i_category = 'Electronics' and (p_channel_dmail = 'Y' or p_channel_email = 'Y' or p_channel_tv = 'Y') and s_gmt_offset = -5 and d_year = 2000 and d_moy = 11
81	select * from customer, customer_address where ca_address_sk = c_current_addr_sk and ca_gmt_offset = -6
82	select * from store_sales, date_dim, store, household_demographics where store_sales.ss_sold_date_sk = date_dim.d_date_sk and store_sales.ss_store_sk = store.s_store_sk and store_sales.ss_hdemo_sk = household_demographics.hd_demo_sk and date_dim.d_dom between 1 and 2 and (household_demographics.hd_dep_count = 9 or household_demographics.hd_vehicle_count= 3) and date_dim.d_year in (1998, 1998+1, 1998+2) and store.s_city in ('Midway', 'Fairview')
83	select * from catalog_sales, date_dim where d_month_seq between 1183 and 1183 + 11 and cs_ship_date_sk = d_date_sk
84	select * from catalog_sales, customer_demographics cd1, customer, customer_address, date_dim, item where cs_sold_date_sk = d_date_sk and cs_item_sk = i_item_sk and cs_bill_cdemo_sk = cd1.cd_demo_sk and cs_bill_customer_sk = c_customer_sk and cd1.cd_gender = 'F' and cd1.cd_education_status = '4 yr Degree' and c_current_addr_sk = ca_address_sk and c_birth_month in (6, 5, 12, 4, 3, 7) and d_year = 2001 and ca_state in ('TN', 'IL', 'GA', 'MO', 'CO', 'OH', 'NM')
85	select * from store_sales, date_dim, store, household_demographics where store_sales.ss_sold_date_sk = date_dim.d_date_sk and store_sales.ss_store_sk = store.s_store_sk and store_sales.ss_hdemo_sk = household_demographics.hd_demo_sk and (household_demographics.hd_dep_count = 1 or household_demographics.hd_vehicle_count > -1) and date_dim.d_dow = 1 and date_dim.d_year in (1999, 1999+1, 1999+2) and store.s_number_employees between 200 and 295
86	select * from item, store_sales, date_dim where ss_item_sk = i_item_sk and ss_sold_date_sk = d_date_sk

	and d_month_seq in (1195, 1195+1, 1195+2, 1195+3, 1195+4, 1195+5, 1195+6, 1195+7, 1195+8, 1195+9, 1195+10, 1195+11) and ((i_category in ('Books', 'Children', 'Electronics') and i_class in ('personal', 'portable', 'reference', 'self-help') and i_brand in ('scholaramalgamalg #14', 'scholaramalgamalg #7', 'exportiunivamalg #9', 'scholaramalgamalg #9')) or (i_category in ('Women', 'Music', 'Men') and i_class in ('accessories', 'classical', 'fragrances', 'pants') and i_brand in ('amalgimporto #1', 'edu packscholar #1', 'exportiimporto #1', 'importoamalg #1'))))
87	select * from catalog_sales, customer_demographics, date_dim, item, promotion where cs_sold_date_sk = d_date_sk and cs_item_sk = i_item_sk and cs_bill_cdemo_sk = cd_demo_sk and cs_promo_sk = p_promo_sk and cd_gender = 'M' and cd_marital_status = 'D' and cd_education_status = 'College' and (p_channel_email = 'N' or p_channel_event = 'N') and d_year = 2001
88	select i_item_id from item where i_color in ('almond', 'dodger', 'dim')
89	select * from store_sales, date_dim, customer_address, item where ss_item_sk = i_item_sk and ss_sold_date_sk = d_date_sk and d_year = 1998 and d_moy = 1 and ss_addr_sk = ca_address_sk and ca_gmt_offset = -5
90	select * from catalog_sales, date_dim, customer_address, item where cs_item_sk = i_item_sk and cs_sold_date_sk = d_date_sk and d_year = 1998 and d_moy = 1 and cs_bill_addr_sk = ca_address_sk and ca_gmt_offset = -5
91	select * from web_sales, date_dim, customer_address, item where ws_item_sk = i_item_sk and ws_sold_date_sk = d_date_sk and d_year = 1998 and d_moy = 1 and ws_bill_addr_sk = ca_address_sk and ca_gmt_offset = -5
92	select * from date_dim, store_sales, item where d_date_sk = ss_sold_date_sk and ss_item_sk = i_item_sk and i_manager_id=40 and d_moy=12 and d_year=2001
93	select * from customer, customer_address, household_demographics, income_band

