

Efficient Structure-aware OLAP Query Processing over Large Property Graphs

by

Yan Zhang

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Mathematics
in
Computer Science

Waterloo, Ontario, Canada, 2017

© Yan Zhang 2017

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

Property graph model is a semantically rich model for real-world applications that represent their data as graphs, e.g., communication networks, social networks, financial transaction networks. On-Line Analytical Processing (OLAP) provides an important tool for data analysis by allowing users to perform data aggregation through different combinations of dimensions. For example, given a Q&A forum dataset, in order to study if there is a correlation between a poster’s age and his or her post quality, one may ask what is the average age of users grouped by the post score. Another example is that, in the field of music industry, it may be interesting to ask what total sales of records are with respect to different music companies and years so as to conduct a market activity analysis.

Surprisingly, current graph databases do not efficiently support OLAP aggregation queries. In most cases, such queries are transformed to a sequence of join operations, and the system computes everything from scratch. For example, Neo4j, a state-of-art graph database system, processes each OLAP query in two steps. First, it expands the nodes and edges that satisfy the given query constraint. Then it performs the aggregation over all the valid substructures returned from the first step. However, in data warehousing workloads, it is common to have repeated queries from time to time. Computing everything from scratch would be highly inefficient.

Materialization and view maintenance techniques developed in traditional RDBMS have proved to be efficient for processing OLAP workloads. Following the generic materialization methodology, in this thesis we develop a *structure-aware* cuboid caching solution to efficiently support OLAP aggregation queries over property graphs. Structure-aware means that our solution takes both heterogeneous attributes and graph topological information into consideration. The essential idea is to precompute and materialize some views based on statistics of history workload, such that future query processing can be accelerated.

We implement a prototype system on top of Neo4j. Empirical studies over real-world property graphs show that, with a reasonable space cost constraint, our solution on average achieves 15-30x speedup over native Neo4j in time efficiency.

Acknowledgements

I would like to thank Professor M. Tamer Özsü, Dr. Xiaofei Zhang, Professor Semih Salihoglu and Professor Khuzaima Daudjee who made this thesis possible.

Dedication

This is dedicated to my mother Limei Leng whom I love.

Table of Contents

List of Tables	ix
List of Figures	x
1 Introduction	1
1.1 Property Graph Model	1
1.2 OLAP over Property Graph	4
1.3 Challenges of Graph OLAP	4
1.4 Our Solution and Contributions	5
2 Background and Related Work	8
2.1 OLAP over Property Graph Model	8
2.1.1 Graph OLAP Examples	9
2.1.2 Structure, Dimension, and Measure	11
2.2 Graph Databases and Neo4j	13
2.3 Related Work	13
3 Problem Definition	17
3.1 Terminology	17
3.1.1 Definition of Property Graph	17
3.1.2 OLAP Query	18
3.1.3 Materialization: Cuboid & Substructure	18
3.2 Problem Definition	20

4	Solution	22
4.1	Solution Framework Overview	23
4.2	Materialized View Selection	23
4.2.1	Overview of Materialized View Selection	24
4.2.2	Greedy Selection Framework	25
4.2.3	CubePlanner	27
4.2.4	Structure Planner	32
4.2.5	ID and Property Selection	34
4.2.6	Update on Materialized Views	35
4.3	Query Processing	36
4.3.1	Substructure Selection	37
4.3.2	Decomposition and Join	38
5	Evaluation	43
5.1	Experiment Setup	43
5.1.1	Datasets	43
5.1.2	Query Workloads	44
5.1.3	System Setting	45
5.2	Aspects of Interest	46
5.3	Results and Discussion	47
5.3.1	Our System vs. Neo4j	47
5.3.2	Frequency Threshold	49
5.3.3	Space Cost Limit	50
5.3.4	Storage Level for Materialized Views	51
5.3.5	CubePlanner vs PMA	52
5.3.6	StructurePlanner vs FPM	54
5.3.7	Substructure Selection	56
5.3.8	Decompose_Join	56
5.3.9	Reflections on Neo4j	59

6 Conclusion	61
References	63
APPENDICES	66
.1 Previous Workload	66
.2 Future Workload	67

List of Tables

2.1	A summary of Graph OLAP literature	16
3.1	Comparisons between Cuboid and Substructure.	20

List of Figures

1.1	A simple property graph modeling “users post posts” (data graph).	2
1.2	Meta graph containing User, Post and Tag.	3
1.3	A snapshot of data graph containing User, Post and Tag.	3
1.4	Neo4j’s execution plans for executing “getting the average post score grouped by users upvotes” for the first time and fifth time.	6
2.1	<i>Structure of Q1</i>	11
2.2	Cube of properties {A,B,C,D}.	12
2.3	<i>Structure of Q3</i>	12
2.4	An example cube in Cube-based [12].	14
4.1	Solution framework.	22
4.2	Neo4j’s execution plan for query <i>User-Badge, User-Post, Post-Tag: Tag.TagName</i> .	31
4.3	A substructure lattice with <i>Badge-User, User-Post, Post-Tag</i> as its root node.	32
4.4	“Border nodes” of structure <i>User-Post, Post-Tag</i> .	35
5.1	The meta graph of StackOverFlow used in experiments.	44
5.2	Time efficiency on the new workload: our solution vs Neo4j.	47
5.3	Substructure selected by StructurePlanner.	48
5.4	Total processing time under different settings of ω .	50
5.5	Total execution time under different space limits.	51
5.6	Main memory storage vs hard disk storage.	52

5.7	Time: CubePlanner vs PMA.	53
5.8	Total cuboid space cost: CubePlanner vs PMA.	53
5.9	Total processing time for new workload: StructurePlanner vs FPM.	54
5.10	Space cost: StructurePlanner vs FPM.	55
5.11	Processing time for each query: StructurePlanner vs FPM.	56
5.12	Processing time for Q10 - Q12 by three approaches.	57
5.13	Joining time in processing Q10 - Q12 by three approaches.	58
5.14	Total processing time for Q10 - Q12 by three approaches.	58
5.15	Total processing time vs trial query processing time.	59
5.16	Execution plans for <i>User-Post: User.Age</i> and <i>User-Post: ID(User), ID(Post)</i>	60

Chapter 1

Introduction

Being a flexible and semantic rich model for graph structured data, the property graph model has been widely adopted and graph database systems have emerged supporting this model, such as Neo4j [25], and PGX [11]. Supporting OLAP (On-Line Analytic Processing) is one critical feature of modern database systems, because efficient OLAP processing is fundamental to many decision-making applications, e.g., business intelligence [21], risk management [28], and trend monitoring [7]. However, our empirical study shows that existing graph database systems do not efficiently support OLAP workloads, especially structure-oriented aggregation queries. Moreover, current graph database systems do not support view-based queries or materialization of some “hot” intermediate results to serve new queries. In this thesis, we study the efficient processing of OLAP queries over property graphs using view materialization.

1.1 Property Graph Model

We are living in an age with exponential growth of data, and a world that is more and more connected. With the fast development of Web2.0 and Internet of Things (IoT) [26], numerous connections of various kinds are being created, producing massive amount of graph structured data. For example, the moment a user creates a new post on an online forum, not only a post is created, a “*creates*” connection between the user and the post is established as well; when a user tags a post, a “*has_tag*” connection is created between certain tag strings and the post; or in a banking scenario, when a transfer happens, a “*transfers*” connection between two accounts is created.

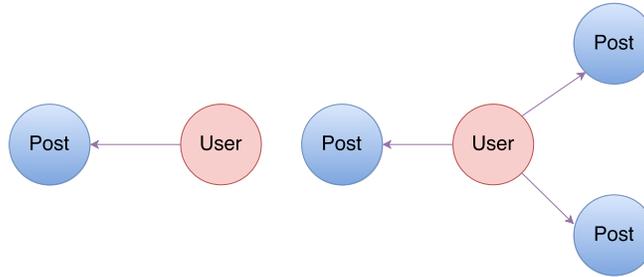


Figure 1.1: A simple property graph modeling “users post posts” (data graph).

To capture the rich semantics of connected real-world entities, property graph model [23] is becoming increasingly popular due to its flexibility for graph structured data. In a property graph nodes and edges can have any number of properties, or attributes, of any type. For example, Figure 1.1 shows a simple property graph of an online Q&A forum. It shows the connections among users (represented by red nodes) and posts (represented by blue nodes). Each edge from a user node to a post node represents a “creates” relation. From the graph, we can clearly see that there is one user who has created one post while the other user has created 3 posts. In addition, a User node can have properties like the user’s Age, Views, UpVotes etc. For presentation purposes, we shall use a property graph which contains information of an online Q&A forum www.stackoverflow.com through this thesis. We name this graph “StackOverFlow graph”.

Note that although the property graph model does not enforce any restrictions on what properties a node or edge can have, a high-level abstraction describing the property relations, named the “meta graph”, is often defined in practice. Meta graph demonstrates the information of entities and entity correlations at a schema level, while data graph refers to the actual graph populated from the meta graph. Figures 1.2 and 1.3 are the meta graph and a partial snapshot of the data graph of the StackOverFlow graph, respectively. As shown in Figure 1.2, there are three types of entities: User, Post, and Tag, which are colored in red, blue and green in the data graph. Each user has a property named “Age”, each post has a property named “Score”, and a property “TagName” associated with each tag. There are two types of edges: “creates” and “has_tag”. Note that figure 1.3 is a partial snapshot of the StackOverFlow data graph. The actual data graph contains millions of nodes and edges.

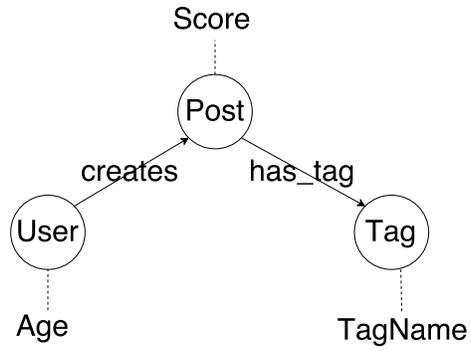


Figure 1.2: Meta graph containing User, Post and Tag.

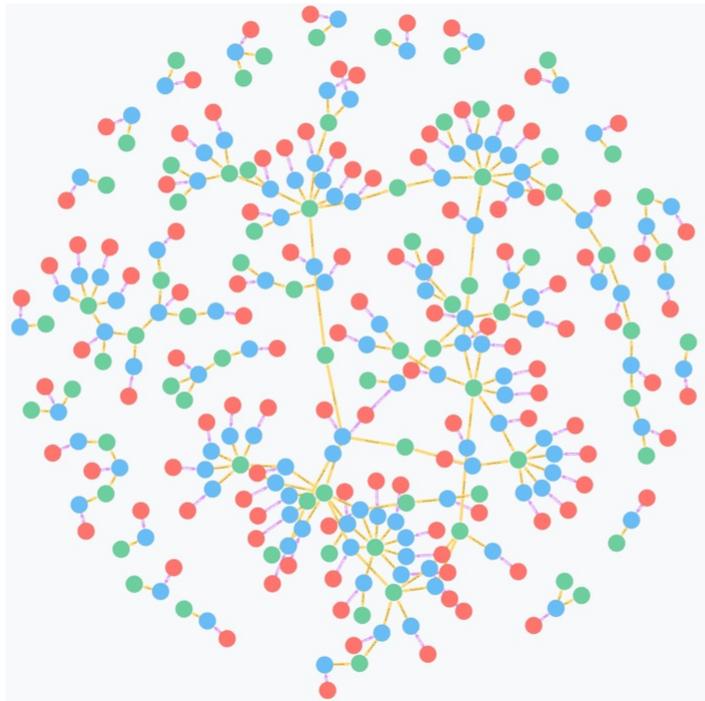


Figure 1.3: A snapshot of data graph containing User, Post and Tag.

1.2 OLAP over Property Graph

In traditional databases and data warehousing, OLAP queries enable users to interactively perform aggregations on underlying data from different perspectives (combinations of dimensions). There are three typical operations in OLAP. Drill-down operation allows users to view data in more detail, while roll-up operation does the opposite. Slicing enables filtering on data.

Supporting efficient OLAP processing on property graphs grants users the power to perform insightful analysis over graph structured data. For example, on the StackOverFlow graph, users can study the correlation between the number of UpVotes and a post’s score by using the following query:

Get the average post score grouped by users upvotes.

If the result shows a tight correlation, it suggests that an author’s upvotes can be used to estimate the quality of his or her post when a post is freshly posted and score of the post has not been settled.

Consider another example, using a property graph dataset on music industry, one can issue the following query to evaluate a company’s strategy to increase the share of young people’s market.

Get the total sum of music purchases by buyers at age 18-25 grouped by music company and month.

For simplicity, we call such OLAP query workloads over property graphs as “Graph OLAP”. Graph OLAP has already been applied in various scenarios like business analytics and decision making and it is attracting increasing research interests in the database community. In this thesis, we focus on supporting Graph OLAP by view materialization.

1.3 Challenges of Graph OLAP

Supporting efficient OLAP in traditional relational database management systems (RDBMS for short) [17] or warehousing applications is a well-studied topic. There is abundant literature attacking this problem from various different perspectives, e.g. data partition [5], view selection [14], partial materialization [6]. However, there are very few research efforts on Graph OLAP. Existing literature concerning OLAP workload over graph data either targets accelerating Graph OLAP over a special subset of property graphs [29], or addresses

generic high-level topics, such as an expressive graph aggregation operator for RDF graphs [18].

Our empirical study (reported later) shows that existing graph databases do not provide efficient support for Graph OLAP, especially when the graph size scales to real-world graph sizes, which usually contains over millions of nodes and edges. To elaborate, Neo4j, a state-of-art graph database, processes OLAP queries in a rather straightforward manner: computing each query from scratch without being aware of the previous workloads. Figure 1.4 reveals that even if we repeatedly execute the same query using Neo4j, the execution plan always stays the same and yields no execution time improvement.

Valuable information extracted from workload history can be helpful to accelerate incoming queries. For example, the above example OLAP query on StackOverFlow graph dataset (of roughly 45GB in size) takes Neo4j more than 2 hours to process. It is frustrating for users to wait that long for the result of one single OLAP query, as it undermines interactivity which is one of the most distinctive features of real-time analytics.

Previous workload provides useful information for new workloads. This is because users do not generally generate OLAP queries randomly. Instead they often tend to be interested in some specific “hot” structures on a meta graph level and some “hot” properties. Such interest is contained in workload history and can serve as a hint on new workloads. Our claim is that by sacrificing some memory space to materialize “hot” structures and properties, new queries can be processed faster.

The real challenge is how to design a score function to evaluate the trade-off between such benefits and the cost of materializing views so that we can select the best views for materialization. Here “best materialization” means that we achieve the maximum speedup in processing new queries within a given memory constraint for materialization.

1.4 Our Solution and Contributions

To address the challenges discussed above, we propose an end-to-end solution to support efficient Graph OLAP.

The essence of our solution is to precompute and materialize popular intermediate results that can be reused by new queries. We assume that most OLAP queries from the same client tend to reside in particular structures and properties (usually closely related with the topics that the client is interested in). Within a particular period of time, there are “hot” structures that the client tends to repeatedly investigate from different dimensions.

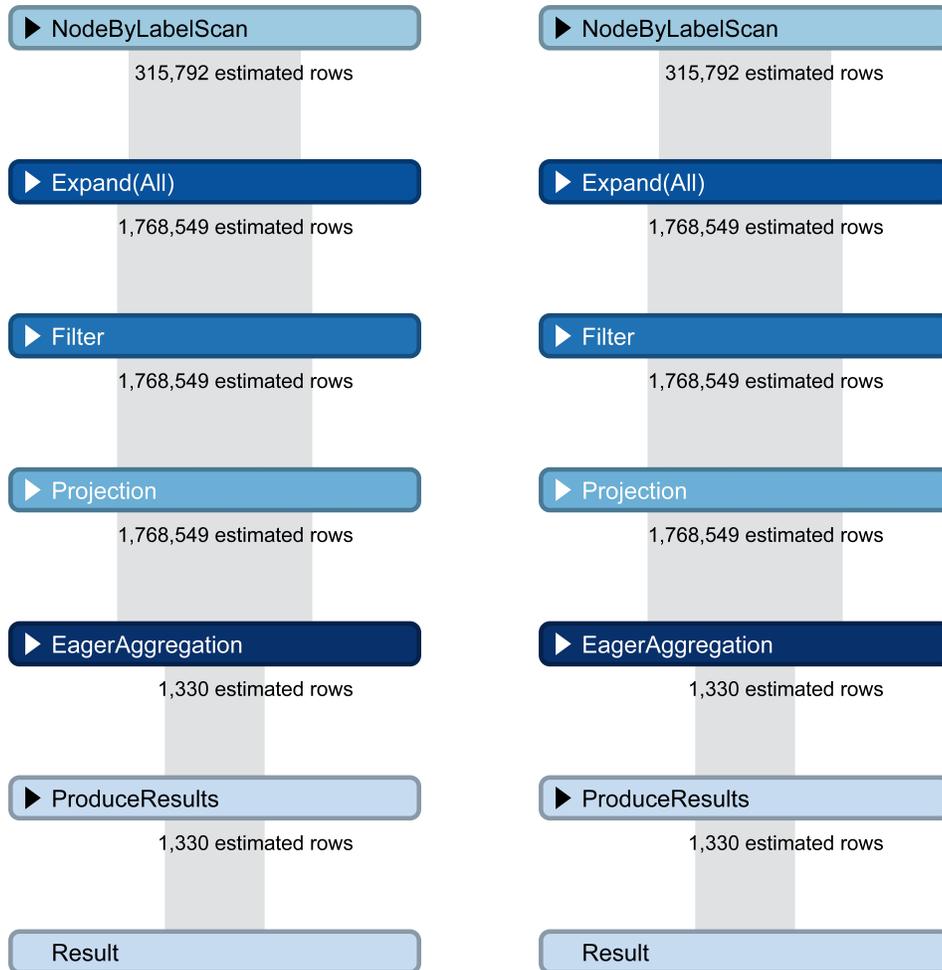


Figure 1.4: Neo4j’s execution plans for executing “getting the average post score grouped by users upvotes” for the first time and fifth time.

