

DistGALO: Distributed System for Query Problem Determination

by

Alexandar Mihaylov

A thesis submitted to the
School of Graduate and Postdoctoral Studies in partial
fulfillment of the requirements for the degree of

Master of Science

in

Computer Science

Ontario Tech University
Oshawa, Ontario, Canada

December 2019

© Alexandar Mihaylov, 2019

Thesis Examination Information

Submitted by: **Alexandar Mihaylov**

Master of Science in Computer Science

Thesis title: DistGALO: Distributed System for Query Problem Determination

An oral defense of this thesis took place on December 5, 2019 in front of the following examining committee:

Examining Committee:

Chair of Examining Committee

Dr. Jeremy Bradbury

Research Supervisor

Dr. Jarek Szlichta

Examining Committee Member

Dr. Heidar Davoudi

Thesis Examiner

Dr. Pejman Mirza-Babaei

The above committee determined that the thesis is acceptable in form and content and that a satisfactory knowledge of the field covered by the thesis was demonstrated by the candidate during an oral examination. A signed copy of the Certificate of Approval is available from the School of Graduate and Postdoctoral Studies.

Abstract

Query optimization is a quintessential element of modern Database Management Systems(DBMSs). Compile-time driven estimates and heuristics aid the compiler in selecting what is deemed the lowest cost Access Plan for a given query. These access plans are seldom optimal, and can oftentimes lead to under-performing query runtimes, with varying severity. Traditionally, domain experts painstakingly examine the access plans to detect and fix problem patterns. DistGALO, the successor to the previous GALO system, was developed to remedy this manual labour by incorporating a cluster of nodes to learn problem patterns in a distributed fashion and apply the fixes automatically. Several partitioning and pruning strategies are employed, including the RSACE module which gives user fine-grained control for trading off runtime versus template creation. In the experimental validation, DistGALO demonstrates the efficiency boost over our previous system using the synthetic TPC-DS benchmark and the effectiveness of the various pruning strategies.

Keywords: distributed computing; partitioning; problem determination; database optimizer;

Co-authorship Statement

The contribution outlined in this work is part of a bigger ongoing project developed in collaboration with IBM. Two previous projects, OptImatch[1][2] and GALO[3][4], have built some of the foundations for the current iteration that I present here. I have previous involvement in the GALO development, but do not claim it as a contribution in this work. The *Matching Engine* and *Transformation Engine* were developed by Guilherme Damasio and only utilized since they are critical in the system's operation. The contribution presented in this work has been developed primarily by myself, with some assistance from the undergraduate student Spencer Bryson during his summer internship in the data science lab. His assistance includes help with the implementation of the distribution of the sub-query generation process (Section 3.3) and dynamic sampling rate. York University collaborator Parke Godfrey has also been involved in this work, by providing orientation, advice, and guidance throughout the whole process. The remainder of the development and evaluation discussed throughout is solely contributed by myself.

Author's Declaration

I hereby declare that this thesis consists of original work of which I have authored. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I authorize the University of Ontario Institute of Technology to lend this thesis to other institutions or individuals for the purpose of scholarly research. I further authorize University of Ontario Institute of Technology to reproduce this thesis by photocopying or by other means, in total or in part, at the request of other institutions or individuals for the purpose of scholarly research. I understand that my thesis will be made electronically available to the public.

Alexandar Mihaylov

Statement of Contributions

All mention and work related to the OptImatch[1] and GALO[4] systems is not considered a contribution in the work presented here. Some modules from GALO, mainly the *Transformation Engine* and *Matching Engine* remain unchanged, but are utilized in DistGALO for system operational purposes. In addition, the GALO *ranker* and *sub-query generator*, excluding the distributed sub-query generation addition in this work, are also utilized with little to no modifications. These aforementioned components though integral in DistGALO's operation, are not considered as a contribution throughout the work presented here.

The development process of DistGALO has in large part been done in IBM's *Centre for Advanced Studies*¹ in Markham, Canada. The development of the system was in whole done by myself with some aid from an undergraduate research student, described in the *Co-authorship Statement*. IBM query compiler domain experts have provided me with guidance and advice throughout the development.

This work is subject to publication at a future date but is not currently submitted, and the conference remains unknown.

¹<https://www.research.ibm.com/university/cas/>

Acknowledgements

In the first place, I would like to acknowledge my parents, Lubomir and Reneta, my sister Iva, and my girlfriend, Luisa, for providing me with the encouragement and assurance throughout my university journey. You have showered me with unconditional love and support during this path that I would otherwise not be on without you. Thank you from the bottom of my heart. I am indebted to my supervisor Jarek Szlichta for the continual guidance during my undergraduate and graduate career both as a professor and mentor. Special thank you to Parke Godfrey, collaborator from York University, for his input on every pivotal decision in this project. I would also like to thank IBM collaborators, Calisto Zuzarte, Vincent Corvinelli and Piotr Mierzejewski, for the insight and technical knowledge that has resulted in countless hours spared on my end. My sincere gratitude goes out to the University of Ontario Institute of Technology and the entire faculty, who have provided the tools and environment to make all my ambitions come to light. Finally, to all my friends that have stuck by me throughout the years, you have grown to be an integral part of my life, and know that I cherish each and every one of you.

Contents

Abstract	iii
Co-authorship Statement	iv
Author’s Declaration	v
Statement of Contributions	vi
Acknowledgements	vii
Table of Contents	1
List of Tables	1
List of Figures	1
List of Abbreviations and Symbols	2
1 Introduction	4
1.1 Motivation	4
1.2 Goals	6
1.3 Real World Example	11

1.4	Contributions	13
2	Background	17
2.1	Preliminaries	17
2.2	System Overview	25
3	Distributed Learning	29
3.1	Cluster and Spark Configuration	31
3.2	Partitioning	34
3.3	Sub-query Generation	39
3.4	Optimizer Sub-query Execution	41
3.5	Random Sub-query Execution	43
3.6	RSACE	50
4	Experimental Evaluation	56
4.1	Distributed Learning Engine Evaluation	57
4.2	Distributed Performance	61
4.3	Pruning Effectiveness	64
5	Related Work	67
6	Conclusion	70

List of Tables

3.1	Apache Spark Cluster Configuration	33
3.2	Sub-query post execution ranking properties considered.	43

List of Figures

1.1	IBM customer problem query and applied fix.	11
2.1	Problem pattern with expensive hash join and DistGALO discovered fix.	20
2.2	Plan optimization guidelines applied to access plan in Figure 2.1b. . .	23
2.3	System architecture of DistGALO's Distributed Learning Engine. . .	26
3.1	Problem pattern with expensive nested-loop join and DistGALO dis- covered fix.	30
3.2	Partitioning process of random sub-queries using one of three available partitioners.	45
3.3	Random Sub-query Abridged Clustering Executor (RSACE) module.	50
3.4	Random sub-query vectorization and grouping in RSACE.	54
4.1	DistGALO's learning elapsed time (seconds) versus DistGALO using TPC-DS benchmark queries.	58
4.2	Distributed Learning with 1G and 10G TPC-DS databases.	59
4.3	Random sub-query execution with different number executors.	61

4.4	Random sub-query execution using Hash, Cost Estimate, and Runtime Partitioners.	62
4.5	RSACE random sub-query reduction and its effect on number of templates found.	65

List of Abbreviations and Symbols

API	Application Programming Interface.
DBA	Database Administrator.
DBMS	Database Management System.
DBSCAN	Density Based Spatial Clustering of Applications with Noise.
DLC	Distributed Learning Component.
FETCH	-IXSCAN Fetch Index Scan.
GALO	Guided Automated Learning for query workload re-Optimization.
GUI	Graphical User Interface.
HDFS	Hadoop Distributed File System.
HSJOIN	Hash Join.
IaaS	Infrastructure as a Service.
IBM	International Business Machines Corporation.
IoT	Internet Of Things.
IXSCAN	Index Scan.
KB	Knowledge Base.
LOLEPOP	LOW LEVEL Plan OPERators.
MSJOIN	Merge Join.

NLJOINNested Loop Join.
QEP Query Execution Plan.
QGM Query Graph Model.
RDD Resilient Distributed Datasets.
RDF Resource Script Framework.
RSACE Random Sub-query Abridged Clustering Executor.
SPARQL SPARQL Protocol and RDF Query Language.
SQL Structured Query Language.
TBSCAN Table Scan.
URI Uniform Resource Identifier.
W3C World Wide Web Consortium.
XML Extensible Markup Language.
Φ Optimizer Sub-queries.
Ψ Random Sub-queries.

Chapter 1

Introduction

1.1 Motivation

The volume of data being stored and processed is typically domain dependant, but recently there's been an immeasurable increase across all domains [5][6]. This trend can be attributed to various factors and technologies, including the widespread availability of IaaS (Infrastructure as a Service), Internet Of Things (IoT), smart devices, and consumer demand for global connectivity. The increase in data posed many challenges for database management systems and their vendors. One consequence was that SQL queries became more complex since middleware tools became capable of automatically generating them [7]. The convenience these tools brought from a user perspective, they took away from a problem determination one. This increase in complexity has led to queries with essentially no limit on the number of algebraic operators they contained, potentially spanning hundreds of lines. This though indifferent to database users has proven increasingly more difficult for domain experts conducting troubleshooting over the queries.

From an automation standpoint, database query optimizers absolve some of the manual labour experts would otherwise have to incur. The query optimization process has been researched thoroughly[8][9] and new techniques are continually being developed[10]. The underlying process relies heavily on cardinality estimates to determine a cost model for varying access plans that express the path of execution. A varying set of access plans are considered, with alternate join orders and operators, with the lowest cost plan ultimately being selected for execution. The selected and optimal access plan would ideally be identical, but is not always the case, since there is a large dependence on estimates and timely expectancy. When cardinality estimations prove inaccurate, the optimizer tends to pick sub-optimal access plans, and under-performance may arise.

In effort to aid database administrators (DBAs), several tools have been developed[11][12] by providing suggestions to the optimizer as to what decisions(operators) are to be chosen in the final access plan. These are denoted as *pragma* in the Oracle Database, and *hints* by Microsoft SQL Server. These suggestions can be embedded into the SQL but may become dated over time as the database data and subsequently, statistics change. DB2 offers an XML based *guideline document* that provides recommendations to the optimizer during the cost-based optimization stage. This document can be submitted alongside the query and provide seamless alteration from a user perspective.

It is at this stage that domain experts must intervene to address the under-performance, by manually analyzing and adjusting the access plan through hints, pragmas, or guidelines, depending on the database vendor. This process, however, is quite laborious and requires a high degree of domain expertise. More so, with the

increasing complexity of automatically generated SQL queries, the manual process is starting to prove unbecoming. The manual problem determination is also done *ad hoc*, and thus recurring patterns are forgotten and redundant work must be performed.

Domain experts are scarce resources and can seldom devote time to the arduous troubleshooting. A more automated approach is demanded, but recent approaches to do so failed to adjust a miss-performing access plan[13][14]. In response, a semi-automated approach OpImatch [1][2] was devised. This approach allowed domain experts to graphically create problem patterns, to later save in a *knowledge base*, and share with other domain experts. The system proved to aid the troubleshooting, but still required some manual input from experts in order to populate the knowledge base. As its successor, GALO[4][3] was developed as a fully automated solution. The system required no domain expert intervention and functioned in a two-step process. First the *Learning Phase* was responsible for learning problem patterns offline, and populate a knowledge base. The *Matching Phase* would then query the knowledge base online and automatically apply access plan repair on problem workloads. The system was well received, and in this work I extend it to DistGALO, a distributed modern-day solution to the growing demand for modern-day scalable systems.

1.2 Goals

The goals for DistGALO subsume the goals of the previous GALO and OptImatch systems, but also extend them as follows:

1. *automatic query problem determination*;
2. *query re-optimization*;
3. *optimization evolution*; and

4. scalability.

DistGALO's goal 1 is in direct response to the manual effort optimizer experts exert in finding problems within large query access plans. To address this, OptImatch was first devised so as to provide some relief in the process by allowing experts to create, save, and apply recurring access plan problem patterns from a shared knowledge base. We further enhanced this process with GALO, a system capable of fully automating the knowledge base population stage, thus eliminating any need for manual expert intervention. The knowledge base was composed of Resource Description Framework(RDF) graph templates and later applied to ill-performing SQL queries in order to improve performance. RDF provides a natural mapping to graph-based representations since execution plans are also graphs, and the RDF representation could be queried using the SPARQL language. The problem patterns, or templates saved in the knowledge base are table and attribute independent, using the selected predicates' cardinalities to create a range which can later be matched. These ranges are created for every LOLEPOP and henceforth any QEPs that fit within that range and structure of the template, will be applied during the matching phase.

Queries must undergo several steps before an access plan is executed, and results are returned to the user. One of these includes the query re-write, where the graphical representation of the query is mapped to a semantically equivalent, but more efficient graph [15]. The graph-based representation is then passed onto the optimizer, where various *Query Execution Plan(QEP)s* are generated, each with an associated estimated cost. Ultimately the QEP with the lowest cost is selected and goes on to be executed by the database manager. QEP costs are estimates calculated at compile-time, and may not always adequately represent the effectiveness of the access plan.

Inhere lies the problem, and is the main motive behind GALO's goal 1. Sub-optimal access plans are typically selected due to the misguided estimated costs, or because the optimal QEP was never considered from the large search space of plans. GALO was able to address this issue by analyzing a large number of query fragments, or sub-queries, to create rules of problem patterns then used to populate its knowledge base. It did so quite successfully, but it came with some inherent shortcomings. The search space of sub-queries generated proved to be quite large, and with a lack of pruning rules, sequentially processing each proved to be a time-consuming process. This deficiency we address in our current work with DistGALO, by distributing the most costly and expensive stage of the process. Also, several pruning rules allow for an even greater speedup by reducing the search space. As a whole, the goal remains the same, but the methodology to achieve it has been greatly improved upon.

We have thus far explored two major steps in the compilation process, mainly the query rewrite, and the cost-based optimization stages. Goal 2 provides a third-tier absent from the compiler that aims to rewrite the final QEP chosen by the optimizer. This is because the chosen QEP is sub-optimal, possibly exhibiting a higher estimated cost than the actual, or a better performing QEP' exists, but remains undiscovered. Stored rules within the knowledge base can be queried and applied to transform an under-performing QEP, to a QEP' that is unhindered by poorly selected operators that experts might otherwise toil over. All the acquired wisdom is aggregated within a single knowledge base, rather than throughout numerous experts, who can only rely on their ad hoc observations to trace the problematic section.