	where ca_city = 'Antioch' and c_current_addr_sk = ca_address_sk and ib_lower_bound >= 9901 and ib_upper_bound <= 9901 + 50000 and ib_income_band_sk = hd_income_band_sk and hd_demo_sk = c_current_hdemo_sk
94	select * from store_sales, date_dim, store, household_demographics where store_sales.ss_sold_date_sk = date_dim.d_date_sk and store_sales.ss_store_sk = store.s_store_sk and store_sales.ss_hdemo_sk = household_demographics.hd_demo_sk and date_dim.d_dom between 1 and 2 and (household_demographics.hd_buy_potential = 'gt10000' or household_demographics.hd_buy_potential = 'from5001to10000') and household_demographics.hd_vehicle_count > 0 and date_dim.d_year in (1998, 1998+1, 1998+2) and store.s_county in ('Williamson County')
95	select * from store_sales, date_dim where ss_sold_date_sk = d_date_sk and d_date between '1998-08-14' and '1998-09-14'
96	select * from store_returns, date_dim where sr_returned_date_sk = d_date_sk and d_date between '1998-08-14' and '1998-09-14'
97	select * from catalog_sales, date_dim where cs_sold_date_sk = d_date_sk and d_date between '1998-08-14' and '1998-09-14'
98	select * from catalog_returns, date_dim where cr_returned_date_sk = d_date_sk and d_date between '1998-08-14' and '1998-09-14'
99	select * from web_sales, date_dim where ws_sold_date_sk = d_date_sk and d_date between '1998-08-14' and '1998-09-14'
100	select * from web_returns, date_dim where wr_returned_date_sk = d_date_sk and d_date between '1998-08-14' and '1998-09-14'
101	select * from store_sales, date_dim where store_sales.ss_sold_date_sk = date_dim.d_date_sk and d_month_seq between 1214 and 1214+11
102	select * from catalog_sales, date_dim where catalog_sales.cs_sold_date_sk = date_dim.d_date_sk and d_month_seq between 1214 and 1214+11
103	select * from web_sales, date_dim where web_sales.ws_sold_date_sk = date_dim.d_date_sk and d_month_seq between 1214 and 1214+11
104	select * from store_sales, date_dim where ss_sold_date_sk = d_date_sk and d_year in (1998, 1998+1)
105	select * from web_sales, date_dim where ws_sold_date_sk = d_date_sk and d_year in (1998, 1998+1)
106	select * from web_returns, date_dim

	where wr_returned_date_sk = d_date_sk and d_year = 2000
107	select * from store_sales, date_dim, store, household_demographics where store_sales.ss_sold_date_sk = date_dim.d_date_sk and store_sales.ss_store_sk = store.s_store_sk and store_sales.ss_hdemo_sk = household_demographics.hd_demo_sk and (date_dim.d_dom between 1 and 3 or date_dim.d_dom between 25 and 28) and (household_demographics.hd_buy_potential = 'from500to1000' or household_demographics.hd_buy_potential = 'from0to500') and household_demographics.hd_vehicle_count > 0 and date_dim.d_year in (1998, 1998+1, 1998+2) and store.s_county in ('Williamson County')
108	select * from web_sales, date_dim where d_date_sk = ws_sold_date_sk and d_moy=12 and d_year=2002
109	select * from catalog_sales, date_dim where d_date_sk = cs_sold_date_sk and d_moy=12 and d_year=2002
110	select * from store_sales, date_dim where d_date_sk = ss_sold_date_sk and d_moy=12 and d_year=2002
111	select * from store_sales, date_dim where ss_sold_date_sk = d_date_sk and d_month_seq between 1223 and 1223+11
112	select * from item, catalog_sales, date_dim where cs_item_sk = i_item_sk and cs_sold_date_sk = d_date_sk and (d_year = 1999 or (d_year = 1999-1 and d_moy = 12) or (d_year = 1999+1 and d_moy = 1))
113	select * from store_sales, date_dim, item where ss_sold_date_sk = d_date_sk and ss_item_sk = i_item_sk and d_year in (1999, 1999+1, 1999+2, 1999+3)
114	select * from catalog_sales, date_dim where d_year = 1999 and d_moy = 6 and cs_sold_date_sk = d_date_sk
115	select * from web_sales, date_dim where d_year = 1999 and d_moy = 6 and ws_sold_date_sk = d_date_sk
116	select * from store_sales, item iss, date_dim d1 where ss_item_sk = iss.i_item_sk