Therefore, previous queries can be used as a good reference to discover the structures and properties in which the client is particularly interested.

A good analogy of this is the establishment of materialized views in relational databases for query processing. In relational databases, we are allowed to build materialized views on structures and attributes that we are interested in. Hopefully when new queries come, we can process them much faster using materialized views. Unfortunately, current graph databases do not support similar operations.

There are two important problems that we need to solve. One key issue is to select the proper set of views to materialize. We need to select and pre-compute those that are most beneficial for new queries. Another key issue is how to come up with a better execution plan for answering a new query efficiently using the materialized views. To address the first issue, we develop a score function to evaluate cost/performance ratio of a materialization. We propose a greedy algorithm to select candidates based on their scores. For the second challenge, if a new query result can be directly answered using a materialization we simply do it. For other cases, we propose a scheduling policy to decompose an incoming query into substructures and join these substructures to produce the final result.

To highlight, we summarize our contributions in this thesis as follows:

- We design an end-to-end solution that realizes structure-aware OLAP query processing on graph databases using precomputation based on previous workloads.
- We proposed our algorithm for smart selection of structures and cuboids to be pre-computed.
- We investigate different query execution plans.
- We implemented our solution on top of Neo4j to test its effectiveness and efficiency.

The following contents are organized as follows. We discuss the preliminaries and related work in Chapter 2, which is followed by the background knowledge about OLAP, graph databases, and Neo4j. We also give a summarization of existing literature concerning OLAP queries over graph data. In Chapter 3 we explain our solution framework and system design in details. We present the experiment design and result discussion in Chapter 4. Chapter 5 concludes this thesis highlighting opening questions and future work.

Chapter 2

Background and Related Work

In this chapter, we first explain Graph OLAP with real examples. Then we briefly introduce Neo4j, a state-of-art graph database system, which is employed as the back end of our proposed solution. In addition, we review and summarize the most recent relevant works on Graph OLAP processing.

2.1 OLAP over Property Graph Model

As noted in the previous chapter, in the property graph model, each node and edge could have arbitrary number and type of properties. A type of property is represented as follows:

$$[NodeType].[PropertyType]$$

For example, `User.Age` denotes an “Age” attribute associated with a node of type “User”. In order to identify a node or edge, a unique ID is assigned to each node and edge. For simplicity, in this thesis we represent IDs of nodes or edges of a certain type as `ID([NodeType])` or `ID([EdgeType])`. For example, `ID(User)` refers to unique IDs of “User” nodes.

OLAP [2, 15, 1] performs aggregation over multiple table columns (or attributes). This is known as the cube [1, 15, 12], and it provides users a multi-dimensional and multi-level view for effective data analysis from different perspectives and with multiple granularities. The key operations in an OLAP framework are slice/dice and roll-up/drill-down, with slice/dice focusing on a particular aspect of the data, roll-up performing generalization if

users only want to see a concise overview, and drill-down performing specialization if more details are needed. For example, we can perform OLAP to analyze the profit performance of an international company by different branches. We can perform a drill-down operation by adding season as an extra dimension to branch to take a closer look at the profit performance of different branches in different seasons. Then we can perform a roll-up operation by discarding the branch dimension and only look into profit performance by seasons. Finally, we can perform a slicing operation by adding a condition, say summer, on season dimension, such that we only focus on profits in the summer season. In this case, OLAP serves as a tool for managers to better understand company performance.

2.1.1 Graph OLAP Examples

Graph OLAP was first proposed by Graph Cube [29] to refer to OLAP over graphs, although no formal definition of the notion Graph OLAP was given. In order to better elaborate how Graph OLAP is interpreted in this thesis, consider the following three example scenarios, where we perform OLAP queries over the StackOverFlow graph.

Example 1. Does the number of high upvotes of a user indicate a high-quality post?

Q1: Get average post score grouped by user’s upvotes.

Sample query result:

User.UpVotes	AVG(Post.Score)
0	1.33
1	2.23
2	2.34
3	2.77
4	3.43

From the query result we can see that upvotes can be used as a good indicator of a user’s post quality. Suppose we would like to propose suggested posts based on scores. When a post is freshly posted and the score of the post has not been voted, we may use the author’s upvotes as a factor to estimate the quality of his or her post.

Example 2. Following the context of Q1, suppose we want to take a closer look at how upvotes are correlated with post scores for different types of posts. It is reasonable to guess that upvotes make more difference in answers’ (posts with PostTypeId 2) scores than the questions’ (posts with PostTypeId 1), because a user’s level of expertise is usually better reflected by his or her answers.

Q2: Get average post score grouped by user's upvotes and post's post types.

Sample query result:

User.Upvotes	Post.PostTypeId	AVG(Post.Score)
0	1	2.14
1	1	2.26
2	1	2.83
3	1	3.04
4	1	3.46
0	2	1.54
1	2	2.21
2	2	2.18
3	2	2.72
4	2	3.58

The query results reveal that users with higher upvotes not only provide good answers but ask valuable questions as well. However, there is a subtle difference on how upvotes are correlated with scores of questions and answers. For example, a really low upvote level indicates a low-quality answer more than a low-quality question. This is probably because people tend to be tolerant of a naive question but not a wrong answer.

Q1 and Q2 simply focus on relationship between User and Post. We may switch our attention to a slightly more complicated structure by adding the Tag.

Example 3. In year 2017, what is the weighted average age of users? Is Python more trendy than C among young users?

Q3: For each tag, get the weighted average age of users who have posted posts in 2017 on the tag. In other words, group all the posts that were posted in 2017 by their tag names and calculate the weighted average of users' age under each tag.

Sample query result:

TagName	AVG(Age)
Router	19.6
Python	24.1
Internet	26.8
C	30.2
programmer	31.4
software	29.8

From the results, one can tell the average users’ age for each tag and compare them easily. It reveals some interesting insights. We know that Python may have a relatively young user group, compared to C. Similarly, “Router” is a topic that is more frequently discussed among younger users than the general topic of “Internet”. Note that we are led to the above findings based on the recent posts (posted in year 2017).

From the above OLAP query examples we can see that OLAP over property graphs provides an interactive and informative way to analyze property graphs from multiple dimensions, and thus helps people find the hidden correlations, aggregated effects, regularities, tendencies and so on.

2.1.2 Structure, Dimension, and Measure

We now explain the three key elements of a Graph OLAP: *structure*, *dimension*, and *measure* using Q1 as an example.

Q1 concerns the structure (colored in purple) shown in Figure 2.1.

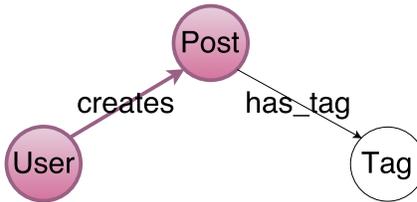


Figure 2.1: *Structure* of Q1

We say that (User)-[creates]->(Post) is the structure of Q1. The query is first aggregated on users’ upvotes. We say that User.Upvotes is the dimension of Q1, and the output of the query is an aggregation function on post’s score. We say that AVG(Post.Score) is the measure of Q1.

Similarly, consider Example 2, which shares the same structure as shown in Figure 2.1. The dimensions of Q2 is User.Upvotes, Post.PostTypeId, and the measure is AVG(Post.Score). Note that Q2 adds Post.PostTypeId to Q1’s dimensions. In other words, Q2 asks for more detailed partitions over dimensions. We call Q2 a drill-down from Q1, and Q1 is a roll-up from Q2. Note that possible property combinations can be modeled as a lattice-structured cube. Figure 2.2 shows what a cube is like for properties {A,B,C}. We can see that roll-up and drill-down operations allow us to navigate up and down on a cube.

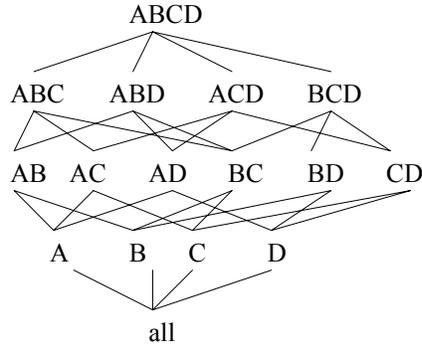


Figure 2.2: Cube of properties $\{A,B,C,D\}$.

We summarize Q3 as follows.

Structure: (User)-[creates]-(Post)-[has_tag]-(Tag)

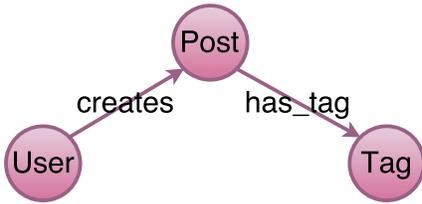


Figure 2.3: *Structure* of Q3

Dimensions: Tag.Tagname

Measure: AVG(User.Age)

Note that Q3 has a different *structure* than Q1 and Q2, as shown in Figure 2.3. Q3 enforces a requirement that posts must be created in year 2017, which picks out a particular subset of the posts. In OLAP this is called the “slicing” operation. Slicing operation allows users to view the data with constraints on selected properties. In this thesis we call these constraints (e.g., $\{\text{Post.Year}=2017\}$ in Q3) “*slicing conditions*”.

To summarize, Graph OLAP allows clients to aggregate different *structures*, over different *dimensions*, on different *measures*, and optionally slice the aggregation result by different *slicing conditions*. Clients can change their views by performing roll-up, drill-down, and slicing freely and interactively.

2.2 Graph Databases and Neo4j

Emerging online applications concerning graph processing has motivated the relational database community to support efficient graph management [27] [13]. However, there has been active debate about the efficiency of using traditional RDBMS for graph computing considering the unique query workloads against graph data [9] [10]. Relational and graph database systems both have their own strengths in terms of query processing. In this thesis, we work with graph database systems.

In graph databases, the storage structure that is most commonly used is an adjacency list. An adjacency list is a list that stores all the edges and neighbor nodes for a given node. A popular graph database system is Neo4j. Like most traditional RDBMS, Neo4j enforces atomicity, consistency, isolation and durability (ACID) properties. Neo4j’s native graph processing engine enables SQL like queries over property graphs which are managed by its own graph storage. Common features in RDBMS such as indexes are also supported. REST API plus official drivers for programming languages such as Java are provided for the sake of easy interactions with other programs. One most special feature about Neo4j is that it allows multiple-labeled nodes and edges; for example, a node referring to a male student could have various labels such as “student” and “male” and etc.

Cypher, Neo4j’s query language, is a powerful SQL like query language, which is highly expressive and simple. For example, consider the following query: what are the number of upvotes and associated average score of answers (posts with PostTypeID of 2)? A Cypher query would be written as follows:

```
match (u:User)-[r:creates]->(p:Post)
where p.PostTypeID='2'
return u.Upvotes, AVG(p.Score)
```

In the above Cypher query, “User” and “Post” are node labels, PostTypeID and Score are properties of “Post”, “UpVotes” is a property of “User”.

2.3 Related Work

There have been a few works discussing efficient Graph OLAP processing on property graphs or RDF graphs.

Cube-based [12] proposes a conceptual graph model to support OLAP on bibliographic networks. It proposes the concept of graphs enriched by cubes where each node and edge of the considered network are described by a cube. Figure 2.4 is an example cube in Cube-based [12]. It summarizes counts of co-authorships between two authors over dimensions of year, venue and keywords. Cube allows the user to quickly analyze the information summarized into cubes.

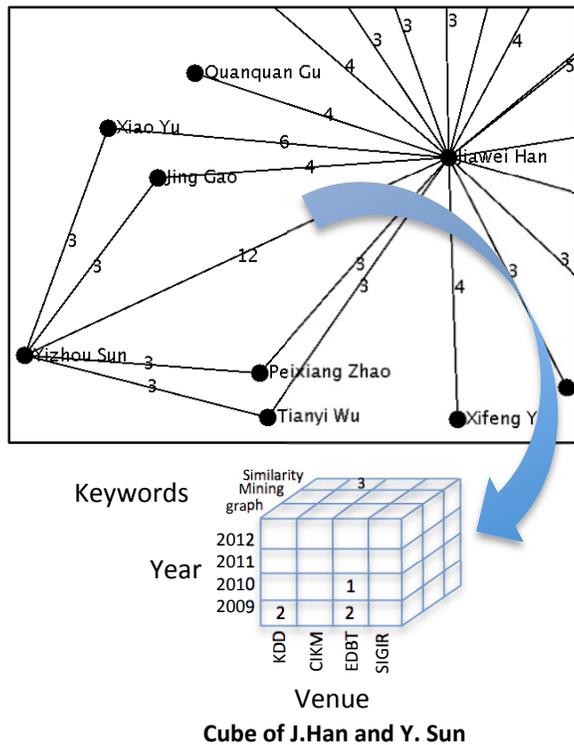


Figure 2.4: An example cube in Cube-based [12].

Gagg [18] introduces an RDF graph aggregation operator that is both expressive and flexible. It defines its operational semantics on top of SPARQL algebra, and proposes an algorithm to answer graph aggregation queries. Gagg achieves significant improvements in performance compared to plain-SPARQL graph aggregation.

Pagrol [24] provides an efficient MapReduce-based parallel graph cubing algorithm, MRGraph-Cubing, to compute cubes for an attributed graph.

Graph Cube [29] introduces a data warehousing model that supports OLAP queries

effectively on large multidimensional networks. It takes account of both attribute aggregation and structure summarization of the networks. In order to deal with “dimensionality curse”, a greedy algorithm is introduced for partial materialization of cuboids. Moreover, it addresses and defines the two most important notions in Graph OLAP scenarios as *dimension* and *measure*. Note that in our work, *structure* is addressed as a third important notion (as discussed in 2.1.2). Graph Cube [29] studies graphs with homogenous types of nodes and edges. Graph Cube [29] focuses on OLAP scenarios over a fixed *structure*, with *dimension* and *measure* varied, whereas our work deals with OLAP workloads over various *structures*.

Graph OLAP Framework [4] studies dimensions and measures in the Graph OLAP scenario and furthermore develops a conceptual framework for data cubes on graphs. It differentiates different types of measures (e.g., distributive measures and holistic measures) by their properties during aggregation. It looks into different semantics of OLAP operations, and classifies the framework into two major subcases: informational OLAP and topological OLAP. It points out a graph cube can be fully or partially materialized by calculating a special kind of measure called aggregated graph.

In Graph Cube [29], concepts of graph cube is introduced. Given a particular structure S , a property set P , and measure set M , we can aggregate over S on $2^{|P|}$ different combinations of dimensions. These $2^{|P|}$ queries can be mapped as a lattice structure, where each combination of dimensions corresponds to a cuboid in the lattice. We call the lattice structure of these $2^{|P|}$ queries a graph cube.

It has been pointed out in Graph OLAP Framework [4] that as long as the domain of measure is a subset of $\{\text{COUNT}, \text{SUM}, \text{AVERAGE}\}$ and M contains $\text{COUNT}(\ast)$, the following feature holds: given any two cuboids C_1 and C_2 from the same graph cube, as long as $\text{dimension}(C_2)$ is a subset of $\text{dimension}(C_1)$, result of C_1 can be used to generate result of C_2 . This means that once a cuboid is materialized, all roll-up operations from this cuboid could be processed simply by scanning the materialized cuboid result. This will dramatically decrease roll-up operation time compared to aggregation from data graph (often of larger size, disk I/O), scanning materialized cuboid result (often of smaller size) is often much faster.

Ideally we can materialize all cuboids. But when the number of dimension is large, number of cuboids grows exponentially, making total materialization hard due to overwhelming space cost. To solve this, Graph Cube [29] proposes a partial materialization algorithm on graph cube. It is a greedy algorithm and the score function is based on benefits of deduction of total computation cost.

We categorize the existing work into two lines. The first focuses on a simple subset

	G. Type	Q. Pattern	Feature
Cube-based [12]	Property	Simple relation	Cubes on edges and nodes
Gagg [18]	Property	SPARQL	Aggregation operator
Pagrol [24]	Property	edge & node attributes	Map-Reduce computing
Graph Cube [29]	Homogenous	node attributes	Partial materialization
Graph OLAP Framework [4]	Property	edge & node attributes	Distributive and holistic measures
Graph Cube [29]	Homogenous	node attributes	Partial materialization
Our Work	Property	edge & node attributes	Materialized views

Table 2.1: A summary of Graph OLAP literature

of property graphs (e.g., graphs with only homogenous nodes and edges), and proposes optimizations in order to accelerate OLAP query processing (e.g., Graph Cube [29]). Although attribute-aware optimizations are proposed in some literature, structure-aware optimizations are not studied. The second focuses on a relatively high-level framework that processes generic queries over generic property graphs (e.g., Gagg [18]). However, query processing efficiency is not the main focus of these works.

To conclude, there is no work on structure-aware optimizations for efficient Graph OLAP. As mentioned in Section 1.3, efficiency issue is one of the most challenging issues on Graph OLAP. Therefore, we investigate faster structure-aware OLAP processing over general property graphs.

Chapter 3

Problem Definition

In this chapter, we first introduce the terminology and notations used in this thesis. Then we formally define the efficient Graph OLAP problem.

