



NATIONAL AND KAPODISTRIAN UNIVERSITY OF ATHENS

**SCHOOL OF SCIENCE
DEPARTMENT OF INFORMATICS AND TELECOMMUNICATIONS**

POSTGRADUATE PROGRAM

“Data Science and Information Technologies, Big Data and Artificial Intelligence”

MASTER'S THESIS

**Using Graph Databases for Query Lineage Tracking in
Relational Data**

Panagiotis X.Dimakopoulos

Supervisor: Katerina Doka, Researcher

ATHENS

MARCH 2025

MASTER'S THESIS

Using Graph Databases for Query Lineage Tracking in Relational Data

Panagiotis X.Dimakopoulos
A.M.: 7115152300022

Supervisor: **Katerina Doka**, Researcher

ΕΞΕΤΑΣΤΙΚΗ ΕΠΙΤΡΟΠΗ: **Delis Alexis**, Professor
Ntoulas Alexandros, Assistant Professor

March 2025

ΠΕΡΙΛΗΨΗ

Η παρακολούθηση της γενεαλογίας ερωτημάτων (query lineage tracking) είναι κρίσιμη στην ανάλυση δεδομένων, καθώς επιτρέπει τον εντοπισμό των πηγών δεδομένων, των μετασχηματισμών και των εξαρτήσεων. Σε ένα περιβάλλον σχεσιακών βάσεων δεδομένων, η ικανότητα ιχνηλάτησης της προέλευσης των δεδομένων είναι απαραίτητη για τη διασφάλιση της ποιότητας των δεδομένων, την αποσφαλμάτωση σφαλμάτων και τη βελτιστοποίηση των αναλυτικών διαδικασιών. Η παρούσα διπλωματική εργασία επικεντρώνεται στην αξιοποίηση των γραφηματικών βάσεων δεδομένων για την αποθήκευση και ανάλυση της γενεαλογίας των ερωτημάτων που εκτελούνται σε σχεσιακές βάσεις δεδομένων. Συγκεκριμένα, το **PostgreSQL** χρησιμοποιείται για την εκτέλεση ερωτημάτων και την παρακολούθηση της γενεαλογίας, ενώ το **Neo4j** αξιοποιείται για την αποθήκευση και ανάλυση των σχέσεων μεταξύ των δεδομένων. Η αξιολόγηση απόδοσης πραγματοποιείται για να εκτιμηθεί η επίδραση της παρακολούθησης γενεαλογίας στην εκτέλεση των ερωτημάτων. Για την αξιολόγηση, χρησιμοποιείται το σύνολο δεδομένων **TPC-H**, το οποίο αποτελεί ευρέως αναγνωρισμένο benchmark για την αποτίμηση της απόδοσης των βάσεων δεδομένων. Τα αποτελέσματα δείχνουν ότι οι γραφηματικές βάσεις δεδομένων προσφέρουν αποδοτική ανάλυση της γενεαλογίας των ερωτημάτων, χωρίς να επηρεάζουν σημαντικά την απόδοση των βασικών λειτουργιών των σχεσιακών βάσεων δεδομένων.

ΘΕΜΑΤΙΚΗ ΠΕΡΙΟΧΗ: Βάσεις Δεδομένων, Ιχνηλάτηση Δεδομένων, Διαχείριση Δεδομένων

ΛΕΞΕΙΣ ΚΛΕΙΔΙΑ: Ιχνηλάτηση Γραμμής Ερωτημάτων, Σχεσιακές Βάσεις Δεδομένων, Γραφοβάσεις Δεδομένων, Γραφο-βασισμένη Ιχνηλάτηση, Αποθήκευση Μεταδεδομένων, Πρότυπο TPC-H

ABSTRACT

Query lineage tracking is crucial in data analytics as it enables the identification of data sources, transformations, and dependencies. In a relational database environment, the ability to trace the origin of data is essential for ensuring data quality, debugging errors, and optimizing analytical processes. This thesis focuses on leveraging graph databases for storing and analyzing the lineage of queries executed in relational databases. Specifically, PostgreSQL is used for executing queries and tracking lineage, while Neo4j is utilized for storing and analyzing relationships between data elements. Performance evaluation is conducted to assess the impact of lineage tracking on query execution. For evaluation, the TPC-H benchmark dataset is used, a widely recognized dataset for database performance assessment. The results demonstrate that graph databases provide efficient query lineage analysis without significantly affecting the performance of core relational database operations.