	and ss_sold_date_sk = d1.d_date_sk and d1.d_year between 1999 and 1999 + 2
117	select * from catalog_sales, item ics, date_dim d2 where cs_item_sk = ics.i_item_sk and cs_sold_date_sk = d2.d_date_sk and d2.d_year between 1999 and 1999 + 2
118	select * from web_sales, item iws, date_dim d3 where ws_item_sk = iws.i_item_sk and ws_sold_date_sk = d3.d_date_sk and d3.d_year between 1999 and 1999 + 2
119	select * from store_sales, date_dim where ss_sold_date_sk = d_date_sk and d_year between 2000 and 2000 + 2
120	select * from catalog_sales, date_dim where cs_sold_date_sk = d_date_sk and d_year between 2000 and 2000 + 2
121	select * from web_sales, date_dim where ws_sold_date_sk = d_date_sk and d_year between 2000 and 2000 + 2
122	select * from store_sales, item, date_dim where ss_item_sk = i_item_sk and ss_sold_date_sk = d_date_sk and d_year = 2000+2 and d_moy = 11
123	select * from catalog_sales, item, date_dim where cs_item_sk = i_item_sk and cs_sold_date_sk = d_date_sk and d_year = 2000+2 and d_moy = 11
124	select * from web_sales, item, date_dim where ws_item_sk = i_item_sk and ws_sold_date_sk = d_date_sk and d_year = 2000+2 and d_moy = 11
125	select * from date_dim where d_year = 2000 + 1 and d_moy = 12 and d_dom = 16
126	select * from date_dim where d_year = 2000 and d_moy = 12 and d_dom = 16
127	select * from catalog_sales, customer_demographics, customer, customer_address, promotion where cs_bill_cdemo_sk = cd_demo_sk and cs_bill_customer_sk = c_customer_sk and cs_promo_sk = p_promo_sk and cd_gender = 'F' and cd_marital_status = 'D' and cd_education_status = '4 yr Degree' and c_current_addr_sk = ca_address_sk

	and c.birth_month in (6, 5, 12, 4, 3, 7) and (p_channel_email = 'N' or p_channel_event = 'N') and ca_state in ('TN', 'IL', 'GA', 'MO', 'CO', 'OH', 'NM')
128	select * from catalog_sales, customer_demographics, date_dim, item, promotion where cs_sold_date_sk = d_date_sk and cs_item_sk = i_item_sk and cs_bill_cdemo_sk = cd_demo_sk and cs_promo_sk = p_promo_sk and cd_gender = 'M' and cd_marital_status = 'M' and cd_education_status = '4 yr Degree' and (p_channel_email = 'N' or p_channel_event = 'N') and ((d_date between '2000-03-22' and '2000-06-22') or d_year = 2001)
129	select * from web_sales, customer_address, date_dim, time_dim , ship_mode where ws_sold_date_sk = d_date_sk and ws_sold_time_sk = t_time_sk and ws_ship_mode_sk = sm_ship_mode_sk and ws_ship_addr_sk = ca_address_sk and (d_year = 1999 or d_qoy = 2 or d_moy = 12) and t_time between 46185 and 46185+28800 and sm_carrier in ('DIAMOND', 'ZOUROS') and ca_gmt_offset = -6
130	select * from item where ((i_category = 'Women' and (i_color = 'olive' or i_color = 'grey') and (i_units = 'Bundle' or i_units = 'Cup') and (i_size = 'petite' or i_size = 'small')) or (i_category = 'Women' and (i_color = 'hot' or i_color = 'thistle') and (i_units = 'Box' or i_units = 'Each') and (i_size = 'large' or i_size = 'medium')) or (i_category = 'Men' and (i_color = 'chiffon' or i_color = 'yellow') and (i_units = 'Carton' or i_units = 'Dozen') and (i_size = 'NA' or i_size = 'extra_large')) or (i_category = 'Men'

	and (i_color = 'bisque' or i_color = 'turquoise') and (i_units = 'Case' or i_units = 'Tsp') and (i_size = 'petite' or i_size = 'small')))) or ((i_category in ('Books', 'Children', 'Electronics') and i_class in ('personal', 'portable', 'reference', 'self-help') and i_brand in ('scholaramalgamalg #14', 'scholaramalgamalg #7', 'exportiunivamalg #9', 'scholaramalgamalg #9')) or (i_category in ('Women', 'Music', 'Men') and i_class in ('accessories', 'classical', 'fragrances', 'pants') and i_brand in ('amalgimporto #1', 'edu packscholar #1', 'exportiimporto #1', 'importoa- malg #1')))) or (i_current_price between 0.99 and 1.49) or i_manufact_id in (744, 691, 853, 946) or i_manager_id=91
131	select * from store_sales, date_dim, store, household_demographics where store_sales.ss_sold_date_sk = date_dim.d_date_sk and store_sales.ss_store_sk = store.s_store_sk and store_sales.ss_hdemo_sk = household_demographics.hd_demo_sk and (household_demographics.hd_buy_potential = 'gt10000' or household_demographics.hd_buy_potential = 'from5001to10000') and (household_demographics.hd_vehicle_count > 0 or household_demographics.hd_dep_count = 5) and date_dim.d_dom between 1 and 2 and d_dow = 1 and (store.s_county in ('Williamson County') or store.s_city in ('Midway', 'Fairview') or store.s_number_employees between 200 and 295)