3.1 Terminology

3.1.1 Definition of Property Graph

A property graph G is defined as $(V, Vid, E, Eid, A, L, f)$. V is the set of nodes. Vid is the set of unique IDs of each node in V . E is the set of edges, where $E \subseteq V \times V$. Eid is the set of unique IDs of each edge in E . A is a set of predefined properties. L is a set of predefined labels. $f = \{f_{VA}, f_{VL}, f_{Vid}, f_{EA}, f_{EL}, f_{Eid}\}$ is a set of mapping functions, such that:

- $f_{VA} : v_i \rightarrow A_i, v_i \in V, A_i \subseteq A$, maps each node to its properties;
- $f_{VL} : v_i \rightarrow L_i, v_i \in V, L_i \subseteq L$, maps each node to its labels;
- $f_{Vid} : v_i \rightarrow vid_i, v_i \in V, vid_i \subseteq Vid$, maps each node to its unique ID;
- $f_{EA} : e_i \rightarrow A_i, e_i \in E, A_i \subseteq A$, maps each edge to its properties;
- $f_{EL} : e_i \rightarrow L_i, e_i \in E, L_i \subseteq L$, maps each edge to its labels;
- $f_{Eid} : e_i \rightarrow eid_i, e_i \in E, eid_i \subseteq Eid$, maps each edge to its unique ID.

3.1.2 OLAP Query

As discussed in Section 2.1.2, four elements of a Graph OLAP query are *Structure*, *Dimension*, *Measure*, and *Slicing Condition* (optional). We now define how we represent these four elements in an OLAP query. We will use Q3 in Section 2.1.1 as an example.

Structure : A *structure* consists of *edges*. We write a *structure* by listing all its *edges* separated by comma, where an *edge* is represented by “*Starting Node Label - Edge Label - Ending Node Label*”. For instance, Q3’s *structure* as shown in Figure 2.3 is written as “*User-creates-Post, Post-has_tag-Tag*”.

Dimension: A *Dimension* is written by listing all properties that act as dimensions in an OLAP query. For example, Q3’s *dimension* is written as “*Tag.Tagname*”.

Measure: We focus on three most common types of *measure*: *COUNT*, *SUM* and *AVG*. Q3’s *measure* is written as “*AVG(User.Age)*”.

Slicing Condition: A *Slicing Condition* is written as “*property = value*”. Q3’s *slicing condition* is written as “*Post.Year=2017*”. In our work we only consider equality but it can easily be extended to other conditions such as the less than inequality.

With the four elements defined above, we write an OLAP query in the following format:

<i>Structure : Dimension, Measure, Slicing Condition</i>

Recall that Q3 is written as

<i>User-creates-Post, Post-has_tag-Tag: Tag.Tagname, AVG(User.Age), Post.Year=2017</i>

where *User-creates-Post, Post-has_tag-Tag* refers to *structure*; *Tag.Tagname* refers to *dimension*, *AVG(User.Age)* refers to *measure*; and *Post.Year=2017* refers to *slicing condition*.

We define some notations adopted in this thesis. Given an OLAP query q , we use “ $q.properties$ ” to refer to the set of **all properties** that appear in *Dimension*, *Measure*, and *Slicing Condition* of q . We use “ $q.structure$ ” to refer to the structure of q . For example, $Q3.properties = \{ Tag.Tagname, User.Age, Post.Year \}$, and $Q3.structure = \{ User-creates-Post, Post-has_tag-Tag \}$.

3.1.3 Materialization: Cuboid & Substructure

As previously discussed, a key issue of our work is to find the most useful common substructures and properties from the previous queries, which are assumed to appear in the

workload frequently. For fast access to results of these useful queries, we store their results (preferably in main memory), and this is known as the *materialization* of the query.

As defined above, in a property graph, each node or edge has a unique ID, which can be treated as a special property. Whether a materialization keeps the unique ID is an important issue. This is because keeping unique ID often increases the space cost. We consider two types of materializations, namely the *cuboid* and the *substructure*, based on whether or not unique IDs of nodes (or edges) are kept. To better elaborate the differences between cuboid and substructure, we consider the following example. Suppose we have the following query workload containing two history queries (P1 and P2) and two incoming queries (Q1 and Q2):

P1 User-creates-Post: User.Age
P2 User-creates-Post: User.Age, (AVG)Post.Score
Q1 User-creates-Post: Post.Score, (AVG)User.Age
Q2 User-creates-Post, Post-has_tag-Tag: User.Age, Tag.TagName

We can tell that users are most interested in the *User-creates-Post* structure. $\{User.Age, Post.Score\}$ is the set of properties being involved in queries over *User-creates-Post*. Thus, we can build a cuboid lattice of all combinations of $\{User.Age, Post.Score\}$. Materialization of the base cuboid of the lattice is

$M_1: \$User-creates-Post: User.Age, Post.Score, COUNT(*)$

Note that for the rest of this thesis, we use the \$ symbol followed by structures and dimensions to denote a materialization, represented by M . The above materialized view is useful for Q1 since we can process Q1 by aggregation (with dimension of Post.Score and measure of (AVG)User.Age) over M_1 . We call such a materialization without ID(s) in *dimension* a **cuboid**.

However, M_1 is not useful for Q2 since they have different *structures*. On the contrary, if we add $ID(Post)$ to *dimension* and materialize

$M_2: \$User-creates-Post: User.Age, Post.Score, ID(Post), COUNT(*)$

so that *Post* is “activated” to be able to join with other materializations containing *Post* and produce results for OLAP over more complicated *structures*. For example, Q2 can be processed through the following steps:

Step 1: joining $\$User-creates-Post: User.Age, Post.Score, ID(Post) COUNT(*)$ and $\$Post-has-tag-Tag: ID(Post), Tag.TagName, COUNT(*)$ on $ID(Post)$;

Step 2: perform aggregation on $\{User.Age, Tag.TagName\}$.

	Cuboid	Substructure
Dimension	Only properties	Properties and ID(s)
Space Cost	“Low”	“High”
Potential benefit	Aggregation	Aggregation & Joining

Table 3.1: Comparisons between Cuboid and Substructure.

In this case, we only need to fetch $\$Post-has.tag-Tag: ID(Post), Tag.TagName, COUNT(*)$ from database to produce result for Q2. We call such materialization with ID(s) in *dimension* as **substructure**.

Table 3.1 shows a comparison between cuboid and substructure. Note that cuboids can only be used in queries with exactly the same structure. They are good for roll-up and slicing operations but not useful for queries with different structures. Substructures can be used to join with other materializations to help with new queries of various types of structures. The drawback is that substructures are generally more space-costly than cuboids, as IDs are unique keys. The trade-off between cuboids and substructures is the trade-off between space cost and the potential saving of join processing.

3.2 Problem Definition

Intuitively, our goal is to answer OLAP queries efficiently by taking advantage of materialized views, which are constructed based on the knowledge of previous queries. Therefore, our solution needs to follow two steps:

- Materialization step: materialized view selection.
- Query processing step: answer new queries as quickly as possible (using materializations).

The materialization step is in fact a “Materialization Selection” (MS for short) problem, as using materialization is good for query efficiency, but comes with a storage cost. So we want to study the problem of how to best utilize materialization within a space budget σ . The query processing step involves “Execution Planning” (EP for short). We formally define these two problems as follows.

Materialization Selection Problem: Given a property graph G , a set of previous queries P on G , space limit σ , find cuboids C and substructures S , such that:

$$\sum_{c_i \in C} c_i.space + \sum_{s_i \in S} s_i.space \leq \sigma; \text{ and}$$

$$\sum_{p_i \in P} T(G, p_i, C, S) \text{ is minimized.}$$

Here $T(G, p_i, C, S)$ is a function to estimate the query processing time of p_i on G using C and S . “*space*” refers to the estimated space cost of a cuboid or substructure. Note that the real run time of a particular query is hard to estimate. Therefore, we use $T(G, p_i, C, S)$ as a cost function to measure the time cost of query processing.

Execution Planning: Given a property graph dataset G , a new query q , materialized cuboids C and substructures S , find a processing plan $process(G, q, C, S)$, such that the processing time of q , denoted by $process(G, q, C, S).time$, is minimized.

As a matter of fact, execution planning can be further divided into two sub-problems. First, which materialized views in C and S should be used to answer q ? Second, how to answer q as fast as possible using the view selected from the first question. Thus, we define the first question as the *Decomposition Problem*, which decomposes q into two parts, one part covered by the views from C and S , while the other part not covered, which is named as the *remaining views*. Note that the “remaining views” refers to the data that needs to be fetched from the database server on the fly. We define the second question as the *Composition Problem*, which performs basic relational operations such as join, projection and selection over views in order to get the result of q .

Composition Problem: Given a property graph G , a new query q , materialized cuboids C' and substructures S' , and remaining views R , find a composition plan $compose(G, q, C', S', R)$, so that estimated composition time $compose(G, q, C', S', R).time$ is minimized. Here $compose(G, q, C', S', R)$ returns result of query q by performing operations (join, selection, projection etc.) over C', S', R .

Decomposition Problem: Given a property graph G , a new query q , materialized cuboids C and substructures S , a composition plan $compose(G, q, C, S, R)$, find $C' \subseteq C, S' \subseteq S$, and remaining views R , so that $compose(G, q, C', S', R).time$ is minimized. Here $compose(G, q, C, S, R)$ is the solution of Composition Problem.

Note that we define Composition Problem before Decomposition Problem. The reason is that we need to consider a composition plan $compose(G, q, C', S', R)$ when making our selection policy of C', S' and R . In other words, these two problems are closely related.

Chapter 4

Solution

In this chapter, we present our complete solution towards efficient OLAP query processing over property graphs. For comprehensive presentation, we first illustrate the overall solution framework in Section 4.1. Then we present our strategy for materialized view selection, as well as the execution planning for query processing in Section 4.2 and 4.3, respectively.

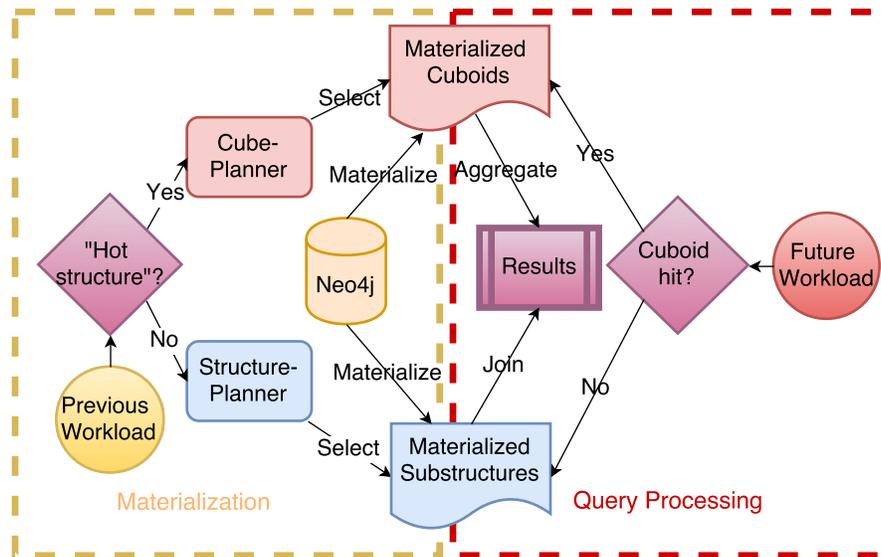


Figure 4.1: Solution framework.

4.1 Solution Framework Overview

Figure 4.1 describes the overall solution framework. Two dashed line rectangles represents the major components of our solution: materialization and query processing. Materialization takes previous workload as input and materializes cuboids and substructures which are beneficial for new workload processing. We adopt a straightforward best effort approach for materialization. Intuitively, we first partition previous queries into “hot” queries and “less hot” queries based on the frequency count of their structures. Modules CubePlanner and StructurePlanner take “hot” queries and “less hot” queries respectively, as input and produce cuboids and substructures (in the form of tables) for materialization. More details are left to Section 4.2.1, where we explain the intuition of categorization of “hot” and “less hot” queries, as well as the reason for passing them to different planners. Query processing component takes incoming queries as input and returns results. Briefly, the workflow of query processing is the following. If a new query happens to contain a “hot” structure, we consult cuboid materializations to see if it can be directly answered by aggregation over a cuboid materialization. In this case, cuboid materialization will be used. Otherwise, if the query cannot be directly answered by any materialized cuboid, we consider available materialized substructures by decomposing the query into substructures for join. Note that if required substructure is not materialized, on the fly data fetching from the graph database server is mandatory.

4.2 Materialized View Selection

Materialized view selection is a research problem that has attracted significant attention in traditional RDBMS community. However, RDBMSs materialize query results, while in this work we wish to take into account the structure of the graph queries and materialize parts of that structure rather than the entire query results. This is mainly because *structure* plays an important role in a query over a property graph. One materialized view may affect the benefits of other ones (e.g., when their *structures* overlap). Therefore combinational benefit for a set of materialized views is not a simple sum of benefits for each one of them. As briefly illustrated before, we consider two types of materializations for efficient OLAP query processing: cuboids and substructures. In this section, we first elaborate the essential heuristic of selecting cuboids and substructures. Then we detail the approaches taken for different types of materializations.

4.2.1 Overview of Materialized View Selection

In Section 3.1.3, we discussed the trade-off between cuboids and substructures. As noted, utilization of a cuboid materialization requires new queries to have exactly the same structure as the materialized cuboid. Therefore, it is only reasonable to materialize a cuboid when we are confident that the same structure is likely to be “hit” by new queries. Otherwise, it is simply a waste of space to materialize cuboids that would be rarely “hit”. On the contrary, substructures do not have such strict structural match requirement. A substructure can be used as long as it appears in a new query.

Considering the different features of cuboids and substructures, we follow the following strategy for materialized view selection. We first perform a frequency count of previous queries. If more than ω queries shares the same structure, where ω is a predefined frequency threshold, this structure is considered to be a *hot structure* and would be passed on to the *CubePlanner* module for cuboid selection. Queries that do not have *hot structures* are passed to *StructurePlanner* for the substructure selection. Algorithm 1 describes the overall framework of the materialized view selection process.

Algorithm 1: Materialization Overview

System setting: ω : frequency threshold for hot structures
Input: Q : a set of previous queries
Output: C : a set of materialized cuboids
 S : a set of materialized substructures

```

1  $CInput \leftarrow \emptyset$ ;
2  $SInput \leftarrow \emptyset$ ;
3 foreach  $p \in P$  do
4   | if  $structureFreq(P, p) > \omega$  then
5   |   |  $CInput \leftarrow CInput \cup \{p\}$ ;
6   |   end
7   |  $SInput \leftarrow SInput \cup \{p\}$ ;
8 end
9  $C \leftarrow CubePlanner(CInput)$ ;
10  $S \leftarrow StructurePlanner(SInput)$ ;

```

As shown in Algorithm 1, function $structureFreq(P, p)$ returns the frequency count of all query structures in P . After *hot structures* are selected, two functions *CubePlanner* and *StructurePlanner* are called to select cuboids and substructures for materialization. Consider the following example. Assume we are aware of the six previous queries as shown below. We can group queries by structure and count the structure frequency.

Previous Workload:

P1 Badge-User, User-Post:Badge.Name,Post.Score,Post.PostTypeId=2
P2 User-Comment, Comment-Post: User.UpVotes, Comment.Score, AVG(Post.Score), Post.PostTypeId=1
P3 User-Post, Post-Vote: User.UpVotes, Vote.VoteTypeId
P4 User-Post, Post-Tag: AVG(User.CreationDate_Year), Tag.TagName
P5 User-Comment, Comment-Post: User.ActiveMonth, Post.CreationDate_Year=2016
P6 User-Comment, Comment-Post: User.Age, AVG(Comment.Score), Post.PostTypeId=2

new workload:

Q1 User-Comment, Comment-Post: User.UpVotes, AVG(Post.Score), Post.PostTypeId
Q2 User-Comment, Comment-Post: User.Age, Post.PostTypeId
Q3 User-Post, Post-PostHistory: User.UpVotes, PostHistory.PostHistoryTypeId
Q4 Badge-User, User-Post: AVG(Post.Score),Post.PostTypeId=2

Structure	Frequency
User-Comment, Comment-Post	3
User-Post, Post-Tag	1
User-Post, Post-Vote	1

Obviously, *User-Comment, Comment-Post* is a *hot structure*. We materialize cuboids over structure *User-Comment, Comment-Post* by passing previous query P2, P5 and P6 to *CubePlanner*. *CubePlanner* will materialize cuboids that benefit processing of new queries Q1 and Q2 (which have *User-Comment, Comment-Post* structure). Then, we pass the three remaining queries of less hot structures, P1, P3, and P4 to *StructurePlanner*, which will discover and materialize most useful substructures. In this case *StructurePlanner* is likely to find *User-Post* as a useful substructure it can be used in joining the result of new queries Q3 and Q4.

4.2.2 Greedy Selection Framework

Before diving into the details of the *CubePlanner* and *StructurePlanner* modules, we first illustrate the essential greedy heuristic employed for view selection.

In our solution framework, *CubePlanner* and *StructurePlanner* are responsible for materialized view selection (over cuboids and substructures, respectively). They both adopt the same greedy selection framework. In Section 3.2, we introduced the “Materialization Selection” problem, which aims at finding best materializations under a space limit σ . Materialization selection is a known NP-complete problem [16]. The difficulty lies in that

the overall benefit of materialized views is not a simple sum of the individual benefits of each view. A materialized view’s marginal benefit may be deducted when another view is selected. For example, suppose we have previous queries over “*User-Post, Post-Tag*”, then the marginal benefit of a substructure over “*User-Post*” will be affected by selecting substructures over “*User-Post, Post-Tag*”. This is because we do not need substructure “*User-Post*” in joining the results of previous queries over “*User-Post, Post-Tag*” when we already have a materialized substructure over “*User-Post, Post-Tag*”. A straightforward approach to solve the materialization selection problem is to enumerate over all possible combinations of cuboids C and substructures S within the space limit σ and find the best combination. But such a brute-force solution is infeasible in practice. In addition, assume that we obtain the optimal C' and S' in some way, it is not guaranteed that the actual total space cost of C' and S' is strictly lower than σ as we only made estimations in our calculation. Therefore, we turn to a greedy algorithm which is better than straightforward approach in terms of efficiency. Besides, it allows materializations to be done one by one while the space limit is better respected.