SUBJECT AREA: Databases, Data Lineage, Data Management

KEYWORDS: Query Lineage Tracking, Relational Databases, Graph Databases, Graph-Based Lineage, Metadata Storage, TPC-H Benchmark

Table of Contents

1. INTRODUCTION	12
2. QUERY LINEAGE THEORY	12
2.1 Definition and Importance of Data Lineage	12
2.1.1 Understanding Data Lineage	Error! Bookmark not defined.
2.2 Why Data Lineage Matters	13
2.2.1 Ensuring Data Accuracy and Integrity	13
2.2.2 Error Detection and Resolution	13
2.2.3 Regulatory Compliance and Governance	13
2.2.4 Business Intelligence and Decision-Making	14
2.2.5 Impact Analysis and Change Management	14
2.3 Challenges in Data Lineage Implementation	14
2.4 Applications of Data Lineage in Various Domains	14
2.4.1 ETL Processes and Data Warehousing	14
2.4.2 Data Analytics and Business Intelligence	15
2.4.3 Regulatory Compliance and Risk Management	16
2.4.4 Machine Learning and AI Pipelines	16
2.5 Granularity in Data Lineage Tracking	16
2.5.1 Fine-Grained Lineage: Precision and Deep Traceability	16
2.5.2 Coarse-Grained Lineage: Scalability and High-Level Insights	17
2.5.3 Finding the Right Balance: The Hybrid Approach	17
3. EXISTING DATA LINEAGE TOOLS AND FRAMEWORKS	17
3.1 OpenLineage: A Standardized Approach to Data Lineage Tracking	18
3.2 Apache Atlas: Enterprise Metadata Management for Hadoop Ecosystems	18
3.3 Egeria: A Universal Metadata Exchange Framework	18
3.4 Spline: Lineage Tracking for Apache Spark Pipelines	19
3.5 OpenMetadata: Unified Metadata and Lineage Management	19

4. SYSTEM ARCHITECTURE	20
4.1 PostgreSQL: The Relational Database for Query Execution and Lineage Logging.....	20
4.1.1 Query Execution and Lineage Logging.....	21
4.1.2 Optimizing Performance for Lineage Tracking.....	21
4.2 Neo4j: Graph-Based Lineage Storage and Analysis	21
4.2.1 Graph Structure and Node Types	22
4.2.2 Relationships and Lineage Tracking.....	22
4.2.3 Enhanced Querying and Metadata Utilization.....	22
4.3 Python Pipeline: Synchronization Between PostgreSQL and Neo4j	23
4.4 Performance Evaluation with TPC-H Benchmark	23
4.5 End-to-End System Flow	25
5. DATA DESCRIPTION	26
5.1 Overview of the Dataset.....	26
5.2 Structure of the TPC-H Database.....	27
5.3 Metadata Storage for Query Lineage Tracking	28
5.3.1 Query Execution Log (query_log)	28
5.3.2 Query Metadata (query_metadata).....	29
5.3.3 Query Results (query_results)	29
5.3.4 Query Cell History (query_cell_history)	30
5.4 Relationships Between Metadata Tables	30
5.5 Integration with Query Lineage Tracking.....	31
6. IMPLEMENTATION	33
6.1 System Setup and Environment	33
6.2 Query Execution and Metadata Logging.....	33
6.3 Python Synchronization Pipeline	34
6.4 Query Performance with Lineage Tracking	34
6.5 Summary of Implementation	35
7. EVALUATION & BENCHMARKS.....	35

7.1	Overview.....	35
7.2	Query Execution Performance.....	36
7.3	Neo4j: Graph-Based Lineage Storage and Analysis	37
7.3.1	PostgreSQL Storage Analysis	37
7.3.2	Neo4j Storage and Graph Growth	38
7.4	System Resource Usage	39
7.5	Scalability Considerations	40
7.6	Summary of Evaluation.....	40
8.	CONCLUSIONS & FUTURE WORK.....	41
8.1	Conclusions	41
8.2	Limitations	42
8.3	Future Work	42
8.3.1	Performance Optimization.....	42
8.3.2	Scalable Storage Management.....	43
8.3.3	Integration with Machine Learning	43
8.3.4	Expanding to Multi-Platform Environments.....	43
8.3.5	Enhancing Regulatory Compliance Features	43
8.4	Final Remarks.....	43
	REFERENCES	44

List of Figures

Figure 2.1: Architecture of the Query Lineage Tracking Framework	20
Figure 5.1: Query Execution Flow for Lineage Tracking in PostgreSQL.....	27
Figure 5.2: Metadata Storage Schema for Query Lineage Tracking.....	32
Figure 6.1: Graph Schema for Query Lineage in Neo4j.....	34
Figure 7.1: Query Execution Time Comparison (With and Without Lineage Tracking)..	37

1. Introduction

In today's data-driven landscape, organizations face mounting pressure to ensure data quality, compliance, and operational efficiency. As data ecosystems grow increasingly complex—spanning hybrid cloud architectures, machine learning pipelines, and heterogeneous sources—the need for transparent data lineage has become critical. Data lineage, the process of tracking data's origin, transformations, and dependencies, is foundational to troubleshooting errors, optimizing workflows, and meeting regulatory mandates like GDPR and CCPA. Yet, traditional lineage methods, reliant on manual documentation or siloed metadata, struggle to scale with modern demands.