Upon locating a problematic segment, one approach would be to apply the matched rewrites directly to the QEP supplied by the optimizer, but this could ultimately re-

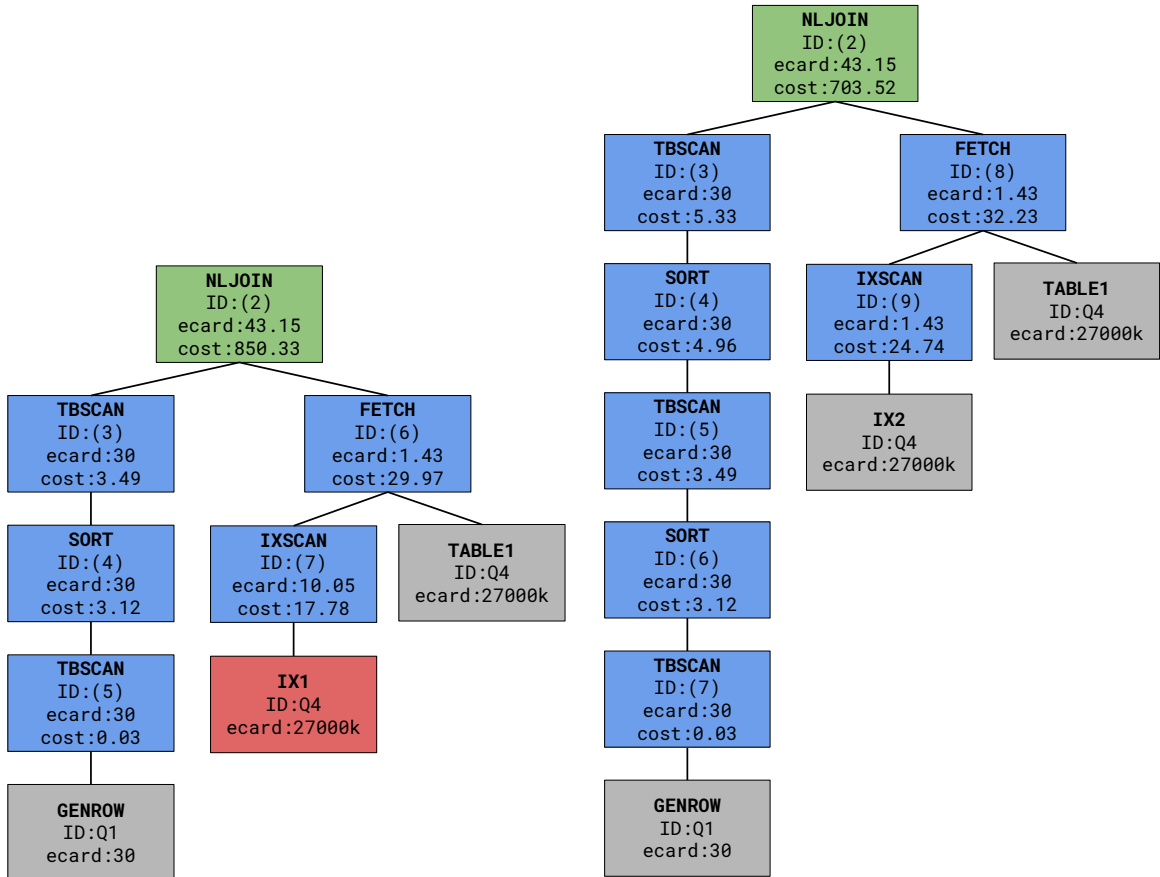
sult in incompatibilities in the overall QEP. Instead, DistGALO generates a guideline document that can be directly appended to the original query, without any user intervention. The query can then be passed through the compiler’s pipeline, including the query rewrite and cost-based optimization stages, ensuring that all intermediate LOLEPOPs are re-evaluated and correctly estimated. The optimizer may then chose to honor the guideline, or in some cases discard it due to incompatibility issues. Ultimately if the guideline proves to be applicable, a potentially never-before-seen QEP’ may be selected, thus providing a faster execution time. This process we term *plan re-optimization*.

Goal 3 persists through GALO and DistGALO as they share many fundamental values. Both provide a valuable supply of problem patterns that can be analyzed by the performance optimization team. This analysis could potentially lead to uncovering of unknown issues lying within the optimizer, enhancing the query rewrite rules, or even applying new heuristics in selecting a more cost-optimal QEP. These improvements are not constrained to academic and benchmark synthetic data, but also can be applied to real-world customer workloads.

Goal 4 is unique to DistGALO, and was the primary motive to bring the system up to par with the vastly growing complexity of queries, and the size of data. We observed a limitation in GALO’s ability to conform to more complex queries and larger datasets. This, though to some degree can be remedied by vertical scaling of the server, was not an acceptable solution. What DistGALO offers instead is a way to horizontally scale in response to the workload complexity and size. Given the maturity of the modern-day cloud systems and their abundant availability, it is not unreasonable to rent nodes on a per need basis. If a larger workload is required,

a business or even an individual is quite reasonably able to acquire more hardware for the duration needed. This task is not only achievable in modern development environments, but can even be automated[16][17]. This can allow for more optimized utilization of resources since they are now also able to scale in accordance with the work required. DistGALO aims to adhere to this growing trend by allowing horizontal scaling in response to the query and workload supplied. It utilizes a cluster of nodes that can very easily be altered, and improve performance if the use case calls for a faster response time.

Two varying degrees of complexity cause a drastic increase in search space for GALO. First, the complexity of the queries being learned causes a larger number of operator combinations, consequently leading to a larger number of sub-queries that are generated. This results in more sub-queries that would require analysis and execution, thus causing a slowdown. The other scenario is one in which the referenced tables are larger, once again causing a slowdown in the execution and probing stages of the system. DistGALO tackles both of these by incorporating multiple machines working harmoniously together, each intelligently delegated its own workload. The individual workloads ensure an overall balance and skew reduction in the system that ultimately fractions the learning runtime. Another objective was to ensure that DistGALO responded well to vertical scaling as well. Pruning rules are presented to drastically reduce the search space, whilst still providing meaningful results. The ability to scale according to the demands of the user or the workload makes DistGALO a very applicable system to utilize in real-world applications.



(a) Plan selected by the optimizer.

(b) Plan chosen by the DistGALO system.

Figure 1.1: IBM customer problem query and applied fix.

1.3 Real World Example

Consider a customer who experiences a slowdown in their daily system metric reports as a consequence of a slow query bottlenecking the whole pipeline. What database vendors, and subsequently domain experts, first have to do is to analyze a problematic query execution plan, like the one in Figure 1.1a. This already arduous task is made even more challenging when considering that the QEP shown is only a subgraph from a much larger access plan (not shown for brevity and to highlight the problem section).

Even the act of identifying the subgraph is a challenging and time-consuming process that, without tools like DistGALO or GALO, must be done manually. Each node in the subgraph can be either a table, index, table functions, or operators such as join methods and scans. Each node we refer to as a low level plan operator (LOLEPOP) hereafter. Each LOLEPOP contains several pieces of information, but for simplicity have only shown a subset of the most relevant ones. Starting from the top, the operator type (ex. NLJOIN, TBSCAN), followed by the ID of the LOLEPOP, unique to each node, typically used for reference. Third down is the optimizer calculated estimated cardinality or *ecard* for short. Lastly, at the very bottom of the LOLEPOP is the optimizer calculated *cost* in timeron units, the estimated total I/O and CPU cost the db2 manager might incur during execution. Values for *ecard* and *cost* might include a k multiplier of 1000 for brevity. Note the leaf nodes are either base tables or indexes and do not follow the formatting scheme outlined above. Rather than the operator name, we denote the name of the table or index being read from. This is followed by the instance name, the table/index referenced by the compiler post-query-rewrite. Finally, at the bottom we include the estimated cardinality as for all LOLEPOPS.

As a concrete example, consider the **NLJOIN** LOLEPOP in Figure 1.1a, which has an ID #2, estimated cardinality of 43 and cost of 850 timerons. Next consider the base table **TABLE1** at the bottom-right-most. We uniquely identify it by its index name *IX1*, which also has a cardinality of $27000k$ or 27 million. Henceforth we will refer to specific LOLEPOPs either by their unique ID or table/index instance name, depending on the type.

The underlying under-performance of the optimizer selected plan in Figure 1.1a

is mainly attributed to the choice of index. The **IXSCAN (#7)** scans a large index *IX1*, using the RIDs (row ids) for the following **FETCH (#6)** above. The fix for the under-performance is a subtle one, attained by swapping the index used in the **IXSCAN (#7)** for a larger index *IX2* (with one more level in the B-Tree). The larger index (*IX2* in Figure 1.1b) is typically more expensive since it needs one extra I/O with each probe in the index. Since that is the only variant in the plan, the optimizer opted for the cheaper index *IX1*. The subtlety, however, lies in the values of the additional attributes present in *IX2* but not *IX1*. These post analysis, turn out to have a large number of distinct values, and predicates on these attributes qualify fewer rows that need to be fetched. The optimizer ultimately selected the cheaper inner (#6 in Figure 1.1a) to minimize the time spent compiling and to reduce the search space. This minute detail leads to an order of magnitude difference in runtime performance, and thus stresses the importance and need for domain experts and automated systems like DistGALO to improve the performance.

1.4 Contributions

Scalability is a critical implicit requirement for modern-day systems and remained the focus when developing DistGALO. The new system improves upon the previous GALO system, whilst still upholding the requirement of providing a third *plan rewrite* optimization stage. We performed a detailed analysis on a per-module basis on GALO and uncovered some bottlenecks. In this work, we have improved upon those bottlenecks by improving or overhauling components of the previous system. DistGALO aims to comply with the demand for scalable systems capable of withstanding the progressively more complex queries and larger databases GALO’s *Transformation*

Engine, *Matching Engine*, and more specifically, *query ranker* and *sub-query generation* (not including the distribution) are all utilized but not considered contributions in this work. All elements relating to distributed computation, including partitioning, parallelization of executors, distributed file system utilization and pruning modules are exclusive to DistGALO. We present our contributions in greater detail as follows:

1. **Distributed Learning Component (DLC)**. First is the *Distributed Learning Component (DLC)* that transforms GALO’s Learning Engine, into the Distributed Learning Engine presented in this work. Solely it is a cohesive system, but can be summarized with three of its largest components including Distributed Sub-query Generator, Query Partitioner, and the Distributed Sub-query Executor.

(a) The *Distributed Sub-query Generator* decomposes workload queries into smaller components or sub-queries. The size of the sub-query depends on the number of joins, and thus a range can be established in the configuration of the system. The process of decomposing the queries is done so in a distributed manner, using heuristics to minimize any load imbalance that may arise. The generated sub-queries can be grouped into two different sets, the optimizer sub-queries, and random sub-queries. These move down the pipeline and are processed separately due to their nature.

(b) The *Query Partitioner* component, is used to mitigate skew throughout the cluster using various strategies. This is a critical component of the distributed environment since it determines how data should be grouped on nodes. We employ three different strategies: hash, cost estimate, and runtime partitioning. Each is utilized in some part of the system since some are dependant on metadata that may or may not be available at a given point in the pipeline.

The choice of partitioning strategies has some flexibility, and where present, can be adjusted through the system configuration file.

- (c) The *Distributed Sub-query Executor* is responsible for taking in already-partitioned sub-queries and processing them on all available nodes. Sub-queries are executed a predetermined amount of times (in effort to reduce noise), metadata from each run is collected and the best performing run is saved into the Distributed File System for later processing. This component processes both the optimizer and random sub-queries separately. Optimizer sub-query runtime information is saved in a lookup table that is later used to aid the partitioner in grouping random sub-queries, and to also provide valuable threshold information to the *distributed sub-query executor*.

2. **Random Sub-query Abridged Clustering Executor.** Second is the *random sub-query abridged clustering executor* (RSACE) which adds a pruning layer to minimize the processing of redundant sub-queries. This module only applies to the random sub-queries since they pose the biggest bottleneck, and also because it relies on some meta-information from the optimizer sub-query execution stage. Clusters of alike sub-queries are created, whereby only representatives from each are executed. Further execution by the *distributed sub-query executor* follows this process in effort to eliminate any false positive representatives.
3. **Experiments.** Finally, we present *experiments* to validate the effectiveness of DistGALO, and its efficiency over its predecessor system, GALO. The effectiveness is measured over the TPC-DS decision support benchmark with synthetically generated data. We demonstrate that DistGALO is just as effective, but is also able to perform the learning process in much less time. We also quantify the

performance of the varying partitioners and the RSACE module.

In Section 2 we describe some of the preliminaries, and provide an overview of the DistGALO system. In Section 3 we detail over the distributed learning process and the varying components involved. In Section 4 we validate experimentally the effectiveness and efficiency of DistGALO, followed by related works in Section 5, finally concluding our work in Section 6.

Chapter 2

Background

2.1 Preliminaries

GALO. There is an inherent problem in which database experts were exerting too much effort and time into finding fixes for problematic customer queries. To aid the laborious task of examining queries to only discover a previously seen problem pattern, OptImatch[1] was developed. It allowed experts to interact with a web interface and input recurring problem patterns they would come across. This proved an invaluable tool but still required some expert intervention. In effort to fully automate the process, GALO[4] was developed. It was a complex system primarily comprised of a *transformation engine*, *learning engine* and *matching engine* that would all interact with a Knowledge Base(KB). The KB would be the central hub of the system, populated by the *learning engine*, and queried by the *matching engine*, using the *transformation engine* as an intermediary. The *learning engine* would profile workloads *offline*, and capture problem patterns that it would then save in a graph-based representation within the KB. A workload refers to a set of SQL queries that require

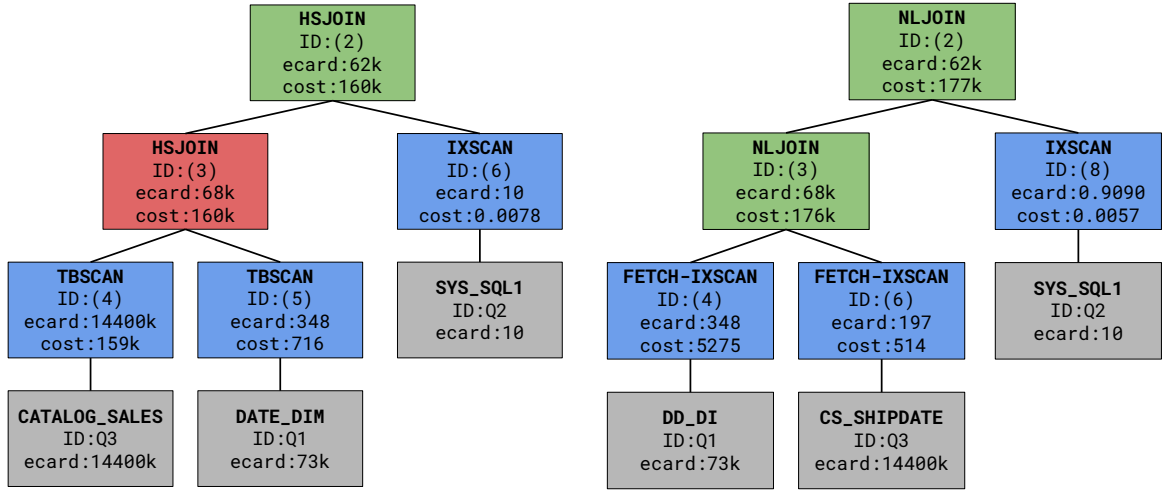
some periodic execution. The *matching engine* processes a separate workload *online* whilst querying the KB for potential plan rewrites overseen by the optimizer. This enables the optimizer to use GALO as the *third tier of optimization*, be re-writing access plans using templates from its previously populated KB. The *transformation engine* acts as a translator between the language the databases uses, and the language of the knowledge base. The decision tree the optimizer selects as the path of execution it must take during runtime is represented as a graph. The KB is also built using a graph based-representation, but using a different format and thus needs to be translated by the *transformation engine*.