We discuss this greedy selection framework first to give a high-level idea of our selection policy. We use a greedy algorithm for both cuboid and substructure selection, as shown in Algorithm 2. The idea is to always pick the next candidate with the highest ratio of marginal benefit against the space cost. After a candidate is picked, we re-evaluate the benefit of remaining candidates. Re-evaluation is mandatory as the marginal benefit of a candidate may be reduced as a result of materialization of a selected candidate.

Algorithm 2: Greedy Selection

System setting: σ : space limit

Input: C : a set of candidates of cuboids or substructures in lattice structure

P : A set of previous queries

Output: R : a list of materialized candidates

```

1 foreach  $c \in C$  do
2    $c.space \leftarrow space(c)$ ;
3    $c.benefit \leftarrow estimateMarginBenefit(c, P, R)$ ;
4    $c.score \leftarrow c.benefit/c.space$ ;
5 end
6 while  $R.totalsize < \sigma$  do
7    $selected \leftarrow$  candidate in  $C$  with the highest  $score$ ;
8    $R.add(materialize(selected))$ ;
9   repeat Lines 1-5;
10 end

```

Lines 1-5 estimate the space cost, the marginal benefit for new workload, as well as the score for each candidate. Functions $space(c)$ and $estimateMarginBenefit(c, P, R)$ estimate the space cost and the marginal benefit of materializing candidate c , respectively. We call this phase **score calculation**. Lines 6-10 keeps picking up candidates with highest score one by one until space limit is hit. Notice that each time a candidate is selected, Line 9 refreshes scores for all candidates by repeating 1-5. We call this phase **pick-and-update**. Note that we use $materialize()$ as a function to denote the materialization of selected cuboids and substructures.

CubePlanner and *StructurePlanner* apply this greedy selection framework with different implementations of score calculation and pick-and-update. Users can adjust the behavior of *CubePlanner* and *StructurePlanner* by plugging in their own implementation of the score calculation function that may consider different database features. Furthermore in Algorithm 2 we describe a scenario where space is the limit (Line 6). Line 6 can be modified to deal with situations of other limits. For example, in our implementation of *CubePlanner*, we select at most n cuboids, where a maximum number of materialized candidates is the limit (which will be discussed in 4.2.3). For the rest of this chapter, we will focus on our implementation of *CubePlanner* and *StructurePlanner* in Neo4j.

4.2.3 CubePlanner

CubePlanner takes previous queries with *hot structures* as input and returns the selected cuboids for materialization. As mentioned in Section 3.1.3, a cuboid is only useful for queries sharing the exact identical structure. To put it another way, cuboids of different structures do not affect each other at all in terms of benefits for new queries. Since the input queries for *CubePlanner* may have different structures, we can group queries by structure and treat each group separately. For each group of input queries, we propose an algorithm named *SingleCubePlanner* to select top- k cuboids. After all groups are finished, we materialize the final top- n cuboids by searching across top- k cuboids from all groups. A good analogy for such a process is to first hold regional competitions and then select national winners from regional winners. Next we will explain *CubePlanner* and *SingleCubePlanner* in detail.

CubePlanner

As we mentioned above, *CubePlanner* performs cuboid selection in a holistic manner by selecting cuboids one-by-one from results of *SingleCubePlanner*. We first explain the workflow of *CubePlanner*, as shown in Algorithm 3. Intuitively, *CubePlanner* first groups P

by structure using the function $group(P)$. Lines 2-4 perform cuboid selection in each partition using *SingleCubePlanner*. A queue of ordered candidates is generated within each group of queries. Lines 5-8 repeatedly check the current top candidate of each partition to select the best candidate among them. n is a user defined parameter, denoting the most number of cuboids for materialization. Note that users may choose other ways, such as a space limit, as the bound for cuboid materialization. In our solution, we use the most number of cuboids n in CubePlanner, and we use a space limit in StructurePlanner. It is worth pointing out that a unified marginal benefit evaluation scheme for cuboids and substructures as a whole needs to be developed to divide memory budget between these two materialized views, which is considered as a future work.

Algorithm 3: CubePlanner

System setting: n : maximum number of cuboids to be precomputed

Input: P : a set of previous queries not necessarily with the same structure

Output: C : a list of materialized cuboids

```

1  $G \leftarrow group(P)$ ;
2 foreach  $group \in G$  do
3   |  $group.results \leftarrow SingleCubePlanner(group)$ ;
4 end
5 for  $i=1$  to  $n$  do
6   |  $g \leftarrow$  the group in  $G$  with the highest  $results.top().score$ ;
7   |  $C.add(materialize(g.results.Dequeue()))$ ;
8 end

```

SingleCubePlanner

Now we elaborate the *SingleCubePlanner* function. As shown in Algorithm 4, *SingleCubePlanner* follows a greedy selection strategy to generate the top- k cuboids.

The algorithm starts with building a lattice over all combinations of dimensions of all attributes that appeared in previous query set P , using a classic lattice construction algorithm described in [19]. Lines 2-4 initialize the best-so-far processing time for each previous query with its estimated naive database processing time. Lines 5-12 perform score calculation following the greedy selection framework presented in Algorithm 1. For each cuboid, we estimate its space (line 6). Lines 8-10 calculate the marginal benefit by iterating over previous queries that can be answered by scanning current cuboid. If the estimated scanning time is less than a previous query’s current best-so-far processing time,

Algorithm 4: SingleCubePlanner

System setting: k : as in “top- k ”

Input: P : a set of previous queries with a same structure

Output: C : a queue of selected cuboids to precompute

```
1  $Lattice \leftarrow buildLattice(P)$ ;  
2 foreach query  $p \in P$  do  
3   |  $p.time \leftarrow time(p)$ ;  
4 end  
5 foreach cuboid  $\in Lattice$  do  
6   |  $cuboid.space \leftarrow space(cuboid)$ ;  
7   |  $cuboid.benefit \leftarrow 0$ ;  
8   | foreach query  $p \in P$  and  $p.properties \subseteq cuboid.properties$  do  
9     |  $cuboid.benefit \leftarrow cuboid.benefit + max(0, p.time - aggreTime(cuboid))$ ;  
10  | end  
11  |  $cuboid.score \leftarrow cuboid.benefit / cuboid.space$ ;  
12 end  
13 for  $i=1$  to  $n$  do  
14   |  $nextBestCube \leftarrow$  the cuboid in  $Lattice$  with the highest score;  
15   | if  $nextBestCube.score < 0$  then  
16     | break;  
17   | end  
18   |  $C.Enqueue(nextBestCube)$ ;  
19   | foreach query  $p \in P$  and  $p.dimension \subseteq nextBestCube.dimension$  do  
20     |  $p.time \leftarrow min(p.time, aggreTime(nextBestCube))$ ;  
21   | end  
22   | Repeat Lines 5-12;  
23 end
```

we add the difference of two times to the cuboid’s total marginal benefit (Line 9). Lines 13-23 perform the pick-and-update, where lines 15-17 terminate the selection process when there is no more extra marginal benefit, and lines 19-22 update the best-so-far processing time for previous queries as a result of the current round of selection.

Now we explain the implementation details of the time estimation function employed in Algorithm 4. Function $time(query)$ estimates the time of naïve processing of a query in a graph database. Implementation of $time(query)$ is database specific as physical storage and execution plans vary among different databases (i.e., not using materialized views). Since Neo4j provides APIs to show the execution plan as well as the estimated intermediate result size, we directly use the total size of intermediate results as an estimation of the time cost. For example, Figure 4.2 is an execution plan provided by Neo4j for query *User-Badge, User-Post, Post-Tag: Tag.TagName*. We can see that the number of estimated rows of intermediate results are provided (e.g., the number of estimated rows after “NodeHashJoin” is 1,721,324). We use the sum of estimated rows of all the intermediate results in the execution plan to estimate the total processing time cost.

For graph databases that do not provide an API to access execution plans and estimated intermediate result sizes, users can construct different estimation functions following the same intuition, which usually depends on specific database implementations. There are many studies on cost estimation for database operations (e.g., join operation). Users may consider joining (expanding) order [3] and estimation of intermediate result sizes [22] as two important factors.

Function $aggreTime(cuboid)$ estimates the time cost for scanning and aggregation over a materialized cuboid. As the actual processing time is difficult to estimate, we alternatively use the space cost of cuboids as the estimation metric. Given a cuboid c , we estimate the space cost of c as follows:

$$spacePerRow = \sum_{p \in c.properties} typeSize(p)$$

Thus,

$$SpaceCost(c) = spacePerRow \times numberOfRows(c)$$

Note that $typeSize(property\ type)$ refers to the standard size of data types. For example, the integer type in “C++” is 2 byte. $numberOfRows(c)$ refers to the number of rows in c . A rough estimation is the size of the Cartesian product of all queried properties:

$$numberOfRows(c) = \prod_{p \in c.properties} |p|$$

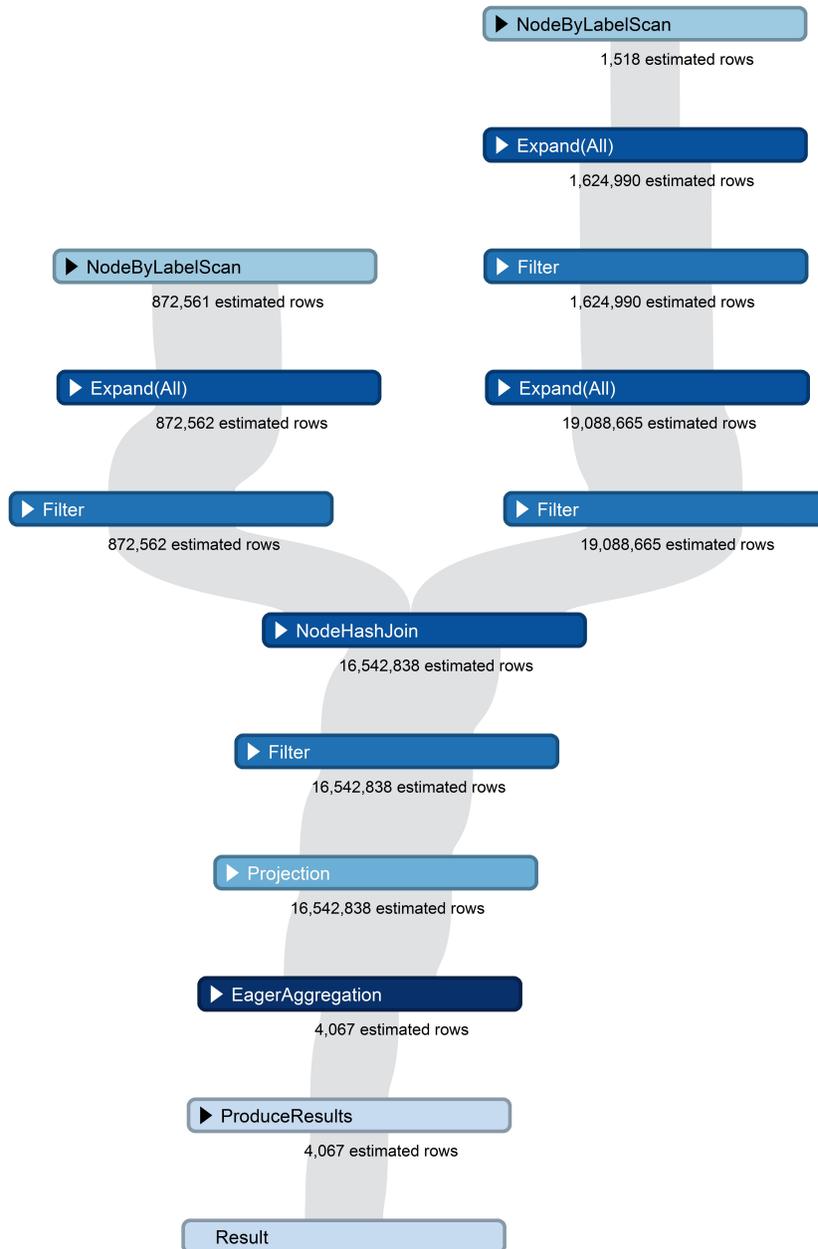


Figure 4.2: Neo4j's execution plan for query *User-Badge, User-Post, Post-Tag: Tag.TagName*.

4.2.4 Structure Planner

As mentioned above, *StructurePlanner* also adopts the same greedy selection strategy described in Algorithm 1. We detail the process of *StructurePlanner* in Algorithm 5. First, we build a lattice over all substructures contained in previous queries P , using the classic lattice construction algorithm (similar to the lattice construction algorithms adopted in *CubePlanner*). Figure 4.3 shows a substructure lattice originating from the root node *Badge-User, User-Post, Post-Tag*. Starting from a union of structures of previous queries as the root node, a lattice can be constructed recursively by populating descendants from parent nodes through edge removals.

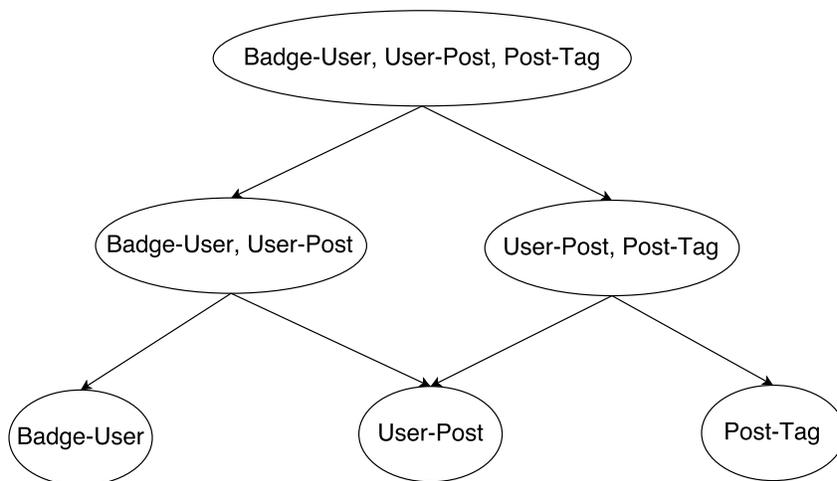


Figure 4.3: A substructure lattice with *Badge-User, User-Post, Post-Tag* as its root node.

Then, we initialize the covered substructures for each previous query as an empty set (lines 2-4). For a previous query, *coveredSubstructure* keeps the information on what substructures have been selected so far. It is updated every time a new substructure is selected. Lines 5-12 perform score calculation. For each substructure, Line 6 estimates its space. Lines 8-10 iterate over all “favored” previous queries (favored by current substructure) and add on the marginal benefit (if any). Here marginal benefit refers to the time saved after adding current substructure to selected substructures (Line 9). Lines 13-23 perform the pick-and-update, and lines 19-22 update the covered substructures for previous queries as a result of current round of selection. Such iteration will be terminated when space limit is exceeded (line 13) or when there is no more marginal benefit (line 15).

We now explain the detailed implementation of the estimation functions employed in Algorithm 5. Note that these are specific to our implementation in Neo4j and may differ in

Algorithm 5: StructurePlanner

System setting: σ : space limit for materialized views

Input: P : a set of previous queries

Output: S : an list of materialized substructures

```
1  $Lattice \leftarrow buildSubstructureLattice(P)$ ;  
2 foreach  $p \in P$  do  
3   |  $p.coveredSubstructures \leftarrow \emptyset$ ;  
4 end  
5 foreach  $substructure \in Lattice$  do  
6   |  $substructure.space \leftarrow space(substructure)$ ;  
7   |  $substructure.benefit \leftarrow 0$ ;  
8   | foreach  $p \in P$  and  $p.structure \subseteq substructure.structure$  do  
9     |  $cuboid.benefit \leftarrow$   
     |    $cuboid.benefit + max(0, benefit(p, substructure, p.coveredSubstructures))$ ;  
10  | end  
11  |  $substructure.score \leftarrow substructure.benefit/substructure.space$ ;  
12 end  
13 while  $System.memoryUsage < \sigma$  do  
14   |  $nextBestSubstructure \leftarrow$  the substructure in  $Lattice$  with the highest  $score$ ;  
15   | if  $nextBestSubstructure.score < 0$  then  
16     | break;  
17   | end  
18   |  $S.add(materialize(nextBestSubstructure))$ ;  
19   | foreach  $p \in P$  and  $p.structure \subseteq nextBestSubstructure.structure$  do  
20     |  $p.coveredSubstructures \leftarrow p.coveredSubstructures \cup \{nextBestSubstructure\}$ ;  
21   | end  
22   | repeat Lines 5-12;  
23 end
```

the case of other systems. Function $space(substructure)$ returns the estimated space cost of a substructure materialization. We use Neo4j’s execution plan API to get an estimated result size of a substructure. Function $benefit(p, substructure, p.coveredSubstructures)$ evaluates the marginal benefit of materializing a substructure for P . We know that execution plan and estimated intermediate result size are provided by Neo4j’s API. But when substructure materialization is used, the actual execution plan (intermediate result) may be different than the naïve processing plan. As a result, estimating the marginal benefit of a substructure is tricky. We estimate the marginal benefit of a substructure using

$$time(p.coveredSubstructures \cup substructure) - time(p.coveredSubstructures),$$

which captures the overall improvement of adding $substructure$ to $coveredSubstructures$ as materializations.

4.2.5 ID and Property Selection

Given a substructure picked by the *StructurePlanner*, we need to decide which IDs and properties need to be stored. Keeping all IDs and attributes makes a substructure materialization more informative but increases the space cost. Thus, the selection of IDs and properties is an important issue. We use substructure *User-Post*, *Post-Tag* as an example and discuss different ID and property selection policies.

For IDs, we consider the following two policies.