This dissertation addresses this gap by proposing a graph-based lineage tracking system that integrates PostgreSQL, a relational database, with Neo4j, a graph database, to enhance lineage visualization, querying, and analysis. Unlike existing tools that focus on single-platform tracking (e.g., Spline for Apache Spark) or lack granularity (e.g., coarse-grained lineage in Apache Atlas), our approach leverages Neo4j's graph traversal capabilities to dynamically map fine-grained dependencies across hybrid environments. By automating lineage capture during query execution and enabling real-time impact analysis, this system empowers organizations to balance scalability with traceability—a challenge underscored in prior work (e.g., [cite sources]).

Structure of the Dissertation

Chapter 2 establishes the theoretical foundations of data lineage, including its role in compliance and the trade-offs between granularity and scalability. Chapter 3 critically evaluates existing lineage tools, identifying gaps in cross-platform interoperability that our architecture addresses. Chapters 4 and 5 detail the system's design and implementation, combining PostgreSQL's query logging with Neo4j's graph-based lineage representation. Finally, Chapters 6 and 7 validate the system's efficiency through TPC-H benchmarks, demonstrating practical viability for enterprise-scale deployments.

By bridging relational and graph-based paradigms, this work contributes a novel framework for lineage tracking that prioritizes both transparency and performance—a necessity for data-driven organizations navigating evolving regulatory and technical landscapes.

2. Query Lineage Theory

2.1 Definition and Importance of Data Lineage

2.1.1 Understanding Data Lineage

Data lineage represents the complete lifecycle of data, mapping its origin, transformations, movement, and dependencies across an organization's ecosystem. By visualizing how data flows between systems, applications, and analytical tools, organizations can ensure transparency, maintain data integrity, and optimize decision-making processes.

The ability to track data lineage is fundamental to ensuring accurate reporting, regulatory compliance, and effective data governance. Without a well-defined lineage framework, organizations risk data inconsistencies, errors in analytics, and inefficiencies in operational processes, which can ultimately impact strategic business objectives.

At its core, data lineage helps organizations answer critical questions such as:

- Where does this data originate?
- How has it been transformed?
- Which systems, processes, or users have interacted with it?
- What dependencies exist between different datasets?
- How will changes to a dataset impact downstream processes?

By addressing these concerns, data lineage enables organizations to manage, analyze, and govern their data assets more effectively, fostering a structured approach to data reliability and operational efficiency.

2.2 Why Data Lineage Matters

2.2.1 Ensuring Data Accuracy and Integrity

As organizations collect vast amounts of data, maintaining accuracy and consistency becomes a significant challenge. Data transformations, manual inputs, and integrations between different systems often introduce discrepancies. Lineage tracking helps validate data at every stage, ensuring that transformations align with business rules and expected outputs. This capability is particularly crucial in analytics, where even minor inconsistencies can lead to flawed insights and poor business decisions.

2.2.2 Error Detection and Resolution

Data anomalies, schema changes, and unexpected transformations can disrupt business intelligence (BI) and reporting processes. Without proper tracking mechanisms, identifying the root cause of these issues is challenging. Data lineage allows teams to trace the flow of data, pinpoint errors at their source, and implement corrective measures without affecting downstream applications.

2.2.3 Regulatory Compliance and Governance

Regulatory frameworks such as GDPR, CCPA, and HIPAA mandate stringent accountability for data handling, requiring organizations to maintain auditable records of data provenance, access, and transformations. Data lineage directly supports compliance by providing end-to-end visibility into how sensitive information—such as PII or health records—flows across systems, enabling organizations to trace unauthorized modifications, validate adherence to retention policies, and efficiently respond to audit requests or data subject inquiries. For instance, GDPR’s “right to explanation” obligates enterprises to disclose how automated decisions derive from raw inputs, a process enabled by fine-grained lineage tracking that maps data from source to output. Traditional approaches, which rely on siloed or coarse-grained lineage, struggle to meet these demands in hybrid environments, risking incomplete audits and regulatory penalties. The proposed system addresses this gap through automated metadata logging (Chapter 5) and graph-powered dependency tracing (Chapter 4), which dynamically link queries, datasets, and users to demonstrate compliance. By leveraging Neo4j’s real-time traversal capabilities, the architecture ensures that organizations can swiftly identify impacted systems during schema changes or data breaches—a critical advantage over static, table-centric tools like Apache Atlas. This integration of granular lineage with scalable graph analytics not only fulfills regulatory requirements but also reduces manual oversight, positioning the framework as a robust solution for modern governance challenges.