Query Compiler. The runtime of queries is a quintessential in the successful operation of some applications. In some cases, hundreds or even thousands of queries may be requested within a very small time period, and so retrieving the results promptly becomes even more critical. The query compiler optimizer aims to aid this process through the use of heuristics and statistics. The query must pass through several crucial steps before the optimal access plan is found and executed. Under the hood, the query is represented by a Query Graph Model (QGM) as an in-memory database. The query must first be checked to be syntactically and semantically sound, to ensure that the user did not make an error. The query is then transformed into an easier to optimize format in the query rewrite stage. Tables can be renamed here, and predicates can be pushed down to other levels of execution to improve performance. The QGM is updated and, with the help of various statistics from tables, indexes, columns, and functions, it generates various execution plans, each with an associated estimated cost. Ultimately the execution plan with the lowest cost is selected and used during runtime. As a final step, the compiler creates the executable access plan

for the query ensuring it is devoid of any redundant computations.

DB2 provides an explain facility to shed some light on some of the steps outlined above and provide greater insight into decisions to be made at runtime. The explain facility can provide detailed information on what tables and indexes were accessed, cost information, statistics for all referenced objects, as well as predicates and selectivity estimates for each, to name a few. Also, it also captures the sequence of operations the optimizer selected to process the query. This sequence of operations can also be represented as a graph and is termed the Query Execution Plan (QEP). This is, in essence, a data-flow graph of operators, where edges are the flow of the data itself, and the nodes are operations like joins or sorts[15]. Note that the QEP is available at compile-time, and as such only presents estimated cardinalities and timers, in contrast to the actual values available post-execution. We have already seen a QEP when describing the real-world problem pattern in Figure 1.1. The described problem pattern is much easier identified and fixed through the use of QEPs, since they provide a much more human-readable translation of the optimizer decisions made. The fix applied in Figure 1.1 was done by experts, and is precisely the effort DistGALO aims to automate.

Let us now examine an automatic fix of a problematic QEP discovered and fixed by DistGALO. Figure 2.1a depicts the optimizer selected plan, and the DistGALO selected variant in Figure 2.1b. Let's first examine the problem with the former, and then describe the fix applied in the latter. The optimizer selected plan in Figure 2.1a suffers from a hash join (**HSJOIN (#3)**) with an expensive table scan (**TBSCAN (#4)**) as its outer input. The expensive table scan is considerably I/O intensive and accounts for 99% of the cost of the **HSJOIN (#3)**. The compiler knows it can



(a) Plan selected by the optimizer.

(b) Plan chosen by the DistGALO system.

Figure 2.1: Problem pattern with expensive hash join and DistGALO discovered fix.

utilize big block I/O for the table scan, which tends to be more efficient, and results in the cheapest overall access plan with a total cost of 160797 timerons.

Let us consider the DistGALO chosen QEP shown in Figure 2.1b. Note the reason for this plan not being chosen is solely due to the total cost of 177,056 (**NLJOIN (#2)**) being higher. The cost of the nested-loop join (**NLJOIN (#3)**) makes up for the majority of the cost so perhaps the secret of the performance boost lies within. First, note that the inner and outer input streams of the join are both obtained from **FETCH-IXSCAN** operators #4 and #6. This means that both inputs into the nested-loop join are ordered, and thus have better buffer pool exploitation. Secondly, the cost estimation tends to be quite pessimistic towards the I/O in nested-loop joins, further compounded by the partitioned indexes **DD_DI** and **CS_SHIPDATE**. These oversights result in a significant 3X speedup in the DistGALO selected plan.

QEPs are fundamental in DistGALO as they give the ability to translate queries into graphs, which can be decomposed into subgraphs, traversed, and altered with

ease. The ability to do so hinges on a well defined and established graph framework, and as such, we opted for the Resource Description Framework.

RDF. The Resource Description Framework (RDF) is an XML based standard for describing data on the web, developed by the World Wide Web Consortium (W3C) ¹. RDF uses Uniform Resource Indicators (URIs) to link varying XML objects together without the need to embed them into one another as you would in XML. Due to the verbose nature of URIs, RDF is typically not intended to be read by programmers, but rather by computers. The RDF format is comprised of statement triples, namely the subject, predicate (property), and object (value). These three pieces of information are arguably enough to describe any single bit of knowledge or data. The subject is the who or what of the statement, the predicate is an existing fact about the subject, and the object is the final descriptive element regarding the subject. Objects do not only have to consist of primitive data types but can also be URIs pointing to other subjects. The QEP's LOLEPOPs are translated into subjects, with its metadata values defined by predicates and objects. The parent and children of each LOLEPOP are also defined similarly, except their values reference other LOLEPOPs. For example, each LOLEPOP consists of an inner and outer input stream, represented by the `< http://DistGALO/planDetails/property/hasInnerInputStream >` and `< http://DistGALO/planDetails/property/hasOuterInputStream >` predicates respectively. These are analogous to the child nodes in a binary tree. The parent node, or output stream of the LOLEPOP is defined by the `< http://DistGALO/planDetails/property/hasOutputStream >` predicate, followed by the reference value of the parent subject. Its flexibility allows us to easily define and

¹<https://www.w3.org/RDF/>

store QEPs obtained from the optimizer. All queries and sub-queries used throughout the system are fed through the RDF Transformation Engine which parses the optimizer provided explain files containing the QEP, into an RDF. It allows for easier parsing and manipulation of the graph-based QEPs. This further gives the ability to breaking down queries into sub-queries and creating templates to be saved in the knowledge base.

To query the RDF defined QEP, we utilize the SPARQL Protocol and RDF Query Language (SPARQL). Also supported by WC3, it is widely used to retrieve, and alter data defined by RDF. The queries themselves are similar to SQL in that they contain the **SELECT** and **WHERE** keywords and allow users to declare variables within the query. The main difference, however, is in the **WHERE** clause, which is comprised of a series of triples, similar to RDF. These triples can, however, be used to declare variables, prefixed with the *?* or *\$* character. For instance, *?x pref:diameter ?diam;* would query any triple with the *pref:diameter* predicate, store the resulting subject in *x*, and the object in *diam* which can be used in the projection of the data.

In this work, we utilize the SPARQL Jena Framework² to apply and make use of any required ACID transactions in Java. In addition, we utilize the Apache Jena Fuseki³ web server as the base for the knowledge base responsible for storing and retrieving templates.

Optimization Guidelines. Guidelines are recommendation rules made to the optimizer during the compilation stage of a query. Three types are made available by the DB2 optimizer: *general*, *query rewrite*, and *plan* optimization guidelines. *General* optimization guidelines are used to set general-level optimization parameters. *Query*

²<https://jena.apache.org/index.html>

³<https://jena.apache.org/documentation/fuseki2/index.html>

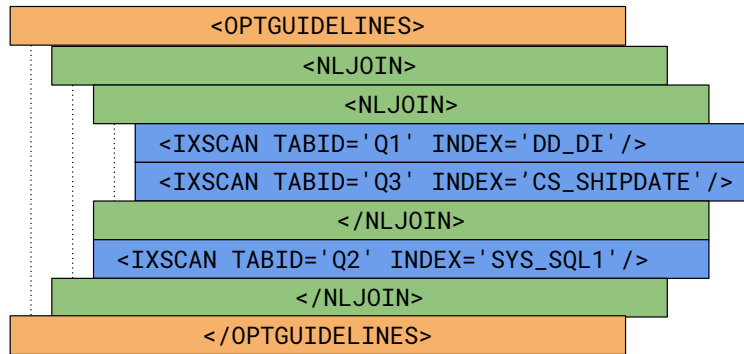


Figure 2.2: Plan optimization guidelines applied to access plan in Figure 2.1b.

rewrite optimization guidelines can apply alterations to the rules during the query rewrite stage. The only guidelines utilized in the DistGALO system are the *plan* optimization guidelines. These are applied during the third tier plan-rewrite stage of the optimization process, and apply recommendations to the access methods or join types that will be chosen. Any invalid or unspecified sections of the query will be determined by the optimizer in the normal cost-based approach. Specifics regarding an access can be requested, such as a table scan, index scan, or list prefetch. Join requests can also be specified, including nested-loop, hash, and merge joins. Optimization guidelines are XML documents with *OPTGUIDELINES* as the root node, all requested operators as intermediate children, and specified base tables/indexes as the leaf nodes. For example, the guideline in Figure 2.2 is used to apply the specified recommendations to map the optimizer selected plan in Figure 2.1a to the DistGALO selected plan in Figure 2.1b. The guideline specifies, read from leaf to root, that *Q1* and *Q2* are to be scanned using the **DD_DI** and **CS_SHIPDATE** indexes respectively. The results are used as inputs to a nested-loop join specified by the `< NLJOIN >< /NLJOIN >` tags, and show as **NLJOIN (#3)** in Figure 2.1b. The output is then also nest-loop joined with *Q2* indexed by **SYS_SQL1** and the

result is finally obtained. This XML document can be appended to the SQL query statement, and applied by the optimizer without any user intervention. This was the method of choice when applying recommendations to user queries since it can be accomplished automatically, and provide a seamless process from a user perspective.

Apache Spark With the growing trend of increased data availability and storage, distributed data processing has also followed the trend. To mend the demand, several distributed frameworks have emerged, with (Apache) Spark[18] as one of the most popular[19]. Spark has many built-in solutions for common-use applications like streaming, machine learning, and SQL queries. It leverages the MapReduce[20] paradigm in combination with resilient distributed datasets (RDDs) that are used to retain read-only data across multiple nodes. Typically all data meant for distributed computation is stored in RDDs in key/value pairs. These RDDs then expose programmers to Spark’s rich API to apply transformations on the data, like *reduceByKey()*, *mapValues(@f)*, and *sortByKey()* to name a few.

One of the bottlenecks in distributed computing is the communication cost between nodes. Great efforts are made to ensure that machines are located in close proximity with fast network speeds between each other, but even despite such efforts, it remains a big obstacle in distributed computing. If alike data can be grouped on a per-node basis, then the amount of data shuffling that will occur between RDD transformations can be reduced. Spark provides the control to specify which key/-value pairs should appear together on a given node. Spark provides seamless fault tolerance since nodes chance to fail, and thus cannot guarantee and consequently allow users to specify a specific node. Spark also provides built-in partitioning strategies that give the user some control over how RDDs should be distributed amongst

the nodes. Two such partitioners exist, including *Hash Partitioners* and *Range Partitioners*. The former using the hash of the key to determine a partition, and the latter allowing buckets with ranges to allocate data distribution. Additionally, Spark gives the ability to create *Custom Partitioners*, allowing programmers to utilize domain-specific information to partition the data. This is something we make heavy use of in this work since we want to ensure there is an overall balance between nodes, so as to not create skew across the cluster. More detail on the domain-specific application of partitioning can be found in Section 3.2.

2.2 System Overview

The previous GALO[4] system contained three major components: *Learning Engine*, *Matching Engine* and the *Transformation Engine*. DistGALO continues to embody similar architecture, but incorporates a revamped *Distributed Learning Engine* consisting of a *Distributed Learning Component* (DLC) and a *Random Subquery Abridged Clustering Executor* (RSACE) module. A high level of the *Distributed Learning Engine* can be seen in Figure 2.3 with emphasis on the DLC. The Ranking Module, Template Creation, and process for populating the RDF Knowledge Base largely remain unchanged from GALO. The DLC can further be broken down into the following elements:

1. *Sub-query generator*;
2. *Optimizer sub-query executor*; and
3. *Random sub-query executor*;

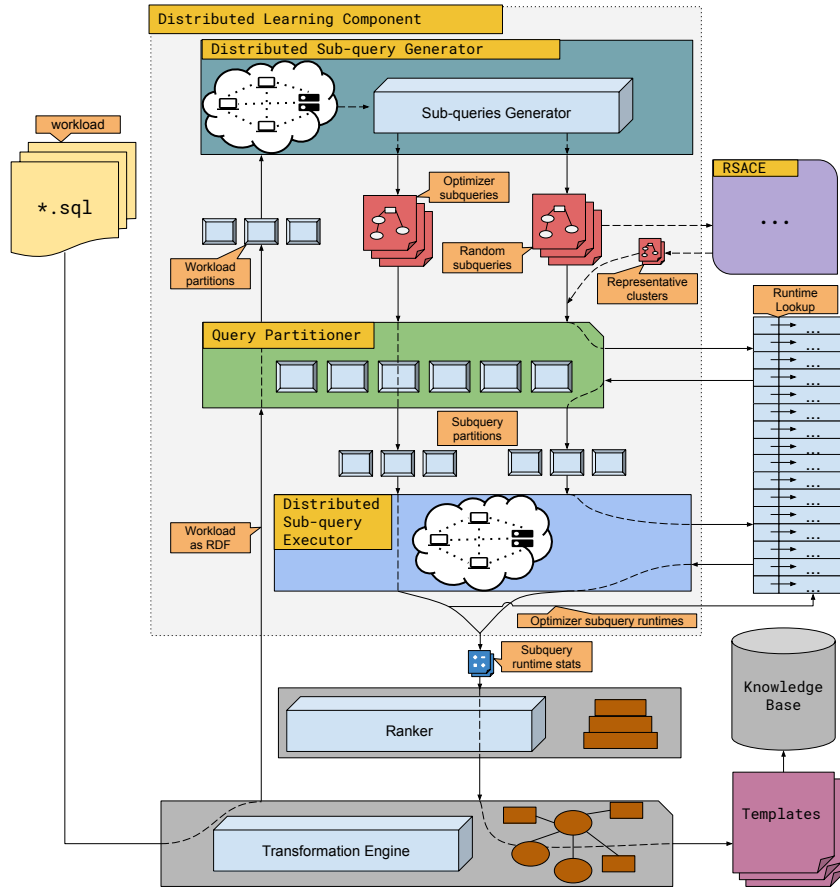


Figure 2.3: System architecture of DistGALO's Distributed Learning Engine.

Sub-query generator The sub-query generator is responsible for processing daily query workloads and generating sub-portions of the original queries, known as sub-queries. After a query is processed, it is executed and the optimizer's chosen QEP is extracted. Using GALO's *Transformation Engine*, the QEP is mapped into its RDF graph representation, parsed, and sub-divided into smaller sub-graphs. The size of the sub-graphs is based on the number of joins, set by the user. The sub-graph we term a sub-query, projects the local and join predicates from the original query. The local predicates are used to generate new values by querying the database with

dynamically sampled probing queries. These newly generated alternate predicate sub-queries we henceforth termed, optimizer sub-queries, or Φ for conciseness. Using the Random Plan Generator (a DB2 built-in tool), a predetermined number of guidelines are generated and attached to each of the optimizer sub-queries to create a new set of sub-queries. These sub-queries are identical to their optimizer sub-query counterpart, with the exception of the attached guideline. The attached guideline provides an alternate execution plan that may otherwise not be considered during the cost-based plan evaluation. This new set of sub-queries are termed random sub-queries, or Ψ , and its size is a factor of the number of random guidelines generated, but typically tends to be much larger than Ψ . DistGALO generates sub-queries in parallel using Spark, in contrast to the sequential generation of GALO. This was done using the workload queries and cost-based partitioner described in Section 3.2. At this stage, both sets Φ and Ψ can be combined and distributed throughout the cluster for processing. This however turns out to be a less effective approach due to the relationship between the sets (more detail in Section 3). Thus, after the sub-query generation step, both the Φ and Ψ sub-query sets are passed along the pipeline to their respective executors.