- “*All nodes preservation*” keeps IDs of all nodes and edges. For *User-Post*, *Post-Tag*, if we keep IDs of all nodes and edges, then we can perform join operation with *Badge-User*, *User-Post*. We call such join an “overlap join” as the two substructures have an overlap part which is *User-Post*. Note that we can join the two substructures only when IDs of nodes (User and Post), and edge (edge between User and Post) are stored in both substructures.
- “*Border nodes preservation*” only keeps IDs of “border nodes” which are on the border of the substructure’s *structure*. Figure 4.4 highlights “border nodes” of structure *User-Post*, *Post-Tag*. In this example we only save IDs of User and Tag. We do not keep IDs of Post as node Post is not located on the border of the *structure*. This saves space compared with “all nodes preservation”, however “border nodes preservation” only enables joins on border nodes. For example we may join *User-Post*, *Post-Tag* with *User-Badge* on their common border node User. But “overlap join” with other

substructures is not enabled because IDs of “inner nodes” are not stored. We cannot join *User-Post*, *Post-Tag* with *Badge-User*, *User-Post* as Post is an “inner node” and IDs of Post is not stored.

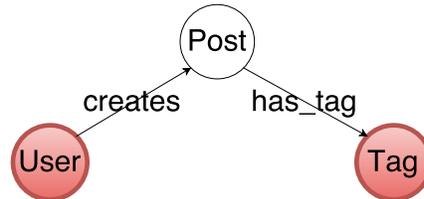


Figure 4.4: “Border nodes” of structure *User-Post*, *Post-Tag*.

In our experiments we choose “all nodes preservation”. But in some applications, the space cost of “all nodes preservation” can be overwhelming, so we can choose “border nodes preservation”.

For properties, we can choose to keep all properties, or only keep properties that were queried in previous workloads. Our suggestion is to consider the proportion of properties that were queried in previous workload over all properties in the data schema. For example, in our experiments only a small proportion of properties were queried. We choose the latter option as it is a waste of space to keep all properties. However, this decision problem can be studied further.

4.2.6 Update on Materialized Views

The solution we have discussed so far is applicable for static scenarios (with fixed previous workload). We now expand our solution to dynamic scenarios: as queries are executed, there could be change of “hot structures” and “properties of interest”. Thus updates on materialized views are necessary. It is worth pointing out that our current solution considers only static data graphs. View maintenance over dynamic data graph is left as a future work. Remember that we have a memory budget, therefore in some cases obsolete materialized views need to be swapped out for new ones. To achieve this, we maintain a sliding window over previous queries. After executing a certain number of queries, we perform materialized view selection over only “recent queries” which are in the sliding window. Old views that need to be swapped out are eliminated by simply releasing their memory. For newly selected views we materialize them. For old views that do not need to be swapped out, we do nothing as they are already materialized in memory. In this way periodical update on materialized views can be realized.

4.3 Query Processing

Query processing aims at processing incoming queries efficiently using materialized substructures and cuboids. When a query q arrives, we first consult materialized cuboids. If q can be answered with an aggregation over any materialized cuboid, we select the cuboid with the minimum space and directly scan over it to produce result of q . If q cannot be answered by any cuboid, we decompose q and use substructures as much as possible to compute the result of q .

Algorithm 6: Query Processing

System: C : a set of materialized cuboids

S : a set of materialized substructures

Input: q : a query

Output: r : result of q

```
1  $minspace \leftarrow \infty$ ;  
2  $mincuboid \leftarrow NULL$ ;  
3 foreach  $cuboid \in C$  do  
4   if  $cuboid.structure = q.structure$  and  $q.dimension \subseteq cuboid.dimension$  then  
5     if  $cuboid.space < minspace$  then  
6        $minspace \leftarrow cuboid.space$ ;  
7        $mincuboid \leftarrow cuboid$ ;  
8     end  
9   end  
10  if  $mincuboid \neq NULL$  then  
11     $r \leftarrow aggregate(mincuboid, q)$ ;  
12  else  
13     $r \leftarrow Decompose\_Join(q, S)$ ;  
14  end  
15 end
```

Algorithm 6 describes query processing. Given an incoming query q , we first look up materialized cuboids and find if any cuboid can be used to answer q (lines 4-9). Note that $cuboid.space$ in line 5 was computed in line 9 in Algorithm 4. Then, we check if q can be answered by cuboid materialization. If so, we perform aggregation operation over the cuboid (line 11). Otherwise, we need to decompose q into substructures and compose the result (line 13). Function $aggregate(mincuboid, q)$ is the classic aggregation operation. We will discuss how function $Decompose_Join(q, S)$ is implemented at the end of this section.

4.3.1 Substructure Selection

Before discussion on $Decompose_Join(q, S)$, we need to first solve a “Substructure Selection” problem. In order to decompose a query q , we need to consider which materialized substructures we need to use. We need to make decision when candidate substructures in S overlap. For example suppose q has structure $Badge-User$, $User-Post$, $Post-Tag$, and S consists of substructures:

- (1) $Badge-User$
- (2) $Badge-User, User-Post$
- (3) $User-Post, Post-Tag$
- (4) $Post-Tag$
- (5) $User-Post$.

We can get structure of q by joining structures (1) and (3). Thus (1) and (3) seem to be a possible combination for substructure selection in this case. Actually we may have at least three possible substructure selections: (1) and (3); (2) and (4); (1), (4) and (5). The key question is which selection will result in fastest processing time on q ? Here are some intuitions to solve this question. First, when we select substructures one by one, we do not select a substructure when it is covered by selected substructures. For example we will not consider (1) if (2) has been selected as (1) is covered by (2). Second, we prefer to minimize total size of selected substructures as we need to at least access each selected view once. Third, we prefer smaller number of selected substructures as intuitively this causes less times of joins.

We propose a greedy algorithm for substructure selection based on user defined heuristics. Users may define heuristic functions based on other intuitions (like the three intuitions mentioned above). The idea of the greedy algorithm is to always pick up next substructure with highest score of user defined heuristic function $h(s)$, which returns heuristic score for a substructure s . Some example heuristics are #edges of substructure, score calculated in StructurePlanner (Line 11 in Algorithm 5), table size, many others.

Lines 1-2 initialize $uncoveredStruc$ and $uncoveredProp$, which keeps track of structures and properties which have not been covered by selected substructures. Such uncovered structures and properties will need to be fetched from the database. Line 3 initializes $coveredStruc$, which keeps union of selected substructures. Line 5 starts iteration over substructures ordered by user-defined heuristics $h(s)$. Line 6 assures that a candidate substructure that is totally covered by selected substructures will be disqualified. In the

above example, suppose we have already selected (2), there is no need to select (1) since (1) is totally covered by (2).

Algorithm 7: SelectSubstructure

System: S : a collection of materialized substructures
 $h(s)$: user defined function. It returns the heuristic score of a substructure s .
Input: q : a new query
Output: V : selected views for future joining
 $uncoveredStruc$: structure not covered by selected views
 $uncoveredProp$: properties not covered by selected views

```

1  $uncoveredStruc \leftarrow q.structure$ ;
2  $uncoveredProp \leftarrow q.properties$ ;
3  $coveredStruc \leftarrow \emptyset$ ;
4  $V \leftarrow \emptyset$ ;
5 foreach  $s \in S$  ordered by  $h(s)$  do
6   | if  $s \subseteq uncoveredStruc$  and  $s \not\subseteq coveredStruc$  then
7   |   |  $V \leftarrow V \cup \{s\}$ ;
8   |   |  $coveredStruc \leftarrow coveredStruc \cup s.structure$ ;
9   |   |  $uncoveredStruc \leftarrow uncoveredStruc - s.structure$ ;
10  |   |  $uncoveredProp \leftarrow uncoveredProp - s.properties$ ;
11  |   end
12 end

```

4.3.2 Decomposition and Join

In this section, we discuss how to implement function $Decompose_Join(q, S)$ (as in Algorithm “Query Processing” in subsection 4.3). Besides $Decompose_Join(q, S)$, we shall discuss two other variations of implementation: $Decompose_Join^*$ and $Decompose_Join^+$.

#1 $Decompose_Join$

Given a query q , we use the previously discussed algorithm “SelectSubstructure” to select a set of substructure materializations V . However, substructures in V may not completely cover the structure of V . If there is any remaining structure ($uncoveredStruc$) and properties ($uncoveredProp$) that V does not cover, we need to retrieve them from the database. We call such remaining structures and properties fetched from the database “complemen-

tary components”. After all these components (both materializations and “complementary components”) are finally ready, we join and aggregate them to produce final results.

Algorithm 8: Decompose_Join

System: S : a collection of materialized substructures

Input: q : a new query

Output: r : result of q

```

1  $\Sigma \leftarrow \emptyset$ ;
2  $V, uncoveredStruc, uncoveredProp \leftarrow SelectSubstructure(q)$ ;
3  $\Sigma \leftarrow \Sigma \cup V$ ;
4  $Splits \leftarrow split(uncoveredStruc, uncoveredProp)$ ;
5 foreach  $s: Splits$  do
6   |  $\Sigma \leftarrow \Sigma \cup \{retrieve(s)\}$ ;
7 end
8  $r \leftarrow join\_aggregate(\Sigma, q)$ ;
```

Line 1 initializes Σ , which maintains a set of all components (materializations and “complementary components”) that are needed. Line 2 selects substructures using `SelectSubstructure` algorithm. `uncoveredStruc` and `uncoveredProp` refer to structures and properties which are not covered by selected substructures. They are “complementary components” and will be retrieved from database servers. Line 4 splits `uncoveredStruc` and `uncoveredProp` into connected components. We retrieve each connected component from the database server. Note that splitting is necessary since `uncoveredStruc` may not be exactly one connected component. Line 8 joins and aggregates all materials together to produce results.

Function `split(uncoveredStruc, uncoveredProp)` is implemented by a classic connected components detection algorithm. It splits `uncoveredStruc` and `uncoveredProp` into connected components (structures). We want to retrieve each connected structure separately from the database because otherwise it may result in unnecessarily large results of Cartesian products of several disconnected structures. Function `materialize(s)` retrieve “complementary components” s from the database server. Function `join(Σ, q)` join tables of Σ together and aggregate over properties based on q . Joins over multiple tables is a well-studied topic. Joining order and join technique (hash join etc.) are two important aspects of this topic. In our implementation we use hash join and our joining order policy is to keep joining two tables that have minimum sum of table sizes and have common column(s). That is, we tend to select two smaller tables to join.

#2 *Decompose_Join**

Decompose_Join retrieve “complementary components” from the database in a naive manner. We adopt the idea of Semi-Join [20] and propose another way of implementation: *Decompose_Join**. Semi-join takes advantage of “selection effect” of natural joins. In *Decompose_Join**, we first perform joins over substructures of V even before fetching “complementary components” from the database. Line 3 performs $join(V)$ before *retrieve* in Line 7. We call this phase “first round of joins”. Note that substructures in V may reside in multiple connected components. Thus V^* may consist of multiple intermediate tables. The purpose for first round of joins is that it provides “candidate” node and edge IDs for future joins (thanks to ‘selection effect’ of natural joins). When fetching “complementary components” from the database server, we inform database server such candidate node and edge IDs so that search space for “complementary components” is narrowed down ($retrieve^*(s, V^*)$ in Line 7). We name this approach *Decompose_Join**.

Algorithm 9: *Decompose_Join**

System: S : a collection of materialized substructures

Input: q : a new query

Output: r : result of q

```
1  $\Sigma \leftarrow \emptyset$ ;  
2  $V, uncoveredStruc, uncoveredProp \leftarrow SelectSubstructure(q)$ ;  
3  $V^* \leftarrow join(V)$ ;  
4  $\Sigma \leftarrow \Sigma \cup V$ ;  
5  $Splits \leftarrow split(uncoveredStruc, uncoveredProp)$ ;  
6 foreach  $s$ :  $Splits$  do  
7 |  $\Sigma \leftarrow \Sigma \cup \{retrieve^*(s, V^*)\}$ ;  
8 end  
9  $r \leftarrow join\_aggregate(\Sigma, q)$ ;
```

We explain implementation of function $retrieve^*(s, V^*)$. It fetches results from the database by passing candidate IDs information (from join result V^*). Syntax to achieve this varies by database. In SQL and Cypher we may pass candidate IDs using a “WHERE” statement.

We compare *Decompose_Join** vs. *Decompose_Join* and summarize following pros and cons of *Decompose_Join**. *Decompose_Join** helps accelerate retrieval process from backend databases in two aspects. First, since screened out candidate IDs are provided, database back end only needs to iterate through a portion of nodes and edges. This saves database processing time. Second, candidate IDs have a “selection” effect thus

size of retrieval results is deducted. Therefore result transmission times will be reduced. The disadvantages are twofold. First, *Decompose_Join** has to transmit IDs. Second, *Decompose_Join** performs two rounds of join; first on V before “complementary components” are ready, followed by second round of joins. In terms of optimization on join orders, *Decompose_Join* is better because the join order can be derived based on full knowledge of all the tables, whereas in *Decompose_Join**, tables are considered separately in two rounds of joins, which narrows down the scope of all possible join orders.

#3 *Decompose_Join*⁺

Algorithm 10: *Decompose_Join*⁺

System: S : a collection of materialized substructures
Input: q : a new query
Output: r : result of q

```

1  $\Sigma \leftarrow \emptyset$ ;
2  $V, uncoveredStruc, uncoveredProp \leftarrow SelectSubstructure(q)$ ;
3  $V^* \leftarrow join(V)$ ;
4  $\Sigma \leftarrow \Sigma \cup V$ ;
5  $Splits \leftarrow split(uncoveredStruc, uncoveredProp)$ ;
6 foreach  $s: Splits$  do
7   | if  $decide(s, V^*)$  then
8   |   |  $\Sigma \leftarrow \Sigma \cup \{retrieve^*(s, V^*)\}$ ;
9   | else
10  |   |  $\Sigma \leftarrow \Sigma \cup \{retrieve(s)\}$ ;
11  |   | end
12 end
13  $r \leftarrow join\_aggregate(\Sigma, q)$ ;

```

We mentioned two advantages of $retrieve^*(s, V^*)$. However a disadvantage of $retrieve^*(s, V^*)$ is an overhead of transmitting candidate IDs to Neo4j. We propose a decisive way to evaluate the trade-off between overhead and benefits of $retrieve^*(s, V^*)$ and choose between $retrieve^*(s, V^*)$ and $retrieve(s)$. *Decompose_Join*⁺ is derived from *Decompose_Join**. The trick is to revise Line 7 by using $decide(s, V^*)$ to choose the better way between $retrieve^*(s, V^*)$ and $retrieve(s)$.

Implementation of function $decide(s, V^*)$ is as follows. We first estimate result sizes of the two retrieval methods: $retrieve^*(s, V^*).estSize$ and $retrieve(s).estSize$. $retrieve(s).estSize$

can be returned by $space(substructure)$ in Algorithm 5. The tricky one is $retrieve^*(s, V^*).estSize$, which we estimate as follows:

1. Randomly sample a small number of candidate IDs.
2. Execute $retrieve^*$ but pass only sampled candidate IDs in (1). We call this a “trial query”. We want to use the “trial query” to estimate result length of actual $retrieve^*(s, V^*)$. Since we only pass a small number of IDs, time cost of “trial query” is acceptably small.
3. Calculate $retrieve^*(s, V^*).estSize$ proportionally using the result of “trial query”.

We compare $Decompose_Join^+$ vs $Decompose_Join^*$: We see that $Decompose_Join^+$ performs two rounds of joins like $Decompose_Join^*$. The major difference is that $Decompose_Join^+$ plays “trial query”. The principle behind “trial query” is to pay an acceptable price to make a wise decision on “complementary components” retrieval. Our experiments show that a good decision making on “complementary components” retrieval saves much more time than the time cost of “trial queries”, results and discussion of which will be presented in next chapter.

Chapter 5

Evaluation

In this chapter, we validate our proposed solution with a set of meaningful queries over a real world dataset. To demonstrate the time efficiency of our approach, we select Neo4j community version 3.1.3 as the baseline. In the experiments, we test query processing time and space cost by running the queries using both our approach and native Neo4j implementation. The results show that our approach is about 20 times faster than the Neo4j implementation on average under the default settings (to be covered in Section 5.2). Moreover, we assess and explain how each aspect of our system affects processing efficiency. Finally, we discuss our reflections on the Neo4j system.

5.1 Experiment Setup

There are two purposes of our experiments. First is to demonstrate the query processing efficiency of our proposed solution by comparing against the native Neo4j system. Second, we would like to test how different settings (of parameters and choices of strategies) in our solution affect query processing efficiency.

5.1.1 Datasets

We use the StackOverFlow dataset in our experiments that is generated by using raw information from <https://archive.org/details/stackexchange>. The dataset contains user-contributed content (such as user information, posts, etc.) on www.stackoverflow.com. The dataset contains 10 different node labels and 12 different edge labels. Figure 5.1 shows

the meta graph of 8 node labels and 8 edge labels that are involved in our experiments. The data graph contains over 300 million vertices and more than 400 million edges and takes 44.5GB of storage.

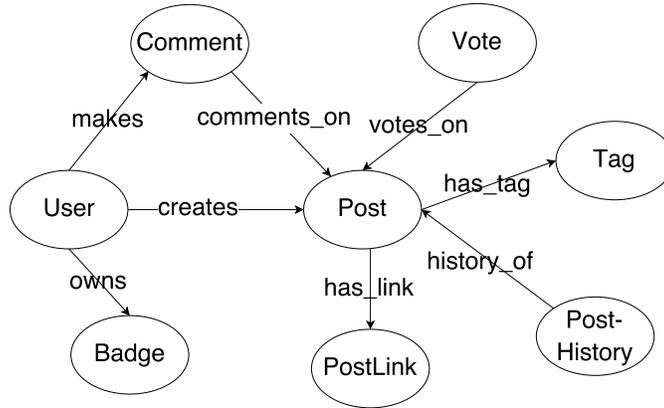


Figure 5.1: The meta graph of StackOverFlow used in experiments.

5.1.2 Query Workloads

We design 24 queries against the StackOverFlow dataset given below. We randomly select 12 queries as the previous workload, leaving the rest 12 ones as the new workload. With these previous and new workloads, three senerios of different cuboid hit and structure hit (to be discussed in Section 5.3.1) happen when processing new queires.