2.2.4 Business Intelligence and Decision-Making

Data-driven decision-making relies on accurate and well-understood data. Without lineage tracking, organizations may struggle to validate the quality of their data sources, leading to uncertainty in analytics. By providing end-to-end visibility into data transformations, lineage tracking enhances trust in reports, dashboards, and predictive models, ensuring that business leaders base their decisions on reliable insights.

2.2.5 Impact Analysis and Change Management

When modifying database schemas, updating ETL workflows, or adjusting reporting logic, organizations need to understand the potential impact of these changes. Data lineage provides a comprehensive view of data dependencies, allowing teams to assess how modifications will affect downstream systems before implementation. This proactive approach minimizes disruptions and prevents data inconsistencies from propagating across the enterprise.

2.3 Challenges in Data Lineage Implementation

Despite its advantages, data lineage tracking presents several implementation challenges.

One of the primary difficulties is scalability. As data ecosystems grow, tracking every transformation and dependency becomes computationally expensive. Storing lineage metadata for vast datasets requires careful optimization to balance storage costs and processing efficiency.

Another major challenge is interoperability across diverse architectures. Many organizations operate in hybrid environments that combine on-premises databases, cloud platforms, and third-party services. Standardizing lineage tracking across these systems requires robust integrations and well-defined metadata standards.

Additionally, real-time lineage tracking introduces performance concerns, particularly in big data environments where thousands of transformations occur per second. Ensuring that lineage tracking does not introduce latency while still capturing detailed metadata is a significant technical challenge.

Finally, security and privacy considerations must be addressed. Organizations must ensure that sensitive data is not exposed through lineage tracking, especially when dealing with personally identifiable information (PII) and confidential business data.

Addressing these challenges requires a combination of automated lineage tracking tools, graph-based metadata storage, and efficient data governance frameworks to ensure seamless integration into existing workflows.

2.4 Applications of Data Lineage in Various Domains

Data lineage plays a fundamental role across various industries and technological domains, ensuring transparency, reliability, and regulatory compliance in data management. By tracking data as it moves through systems, lineage provides organizations with the ability to analyze dependencies, detect inconsistencies, and optimize workflows. From ETL processes to advanced machine learning pipelines, lineage tracking serves as a critical tool for both operational efficiency and strategic decision-making.

2.4.1 ETL Processes and Data Warehousing

Extract, Transform, Load (ETL) pipelines are at the core of modern data architectures, enabling organizations to integrate data from multiple sources, clean and transform it

based on business rules, and store it in data warehouses or data lakes. These processes are critical in ensuring that data is structured and ready for analytics and reporting. However, given the complex nature of ETL workflows, data lineage tracking is essential to ensure accuracy, optimize performance, and enable debugging when data inconsistencies arise.

One of the biggest challenges in ETL workflows is identifying the root cause of errors. When data discrepancies occur, tracing the lineage of data transformations allows organizations to pinpoint whether the issue originated from an incorrect data source, a faulty transformation logic, or an incomplete data load. By continuously monitoring lineage in ETL workflows, organizations can minimize data corruption, reduce operational downtime, and ensure that analytics are built on reliable datasets.

Another major benefit of lineage tracking in ETL pipelines is performance optimization. ETL processes often involve multiple transformations, joins, aggregations, and schema modifications, each of which affects processing time and resource consumption. Lineage tracking enables data engineers to detect unnecessary transformations, identify redundant operations, and optimize ETL logic, leading to faster processing times and lower infrastructure costs.

Additionally, data lineage in ETL workflows ensures compliance and auditability by maintaining a historical record of how data has been processed. This is especially important in regulated industries, where businesses need to demonstrate how data moves through their systems, from ingestion to storage. As organizations increasingly transition to cloud-based data pipelines, the ability to maintain accurate lineage records across hybrid infrastructures has become even more critical.

2.4.2 Data Analytics and Business Intelligence

Beyond ETL, data lineage is crucial in business intelligence (BI) and analytics, where accurate insights drive decision-making. The quality of reports, dashboards, and predictive models depends on the reliability of the underlying data. Without a clear understanding of data lineage, organizations may struggle to verify the accuracy of their datasets, leading to inconsistencies, incorrect trends, and flawed strategic choices.