Optimizer sub-query executor The Optimizer sub-query executor is responsible for processing all the Φ sub-queries previously generated. This is done in a distributed fashion throughout the cluster. The sub-queries are intelligently partitioned throughout all nodes, and processed. The processing of each ϕ sub-query involves generating their QEPs, and executing each in order to obtain the runtime and runtime statistics. Based on the metrics obtained, all runs are ranked by the *Ranking Module*, with the winner being saved in the Hadoop DFS for future reference. A time lookup table with all the best ranked execution of each ϕ is generated

to provide upper-bound thresholds for the processing of Ψ .

Random sub-query executor The Random sub-query executor, as the name suggests, is responsible for processing Ψ . This is critical step, previously proving to be a bottleneck in GALO, due to the large number of sub-queries generated. DistGALO now tackles this stage very strategically, from the partitioning strategies, to pruning approaches to remedy the massive search space. The time lookup table from the optimizer sub-query executor is also heavily utilized to ensure that ψ sub-queries are not run for longer than they need to be. Several partitioning strategies are also made available, each with its strengths and trade-offs. Similarly, sub-queries are processed in a similar fashion to those from the previous executor, and so all top-ranking sub-queries are ultimately saved in the DFS.

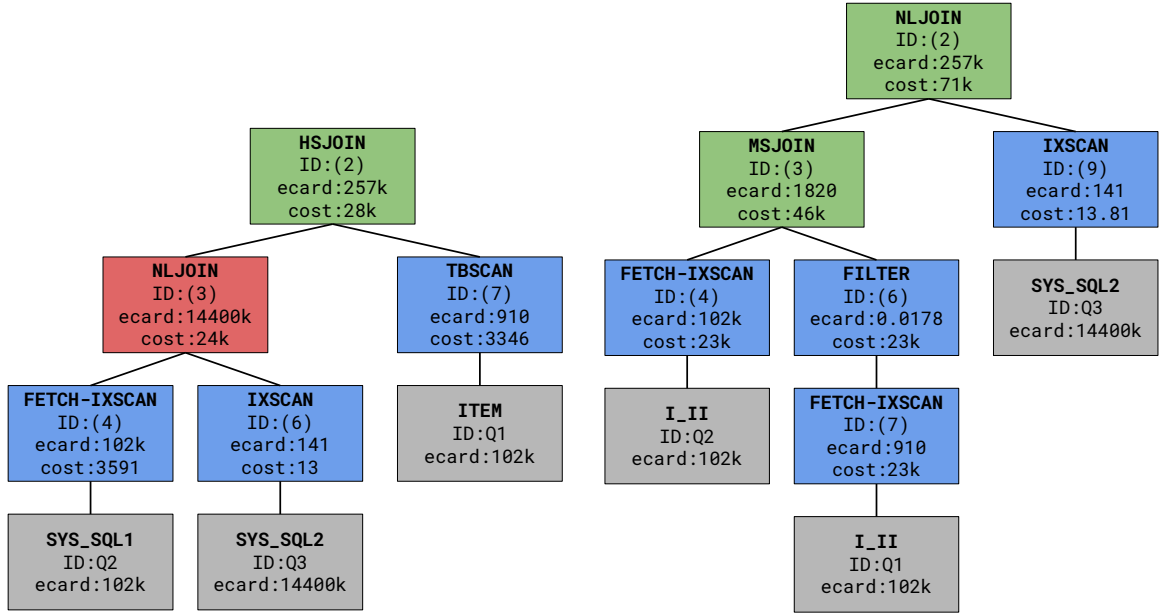
Random sub-query abridged clustering executor (RSACE) RSACE was devised as a pruning strategy to provide further speedups with minimal impairment to the objective of the Learning Engine. The module is a feature optional to the user, that can be toggled on or off on demand. This strategy is also only applicable to the Random sub-query executor since it has the prerequisite of a time lookup table from the optimizer sub-query executor. The process involves vectorizing the predicates of each ψ , grouping by their respective QEP structures, applying clustering techniques within each group, and executing representatives from each cluster. The module can toggle to account of any false positive representatives that may have arisen. Potential false positive representatives' clusters are processed using the Random sub-query executor to ensure the *Distributed Learning Engine* is still effective in the templates it saves in the Knowledge Base.

Chapter 3

Distributed Learning

The Learning Engine is perhaps the more complex component of the DistGALO architecture as a whole. It is further subdivided into several stages: *sub-query generation*, *optimizer sub-query execution*, *random sub-query execution*, *sub-query ranking*, and *template creation*. Though intermediate steps are involved, like creating generating QEPs with actuals, moving files to/from the HDFS, and preparing queries, they are omitted from the discussion since they are significantly less computationally intensive. The main focus will be limited to the 5 processes mentioned above as they are the most computationally and resource heavy.

The output of the Distributed Learning Engine in DistGALO remains the same as that of GALO. It aims to discover the more optimal, less resource consuming, faster executing, random sub-queries, in comparison to their optimizer sub-query counterpart. Let us consider one such case depicted in Figure 3.1 where we observe a problematic QEP chosen by the optimizer, and the appropriate fix selected by DistGALO; let's examine each independently. The optimizer selected plan in Figure 3.1a first joins Q_2 with Q_3 using nested-loop join **NLJOIN (#3)**. It then joins



(a) Plan selected by the optimizer.

(b) Plan chosen by the DistGALO system.

Figure 3.1: Problem pattern with expensive nested-loop join and DistGALO discovered fix.

the results with $Q1$ using the hash join **HSJOIN** (#2), for a total cost of 28860 timerons. Note that by joining $Q2$ with $Q3$ first, it leads to an expanding join, with cardinality of roughly 14million. This expanding join results in the majority of cost, with 24902 timerons.

Consider the DistGALO selected plan in Figure 3.1b. Its total cost of 63682 is more than double than that of the optimizer, yet exhibits a 4X performance speedup. Let us proceed by analyzing the QEP inner-workings and determine the optimizer oversight. First note the join order difference, whereby $Q2$ is first joined with $Q1$, then finally with $Q3$. The former join is performed by what seems an expensive **MSJOIN** (#3) with a cost of 46019 timerons. Note however, that this join, despite being the most costly, reduces its inputs to an estimated cardinality of 1020 in contrast to

the expanding **NLJOIN** with 14 million in the optimizer selected plan. The inner and outer inputs utilize **FETCH (#4,#7)** operators, each with a cost of 23000 timerons. A crucial observation is that the index, **I.II** used by both to fetch row ids, is the exact same. This means that during the merge join operation, whilst the outer is being read, the buffer pool is heavily utilized for the inner with minimal I/O. Effectively the cost of the **MSJOIN (#3)** is halved, propagating the speedup to the following **NLJOIN (#2)** above. This oversight is largely attributed to the cost model, since it assumes I/O must be performed. These types of modification provide insight into the ever demanding need for domain experts, and more-so automated systems like DistGALO.

3.1 Cluster and Spark Configuration

Cluster To remedy the lengthy runtime of the GALO Learning Engine, a distributed cluster was used to improve scalability, runtime, and to fully utilize the hardware available. Computations were performed on the SOSCIP Consortium’s Cloud Data Analytics computing platform(s). SOSCIP is funded by the Federal Economic Development Agency of Southern Ontario, the Province of Ontario, IBM Canada Ltd., Ontario Centres of Excellence, Mitacs and 15 Ontario academic member institutions [21]. The SOSCIP platform ensures that nodes are free of user traffic, thus minimizing any noise that would otherwise be present in multi user nodes. All nodes are located on the same internal network, so communication overhead between them is minimized and limited to the partitioning strategy employed by the system. The cluster is comprised of 9 nodes, including a master, each with 4 virtual CPUs, 16GB of RAM, and 120GB of disk space. To effectively use the cluster, we used the general-purpose

cluster computing platform, Apache Spark. The open-source platform extends the MapReduce model to support a large variety of computations including SQL queries, which fit the DistGALO system requirements closely. It has also proved to be quite effective in the community and has withstood the test of time, and was therefore selected as the tool of choice. We also utilized the Hadoop Distributed File System (HDFS) to seamlessly share data between all the nodes. Great effort was made to minimize the amount of reads/writes to the HDFS since it could pose as a potential bottleneck due to the large amount of files and data involved. At the time of writing, and in the latest iteration of DistGALO, Apache Spark 2.4.0 and Apache Hadoop 2.7.7 were used.

Spark Configuration Apache Spark’s configuration is quite extensive and allows for very flexible and versatile tuning of the cluster. Some of the available configurations available include setting: number of cores per executor, memory per executor, total cores available to the cluster and so on, see¹ for a comprehensive list of Spark’s configuration options. The cluster used includes 9 nodes, one acting as the master and worker, and the remaining 8 just as workers. The cluster has 36 virtual CPU cores, and 144GB of RAM available. Since the system is very precise in how it must be executed, no automatic methods exists for finding and setting the most optimal configuration. There are also various variables that such a system would not be able to tune; including the database size, database configuration, and query workload, just to name a few. We have thus opted with the arduous task of experimentally finding the most optimal configuration. Table 3.1 displays the Spark configuration settings experimentally found to be best. The philosophy behind each is as follows.

¹<https://spark.apache.org/docs/latest/configuration.html>

Table 3.1: Apache Spark Cluster Configuration

Spark Property	Value
spark.cores.max	36
spark.executor.memory	4G
spark.executor.instances	36
spark.executor.cores	1

We limit the number of cores utilized per executor (*spark.executor.cores*) to 1 to maximize concurrency on each node. Since each node has 4 cores available we expect 4 executors to be running concurrently, thus limiting the RAM to 4GB per core (*spark.executor.memoy*). Finally since we have 9 available machines with 4 cores each, we limit both the number of executors (*spark.executor.instances*), and the total number of cores (*spark.cores.max*), to 36, since those are the total available VCPUs available in the cluster. These configuration settings are based on several properties that were measured in-house, primarily runtime and resource utilization. A common symptom and bottleneck of distributed systems is skew. This typically occurs when a small subset of the nodes available take longer to complete the task, thus halting the whole work flow including the nodes that have already completed. This is something that Spark tries to address under the hood, but when skew continues to be evident with the default settings, more customized setting are required. Due to the non-deterministic nature of executing randomly generated execution plans, skew is quite prominent and must be addressed directly. We thus make use of several characteristics and heuristics from the generated sub-queries in effort to distribute the work load evenly throughout the cluster.

3.2 Partitioning

Partitioning in distributed systems is a crucial element in ensuring the cluster is working at its optimum. In most cases of distributed computing, the main bottleneck tends to be the communication cost between nodes, or network traffic. When performing transformations on data, shuffles of data are triggered across the cluster, incurring communication cost and slowing down overall runtime. If alike data is however grouped, transformations would cause less shuffling across the network, thus reducing communication cost. Partitioning provides a grouping of keys based on some function, either built into, or defined by the user (in the context of Apache Spark). In other words, partitioning gives the ability to preemptively designate how keys are to be grouped when allocated to nodes. So a user has the control to determine if two keys should end up on the same node, though cannot guarantee that they will not. Typically the effectiveness of partitioning comes from reducing data shuffle during the transformations applied, but DistGALO only does a single mapping followed by a single reduction. How then do we benefit from partitioning, and is it significant enough to warrant a discussion? DistGALO main goal is to execute a large number of sub-queries with varying complexities. More complex queries usually take a longer time to execute, and so grouping many complex queries on a single node would cause some skew. So the benefit of partitioning comes from the initial sub-query allocation, as opposed to the reduction of data shuffling. The initial distribution dictates which nodes will be overloaded, and which will sit idle waiting for the rest to finish. This can be achieved through partitioning and careful analysis on how to partition sub-queries together. Let us explore the varying partitioning strategies made available through Apache Spark, and others that are proposed in this work.

Apache Spark comes equipped with three options when selecting a partitioner. The first is a Hash Partitioner, which groups data based on the hashed key values. This is a viable option in some scenarios so it is worth further discussion, which can be found in the section to follow. The second partitioning option is the Range Partitioner, which partitions based on the natural ordering of keys, and places keys in predetermined ranges. This could allow us to possibly group similarly complex queries together, which provides the opposite of the desired effect, which is to balance the overall complexity among the partitions. This we deemed not a viable option and is further excluded from the discussion. The last option that Spark provides, is the ability to write a Custom Partitioner. This gives the user the ability to customize precisely how keys are grouped within the context of their application. Consequently we have devised two different partitioning methods, a Cost Estimate Partitioner, and a Runtime Partitioner. Of the three detailed strategies, there isn't one that triumphs over the others, but rather, each has its own unique scenario in which it proves the most optimal choice.

Hash Partitioner The Hash Partitioner groups keys based on the hash value of each key, so keys with the same hash values end up in the same partition. More formally we can write this as: $partition = getHash(key) \% numPartitions$, where *getHash* is a built-in function which generates a hash value, and where *numPartitions* are the number of partitions set by the user or automatically by Spark. Given a hash function adhering to uniformity, we expect that the resulting partitions would have an equal number of key/value pairs assigned to each. This is a desired quality in some parts of DistGALO, like in the optimizer query execution phase of the system. This approach is advantageous in that it requires no prior knowledge about the key/value

pairs and therefore has little overhead. There are scenarios in which an even distribution is not as desirable, and may hinder the runtime. In the context of SQL queries, the hash partitioner doesn't have any information regarding the complexity of the queries. This oversight may result in complex queries being grouped on the same node, leaving the less complex and completed nodes to sit idle. We thus want a more quantified way of assessing the complexity of a query to create an even distribution of workload complexity amongst all nodes.

Cost Estimate Partitioner We turn to a heuristic based approach to partition the sub-queries based on some estimated metric of complexity. DB2 offers such a heuristic in the form of an estimated cost in timeron units. A timeron is an estimate of the total I/O and CPU cost the DB2 manager might incur during execution. The estimate is derived from table statistics, indexes, predicates involved, cardinalities and other variables. A great advantage this approach offers is that the cost estimate values are available at compile time. That is they do not require execution of the sub-query, which is a timely and costly procedure. Timeron values are obtained from the QEP that the DB2 optimizer generates through the explain facility. It does however have a slight overhead since a connection to the database manager must be established, and internal commands must be made to generate the QEP.