Previous Workload:

- P1 User-Post: User.UpVotes, Post.Score=10
- P2 User-Post: User.UpVotes, AVG(Post.Score)
- P3 User-Post: User.Age, SUM(Post.ActiveMonth)
- P4 User-Post: User.CreationDate_Year, Post.PostTypeId=1
- P5 Badge-User, User-Post, Post-Tag: Tag.TagName, User.CreationDate_Year=2017
- P6 Badge-User, User-Post, Post-Tag: Tag.TagName, Badge.Name
- P7 Badge-User, User-Post:Badge.Date_Year, AVG(Post.Score)
- P8 Badge-User, User-Post:Badge.Class, AVG(Post.ActiveMonth)
- P9 User-Post, Post-Tag: AVG(User.Age), Tag.TagName

- P10 User-Post, Post-Tag: AVG(User.UpVotes), Tag.TagName=Java
- P11 User-Post, Post-Vote: AVG(User.UpVotes), Vote.VoteTypeId
- P12 Post-Comment, Post-PostLink: PostLink-LinkTypeId, AVG(Comment-Score)

New Workload:

- Q1 User-Post: User.CreationDate_Year=2017, Post.PostTypeId
- Q2 User-Post: AVG(User.UpVotes),Post.Score
- Q3 User-Post: Post.ActiveMonth, AVG(User.Age)
- Q4 User-Post: User.CreationDate_Year
- Q5 Badge-User, User-Post, Post-Tag: Tag.TagName, Badge.Class
- Q6 Badge-User, User-Post, Post-Tag: Tag.TagName, Badge.Date_Year
- Q7 Badge-User, User-Post:Badge.Name, Post.PostTypeId
- Q8 User-Post, Post-Tag: User.UpVotes, Tag.TagName, Post.PostTypeId=2
- Q9 User-Post, Post-Tag:User.CreationDate_Year, Tag.TagName
- Q10 Badge-User, User-Comment: Badge-Class, AVG(Comment-Score)
- Q11 Badge-User, User-Comment: Badge-Name, AVG(Comment-Score)
- Q12 Post-PostHistory, Post-Tag: Tag-TagName, PostHistory-PostHistoryTypeId

5.1.3 System Setting

We run the experiments on a Linux server with 256GB main memory. Our system is implemented in Java. We set the initial JVM memory to 100GB, and the maximum JVM memory to 200GB.

Our solution is developed on top of Neo4j Community v3.1.3. In the experiments, we set Neo4j’s initial memory to 60GB and 200GB for the maximum usage. We use Neo4j’s official Java driver (<https://neo4j.com/developer/java/#neo4j-java-driver>) to interact with Neo4j server.

5.2 Aspects of Interest

As we mentioned, there are two purposes of our experiments. In addition to the efficiency test, we would like to study how different settings in our system could affect query processing efficiency. We list the following aspects of interest to be tested. Default setting for each is also presented.

Materialization

- Space cost limit (σ in Algorithm 5 in Section 4.2.4) is set to 6GB by default. Results and discussion are presented in Section 5.3.3.
- Algorithms in materialized view selection described in Section 4.2. For cuboid selection, we compare CubePlanner (in Section 4.2.3) with the “Partial Materialization” algorithm (PMA) proposed in Graph Cube [29]. We materialize the best 3 cuboids in both approaches (n in Algorithm 3 is set to 3). For substructure selection we compare StructurePlanner (in Section 4.2.4) with the well-known maximal frequent pattern mining algorithm (FPM) [8]. In other tests, CubePlanner and StructurePlanner are used by default. Results are discussed in Sections 5.3.5 and 5.3.6, respectively.
- Frequency threshold for identification of hot structures (ω in Algorithm 1). We set the default value to be 4; the intuition of default value, together with results and discussion are elaborated in Section 5.3.2.
- Storage level for materialized views. We compare main memory storage vs hard disk storage. Note that memory based materialization is set as the default in all other tests. Detailed discussion is provided in Section 5.3.4.

Processing of New Queries

- Choice of score function in ranking substructures during the “Substructure Selection” ($h(s)$ in Algorithm 7 discussed in Section 4.3.1). By default, $h(s)$ returns the score of s calculated in StructurePlanner (as result of line 11 in Algorithm 5). Results and explanations can be found in Section 5.3.7.
- Choice among using *Decompose_Join*, *Decompose_Join** and *Decompose_Join+* in “Decomposition and Join” in Section 4.3.2. We use *Decompose_Join* as the default method. Detailed discussion is in Section 5.3.8.

5.3 Results and Discussion

5.3.1 Our System vs. Neo4j

We first show the time efficiency of our solution running new queries in default settings against the native Neo4j implementation (Figure 5.2).

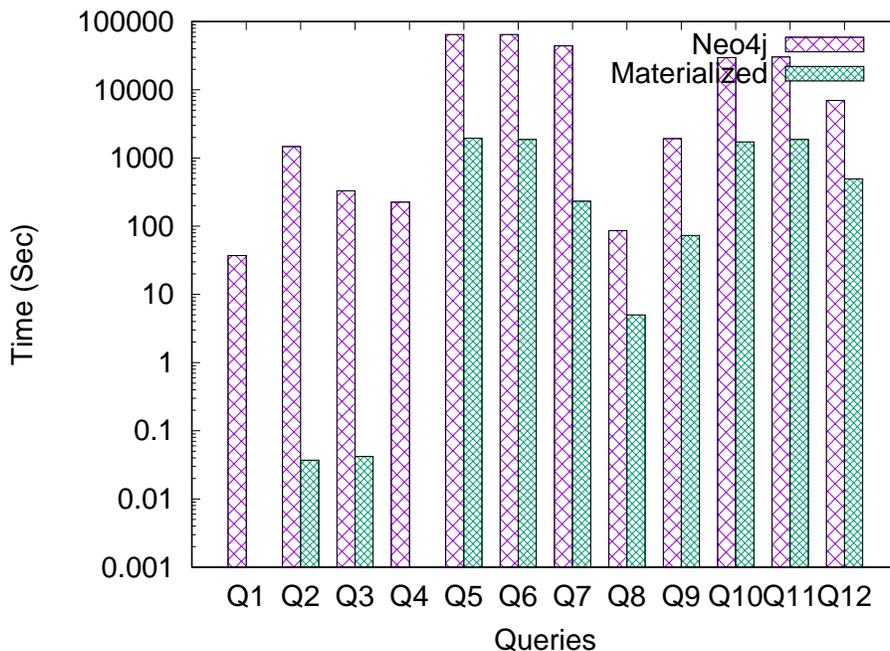


Figure 5.2: Time efficiency on the new workload: our solution vs Neo4j.

Note that the execution time for Q1 and Q4 in our system is 0.001s, which is negligible to show in the figure. Neo4j does not finish Q5 and Q6 in 18 hours so we terminate execution and record the processing time as 18 hours. It is worth pointing out that with an extra 4,045s time cost and 5.7GB space cost for materialized views, our solution achieves remarkable efficiency improvement (29.8 times speedup and 236,514s of total time saving). Given the size of our dataset, this extra space is acceptable.

We now detail how our system works over the 12 queries. First, *User-Post* is identified as a hot structure as its frequency count is larger than or equal to the frequency count threshold. As a result, P1 - P4 are passed to the CubePlanner and P5 - P12 are passed to the StructurePlanner. Cuboids selected by CubePlanner are {User.Age, Post.ActiveMonth},

$\{\text{User.CreationDate_Year}, \text{Post.PostTypeId}\}$, and $\{\text{User.UpVotes}, \text{Post.Score}\}$. Figure 5.3 highlights substructures S that StructurePlanner selects: *User-Post* (s_1), *Post-Tag* (s_2) and *Badge-User* (s_3). From the observation of the previous workload, we can tell that the StructurePlanner makes a good decision as these three substructures are able to cover most of previous queries.

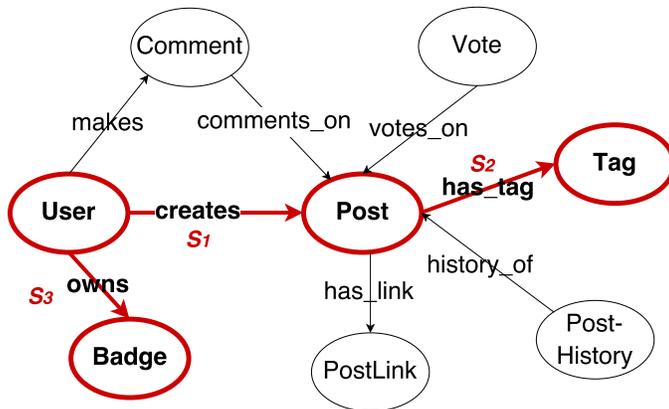


Figure 5.3: Substructure selected by StructurePlanner.

We further explain the performance variance of the new workload. Overall improvement rate for Q1 - Q4 is more than 20000. This is because Q1 - Q4 are answered with a “cuboid hit”. As a result, the time savings for these 4 queries are much greater than for Q5 - Q12 (which do not get a “cuboid hit”). It is worth pointing out that time complexities of cuboid aggregation and substructure joins are much different. The time complexity of a cuboid aggregation is bounded by the size of the Cartesian product of its dimensions. When it comes to substructure joins, however, the time complexity is related to the actual data size. Note that Q5 - Q9 are totally covered by S . In these cases, no complementary components are fetched from the Neo4j database. Q10 - Q12 are partially covered by S . Therefore, complementary components are fetched from the database, where a series of I/O operations increases the total time. Overall, the overall improvement rate for Q5 - Q9 is around 40 times, while improvement rate for Q10 - Q12 ranges around 15 times. Clearly, our system could greatly improve query processing efficiency with an acceptable incremental space cost. While the improvement ratio varies in different scenarios of cuboid and structure “hits”.

5.3.2 Frequency Threshold

Frequency threshold ω can significantly affect the query processing efficiency. A change of ω may result in different hot structures, followed by varied inputs for the CubePlanner and the StructurePlanner. It would further lead to different materialized view selection and overall query processing performance. As indicated in Section 4.2.1, ω serves as a minimum threshold for building a cuboid over a hot structure (“cuboid hit” assured). We consider 4 an appropriate choice for ω in our test case considering the frequency count of structures in previous workload (as shown in the table below).

Structure	Frequency
User-Post	4
Badge-User, User-Post, Post-Tag	2
Badge-User, User-Post	2
User-Post, Post-Tag	2
User-Post, Post-Vote	1
Post-Comment, Post-PostLink	1

Suppose we lower ω to 2. *Badge-User, User-Post, Post-Tag* will be hot structure. As a result, cuboid selection over *Badge-User, User-Post, Post-Tag* will be considered based on merely two queries (P5 and P6). Notice that Q5 and Q6, which have structure *Badge-User, User-Post, Post-Tag*, will never get cuboid hit. This is because properties *Badge-Class* in Q5 and *Badge-Date_Year* in Q6 do not even appear in P5 and P6. That is to say, in this case any cuboid selected over *Badge-User, User-Post, Post-Tag* will be useless. In addition, when ω is set to 2, input for StructurePlanner will be only two queries (Q11 and Q12). Note that structure frequency counts of these two queries are both 1. In other words, these are the most “random” queries. Note that the idea of StructurePlanner is to discover of useful substructures based on a sufficient number of “less hot” queries. In this case, two random queries are not the ideal input for the StructurePlanner.

Suppose we increase ω to 5; then there will be no cuboid that is materialized in our test case. As a result Q1 - Q4 will be processed using substructure materialization (s_1). Although it is still faster than the native Neo4j system, the outstanding improvement ratio of cuboid hit cannot be achieved.

Figure 5.4 presents the total processing time under different settings of ω . Note that no structure has frequency count of 3, therefore setting ω as 3 and 4 would categorize the same set of hot structures and thus yield the same materialized views. We see that ω does

have a significant effect over the system performance. In order to determine the optimal value of ω (which leads to best processing efficiency), we need to evaluate the tradeoff between inputting a previous query to CubePlanner and StructurePlanner. We leave this as future work.

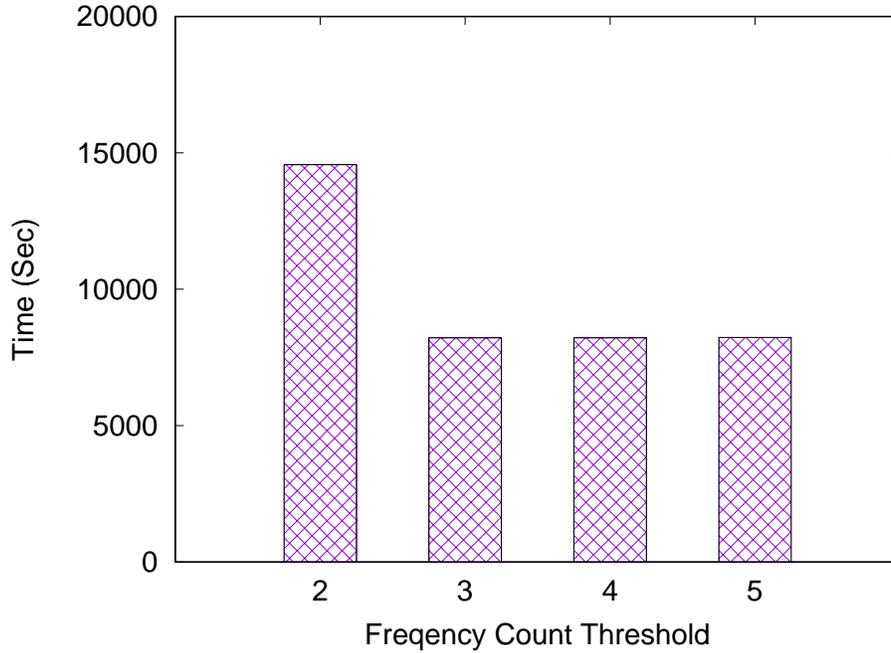


Figure 5.4: Total processing time under different settings of ω .

5.3.3 Space Cost Limit

As pointed out in Section 5.3.1, the space cost of our materialized views is 5.7GB, while the default space cost, σ , is set to 6GB. In this section, we study the effect of σ on the query processing efficiency. Figure 5.5 shows how the total processing time varies with different space costs (which were caused by setting σ to 6GB, 4GB, and 2.5GB, respectively). As more views being materialized, the total processing time decreases.

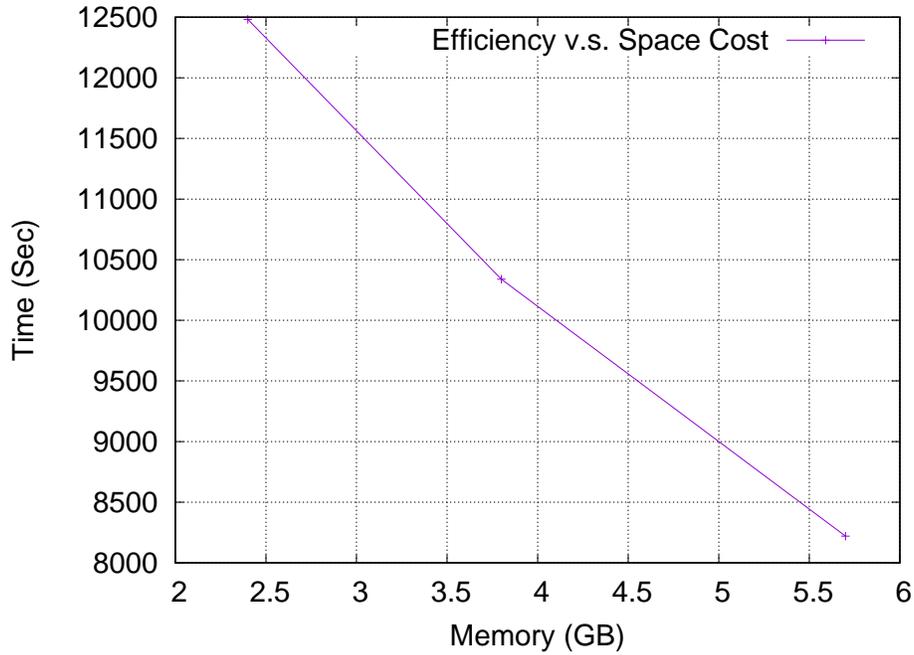


Figure 5.5: Total execution time under different space limits.

5.3.4 Storage Level for Materialized Views

As discussed in Section 5.2, by default, materialized views are stored as objects in main memory. This guarantees fast data access on the cost of extra memory consumption. Alternatively, materialized views can be serialized and stored as files on hard disks. Figure 5.3.4 shows a comparison of the total query processing time using memory-based views and disk-based views. As expected, hard disk storage does not perform as fast as main memory storage. However, it is worth noting that the drop in efficiency is acceptable.

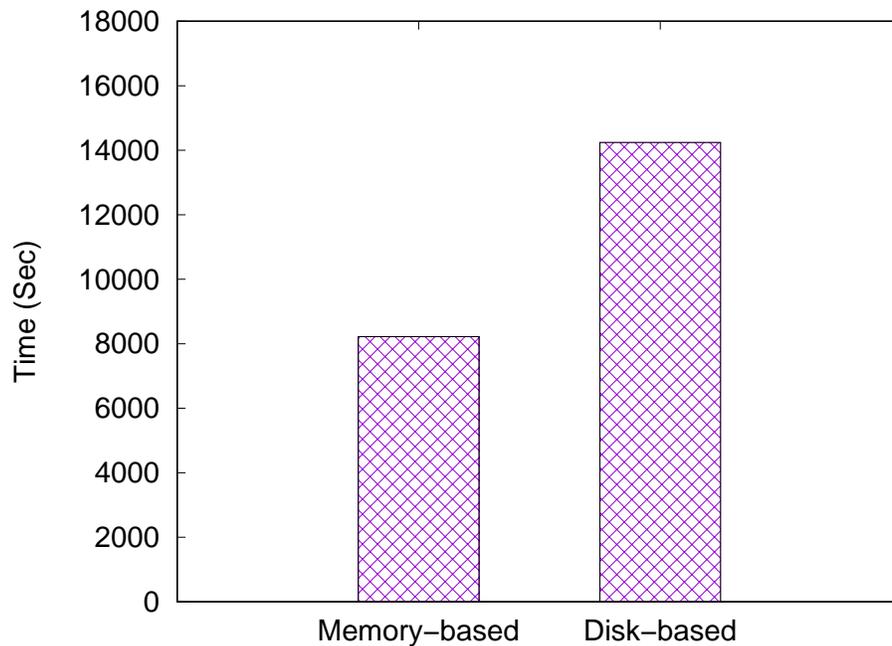


Figure 5.6: Main memory storage vs hard disk storage.