A major issue faced by analytics teams is ensuring trust in data sources. Many organizations pull data from multiple departments, databases, and external APIs, making it difficult to assess which data is the most accurate and up-to-date. Lineage tracking helps analysts trace the origins of datasets, verify transformations, and confirm that insights are based on validated and well-structured information.

Additionally, lineage tracking improves metric reproducibility. Analysts frequently need to replicate past analyses, compare current data with historical trends, or validate KPI calculations. Without lineage, these tasks become challenging, as changes in underlying data sources or transformations can alter results. By implementing lineage tracking, organizations can maintain a transparent workflow, ensuring that analyses remain consistent over time.

In environments where companies operate across multiple data storage solutions—such as cloud platforms, on-premise databases, and third-party services—lineage tracking helps maintain data consistency across BI tools. This prevents discrepancies in reports, ensures alignment across analytical platforms, and allows business leaders to make confident, data-driven decisions.

2.4.3 Regulatory Compliance and Risk Management

For organizations operating in highly regulated sectors such as finance, healthcare, and government, compliance with data protection laws is non-negotiable. Regulatory frameworks such as GDPR, CCPA, HIPAA, and SOX require businesses to maintain clear records of how data is processed, stored, and shared. Data lineage tracking supports compliance by providing an auditable history of data movements, allowing organizations to demonstrate transparency in their data handling practices.

Beyond regulatory compliance, lineage tracking enhances risk management by ensuring that access to sensitive data is properly controlled. Unauthorized modifications to data, whether intentional or accidental, can lead to financial losses, legal consequences, and reputational damage. By implementing lineage tracking, organizations can monitor who has accessed or altered specific datasets, enforce data governance policies, and prevent potential security breaches.

Additionally, regulatory audits and legal inquiries often require companies to provide documentation on historical data transformations. With lineage tracking in place, organizations can efficiently generate reports that showcase how data has been handled over time, reducing the complexity and time required for audit preparation.

2.4.4 Machine Learning and AI Pipelines

With the rise of machine learning (ML) and artificial intelligence (AI), data lineage has become an essential component in model governance. The quality, fairness, and accuracy of AI models are directly influenced by the datasets used for training, making lineage tracking indispensable for ensuring data transparency and accountability.

Lineage tracking in AI/ML environments allows organizations to:

- Trace training data sources to ensure that models are built on unbiased, representative datasets.
- Monitor feature engineering processes, ensuring that transformations applied to data do not introduce unintended biases.
- Reproduce experiments, allowing researchers to validate model performance using the same preprocessing and data preparation steps.
- Implement version control, tracking how data changes impact model predictions over time.

As AI adoption increases in high-stakes applications—such as healthcare diagnostics, financial risk assessments, and automated decision-making—explainability and trust in AI models become more critical. Data lineage provides an essential framework for AI ethics and fairness, ensuring that models are trained on transparent, well-documented datasets.

2.5 Granularity in Data Lineage Tracking

The level of granularity in data lineage tracking determines how much detail is captured about a dataset's journey through an organization's systems. Organizations must carefully consider whether to implement fine-grained or coarse-grained lineage based on their analytical needs, compliance requirements, and resource constraints.

2.5.1 Fine-Grained Lineage: Precision and Deep Traceability

Fine-grained lineage captures data transformations at an extremely detailed level, tracking individual records, attributes, or even specific cell values. This level of precision

allows organizations to audit data transformations with high accuracy, making it invaluable for compliance, error detection, and root cause analysis.

One of the key advantages of fine-grained lineage is its ability to provide high-resolution impact analysis. When a dataset is modified, organizations can determine exactly which reports, machine learning models, or business workflows will be affected. This prevents unexpected disruptions, ensuring that data consumers are aware of potential changes before they impact critical systems.

However, maintaining fine-grained lineage requires significant storage and processing resources. Capturing every individual data change, especially in high-volume environments, can introduce computational overhead, slowing down query performance. To mitigate this, organizations must implement efficient metadata storage solutions and leverage graph databases to handle lineage at scale.

2.5.2 Coarse-Grained Lineage: Scalability and High-Level Insights

Unlike fine-grained lineage, coarse-grained lineage tracks data movement at a higher level of abstraction, such as table-level, dataset-level, or workflow-level tracking. This approach is significantly more scalable and computationally efficient, making it well-suited for big data environments, cloud platforms, and distributed architectures.

While coarse-grained lineage is beneficial for understanding data dependencies and monitoring system-level changes, it lacks the precision needed for detailed auditing or troubleshooting. Organizations that rely solely on coarse-grained lineage may find it difficult to trace the exact source of errors, making debugging more challenging.