The Cost Estimate Partitioner is a Custom Partitioner that groups queries based on their their total estimated cost timerons. The prerequisite is that all cost estimate values must be available for each sub-query to be partitioned. This means that every sub-query in question must have their QEP generated by the DB2 manager. The initial partitioning process works by creating a bucket for each partition, an attribute set by the user. The queries are then sequentially evaluated, and their cost estimated

input : The set of sub-queries to be distributed \mathbb{S}
output: The partitions with sub-queries assigned **filePartitions**

```

1 Let filePartitions = Map(query, metric) ;
2 Let partitionBuckets = [0] * numPartitions ;
3 for  $s \in \mathbb{S}$  do
4   | minIndex = getMinBucketIndex(partitionBuckets) ;
5   | partitionBuckets[minIndex] += s.metric ;
6   | filePartitions.put(s, minIndex) ;
7 end
8 return filePartitions ;
```

Algorithm 1: Algorithm for partitioning of sub-queries

values are used to populate the buckets. The assignment is as simple as assigning the current cost value to the lowest bucket. Algorithm 1 shows the trivial, yet effective assignment of sub-queries to partitions. First the *filePartitions* are initialized (line 1) with an empty map, whereby the keys are the sub-queries, and the values, the corresponding partition the query is to be put in. The empty set of buckets are then initialized (line 2) and will keep track of the total cost for each partition assigned thus far. For each sub-query, the current smallest bucket is obtained (line 4), updated with the query metric (line 5), and set to the corresponding partition (line 6). The query metric can be any characteristic of a given sub-query. In the case of the Hash Partitioner, the metric is the estimated total cost obtained from the optimizer.

Though there is quite a large variation of cost values, when dealing with a large number of sub-queries, the partitions are eventually able to flatten out and more or less have the same values. Optimally, each partition would have the same cost values, but in practice is not the case, though relatively close. Since the estimated costs relate to the complexity of the queries, having partitions with an overall equal level of complexity, translates to partitions that must finish execution at roughly the same

time. This in practice is not always the case since we are relying on estimates, but in general is quite effective in its predictions. The accuracy of the Cost Estimate Partitioner is largely dependent on the accuracy of the estimated cost values. The estimates are constantly being improved by the DB2 optimizer team, and with more and more accurate cost values, the Cost Estimate Partitioner benefits just as equally.

Runtime Partitioner Instead of using estimated timeron values, a more reliable and accurate approach is to use the runtime values. Runtime values refer to the exact CPU time taken to execute a given query. We must emphasize that timerons are not the same as runtimes, since the former is an estimate, and the latter an exact value. The algorithm used is identical to the Cost Estimate Partitioner (Algorithm 1), but with a different query metric. Instead of using the timeron values, we simply swap them out for the actual runtime values. The process of filling the buckets works identically, by sequentially iterating through the queries' runtime values, and assigning each to the minimum valued bucket at the time of execution. The result is a set of partitions each assigned their own set of queries, and a partition value depicting the exact runtime it will take for it to complete execution of its assigned queries. The value of all buckets should be as close as possible to ensure partitions are completed in as similar time as possible. Partitions are later subdivided into multiple executors running concurrently on each node. It should be noted that the partitioning stage precedes the query execution stage, so using runtimes to partition seems paradoxical. This is indeed true if we are partitioning the same queries being executed, but in general is not the approach taken when using the Runtime Partitioner. This requires a more in depth analysis into the queries being executed, this is further elaborated upon in Section 3.5. To show the effectiveness of the Runtime Partitioner, an experiment

comparing all partitioners was conducted and further described in Section 4.2.

3.3 Sub-query Generation

The sub-query generation process takes an SQL query as input and creates various sub-queries, including ones with optimizer chosen execution plans, and others with a randomly generated execution plan. Once the two batches of optimizer and random sub-queries has been generated, they are passed on to the execution step of the pipeline. Let us delve further into the sub-query generation pipeline. The process begins with a set of queries(workload) to be learned from, $Q = \{q_1, q_2, \dots, q_n\}$ for n input queries. Each query q_i is further broken down into a set of sub-queries $S_i = \{s_{i1}, s_{i2}, \dots, s_{im}\}$ where m sub-queries are generated for a given query q_i . The generation of sub-queries for a given query is based on the execution plan generated by the DB2 optimizer. Given a sub-query of size J joins, every combination of J join sub-graphs are generated by $subgen(q_i, J)$ [4]. The sub-query s_{ij} thus represents the j^{th} query generated from the original q_i^{th} query. These set of sub-queries can be more compactly be represented by $\mathbb{S} = \{subgen(q, J) \mid q \in Q\}$. For each sub-query s_{ij} a set of k alternate predicate sub-queries $\{a_{ij1}, a_{ij2}, \dots, a_{ijk}\}$ are generated by the algorithm $altgen(s)$ [4]. These are all modifications of the original sub-query s_{ij} but where, through sampling of the attributes, have different local predicates. For instance, a_{ij1} and a_{ij2} would be the exact same query but would have different predicate values. These sets of alternate predicates can too more succinctly be represented by $A = \{altgen(s) \mid s \in \mathbb{S}\}$. To summarize, the alternate predicate sub-query a_{ijk} thus represents a sub-query formed from query q_i , with the j^{th} sub-query, and the k^{th} alternate predicate sub-query. Each alternate predicate sub-query then forms a

group $g_{ijk} = \{ \phi, \psi_1, \psi_2, \dots, \psi_q \}$ where ϕ represents the optimizer chosen query execution plan, and ψ_q is the q^{th} randomly generated query execution plan for any given alternate predicate sub-query a_{ijk} . To simplify the nomenclature, if G were the set of all groups g_{ijk} and a flat-map like function were applied, then grouping all optimizer chosen plan sub-queries into a set Φ and grouping all the random plan sub-queries into the set Ψ . If the sub-query generation process were observed to be a black box, the input would be the set of queries Q , and the output would consist of the sets Φ and Ψ .

The sub-query process is fairly computationally intensive since tables must be sampled, and a fair number of files are created as a byproduct. The process was therefore distributed among the nodes in the cluster to fully utilize resources. Each query in Q is first executed to obtain the actual runtime and cardinalities as opposed to the just the estimates. The actual statistics served as a fairly accurate heuristic for partitioning the queries throughout the nodes. The heuristic we follow is as follows: a higher total runtime typically corresponds to a more complex and thus larger query execution plan. Larger QEPs would naturally translate to a larger number of joins, resulting in a larger number of possible combinations between the joins, and thus a larger number of sub-queries. This meant that a longer runtime typically translated to a larger number of sub-queries that could be expected from the given query. Consequently, partitioning by cost estimate provided a balance of complexity among the executors and therefore individual nodes. The generation process requires the QEP to be generated, so cost information is already available. Therefore the Cost Estimate Partitioner seemed a natural choice since it incurs no extra cost and has proven quite effective in our experiments. When all executors complete their workloads, the HDFS

is populated with both the optimizer selected QEP sub-queries Φ and the randomly generated plan sub-queries Ψ . These are then passed down into the execution stage of the pipeline.

3.4 Optimizer Sub-query Execution

The next stage involves running the optimizer chosen query execution plan sub-queries described by Φ in Section 3.3. This means that each sub-query must be executed by DB2 to obtain the actual runtimes. As this is a lengthy and computationally heavy step, the choice of distributing the work among the cluster was made.

The next and arguably most important step is the choice of partitioning strategies described in Section 3.2. The options, however, are limited to only two, as runtime partitioning is not available pre-execution of sub-queries. The objective is to obtain the runtimes from each sub-query, so the partitioning strategies are narrowed down to either cost estimate, or hash partitioning. As previously observed, the hash partitioning strategy does not offer much intelligence in the distribution of sub-queries between partitions. It does not take into account any heuristics and instead simply distributes the sub-queries as evenly as possible throughout the partitions. Conversely, the cost-based partitioning strategy offers insight into the complexity of the query and is a decent heuristic to group low cost, fast-running queries together, while allowing the larger ones to run independently, and ultimately avoid skew. This makes a strong case for the latter, however as previously mentioned, there is a prerequisite to running using a Cost Estimate Partitioner. The estimates must be available, which requires the execution of program *db2exfmt* on each query. This ensures that the QEP for the query is generated, and though it is not as costly as actually running the query,

it does add some overhead. In our experiments, we observed that since most of the optimizer selected sub-queries ran in a reasonable amount of time, the strategy of distributing them using the hash partitioner seemed fairly effective. This approach negates the overhead running each sub-query through the explain facility. There are scenarios in which the sub-queries do not run in a respectable amount of time, and so a global threshold was placed to ensure that if met, the query execution was cut off. This avoided some edge cases where a sub-query would run indefinitely thus halting the system since all other nodes wait until completion.

Even though the environment used was not shared amongst users, the system is still susceptible to small amounts of noise which propagate and cause variations on query times. In effort to reduce this noise, we have opted to run each query three times and obtain the best execution. We term the best run here as the one that ranks the highest amongst all other runs. This includes an additional ranking process, similar to that of GALO, that must be run after all three runs of each sub-query. The best performing run is kept and saved in the HDFS, while the remaining two are discarded. The ranking takes into account several statistics extracted from the optimizer, all of which are listed in Table 3.2. The ranking is based on weights and each query run is given a score. `ELAP_T` is given a weight of 50% and all the remaining properties are given the remaining combined 50%. The weight favours the elapsed time since that is the most critical element, and most costly in obtaining, therefore has the most significance. The ranking process is described in greater detail in GALO[4].

Table 3.2: Sub-query post execution ranking properties considered.

Abbreviation	Description
ELAP_T	Elapsed Time
BP_D_LR	Buffer pool data logical reads
BP_D_PR	Buffer pool data physical reads
BP_TD_LR	Buffer pool temporary data logical reads
BP_TD_PR	Buffer pool temporary data physical reads
BP_D_W	Buffer pool data writes
BP_I_LR	Buffer pool index logical reads
BP_I_PR	Buffer pool index physical reads
BP_TI_LR	Buffer pool temporary index logical reads
BP_TI_PR	Buffer pool temporary index physical reads
BP_I_W	Buffer pool index writes
U_CPU_T	Total User CPU Time used by agent (s)
SSH_HWM	Shared Sort heap high water mark
ROWS_R	Rows read

3.5 Random Sub-query Execution

After the optimizer selected QEP sub-queries have completed, the random QEP queries are prepared for execution. This is similar to the process of executing the optimizer variant described in Section 3.4. Similarities include the need for execution to obtain runtime statistics and elapsed time, thus prompting the need for a distributed solution. There are some subtle differences, however. The first most significant observation is the sheer amount of random sub-queries in comparison to the optimizer sub-queries generated. The former exceeded the latter by about 7 times, on average. This figure does vary slightly since the queries generated are random and as such vary in number, but also because the queries in the workload also vary. The second observation that can be made is that some valuable information is now accessible from the execution of the optimizer queries, mainly the runtimes. This

may not seem to be much of use, but we must first examine the relationship between Φ and Ψ described in Section 3.3.

Recall the groups $g_{ijk} = \{ \phi, \psi_1, \psi_2, \dots, \psi_q \}$ created during the sub-query generation phase, where i represents the original query number, j the sub-query number, and k the alternate predicate number. Within each group g there exists an optimizer chosen QEP sub-query or ϕ . In addition there exists q number of ψ randomly generated QEP sub-queries. As a recap, the ultimate objective in creating a template is to discover at least two sub-queries with randomly generated QEPs, which are part of the same sub-query, but with varying alternate predicates. In other words, for two groups belonging to the sub-query, there must exist at least two ψ sub-queries that are better than their ϕ sub-query counterpart, in order for a template to form. At the group level, the only concern is for any given group g , to find a ψ that outperforms the ϕ . Any ψ that does not outperform the groups ϕ is discarded as it cannot possibly contribute to the creation of a template. The optimizer chosen QEP ϕ therefore governs the runtime for each ψ in the group thus describing the relationship between Φ and Ψ .

This relationship though not immediately evident provides three possible ways to provide significant speedups in execution of Ψ , without any sacrifice to the number of templates created. The first speedup, described in Section 3.5, involves utilizing the cluster to distribute all of the sub-queries in Ψ with a better partitioning strategy. The second involves using thresholds to limit the execution time of the sub-queries, with more detail in Section 3.5. The last speedup involves a pruning method that removes unnecessary runs, with more detail in Section 3.5.

Partitioning

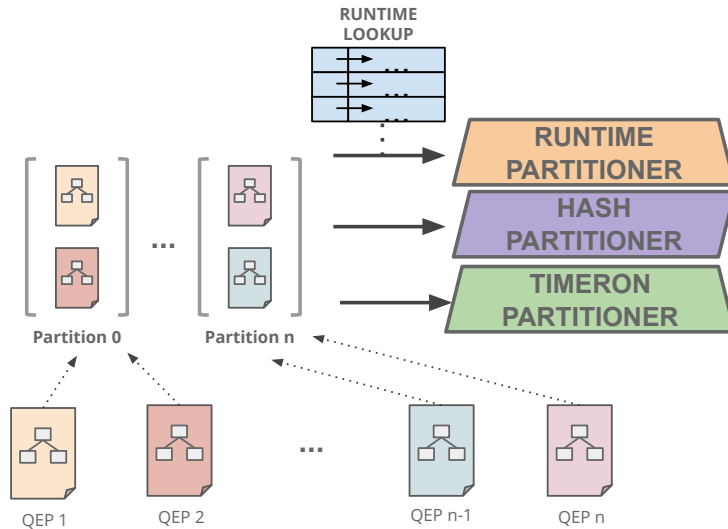


Figure 3.2: Partitioning process of random sub-queries using one of three available partitioners.

The task of executing a large number of sub-queries is substantial, and so there is motivation for utilizing the available cluster to divide the work up into chunks. The main goal in dividing up the work is to reduce skew as much as possible. One of the ways to mitigate skew is through partitioning the work as evenly as possible, so selecting the optimal partitioning strategy becomes the crucial task. Figure 3.2 provides a graphical representation of the partitioning process. The figure only depicts two partitions for simplicity, and the three available partitioners described in Section 3.2. Let us consider each partitioning strategy individually.

A naive solution would be to divide the work up evenly on each node, which can be achieved with ease using the Hash Partitioner. The resulting state of the cluster would be an even distribution of queries amongst all nodes. The limitation to this, however, is that the complexity of the queries is not taken into account. Take for example a scenario where two nodes must each execute 3 queries, but by chance,

the first node got assigned fairly simple queries that run fast and finish quickly. The second node, again by chance, as is often the case with distributed systems, is assigned 3 very complex queries, which take a significantly longer amount of time to complete, and have now introduced skew into our system.

The second approach in partitioning Φ is the Cost Estimate Partitioner. The estimates are obtained through the QEP files, which are generated through the explain table format command, *db2exfmt*. These values tend to be higher with more complex queries, and can, therefore, be used to balance sub-queries throughout the cluster. The caveat is that there is some overhead associated with creating the QEP files. The files do not happen to be available at this stage of the execution process, so they must, therefore, be generated. This overhead though not very significant on a per-query basis when considering thousands of sub-queries, becomes quite significant. The overhead negates some of the speedup advantages the cost-based partitioning strategy provides, making it less effective. To compound this further, the estimates are not always reliable and can sum up to a significant amount of error. Both of these factors surmise to a considerable amount of skew re-introduced into the cluster, nullifying the purpose of a partitioning strategy. We must, therefore, seek an alternate partitioning strategy void of such shortcomings.

The third option for partitioning is using the actual runtimes, in contrast to the estimates. Ideally, the runtimes of the ψ queries would be used, but this is only available post-execution, which is the very task we sought to achieve. Consider however the actual runtime values of Φ that were executed in the previous step. These runtimes, in essence, dictate the upper-bound thresholds of their ψ counterparts. Let us recap, the original goal of partitioning is to ensure that all nodes complete

execution as close as possible to one another. We have now established that we have the upper-bound execution time of each ψ query, and thus if we partition based on the upper-bounds, then we are achieved what we originally set out to do. This strategy doesn't rely on any estimates, so the expectation is that each node will run the exact time it was allotted. This generally holds true, with exception where ψ sub-queries significantly undercut their upper-bound. This in practice is exhibited on each node and tends to balance out overall. The runtime partitioning strategy introduces no overhead since it only utilizes information already available, and thus is the better performing of the three in this stage of the learning process.