5.3.5 CubePlanner vs PMA

We compare CubePlanner in our solution (Section 4.2.3) with PMA in Graph Cube [29]. Figures 5.7 and 5.8 show that CubePlanner outperforms PMA in both query processing efficiency and space cost. Cuboids selected by CubePlanner are {User.Age, Post.ActiveMonth}, {User.CreationDate_Year, Post.PostTypeId}, and {User.UpVotes, Post.Score}. While PMA selected {User.Age, User.UpVotes, User.CreationDate_Year, Post.PostTypeId, Post.Score, Post.ActiveMonth}, {User.Age, User.UpVotes, User.CreationDate_Year, Post.PostTypeId, Post.ActiveMonth}, {User.Age, User.UpVotes, User.CreationDate_Year, Post.PostTypeId, Post.Score}.

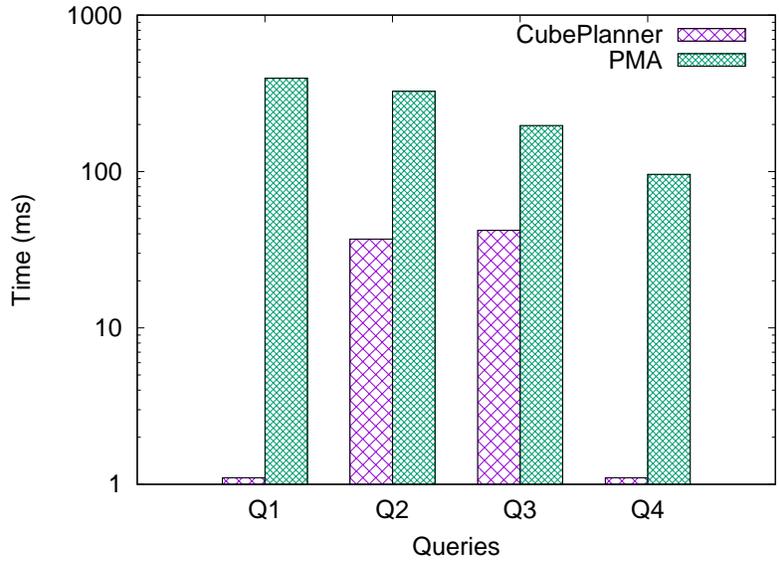


Figure 5.7: Time: CubePlanner vs PMA.

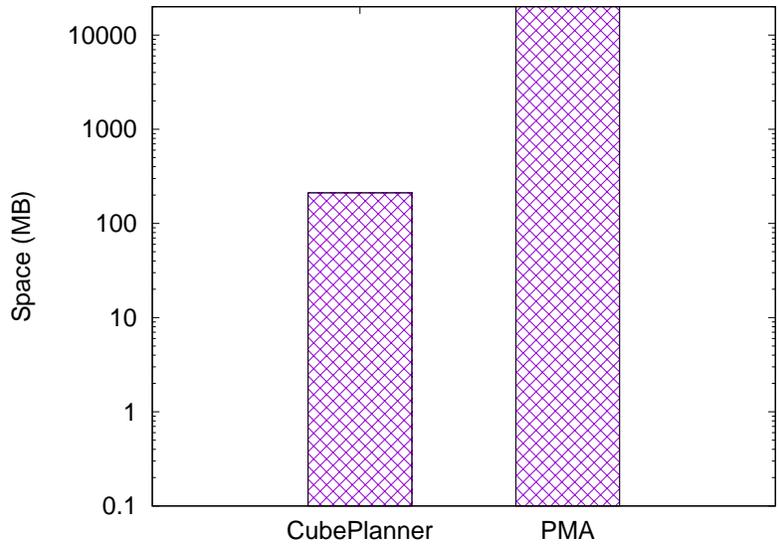


Figure 5.8: Total cuboid space cost: CubePlanner vs PMA.

Both approaches are considered as implementations of the Greedy Selection Framework discussed in Section 4.2.2. The differences between the two implementations are the

following. CubePlanner uses the ratio of marginal benefit over space cost as a score for candidate ranking (line 11 in Algorithm 4), while PMA only considers marginal benefit. That is to say, space cost is not taken into account in PMA. Moreover, PMA treats each combination of properties with an equal weight, regardless of how many times a combination has appeared in previous queries. For example, in our test case, the combination of User.UpVotes, Post.Score appears twice in previous workload. But PMA would treat User.UpVotes, Post.Score with the same weight as those combinations that are not even queried in previous workload ($\{User.UpVotes, Post.PostTypeId\}$ etc). As a result, CubePlanner uses more information from previous workload and thus makes a better selection.

5.3.6 StructurePlanner vs FPM

We compare our Algorithm 5 in StructurePlanner with FPM. In FPM we set the minimum support to 2 considering the frequency count as listed in the table in Section 5.3.2. These two ways provide different substructure selections which lead to different processing efficiency. Figures 5.9 and 5.10 show that our StructurePlanner outperforms FPM in both efficiency and space cost.

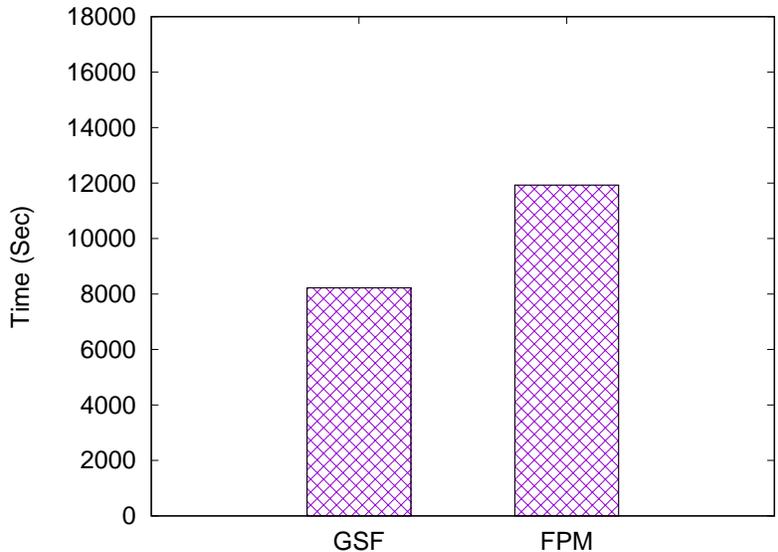


Figure 5.9: Total processing time for new workload: StructurePlanner vs FPM.

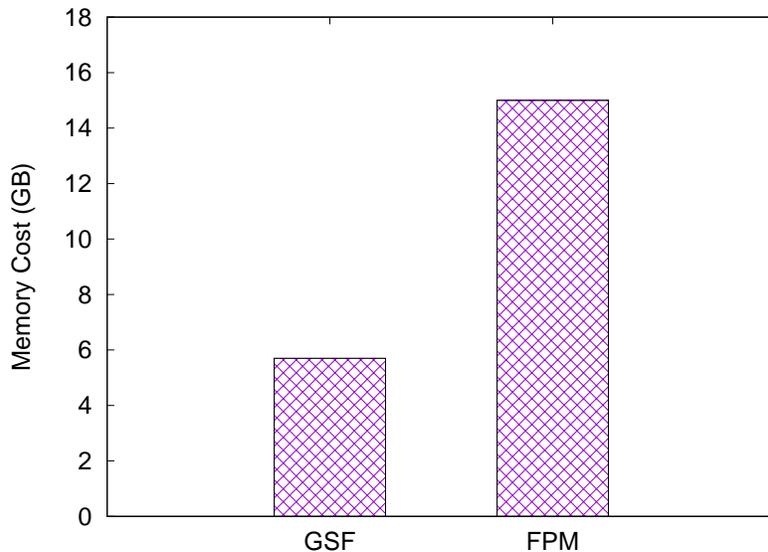


Figure 5.10: Space cost: StructurePlanner vs FPM.

Figure 5.3 highlights three selected substructures by StructurePlanner. As mentioned, it is a good selection as these three substructures are able to cover most of the previous queries. However, FPM selects *Badge-User*, *User-Post*, *Post-Tag*, which is a bad selection because it is useful only for Q5 and Q6. In addition, materialization of *Badge-User*, *User-Post*, *Post-Tag* results in even more space cost than materialization of the three edges separately (as selected by StructurePlanner). Figure 5.11 details processing time for each query. FPM only outperforms StructurePlanner on Q5 and Q6. This is because Q5 and Q6 would be able to perform aggregation over the materialization of *Badge-User*, *User-Post*, *Post-Tag* when FPM is applied. While table joins of s_1 , s_2 and s_3 are required if StructurePlanner is applied, which clearly is more time consuming. However for Q7 - Q12, StructurePlanner is the winner as it gets at least a partial “substructure cover”, while the FPM-selected structure is not helpful at all.

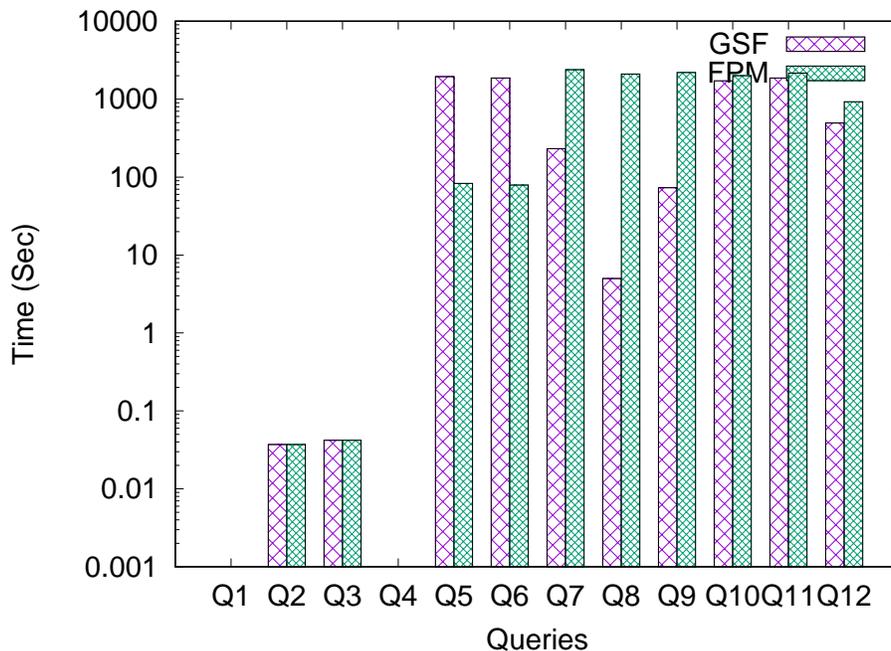


Figure 5.11: Processing time for each query: StructurePlanner vs FPM.

5.3.7 Substructure Selection

In our experiment, $h(s)$ in Algorithm 7 does not make any difference during “Substructure Selection”. This is because the three selected substructures do not share any edges. Note that scenarios in Section 4.3.1, where multiple valid combinations of materialized substructures exist, only happen when materialized substructures have overlaps.

5.3.8 Decompose_Join

We now present the experiments comparing *Decompose_Join*, *Decompose_Join** and *Decompose_Join+* in “Decomposition and Join” (presented in Section 4.3.2). Three different implementations of “Decomposition and Join” would lead to different processing performance for Q10 - Q12, as they are partially covered by S and fetching complementary components from Neo4j is necessary. Figure 5.12 provides the processing time for Q10 - Q12 using the three approaches. *Decompose_Join** performs better than *Decompose_Join* in Q10. This is because *Decompose_Join** passes to Neo4j candidate IDs of users that have

badges, which provides a considerable filtering effect when fetching *User-Comment*. As a result, *Decompose_Join**'s processing time for fetching *User-Comment* is reduced. Besides, its time for joining *Badge-User* and *User-Comment* also decreases because the table size of *User-Comment* is smaller than that of *Decompose_Join*, thanks to the filtering effect. This is reflected in Figure 5.13, where the time for join is saved in *Decompose_Join**. However *Decompose_Join** performs badly for Q12. This is because the filtering effect of *Post-Tag* in fetching *Post-PostHistory* is small as most posts have tags. In addition, *Decompose_Join** has an overhead of scanning the *Post-Tag* table in order to get the set of candidate IDs. It explains why *Decompose_Join** takes longer in processing Q12. Figure 5.14 gives the total processing time for Q10 - Q12 using the three approaches. We see that *Decompose_Join+* has best overall performance. To explain, *Decompose_Join+* is able to choose the faster approach when fetching complementary components from Neo4j in scenarios like Q10 and Q12 with the cost of a cheap trial query (as discussed in Section 4.3.2). Note that the time cost for a trial query is bounded by a constant sample size, and is not proportional to actual data size in the dataset. A too small sample size results in bias in estimation, whereas an unnecessarily large sample size causes too much time. We tried different sample sizes and find 100 to be an appropriate value, which serves estimation properly and saves time. Figure 5.15 shows that the time cost of the trial query is negligible compared to the overall processing time.

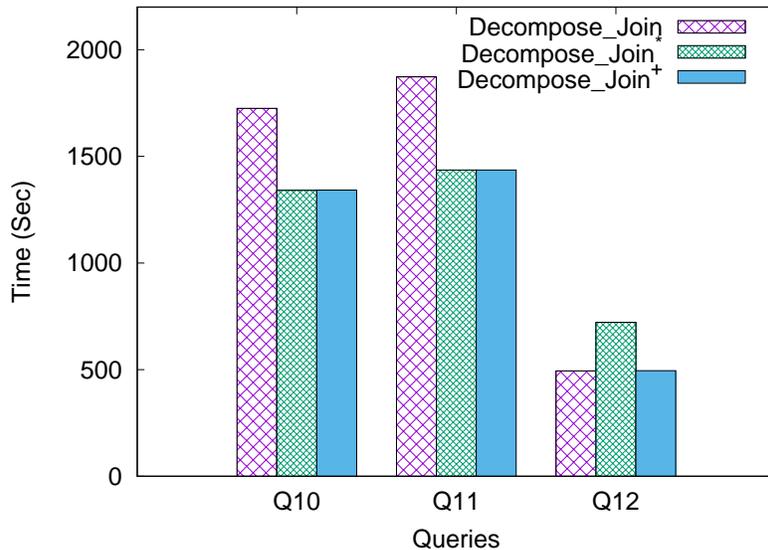


Figure 5.12: Processing time for Q10 - Q12 by three approaches.

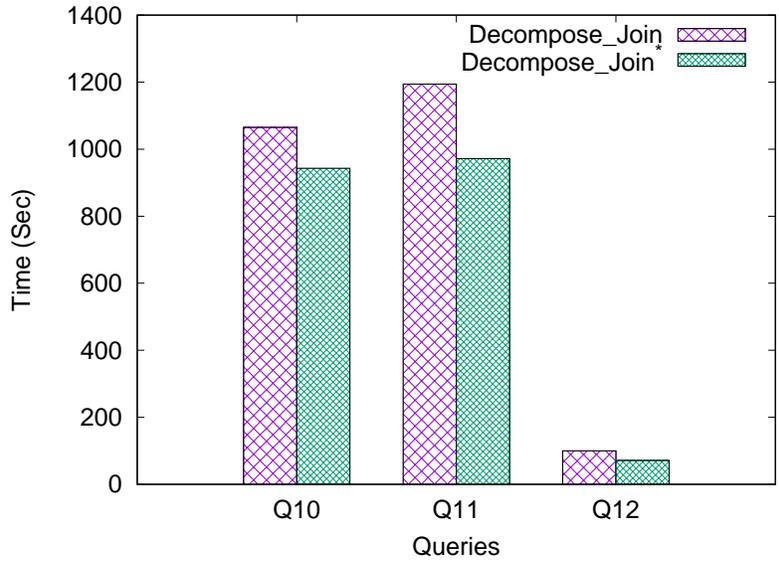


Figure 5.13: Joining time in processing Q10 - Q12 by three approaches.

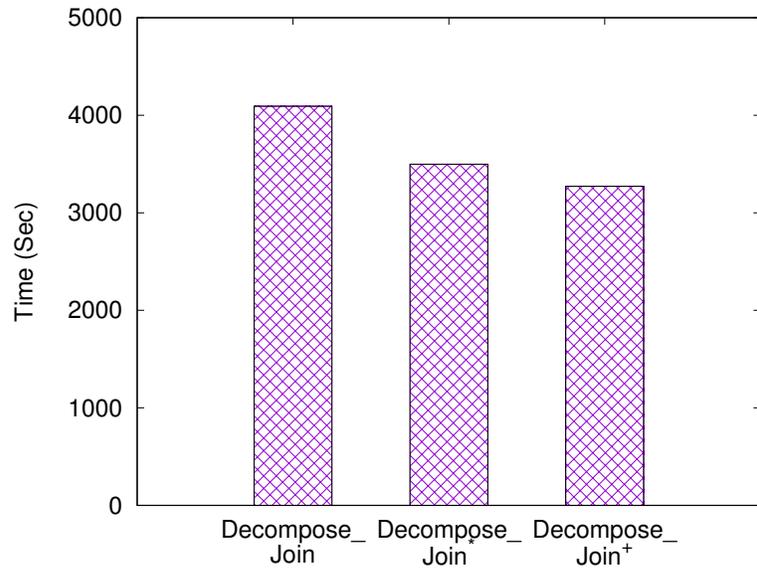


Figure 5.14: Total processing time for Q10 - Q12 by three approaches.