2.5.3 Finding the Right Balance: The Hybrid Approach

Many organizations adopt a hybrid approach to lineage tracking, combining fine-grained and coarse-grained methods based on specific business needs. For example:

- Financial institutions may track transactions at the cell level for compliance but monitor general data flow at the dataset level.
- AI/ML teams may use fine-grained tracking for training datasets but rely on coarse-grained lineage for feature selection and model deployment.
- Regulated industries may enforce fine-grained tracking for sensitive data while maintaining high-level lineage for operational workflows.

As data ecosystems grow more complex, intelligent lineage management becomes critical. Organizations must leverage automated lineage tracking solutions that dynamically adjust granularity based on usage patterns, ensuring efficiency without sacrificing traceability.

3. Existing Data Lineage Tools and Frameworks

The growing complexity of modern data ecosystems has led to the development of various open-source tools aimed at data lineage tracking. These tools help organizations visualize data flow, understand dependencies, and ensure compliance with regulatory frameworks. While some focus on broad metadata management and governance, others specialize in tracking transformations within specific data processing environments. Below is an overview of five prominent data lineage tools: OpenLineage, Apache Atlas, Egeria, Spline, and OpenMetadata.

3.1 OpenLineage: A Standardized Approach to Data Lineage Tracking

OpenLineage is an open standard designed to provide a consistent and interoperable approach to lineage tracking across various data processing frameworks. Unlike traditional metadata management solutions that are often tied to specific platforms, OpenLineage introduces a unified API that allows multiple tools to collect and share lineage metadata in a standardized format.

A key advantage of OpenLineage is its focus on automation—it enables organizations to automatically capture data lineage across batch and streaming workflows, reducing manual effort. The framework is particularly useful for debugging, as it allows engineers to track errors and discrepancies across different processing stages. Additionally, OpenLineage supports compliance and governance, as it helps organizations maintain transparency into how data is accessed and transformed.

With its extensible architecture, OpenLineage integrates seamlessly with multiple data orchestration tools such as Apache Airflow, Apache Spark, dbt (data build tool), and Great Expectations. This makes it a powerful solution for organizations looking to unify lineage tracking across a heterogeneous data stack.

3.2 Apache Atlas: Enterprise Metadata Management for Hadoop Ecosystems

Apache Atlas is an open-source metadata and data governance framework originally developed for the Hadoop ecosystem, though it has since evolved to support a broader range of platforms. It provides comprehensive metadata management, including data classification, policy-based access control, and fine-grained lineage tracking.

One of Apache Atlas's core strengths is its ability to offer graph-based visualization of data lineage, allowing users to explore data transformations and dependencies intuitively. Organizations using Apache Atlas can define business glossaries, enforce security policies, and track data movement across distributed environments.

Apache Atlas is particularly beneficial for enterprises that operate within a big data landscape, as it integrates seamlessly with tools like Apache Hive, Apache HBase, Apache Kafka, and Apache Spark. Additionally, it provides REST APIs that allow easy integration with external metadata stores and governance frameworks.

Despite its strengths, Apache Atlas has a steeper learning curve and is primarily suited for organizations operating large-scale Hadoop-based architectures. While it offers extensive lineage tracking, users looking for lightweight and cloud-native solutions may find other tools more suitable.

3.3 Egeria: A Universal Metadata Exchange Framework

Egeria is an open-source project developed by the ODPi (part of the Linux Foundation) to provide an interoperable metadata exchange framework across diverse technologies. Unlike tools that are limited to a single platform, Egeria is designed for organizations that need to share and govern metadata across multiple data environments, vendor solutions, and cloud platforms.

Egeria enables automated lineage tracking by capturing metadata from data lakes, data warehouses, and analytics tools, offering a unified view of data flow across an enterprise. It supports both fine-grained and coarse-grained lineage tracking, ensuring that organizations can trace how specific attributes evolve while also maintaining high-level oversight.

A key differentiator of Egeria is its collaborative ecosystem approach—it facilitates seamless metadata exchange between different tools and platforms. Organizations can integrate Egeria with existing data catalogs, compliance systems, and governance frameworks, making it a vendor-agnostic solution for metadata management.

As organizations move towards multi-cloud and hybrid environments, Egeria provides a future-proof approach to metadata governance by ensuring that metadata can be easily transferred and shared across disparate systems.

3.4 Spline: Lineage Tracking for Apache Spark Pipelines

Spline is a specialized data lineage tracking tool for Apache Spark applications, designed to automatically capture and visualize lineage information at runtime. Unlike other metadata management tools that rely on static metadata collection, Spline dynamically tracks lineage information as Spark jobs are executed.

One of the primary advantages of Spline is its non-intrusive nature—it does not require modifications to existing Spark code. Instead, it passively collects lineage metadata during job execution, making it an easy-to-adopt solution for organizations running Spark-based ETL pipelines.