Thresholds Once the nodes have been assigned the group of sub-queries that must be executed, a brute force method of simply running the queries does not suffice. There are scenarios in which some of the randomly generated QEPs will not be valid, and will never truly terminate on their own. To mend this, a global threshold can be set, similar to the one outlined in Φ execution in Section 3.4. This though somewhat effective does not perform too well due to the substantial number of ψ sub-queries. The above-mentioned relationship can, however, aid in making a threshold that caters to each group. A threshold can be created for each group $g_{ijk} = \{ \phi, \psi_1, \psi_2, \dots, \psi_q \}$, such that the threshold is the runtime of ϕ which is the upper bound runtime for all ψ sub-queries in the group. The set of thresholds that are created can be formally described by $T = \{ runtime(\phi) \mid \phi \in g \in G \}$ where *runtime* is a function that extracts the execution time of a given query, and G is the set of all groups of sub-queries described in Section 3.3. For every node in the cluster, the set of T thresholds is created, though this may seem excessive, it is a necessity due to the non-deterministic nature of Spark's distribution of sub-queries amongst the nodes. This means that

there is no guarantee that any given ψ will end up on a specific node, and so to cover all ground, every threshold in T must be created on each node on the cluster. This is also done in a distributed fashion, but the time is negligible so we are able to only reap the benefits of the speedup without incurring any added cost.

Once the thresholds have been set up, each node will have a designated batch of ψ queries it must execute, and this is done concurrently using four executors, each utilizing a core from the four available. Similar to Φ query execution process described in Section 3.4, each sub-query is run a total of 3 times, to account for any noise the system may exhibit. It is possible that during any one of the three runs, the threshold limit is reached, and the DB2 manager terminates execution. This does not mean however that the remaining runs can be ignored since it is possible that later runs would result in faster execution. So if a run is terminated due to a threshold limit being reached, we would like to continue the subsequent runs in hope of finding one that performs better than the ϕ runtime.

It is worth noting, as a limitation of the system, the thresholds are only created in 10-second intervals, so any threshold in-between is not possible. For example, the thresholds 10, 20, 30 are valid, but the thresholds 5, 15, 25 are not. This makes up for any queries that may terminate very closely after their ϕ upper-bound, but end up being terminated upon reaching the exact time. This, in essence, acts as a buffer to ensure those edge cases are mitigated, though not completely eliminated. This buffer though somewhat significant with fast-running queries does not have the same effect for slower running queries since 10-second intervals are much less significant. We thus have the motivation to find a method to cater to all queries, and also introduce a pruning ability as a result.

Pruning The number of sub-queries generated per query is roughly a 200 time multiplier. For example, in a sample run with 99 TPC-DS queries, a total of 19845 sub-queries were generated, 2461 of which are ϕ optimizer chosen sub-queries, and the remaining 17384 are ψ randomly generated sub-queries. Each sub-query must also be executed three times to reduce noise, and so from just the 99 TPC-DS queries, we expect 59535 sub-query executions. Therefore, the ability to reduce any number of query executions is of great contribution to the speedup of the learning engine. We, therefore, introduce a pruning rule that aims to remove any runs that are deemed outliers and unable to contribute to templates.

We have mentioned that an upper bound already exists from the Φ sub-query execution, and a threshold slightly above that exists to terminate any slow-performing queries. Upon a sub-query hitting its threshold, it is terminated, but subsequent runs still continue. We make the argument, that any sub-queries executing on their first run and exceed their group's ϕ by 15% should not continue executing the second and third runs. The 15% threshold takes into account for any noise, and also the fact that the runtime is not the only determining factor in the ranking process. This reduces the number of sub-query executions by 1/3 for all sub-queries that meet this pruning rule. The main concern with pruning is potentially removing a ψ in the first run, that could potentially have a second or third run that outperforms ϕ and eventually contributes to a template. This is a scenario that is negated with a higher pruning threshold percentage, at the cost of a higher overall runtime. We have experimentally found that 15% finds that balance perfectly with minimal sacrifice on the number of templates discovered.

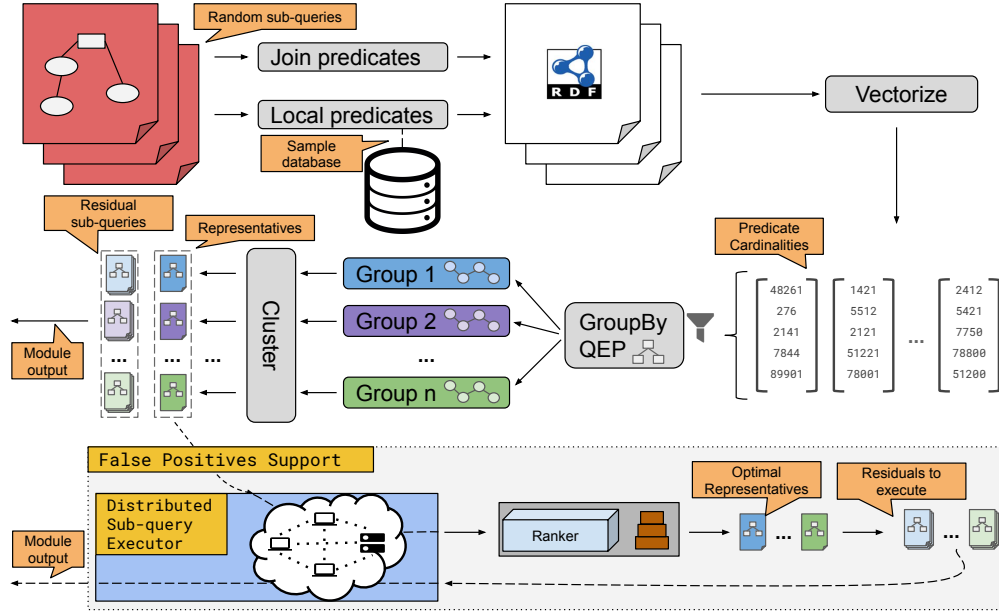


Figure 3.3: Random Sub-query Abridged Clustering Executor (RSACE) module.

3.6 RSACE

Due to the nature of the sub-query generation, we observe a massive search space of random sub-queries that must be executed. In an effort to remedy the large search space, we have devised the Random sub-query Abridged Clustering Execution (RSACE) pruning module. The architecture of this component can be seen in Figure 3.3. This technique is only applied to the random sub-queries Ψ , since the runtimes of the optimizer sub-queries Φ are still required to create thresholds. This RSACE process is thus performed after the Φ execution, but before the Ψ execution, and is performed in several stages: predicate cardinality extraction, vectorization, QEP grouping, clustering, and finally a false positives support stage. We summarize the process of the RSACE module in Algorithm 2 and provide references below with specific line numbers.

input : The set of random subqueries to be executed Ψ
output: Cluster representatives \mathfrak{R} , or set of residuals Ψ'

```

1 for  $\psi \in \Psi$  do
2   |  $\psi.planID = parsePlanID(\psi.QEP)$  ;
3   |  $\psi.predJoin = extractFromActuals(\psi.parentQuery)$  ;
4   |  $\psi.predLocal = parseLocalPred(\psi.QEP)$  ;
5   |  $\psi.predVector = vectorizePredicates(\{\psi.predJoin\} \cup \{\psi.predLocal\})$  ;
6 end
7  $G = \Psi.groupBy(\psi \rightarrow \psi.planID)$  ;
8 Let  $\mathfrak{R} = \{\}$  be the set of cluster representatives ;
9 for  $g \in G$  do
10  |  $\mathfrak{R} = \mathfrak{R} \cup \{hCluster(g)\}$  ;
11 end
12 if not falsePositiveSupport then
13  | return  $\mathfrak{R}$  ;
14 else
15  |  $distExec(\mathfrak{R})$  ;
16  | Let  $\Psi' = \{\}$  be set of optimal representative residuals ;
17  | for  $r \in \mathfrak{R}$  do
18  |   | if  $r.execTime \leq getOptimizerCounterpart(r).execTime$  then
19  |   |   |  $\Psi' = \Psi' \cup flatten(r.cluster)$  ;
20  |   |   end
21  |   end
22  | return  $(\Psi')$  ;
23 end

```

Algorithm 2: Algorithm for RSACE Module.

In the predicate cardinality extraction stage, we extract all relevant information pertaining to predicates, including the values, tables, and cardinalities, all represented in a pseudo-RDF type format. These can be of two different types, local predicates, and join predicates. Since the cardinalities are of interest, we can obtain the actual join predicate cardinalities through the previously executed parent query (line 3). The query execution was the first step in the Learning process and therefore provides, not only the estimated but also the actual cardinalities of all the joins. Since queries subsume sub-queries, a sub-query’s join will always be part of its parent, and thus the cardinality for any given sub-query will also be available. The local predicates were obtained during the sub-query generation phase, as with GALO, and are therefore easily parsed (line 4). Lastly, the *planID*, a table-independent hash value that defines the structure of the QEP, is parsed from the explain file (line 2).

During the vectorization step (line 5), the goal is to create a query to integer vector mapping for all Ψ sub-queries. The parsed information from the previous step is available in the pseudo-RDF format and allows for easy manipulation. The *vectorizePredicates* function extracts the sub-query’s predicate, hashes it, and saves it in a tree map; storing information in a sorted manner according to the natural ordering of the keys. This ensures that predicate cardinalities from different sub-queries are in the same indexed position of the vector. Each vector thus represents an ordered set of predicate cardinalities that can also be compared with other sub-queries with the same predicates.

A set of vector groups, G are created (line 7) to ensure that only sub-queries with the same predicates are compared. The grouping is based on the sub-query’s QEP structure, such that sub-queries with identical QEPs are grouped together. This

is achieved using the *planID* hash previously obtained. Since all Ψ sub-queries are derived from Φ sub-queries with guideline alterations, we observe a grouping of sub-queries with the identical guidelines, but not always derived from the same ϕ group. The process described so far can be summarized in Figure 3.4. In this example the actual cardinalities of join predicates **ws_item = i_item** and **ws_sold_date = d_date** are obtained from the parent query’s QEP. The remaining predicates, **category = 'Jewelry'**, **category = 'Music'**, **d_date = '2016-01-02'**, and **d_date = '2016-12-01'**, have their estimated cardinalities calculated using the probe queries.

For each group in G , clusters of alike vectors are created (line 10), wherefrom within each, a single representative is selected. The sub-query representative should thus closely resemble all sub-queries within its cluster, since each is guaranteed to have the same execution plan, the same predicates, and now as we’ve established, very similar cardinalities. We can conclude that the runtime of the representative should be quite similar to the runtimes of the sub-queries within the cluster. The set of representatives $\mathfrak{R} \subseteq \Psi$ are the only needed sub-queries that require execution since we can interpolate the runtimes of their corresponding clustered sub-queries. The clustering within each group G can be accomplished using various existing clustering techniques, but we observed that most required the number of clusters. We opted for the Density Based Spatial Clustering of Applications with Noise (DBSCAN)[22] algorithm, since it can find clusters of arbitrary size, and most importantly does not require an initial number of clusters. DBSCAN finds the number of clusters starting from the estimated density distribution of corresponding nodes. It requires the neighbourhood of a point ϵ , and *MinPts* which is the minimum number of *density-connected* points required to form a cluster. Two points a and b are *density-connected*

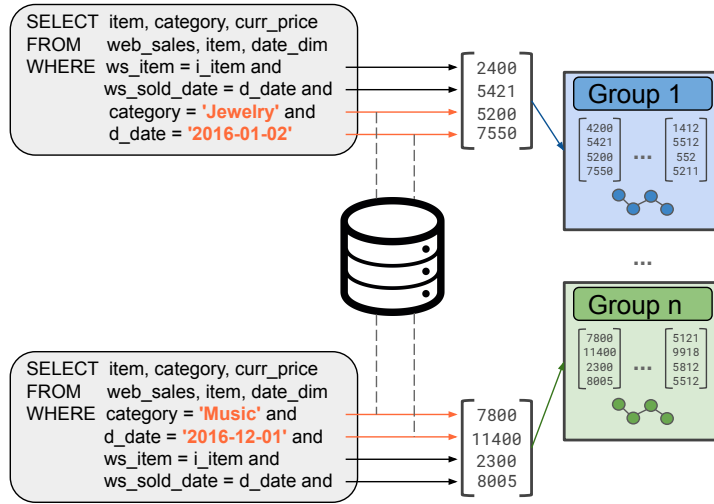


Figure 3.4: Random sub-query vectorization and grouping in RSACE.

if there exist a set of points $\{p_1, p_2, \dots, p_n\}$ such that $p_1 = a$, $p_n = b$, and p_{i+1} is *directly density reachable* from p_i . Two points are *directly density reachable* from one another if they are within each other's ϵ neighbourhoods. The remaining critical parts for the system are to select appropriate values for ϵ and *MinPts*. Since we observed relatively small groups and consequently clusters (varies with the number of random guidelines generated), a value of 0 for *MinPts* allows queries to be in their own cluster. This, in essence, means that outliers will be in the set \mathfrak{R} and as a result will be executed. This is acceptable since it will reduce the false positives/negatives.

In most cases, the representatives \mathfrak{R} adequately capture the runtimes but we must also consider times where they do not. The first of these is the false positive case, in which the ψ representative will **outperform** its ϕ counterpart, but is not the case for some or all ψ sub-queries in the corresponding cluster (residuals). The second is the false-negative case, in which the ψ representative will **under-perform** its ϕ counterpart, but some or all ψ sub-queries in the cluster outperform their ϕ counter-

parts. To remedy either case, the corresponding clusters could be re-run to correctly capture the runtimes of the sub-queries. This is however not a practical approach since this would encompass execution of all Ψ sub-queries, and defeat the purpose of the pruning. We must consider that, in general, ϕ tends to outperform ψ sub-queries, so quantitatively there will be less false positives than negatives. We can thus partially remedy these cases by executing all sub-queries in clusters of ψ representatives that outperformed their ϕ counterparts (line 15 - line 22). This ensures that all false positives are not miss-representing their cluster, and lead to faulty templates. The execution of the residuals is also done in a distributed fashion. False negatives are costly to address and thus we consider this approach a pruning rule, whose outcome depends on the quality of the clusters. This is attainable through parameter tuning, or even different clustering algorithms.

Chapter 4

Experimental Evaluation

In this section, we present the experimental evaluation of DistGALO in the following three categories:

1. *Scalability.* We demonstrate the ability of DistGALO to more efficiently find problem patterns and handle larger databases previously unfeasible in GALO.
2. *Distributed Performance.* We evaluate several cluster parameters and various partitioning strategies.
3. *Pruning Effectiveness.* Finally, we present the effectiveness of the various pruning strategies employed, and what trade-offs they may present to the user.