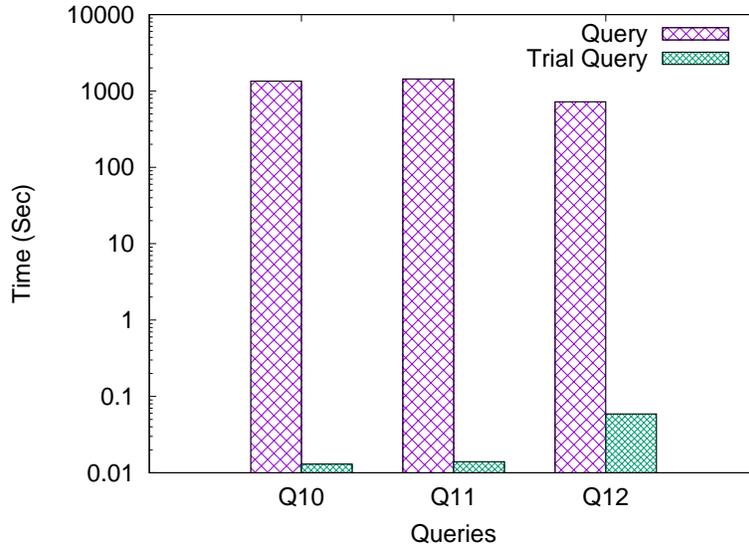


Figure 5.15: Total processing time vs trial query processing time.

To conclude, filtering effect is an important factor in the performance of these three different implementations of “Decomposition and Join”. In general, *Decompose_Join⁺* is the recommended approach as it is able to select the better solution with a small cost of trial query.

5.3.9 Reflections on Neo4j

During the experiments, we found that Neo4j uses a naïve approach to result size estimation for aggregation queries. It simply takes the square root of table size before aggregation as the estimated size for the aggregated result, regardless of which properties are being aggregated.

For example, Figure 5.16 presents Neo4j’s execution plans for queries *User-Post: User.Age* and *User-Post: ID(User), ID(Post)*, respectively. Obviously the latter query should have much larger estimated result size than the former one. However Neo4j returns exactly the same estimated result size by simply taking the square root of table length (6,660,515 estimated rows in “Projection” step) in the previous step. Actual result sizes for the two queries are 85 and 6,660,515. From this example we can tell that Neo4j’s naïve estimation may lead to a huge bias.

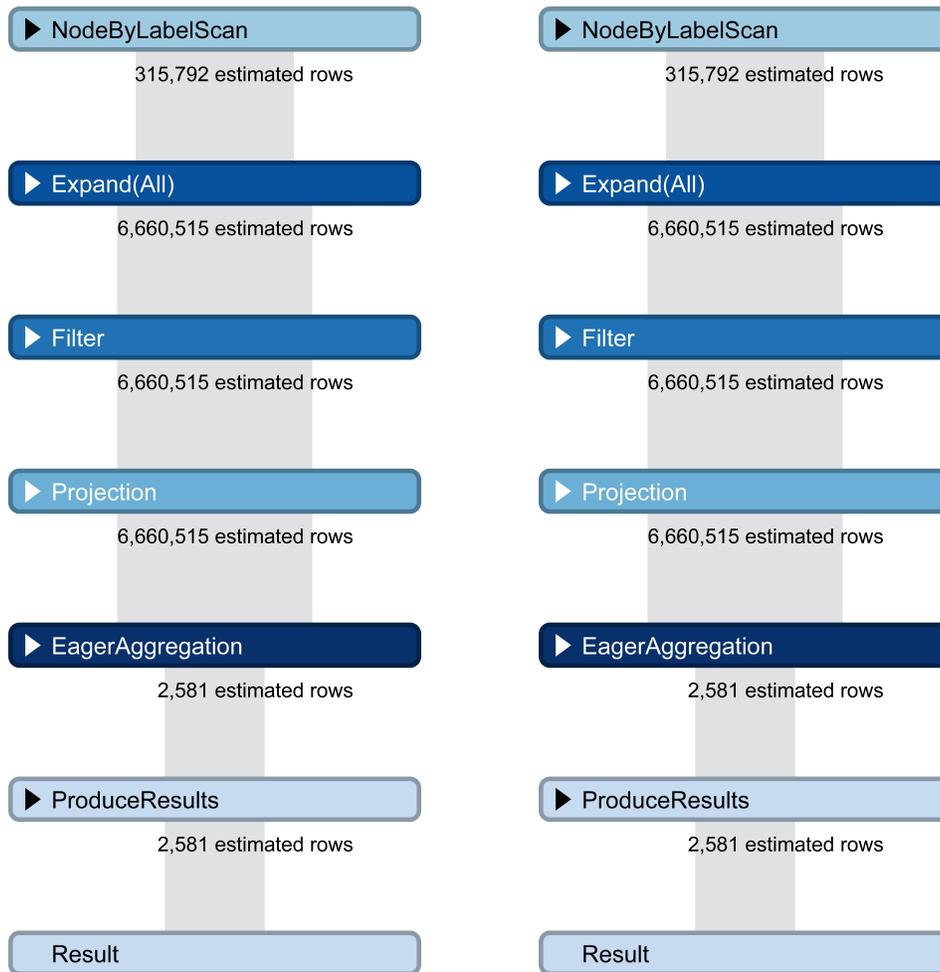


Figure 5.16: Execution plans for *User-Post: User.Age* and *User-Post: ID(User), ID(Post)*.

This is why we use the Cartesian product of dimensions to estimate cuboid sizes in “Single CubePlanner” (covered in Section 4.2.3).

Chapter 6

Conclusion

Our work addresses the urgent need for efficient OLAP processing over large property graphs. We present an end-to-end system to tackle the problem we integrate our proposed solution into Neo4j. The main idea of our solution is to accelerate processing of new queries using materialized views that were selected based on their benefits in previous workload. We emphasize structure as a key feature of graph OLAP queries, which is a complement of previous work [29]. We introduce and distinguish the concepts of cuboid and substructure, which are two important types of materialized views and should be handled differently. Our solution aims at property-aware and structure-aware materialization by taking both cuboids and substructures into account. To achieve this, a greedy selection framework is proposed for materialized view selection. We provide three approaches for substructure join and data retrieval from Neo4j, and analyze their advantages and disadvantages. Our solution works for any SPARQL-like query over the schema graph of any property graph, and it is compatible for other graph databases. We conduct experiments by running a set of queries using our system and native Neo4j. Results show that with acceptable space cost in materialization (12.8% of the total size of the dataset), our system achieves 30x speedup over native Neo4j. The experiments demonstrate our system to be a feasible solution that achieves efficient OLAP queries processing on large graph datasets. In addition, we make reflections on Neo4j and find out a defect in its result size estimation for aggregation queries.

There are a number of directions for future work. First, as reflected in experiments in Section 5.3, performance of our system is affected by system settings, hence investigation on appropriate system settings would be a valuable follow-up to our work. Second, the system we implemented so far supports SPARQL-like queries over schema graph, instead of data graph. It is possible to execute SPARQL-like queries over data graph within our proposed

solution framework. The key issue is to take isomorphism into consideration during query decomposition and composition. Third, future researchers may conduct experiments in which all the views are materialized, which can be used as the optimal situation in efficiency comparison. We did not do such experiments because the space cost would be overwhelming and we did not have a proper memory management system. Besides time for materializing all the views would be too long. In our experiment, simply materializing one substructure of three edges takes more than 5 days, not to mention materializing views of all combinations of structures and attributes. Third, estimation plays an important part in our solution. Better methods for estimation of time cost and space cost can be studied. Finally, greedy selection framework is not a perfect solution for “Materialization Selection” problem. It may cause a local maximal issue. More advanced solutions like simulated annealing could be borrowed to solve this problem.

References

- [1] Deepak Agarwal and Bee-Chung Chen. Latent OLAP: data cubes over latent variables. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 877–888, 2011.
- [2] S. Antony, Ping Wu, D. Agrawal, and A. El Abbadi. Moolap: Towards multi-objective olap. In *Proc. 24th Int. Conf. on Data Engineering*, pages 1394–1396, 2008.
- [3] Surajit Chaudhuri. An overview of query optimization in relational systems. In *Proc. 17th ACM SIGACT-SIGMOD-SIGART Symp. on Principles of Database Systems*, pages 34–43, 1998.
- [4] Chen Chen, Xifeng Yan, Feida Zhu, Jiawei Han, and Philip S. Yu. Graph OLAP: Towards online analytical processing on graphs. In *Proc. 8th IEEE Int. Conf. on Data Mining*, pages 103–112, 2008.
- [5] Alfredo Cuzzocrea and Carson Kai-Sang Leung. Efficiently compressing OLAP data cubes via r-tree based recursive partitions. In *Proc. Foundations of Intelligent Systems 20th Int. Symposium*, pages 455–465, 2012.
- [6] Grzegorz Drzadzewski and Frank Wm. Tompa. Partial materialization for online analytical processing over multi-tagged document collections. *Knowl. Inf. Syst.*, 47(3):697–732, 2016.
- [7] Shifeng Fang, Li Da Xu, Yunqiang Zhu, Jiaerheng Ahati, Huan Pei, Jianwu Yan, and Zhihui Liu. An integrated system for regional environmental monitoring and management based on internet of things. *IEEE Trans. Industrial Informatics*, 10(2):1596–1605, 2014.
- [8] Karam Gouda and Mohammed Javeed Zaki. Efficiently mining maximal frequent itemsets. In *Proc. 1st IEEE Int. Conf. on Data Mining*, pages 163–170, 2001.

- [9] Daniel Hernández, Aidan Hogan, Cristian Riveros, Carlos Rojas, and Enzo Zerega. Querying wikidata: Comparing sparql, relational and graph databases. In *Proc. 15th Int. Semantic Web Conf.*, pages 88–103, 2016.
- [10] Jürgen Hölsch, Tobias Schmidt, and Michael Grossniklaus. On the performance of analytical and pattern matching graph queries in neo4j and a relational database. In *Proc. 2017 Workshops of the EDBT/ICDT Joint Conf.*, 2017.
- [11] Sungpack Hong, Siegfried Depner, Thomas Manhardt, Jan Van Der Lugt, Merijn Verstraaten, and Hassan Chafi. PGX.D: a fast distributed graph processing engine. In *Proc. 2015 ACM/IEEE Conf. on High Performance Computing, Networking, Storage and Analysis*, pages 58:1–58:12, 2015.
- [12] Wararat Jakawat, Cécile Favre, and Sabine Loudcher. OLAP cube-based graph approach for bibliographic data. In *Proc. of Student Research Forum Papers and Posters at 2016 co-located with 42nd Int. Conf. on Current Trends in Theory and Practice of Computer Science*, pages 87–99, 2016.
- [13] Alekh Jindal, Samuel Madden, Malú Castellanos, and Meichun Hsu. Graph analytics using vertica relational database. In *Proc. 2015 IEEE International Conference on Big Data*, pages 1191–1200, 2015.
- [14] Michael Lawrence and Andrew Rau-Chaplin. Dynamic view selection for OLAP. In *Proc. 8th Int. Conf. Data Warehousing and Knowledge Discovery*, pages 33–44, 2006.
- [15] Chuan Li, Philip S. Yu, Lei Zhao, Yan Xie, and Wangqun Lin. InfoNetOLAPer: Integrating InfoNetWarehouse and InfoNetCube with InfoNetOLAP. *Proc. VLDB Endowment*, 4(12):1422–1425, 2011.
- [16] Wen-Yang Lin and I-Chung Kuo. A genetic selection algorithm for OLAP data cubes. *Knowl. Inf. Syst.*, 6(1):83–102, 2004.
- [17] Ling Liu and M. Tamer Özsu, editors. *Encyclopedia of Database Systems*. Springer, 2009.
- [18] Fadi Maali, Stéphane Campinas, and Stefan Decker. Gagg: A graph aggregation operator. In *Proc. 14th Int. Semantic Web Conf.*, pages 491–504, 2015.
- [19] Lhouari Nourine and Olivier Raynaud. A fast algorithm for building lattices. *Inf. Process. Lett.*, 71(5-6):199–204, 1999.

- [20] Z. Meral Özsoyoglu. Review - using semi-joins to solve relational queries. *ACM SIGMOD Digital Review*, 1, 1999.
- [21] André Petermann, Martin Junghanns, Robert Müller, and Erhard Rahm. Graph-based data integration and business intelligence with biiig. *PVLDB*, 7(13):1577–1580, August 2014.
- [22] Arun N. Swami and K. Bernhard Schiefer. On the estimation of join result sizes. In *Advances in Database Technology, Proc. 4th Int. Conf. on Extending Database Technology*, pages 287–300, 1994.
- [23] Dominik Tomaszuk. RDF data in property graph model. In *Proc. Metadata and Semantics Research - 10th Int. Conf.*, pages 104–115, 2016.
- [24] Zhengkui Wang, Qi Fan, Huiju Wang, Kian-Lee Tan, Divyakant Agrawal, and Amr El-Abbadi. Pagrol: PArallel GRaph OLap over lege-scale attributed graphs. In *Proc. 30th Int. Conf. on Data Engineering*, pages 496–507, 2014.
- [25] Jim Webber. A programmatic introduction to neo4j. In *Proc. 2012 Conf. on Systems, Programming, and Applications: Software for Humanity*, pages 217–218, 2012.
- [26] Feng Xia, Laurence T. Yang, Lizhe Wang, and Alexey V. Vinel. Internet of things. *Int. J. Communication Systems*, 25(9):1101–1102, 2012.
- [27] Konstantinos Xirogiannopoulos, Virinchi Srinivas, and Amol Deshpande. Graphgen: Adaptive graph processing using relational databases. In *Proc. 5th Int. Workshop on Graph Data Management Experiences and Systems*, pages 9:1–9:7, 2017.
- [28] Claudio Zaza, Sandro Bimonte, Crescenzo Gallo, Nicola Faccilongo, Piermichele La Sala, and Francesco Contò. OLAP analysis of integrated pest management’s defense rules: Application to olive crop in apulia region. In *Proc. Int. Conf. on Computational Science and Its Applications*, pages 31–44, 2017.
- [29] Peixiang Zhao, Xiaolei Li, Dong Xin, and Jiawei Han. Graph cube: on warehousing and OLAP multidimensional networks. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 853–864, 2011.

APPENDICES

.1 Previous Workload

Meaning for each previous query is listed below.

P1 User-Post: User.UpVotes, Post.Score=9

To see the distribution of users' upvotes for high score (score=9) posts.

P2 User-Post: User.UpVotes, AVG(Post.Score)

To see average post scores by different upvotes.

P3 User-Post: User.Age, SUM(Post.ActiveMonth)

To see each age group's contribution to total posts' active time. What is the main stream users' age range in stackoverflow.com?

P4 User-Post: User.CreationDate_Year, Post.PostTypeId=1

To see numbers of questions (Post.PostTypeId=1) posted by different years when joining the forum.

P5 Badge-User, User-Post, Post-Tag: Tag.TagName, User.CreationDate_Year=2017

In 2017, how many badges are "involved" in each tag?

P6 Badge-User, User-Post, Post-Tag: Tag.TagName, Badge.Name

For each tag, what is the distribution of different types of "involved" badges?

P7 Badge-User, User-Post:Badge.Date_Year, AVG(Post.Score)

How average post score varies by years that badges are honored? Has the "value" of badges changed by time?

P8 Badge-User, User-Post:Badge.Class, AVG(Post.ActiveMonth)

Does high class badge indicate long post active month?

P9 User-Post, Post-Tag: AVG(User.Age), Tag.TagName

Which topics are trendy among youngster users and which ones are popular among middle-aged users?

P10 User-Post, Post-Tag: AVG(User.UpVotes), Tag.TagName=Java

What's the average upvotes (weighted) for users who involve in tag "Java".

P11 User-Post, Post-Vote: AVG(User.UpVotes), Vote.VoteTypeId

For different type of votes, is there a difference in the voters average upvotes?

P12 Post-Comment, Post-PostLink: PostLink-LinkTypeId, AVG(Comment-Score)

Is there a connection between post's link type and post's average comment score?

.2 Future Workload

Intuition for asking each future query is listed as follows.

Q1 User-Post: User.CreationDate_Year=2017, Post.PostTypeId

For users who join recently in year 2017, how many posts are posted for each type of posts? How many are questions and answers?

Q2 User-Post: AVG(User.UpVotes), Post.Score

What is the average users' upvotes for each post score?

Q3 User-Post: Post.ActiveMonth, AVG(User.Age)

For posts of different active timespan, is there a difference in average users' age?

Q4 User-Post: User.CreationDate_Year

How many posts are posted for users who joined in different years?

Q5 Badge-User, User-Post, Post-Tag: Tag.TagName, Badge.Class

Do users of different classes of badges have different topics of interest?

Q6 Badge-User, User-Post, Post-Tag: Tag.TagName, Badge.Date_Year

Is there a major shift in topics of interest for users who receive badges in different years?

Q7 Badge-User, User-Post:Badge.Name, Post.PostTypeId

How many posts are posted for different types of posts and badges?

Q8 User-Post, Post-Tag: User.UpVotes, Tag.TagName, Post.PostTypeId=2

What is the distribution of users' upvotes for each topic of answers (Post.PostTypeId=2)?

Q9 User-Post, Post-Tag:User.CreationDate_Year, Tag.TagName

Do users who join in different years have different topics of interests?

Q10 Badge-User, User-Comment: Badge-Class, AVG(Comment-Score)

Do users of higher classes of badges tend to be more "picky with comments?"

Q11 Badge-User, User-Comment: Badge-Name, AVG(Comment-Score)

Do users of different types of badges give different comment scores? For example, do "masters" tend to give lower comment scores than "students"?

Q12 Post-PostHistory, Post-Tag: Tag-TagName, PostHistory-PostHistoryTypeId

For different tags, is there a difference in post histories related to the tags? For example, posts of which tags are more often re-edited?