Spline provides an interactive web interface that enables users to visualize and analyze data flow dependencies within Spark transformations. This capability is particularly useful for:

- Debugging complex Spark applications
- Optimizing performance by identifying inefficient transformations
- Ensuring compliance by documenting how data is processed

However, Spline's primary limitation is that it is designed exclusively for Apache Spark, meaning it may not be suitable for organizations that require multi-platform lineage tracking. Despite this, for companies operating within big data environments, Spline offers a lightweight and effective approach to lineage tracking within Spark-based workflows.

3.5 OpenMetadata: Unified Metadata and Lineage Management

OpenMetadata is a comprehensive metadata management solution that consolidates multiple aspects of data governance, including data discovery, profiling, and lineage tracking. Unlike standalone lineage tracking tools, OpenMetadata aims to provide a unified metadata repository where organizations can centralize information about their data assets, schemas, quality metrics, and dependencies.

One of OpenMetadata's standout features is its column-level lineage tracking, which enables organizations to trace transformations at a granular level. This makes it particularly well-suited for use cases where understanding attribute-level changes is essential, such as financial reporting and regulatory compliance.

Another key strength of OpenMetadata is its integration with modern data tools such as Snowflake, BigQuery, Redshift, Airflow, and dbt. The platform includes a no-code editor that allows non-technical users to manage metadata, making it a more accessible option compared to developer-focused lineage tracking solutions.

OpenMetadata's holistic approach makes it a strong candidate for organizations that need end-to-end visibility into their data assets. However, for teams that are only

interested in lightweight lineage tracking, OpenMetadata's broad feature set may introduce unnecessary complexity.

4. System Architecture

The proposed system architecture for query lineage tracking and performance evaluation integrates a relational database with a graph-based database to enable efficient tracking, visualization, and analysis of data lineage. The model is built upon three primary components: PostgreSQL, which serves as the main query execution environment and stores structured metadata for lineage tracking; Neo4j, which provides a graph-based representation of lineage data to facilitate analysis; and a Python-based synchronization pipeline, responsible for extracting, transforming, and loading lineage information from PostgreSQL to Neo4j.

The system is designed to support both real-time and historical analysis of query lineage while maintaining a minimal impact on query execution performance. The incorporation of TPC-H benchmarks allows for systematic evaluation of the performance overhead introduced by lineage tracking.

Each component plays a distinct role in ensuring that lineage tracking is both accurate and scalable. PostgreSQL is responsible for executing analytical queries and capturing metadata, while Neo4j enables users to explore relationships between queries, datasets, and transformations efficiently. The Python script functions as a bridge, continuously synchronizing relational metadata with the graph database to ensure that the lineage representation remains up to date.

Architecture of the Query Lineage Tracking Framework




Figure 2.1: Architecture of the Query Lineage Tracking

4.1 PostgreSQL: The Relational Database for Query Execution and Lineage Logging

PostgreSQL functions as the primary query execution engine, handling analytical queries while simultaneously capturing and logging lineage-related metadata. This dual functionality ensures that every executed query is not only processed efficiently but also recorded with sufficient detail to support comprehensive lineage tracking. The database contains both the TPC-H dataset, used for benchmarking and testing, and dedicated metadata tables, which store query execution details, data dependencies, and transformations.

4.1.1 Query Execution and Lineage Logging

Every query executed in PostgreSQL generates results while also triggering a structured metadata logging process. This is facilitated by PL/pgSQL functions, which extend standard query execution by embedding lineage tracking within their logic. These functions ensure that key metadata is collected without requiring external monitoring tools or additional background processes.

The lineage tracking process in PostgreSQL follows a systematic approach:

- **Query Execution Logging:** Each query execution is recorded in the `query_log` table, which stores details such as the query name, execution timestamp, and input parameters. This serves as the foundation for tracking query history.
- **Metadata Capture:** Lineage-specific metadata is stored in the `query_metadata` table, documenting which tables and columns were accessed during execution, how data was transformed, and the dependencies between different datasets.
- **Cell-Level Lineage:** Fine-grained tracking is enabled through the `query_cell_history` table, which maps individual data values that were referenced or modified by the query, providing high precision in data tracking.
- **Result Storage:** The output of each query is saved in the `query_results` table, linking generated results directly to the corresponding query execution for traceability.

By maintaining these structured logs, PostgreSQL ensures a detailed and auditable record of all data transformations, facilitating effective debugging, impact analysis, and regulatory compliance.

4.1.2 Optimizing Performance for Lineage Tracking

Since lineage tracking introduces additional write operations alongside query execution, the system is designed to minimize overhead. The metadata logging process is tightly integrated within the function execution logic, ensuring that additional queries or background processes are not required to capture lineage details.