Upholding the previous system GALO as a benchmark, we demonstrate the ability of our *Distributed Learning Engine* to more efficiently find templates. Furthermore, we display the system’s ability to better handle larger data that was previously inconceivable to run in a reasonable amount of time. Finally, we showcase our pruning strategies and some of the benefits and inherent trade-offs they provide users. All experiments, with exceptions (noted), were conducted on the SOSCIP Cloud Computing Platform[21] using a cluster consisting of 9 nodes (including master). Each

node consists of 4 Intel Core Processor (Haswell, no TSX, IBRS) 2.299 GHz Virtual CPUs, and 16GB of RAM. Experiments were conducted using the synthetic TPC-DS benchmark consisting of 25 tables, and 99 queries. All sub-queries considered in the evaluation consist of single, double, and triple joins, as was previously done in GALO and has shown to be most suitable for the synthetic data, but can otherwise be adjusted to any number of joins.

4.1 Distributed Learning Engine Evaluation

Exp-1: Revised Learning. One of the main motives for developing DistGALO was to address the lengthy learning time observed in GALO. Ultimately the goal was to reduce this learning time as much as possible, so as to permit larger data and more complex queries. In this experiment we compare the elapsed times of DistGALO and GALO from start (workload submission) to the end (template creation) of their respective (distributed) learning engines. The underlying learning process largely remains the same, mainly the decomposition of workload queries into their sub-query counterparts. Parameters across systems were kept similar including, three runs per sub-query for noise reduction, and 2/3/4way join sub-query decomposition.

The results are reported in Figure 4.1. The learning elapsed times on a per-query(Figure 4.1a) and a per-sub-query(Figure 4.1b) basis are reported as per GALO’s *Learning Scalability and Effectiveness* experiment[4]. Though the focus previously was on the time complexity with respect to join-number, here we focus on the raw runtimes of each. First considering the time it takes to learn a single query(Figure 4.1a) we observe a drastic drop from 166218 to 54.2 seconds for 2way join sub-queries, surmounting to a 306442% speedup. We observe consistent results with similarly ex-

System	Table Join Size		
	2way	3way	4way
GALO	166218	315438	729948
DistGALO	54.2	297.8	89.0

(a) Learning time on a per-query basis

System	Table Join Size		
	2way	3way	4way
GALO	27702	54090	88374
DistGALO	1.2	3.0	1.3

(b) Learning time on a per-subquery basis

Figure 4.1: DistGALO’s learning elapsed time (seconds) versus DistGALO using TPC-DS benchmark queries.

treme speedups of 105831% and 819957% for 3way and 4way sub-queries respectively.

Analyzing the runtimes on a per-query basis however, doesn’t accurately depict the overall speedup achieved through DistGALO. Different queries produce varying amounts of sub-queries at the varying levels of join-numbers, thus we only observe an average. Let us consider the amount of time taken to learn a single sub-query (Figure 4.1b). GALO was previously reporting 27702(7.6 hours), 54090(15.0 hours), 88374(24.5 hours) for 2way, 3way, and 4way join sub-queries respectively. All the runtimes were reduced to a sub-5-second range using the current Distributed Learning Engine.

The drastic difference can thus be attributed to the various partitioning strategies, and other optimizations like multi-executor concurrency execution, selective thresholds, and others mentioned throughout this work. The results are devoid of any pruning rules to maximize template creation, resulting in 108 templates found, in comparison to the 98 found in GALO over the TPC-DS 99 queries. The discrepancy we attribute to the randomness of the generated sub-queries since no fundamental

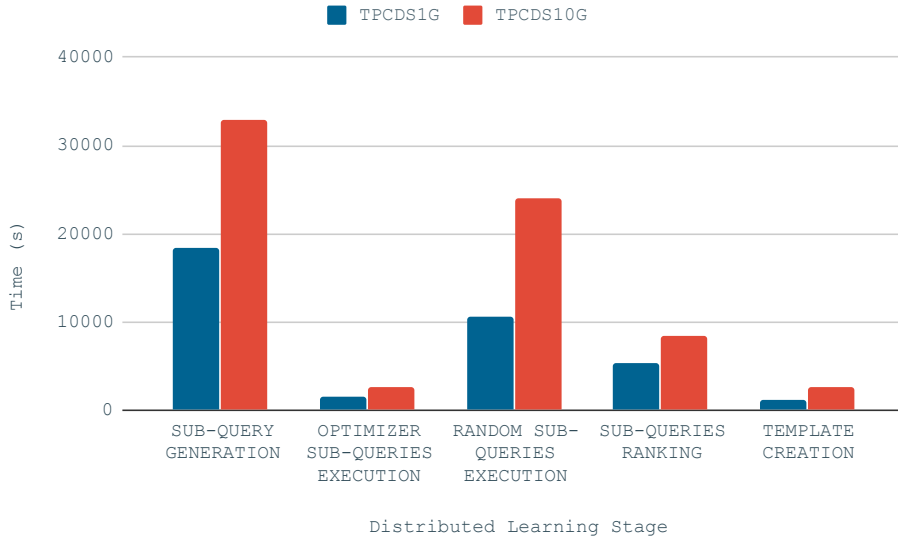


Figure 4.2: Distributed Learning with 1G and 10G TPC-DS databases.

process was changed in how sub-queries are generated, ranked, and used to create templates. The DistGALO average performance improvement over the problematic sub-queries is 35% in comparison to the GALO reported 37%. This we too attribute to the random nature of the system when creating sub-queries as was done for both GALO and DistGALO.

Exp-2: Database Scalability. Next we consider how DistGALO is able to scale in accordance to an increased database size of 10GB. This was an inconceivable task with the previous GALO system, taking upwards of 48 hours before it was terminated by force. We, therefore, are able to showcase a feature beyond just improvement, but rather a whole new ability to tackle large data. This we hope is one step closer to bringing the system to meet modern-day data requirements, with some exhibiting sizes beyond the petabyte range. However, due to limited hardware available, we have opted for a 10x approach, showcasing the ability to handle a database size of

10GB versus the 1GB TPCDS benchmark.

We only consider the Distributed Learning Engine since the Matching Engine remains largely unmodified, and was proven to scale well[4]. We thus split up the Learning Engine process into several smaller stages: *sub-query generation*, *optimizer sub-query execution*, *random sub-query execution*, *sub-query ranking*, and *template creation*. We further analyze the runtime of each of these and discuss the observations made.

Figure 4.2 displays the various stages of the Distributed Learning Engine in a 1GB versus a 10GB environment over the TPC-DS benchmark. The experiment demonstrates how DistGALO is effectively able to handle a database scaling of 10X. Our results reveal a learning runtime increase from approximately 10.4 hours for TPCDS 1G, to 20.2 hours for TPCDS 10G. This 2X scalability experiment shows how effective the system handles an increase in data size, which is much more in line with current-day demands. Two stages in the learning process are worth mentioning, mainly the *sub-query generation* and *random sub-query execution*, since they take a considerably longer time than the remaining. The latter proved to be a major bottleneck in GALO and therefore was the primary area of focus for optimization in DistGALO. This we believe was successfully achieved since it no longer stands to be the bottleneck of the learning process. Consequently, however, we observe a new bottleneck, albeit smaller than previous, emerge. In both 1G and 10G TPC-DS datasets, the *sub-query generation* makes up for 50% of the total learning runtime in both cases. It currently stands to be the main time-sensitive obstacle largely attributed to the probing queries used to obtain local predicate ranges. This step though already distributed will likely be the main focus of our future work.




Figure 4.3: Random sub-query execution with different number executors.

4.2 Distributed Performance

Exp-3: Executor Impact. Despite distributing the work amongst nodes, we are able to leverage the parallelism within nodes in order to further optimize. This is achieved by utilizing distributed workers, or *executors*, in order to make use of all available cores on each node. Each executor, in the context of DistGALO, is responsible for executing a single sub-query, ranking all runs, and storing the results in the distributed file system. We experiment with various executors, starting at 9, the number of nodes, up to 36, the total number of virtual CPUs on the cluster. The results are shown in Figure 4.3.

The results are quite intuitive since with 9 executors, we are only able to achieve serial execution on each node using 1 VCPU, and the remaining 3 on each node are left idle. This highly under-utilizes resources and shows with a higher overall execution time. We then observe a gradual decrease as more executors are added until we reach

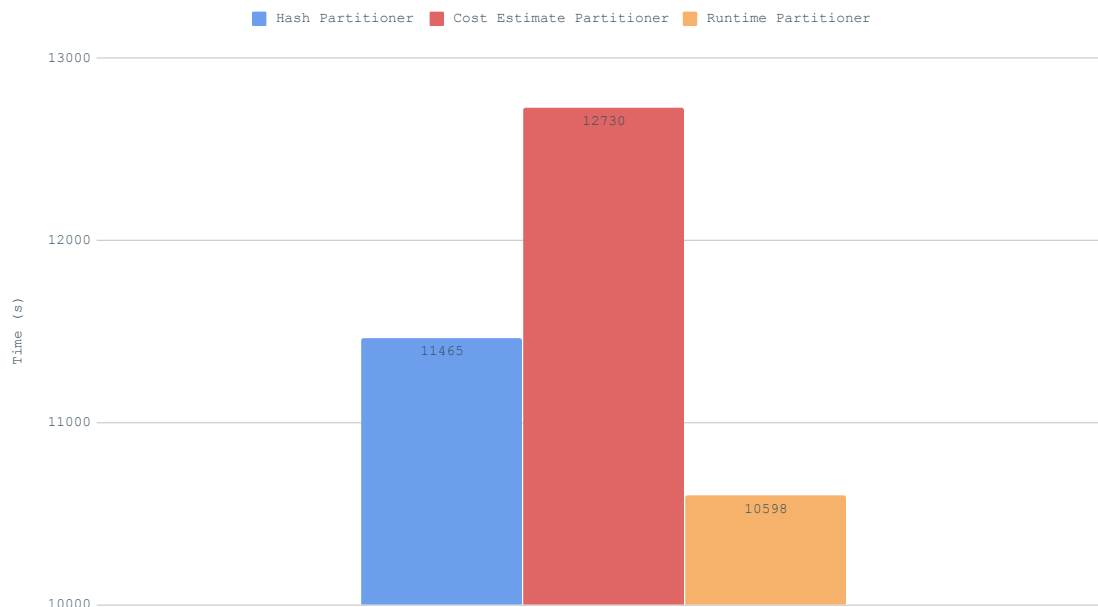


Figure 4.4: Random sub-query execution using Hash, Cost Estimate, and Runtime Partitioners.

36, the total number of cores available. These results though not quite surprising have great importance. What we can deduct is that regardless of how the system scales, be it vertical or horizontal, so long as the number of executors is increased, there will be a definite gain. This gives great flexibility to users in terms of scalability, since some have more powerful nodes but lesser in number, whilst others may have less powerful nodes but in greater quantity. In either case, DistGALO will ensure that all CPUs across the cluster are fully utilized.

Exp-4: Partitioner Preference. In order to minimize the amount of skew exhibited by the cluster, nodes must be balanced as evenly as possible. Skew can be greatly reduced by minimizing communication cost between nodes, and partitioning allows that fine-grained control. Partitioning allows DistGALO to group certain

sub-queries together, selected by the partitioning strategy. We previously discussed the three strategies (Section 3.5) employed by the DLC: Hash, Cost Estimate, and Runtime partitioners. We observe the effect these various partitioners have on the sub-query execution in Figure 4.4.

Note that the differences in partitioner execution times are highlighted in Figure 4.4 by limiting the y-axis minimum runtime to 10000 seconds. The results indicate that the Runtime partitioner is overall best performing with an overall runtime of 10598 seconds, followed by the Hash partitioner with 11465 seconds, and lastly by the Cost Estimate partitioner with 12730 seconds. Intuition dictates that the Cost Estimate partitioner should outperform the Hash partitioner since the former demonstrates some level of intelligence through the use of timeron values, while the latter simply uses the hash values to group sub-queries into their respective partitions.

Analysis of the different strategies reveals an interesting correlation between performance and sub-query distribution among the partitions. The standard deviation, that is the amount of variation between partitions in their number of assigned sub-queries, was observed as follows. The Runtime, Hash, and Cost Estimate partitioners exhibit standard deviation values of 0.48, 15.81, 17.13, respectively. This ranking order is also expressed in the runtime performance of the partitioners, suggesting there is perhaps some relationship. This observation indicates that a more evenly distributed set of partitions, result in less skew, implying that sub-queries are executed in very similar lengths of time.

Further analysis shows that 89% of random sub-queries are cut off by their thresholds, most of which are capped at 10 seconds. This reinforces the correlation between standard deviation and runtime described previously and also describes the

performance results of the petitioners. The Hash partitioner attempts to distribute sub-queries based on their hash values, so in general, the distribution is somewhat balanced, though not perfect. The Cost Estimate partitioner, on the other hand, attempts to balance out the complexity between partitions, and as a result, packs many low complex sub-queries into a single partition, and leaves fewer more complex ones in others. This heuristic is sound and would work well if sub-queries were allowed to complete in their required time, but are instead cut off by the thresholds set in place. It is for this reason that the more significant skew and higher runtime are exhibited, and why it is not the partitioning strategy of choice. The Runtime partitioner, on the other hand, distributes sub-queries based on their optimizer sub-query counterpart thresholds, and as such has the most balanced distribution among partitions, and thus the lowest standard deviation. The runtime, the metric evaluated in this experiment, is consequently the lowest, and thus indicates that the Runtime partitioner is the optimal partitioning strategy in executing the random sub-queries.

4.3 Pruning Effectiveness

Exp-5: RSACE Effectiveness. DistGALO aims to provide more versatile and flexible parameters for users to adjust. The RSACE module prunes sub-queries of similar structure and predicates, so that two very similar sub-queries need not be executed twice. The pruning process relies on clustering groups of sub-queries and selecting representatives \mathfrak{R} to be executed. The better the representative can summarize members of its group, the more can be inferred about those members, and therefore the smaller the search space. Consider that any clustering algorithm can be chosen, but in our experiments, we opted for the DBSCAN clustering algorithm

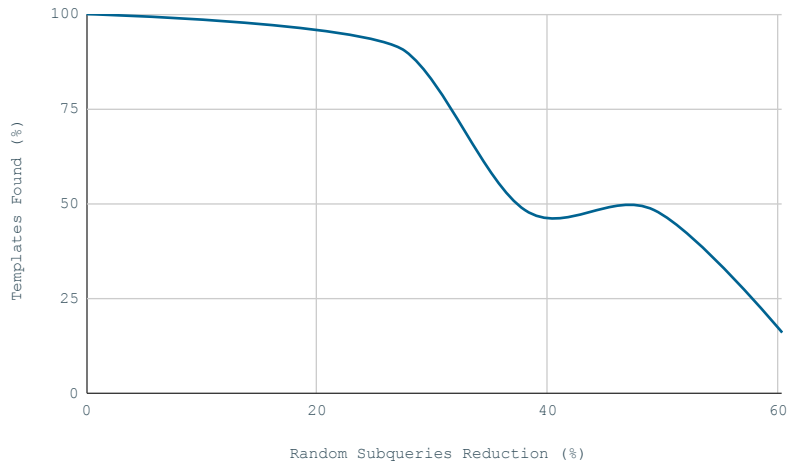


Figure 4.5: RSACE random sub-query reduction and its effect on number of templates found.

[23] since it is able to find clusters of arbitrary sizes, and does not require an initial cluster size. DBSCAN’s aggressiveness can be adjusted with a larger neighbourhood range, controlled by the ϵ parameter. By increasing the ϵ value, we more aggressively group sub-queries, thus resulting in fewer representatives and therefore less random sub-queries that must be executed. The aggressiveness of the pruning will have some effect on the accuracy, or the number of templates discovered in the context of Dist-GALO. With a more aggressive approach intuition dictates a lower runtime, with a compromise on the number of templates discovered. The unknown is how expensive that trade-off is, and at what intervals. The effectiveness of the RSACE pruning strategy with varying aggressive parameters can be seen in Figure 4.5.

The worst-case result for the experiment would be a $-\log$ effect, where minor pruning adjustments cause a drastic loss in accuracy. Our results show an almost inverse effect for the range of (0%, 27%] and (38%, 48%] of pruned random sub-queries. The former indicates that an ideal selection of aggressiveness is at an ϵ value that

prunes 27% of random sub-queries, while still identifying 90% of the templates that would otherwise be discovered without RSACE. Beyond this range, there is a large loss in the number of templates found with only 48% being discovered with 38% of sub-queries pruned. Following that we exhibit an odd behaviour in which higher pruning results in more templates being discovered. This can be attributed to the possibility of a cluster restructure causing different representatives to be selected, marginally more representative of the groups, thus resulting in more templates. This we consider an anomaly as the trend continues past 50% pruning. We conclude that the RSACE module is quite effective as it can provide a 27% speedup for a 10% loss in accuracy.

Chapter 5

Related Work

The classical cost-based method for selecting an efficient query plan, dates back to the the predecessor of all modern day relational databases, System R[24]. This relational approach to database management pioneered the cost-based optimizer methodology of obtaining a low cost means for query execution. The optimizer's cost based metric depends on disk page accesses, including CPU instructions, with an effort to minimize the number of pages being fetched from secondary storage. This early work applied a bottom-up dynamic programming plan enumeration technique to efficiently create a sub-set of plans from a massive search space. Since then, many advancements have been made in query optimization techniques, with modern-day systems regarding a high demand since the Big Data Revolution[25].

One approach to advancing the optimization process is to devise new strategies for pruning the join order search space. Finding the optimal access plan is an NP-hard problem, and thus the classic Dynamic Programming approaches switch to heuristic or randomized methods to resolve complex, high join order, queries. Methods have been devised, so that when a query becomes too complex to be optimized accurately, the

join graph is reduced to a simpler one until it becomes tractable within a given time budget[26]. Alternative methods introduce a top-down join enumeration algorithm, accompanied by various branch-and-bound pruning techniques[23].

Another heavily researched approach has been to improve the accuracy of cardinality estimation, as it is a chief determinant of the overall cost of a plan. The traditional and dominant histogram-based approach independently calculates selectivities of local predicates, while disregarding statistical correlations when calculating the net selectivity. This results in inaccurate cardinality estimates with more complex and often real-world data sets[27][28], since they don't fit the assumptions of uniformity, independence, and inclusion that optimizers make. Several sampling methods[29][30][31] attempt to circumvent this issue since they are better able to capture data correlation. A recent approach uses index-based join sampling, where more accurate cardinality estimates are derived from existing index structures and sampling[32]. This technique takes advantage of the recent advancements of in-memory databases and leverages sampling with a designated operator that utilizes already-existing index structures.

Recently Machine Learning has been a catalyst in many fields within research and development, and query optimization has not been an exception[33]. Neural Networks have been used to accurately estimate the selectivity of queries over highly skewed or correlated data[34]. Other Deep Learning techniques[35][36][10] structure the join ordering as a reinforcement learning problem to obtain the query plans, while others apply supervised learning to solve cardinality estimation in isolation[37]. Machine Learning has proved an effective enough method and could revolutionize the database optimizers of the future.

The discussed methods thus far have been an integral component or modification

of the cost-based optimization stage of the query compiler. The matter of fact is that though there has been considerable improvement, the optimizer will not always be able to pick the optimal access plan, and domain experts will consequently always be needed to troubleshoot such cases. It is the troubleshooting stage for which this work makes an effort to automate, not the query optimization itself. This third-tier optimization step, we term the *access plan rewrite* stage, in which the chosen optimizer plan is rewritten such that subsequent executions suggest a more optimal plan to the optimizer. The first undertaking, OptImatch[1][2] aimed to aid domain experts by letting them build problem patterns, store, and later retrieve to automatically apply to sub-optimal queries. This proved to be a useful tool but still required some manual input in order to populate the knowledge base. To remedy this, we devised GALO[4][3], a system capable of automating the entire process from end to end, without any domain expert intervention. It was capable of automatically discovering and saving problem patterns into a knowledge base, later applied as rewrites to non-optimal access plans. It still filled the requirements of OptImatch in that it acted as a query re-optimization tool, applied post query compiler evaluation. A similar strategy of query re-optimization[38] aims to automatically provide refined, more accurate, sampled cardinalities to sub-optimal queries. This allows the optimizer to take the updated cardinalities and make a new, more informed decisions, resulting in more optimal access plans. This process would be repeated until the refined cardinalities being fed no longer lead to a different access plan. The difference between the latter mentioned approach and this work is that DistGALO relies on real run-time statistics of sub-queries, and is thus able to make objective statements that are not influenced by optimizer heuristics.

Chapter 6

Conclusion

In this work we improve upon the well-received GALO system with DistGALO by re-vamping the Learning Engine into a Distributed Learning Engine. We utilize modern-day scalable technologies[25] like Apache Spark and Hadoop Distributed File System, to effectively and efficiently populate its knowledge base. DistGALO is able to leverage the availability of cloud machines to allow on-demand horizontal scalability to fit the workload domain requirements. Using various partitioning strategies, we effectively minimize the skew throughout nodes. We also introduce pruning strategies that decrease the redundancy throughout executions, whilst maintaining its effectiveness in learning problematic patterns. Despite the numerous structural changes, the technique remains the same, and DistGALO can still be regarded as third, *plan rewrite*, stage of query optimization. New bottlenecks have emerged, mainly the sub-query generation, as we find the predicate probing queries to be expensive. New methods for generating sub-queries and translating them to RDF could be explored, potentially accomplishing it directly from the QGM without any intermediate parsing. Automated scaling could also trivially be applied so that heuristics and metadata

could be obtained from query workloads, to automatically adjust the cluster size as needed. The growing popularity of containerization could aid in the dynamic ability to load balance and scale DistGALO. The system has proved to be very effective and continues to evolve as the technologies it relies on also evolve.

References

- [1] G. Damasio, P. Mierzejewski, J. Szlichta, and C. Zuzarte, “Optimatch: Semantic web system for query problem determination,” in *2016 IEEE 32nd International Conference on Data Engineering (ICDE)*, IEEE, 2016, pp. 1334–1337.
- [2] —, “Query performance problem determination with knowledge base in semantic web system optimatch.,” in *EDBT*, 2016, pp. 515–526.
- [3] G. Damasio, S. Bryson, V. Corvinelli, P. Godfrey, P. Mierzejewski, J. Szlichta, and C. Zuzarte, “Galo: Guided automated learning for re-optimization,” *Proceedings of the VLDB Endowment*, vol. 12, no. 12, pp. 1778–1781, 2019.
- [4] G. Damasio, V. Corvinelli, P. Godfrey, P. Mierzejewski, A. Mihaylov, J. Szlichta, and C. Zuzarte, “Guided automated learning for query workload re-optimization,” *arXiv preprint arXiv:1901.02049*, 2019.
- [5] U. Sivarajah, M. M. Kamal, Z. Irani, and V. Weerakkody, “Critical analysis of big data challenges and analytical methods,” *Journal of Business Research*, vol. 70, pp. 263–286, 2017.
- [6] R. Agarwal and V. Dhar, *Big data, data science, and analytics: The opportunity and challenge for is research*, 2014.

- [7] J. Gryz, Q. Wang, X. Qian, and C. Zuzarte, “Sql queries with case expressions,” in *International Symposium on Methodologies for Intelligent Systems*, Springer, 2008, pp. 351–360.
- [8] M. Jarke and J. Koch, “Query optimization in database systems,” *ACM Computing surveys (CsUR)*, vol. 16, no. 2, pp. 111–152, 1984.
- [9] P. Roy, S. Seshadri, S. Sudarshan, and S. Bhoje, “Efficient and extensible algorithms for multi query optimization,” in *ACM SIGMOD Record*, ACM, vol. 29, 2000, pp. 249–260.
- [10] J. Ortiz, M. Balazinska, J. Gehrke, and S. S. Keerthi, “Learning state representations for query optimization with deep reinforcement learning,” *arXiv preprint arXiv:1803.08604*, 2018.
- [11] N. Bruno, S. Chaudhuri, and R. Ramamurthy, “Power hints for query optimization,” in *2009 IEEE 25th International Conference on Data Engineering*, IEEE, 2009, pp. 469–480.
- [12] N. Bruno, S. Chaudhuri, and R. Ramamurthy, “Interactive plan hints for query optimization,” in *Proceedings of the 2009 ACM SIGMOD International Conference on Management of data*, ACM, 2009, pp. 1043–1046.
- [13] *IBM infosphere optim query workload tuner*. <https://www.ibm.com/support/knowledgecenter>, 2018. [Online]. Available: https://www.ibm.com/support/knowledgecenter/en/SSEPGG_11.1.1.0/com.ibm.db2.luw.idm.tools.doc/doc/c0057033.html.
- [14] B. Dageville, D. Das, K. Dias, K. Yagoub, M. Zait, and M. Ziauddin, “Automatic sql tuning in oracle 10g,” in *Proceedings of the Thirtieth international con-*

- ference on Very large data bases- Volume 30*, VLDB Endowment, 2004, pp. 1098–1109.
- [15] D. Simmen, E. Shekita, and T. Malkemus, “Fundamental techniques for order optimization,” in *ACM SIGMOD Record*, ACM, vol. 25, 1996, pp. 57–67.
- [16] *Openshift: Container application platform by red hat, built on docker and kubernetes*, <https://www.openshift.com/>, (Accessed on 10/21/2019).
- [17] *Production-grade container orchestration - kubernetes*, <https://kubernetes.io/>, (Accessed on 10/21/2019).
- [18] M. Zaharia, M. Chowdhury, M. J. Franklin, S. Shenker, and I. Stoica, “Spark: Cluster computing with working sets,” *HotCloud*, vol. 10, no. 10-10, p. 95, 2010.
- [19] M. Zaharia, R. S. Xin, P. Wendell, T. Das, M. Armbrust, A. Dave, X. Meng, J. Rosen, S. Venkataraman, M. J. Franklin, *et al.*, “Apache spark: A unified engine for big data processing,” *Communications of the ACM*, vol. 59, no. 11, pp. 56–65, 2016.
- [20] J. Dean and S. Ghemawat, “Mapreduce: Simplified data processing on large clusters,” *Communications of the ACM*, vol. 51, no. 1, pp. 107–113, 2008.
- [21] *Soscip — smart computing for innovation*, <https://www.soscip.org/>, (Accessed on 10/21/2019).
- [22] M. Ester, H.-P. Kriegel, J. Sander, X. Xu, *et al.*, “A density-based algorithm for discovering clusters in large spatial databases with noise,” in *Kdd*, vol. 96, 1996, pp. 226–231.

- [23] P. Fender, G. Moerkotte, T. Neumann, and V. Leis, “Effective and robust pruning for top-down join enumeration algorithms,” in *2012 IEEE 28th International Conference on Data Engineering*, IEEE, 2012, pp. 414–425.
- [24] M. M. Astrahan, M. W. Blasgen, D. D. Chamberlin, K. P. Eswaran, J. N. Gray, P. P. Griffiths, W. F. King, R. A. Lorie, P. R. McJones, J. W. Mehl, *et al.*, “System r: Relational approach to database management,” *ACM Transactions on Database Systems (TODS)*, vol. 1, no. 2, pp. 97–137, 1976.
- [25] A. Oussous, F.-Z. Benjelloun, A. A. Lahcen, and S. Belfkih, “Big data technologies: A survey,” *Journal of King Saud University-Computer and Information Sciences*, vol. 30, no. 4, pp. 431–448, 2018.
- [26] T. Neumann, “Query simplification: Graceful degradation for join-order optimization,” in *Proceedings of the 2009 ACM SIGMOD International Conference on Management of data*, ACM, 2009, pp. 403–414.
- [27] G. Lohman, “Is query optimization a “solved” problem,” in *Proc. Workshop on Database Query Optimization*, Oregon Graduate Center Comp. Sci. Tech. Rep, vol. 13, 2014.
- [28] V. Leis, A. Gubichev, A. Mirchev, P. Boncz, A. Kemper, and T. Neumann, “How good are query optimizers, really?” *Proceedings of the VLDB Endowment*, vol. 9, no. 3, pp. 204–215, 2015.
- [29] P. J. Haas, J. F. Naughton, S. Seshadri, and A. N. Swami, “Selectivity and cost estimation for joins based on random sampling,” *Journal of Computer and System Sciences*, vol. 52, no. 3, pp. 550–569, 1996.

- [30] P. J. Haas and A. N. Swami, “Sampling-based selectivity estimation for joins using augmented frequent value statistics,” in *Proceedings of the Eleventh International Conference on Data Engineering*, IEEE, 1995, pp. 522–531.
- [31] R. J. Lipton, J. F. Naughton, and D. A. Schneider, *Practical selectivity estimation through adaptive sampling*, 2. ACM, 1990, vol. 19.
- [32] V. Leis, B. Radke, A. Gubichev, A. Kemper, and T. Neumann, “Cardinality estimation done right: Index-based join sampling,” in *CIDR*, 2017.
- [33] W. Wang, M. Zhang, G. Chen, H. Jagadish, B. C. Ooi, and K.-L. Tan, “Database meets deep learning: Challenges and opportunities,” *ACM SIGMOD Record*, vol. 45, no. 2, pp. 17–22, 2016.
- [34] H. Liu, M. Xu, Z. Yu, V. Corvinelli, and C. Zuzarte, “Cardinality estimation using neural networks,” in *Proceedings of the 25th Annual International Conference on Computer Science and Software Engineering*, IBM Corp., 2015, pp. 53–59.
- [35] S. Krishnan, Z. Yang, K. Goldberg, J. Hellerstein, and I. Stoica, “Learning to optimize join queries with deep reinforcement learning,” *arXiv preprint arXiv:1808.03196*, 2018.
- [36] R. Marcus and O. Papaemmanouil, “Deep reinforcement learning for join order enumeration,” in *Proceedings of the First International Workshop on Exploiting Artificial Intelligence Techniques for Data Management*, ACM, 2018, p. 3.
- [37] A. Kipf, T. Kipf, B. Radke, V. Leis, P. Boncz, and A. Kemper, “Learned cardinalities: Estimating correlated joins with deep learning,” *arXiv preprint arXiv:1809.00677*, 2018.

- [38] W. Wu, J. F. Naughton, and H. Singh, “Sampling-based query re-optimization,” in *Proceedings of the 2016 International Conference on Management of Data*, ACM, 2016, pp. 1721–1736.