However, despite these optimizations, lineage tracking does introduce a computational cost. To quantify this impact, performance evaluation is conducted using the TPC-H benchmark queries, comparing execution times, storage overhead, and system efficiency with and without lineage tracking enabled. These evaluations help determine the trade-offs involved in maintaining high-resolution lineage tracking while preserving database performance.

By structuring lineage metadata within PostgreSQL and efficiently capturing execution details, the system lays a solid foundation for further lineage analysis. However, relational databases are not inherently designed for complex relationship-based queries, making it necessary to integrate Neo4j to enable intuitive graph-based lineage exploration.

4.2 Neo4j: Graph-Based Lineage Storage and Analysis

While PostgreSQL efficiently stores structured metadata, it is not optimized for lineage exploration, particularly when queries span multiple datasets and transformations. Traditional relational databases rely on complex JOIN operations to reconstruct lineage, making real-time exploration inefficient at scale. To address this limitation, Neo4j is integrated into the system, offering a graph-based representation of lineage data that allows for intuitive querying and efficient traversal of data dependencies.

Neo4j models lineage data as a graph, where key entities involved in query execution—such as queries, tables, columns, and individual results—are represented as nodes, while their interdependencies are established through relationships. Each query execution in PostgreSQL is mirrored in Neo4j, forming a structured lineage graph that accurately captures how data flows through the system. This transformation makes it possible to trace data origins, monitor dependencies, and conduct impact analysis with minimal computational overhead.

4.2.1 Graph Structure and Node Types

The lineage graph consists of three primary types of nodes, each serving a specific role in tracking data movement and transformation:

- QueryLog nodes represent individual query executions. Each node records a unique query execution event, linking it to metadata such as execution time and parameters.
- QueryResultEntry nodes store the actual results produced by queries, preserving the data output from each execution. These nodes serve as an essential reference for tracking how query results are derived from raw datasets.
- Cell nodes represent individual values within tables that were accessed or modified during query execution. By tracking values at the data level, fine-grained lineage analysis becomes possible, allowing users to pinpoint how specific values were transformed over time.

4.2.2 Relationships and Lineage Tracking

To establish connections between these entities, Neo4j defines a set of relationships that capture how data flows across queries and datasets:

- HAS_RESULT: Connects a QueryLog node to its corresponding QueryResultEntry nodes, establishing a direct link between query executions and their produced outputs.
- HAS_CELL: Links QueryLog nodes to Cell nodes, recording which specific data points were referenced or modified during execution.
- CONTRIBUTED_FROM: Establishes a relationship between QueryResultEntry nodes and Cell nodes, making it possible to trace how each output was derived from underlying raw data.

These relationships provide a clear lineage path, allowing users to track the transformation of data across multiple queries. Unlike traditional relational databases, where lineage tracking requires expensive joins and subqueries, Neo4j enables near-instantaneous retrieval of lineage data using Cypher queries, making data exploration and debugging significantly more efficient.

4.2.3 Enhanced Querying and Metadata Utilization

Beyond simple lineage tracking, Neo4j enables advanced filtering and contextual analysis by incorporating additional metadata attributes into the graph. When metadata such as query execution timestamps, column dependencies, and transformation types is transferred from PostgreSQL to Neo4j, it enhances the analytical capabilities of lineage queries. Users can refine lineage searches based on execution time, data complexity, or transformation logic, making it easier to investigate performance bottlenecks, troubleshoot errors, and audit data movements.

By leveraging a graph-based approach, Neo4j optimizes the storage and retrieval of lineage data, providing a scalable and interactive model for data tracking. This integration ensures that users can efficiently navigate, analyze, and validate data lineage, making it a powerful tool for compliance, debugging, and performance optimization.

4.3 Python Pipeline: Synchronization Between PostgreSQL and Neo4j

To integrate structured lineage logging from PostgreSQL with the graph-based representation in Neo4j, a Python-based synchronization pipeline is implemented. This pipeline acts as an intermediary layer that monitors query execution activity in PostgreSQL and transfers lineage metadata to Neo4j, ensuring that lineage relationships remain current and accurate. By automating this process, the system maintains a real-time and consistent view of query dependencies, eliminating the need for manual updates while minimizing processing overhead.

The synchronization process follows a systematic and efficient workflow that ensures both performance and accuracy in lineage tracking. The script first establishes a connection with PostgreSQL, retrieving details of newly executed queries from the `query_log` table. This includes execution time, query parameters, and references to affected tables and columns. Once retrieved, the metadata is further enriched by extracting additional lineage information from related tables, such as `query_results`, `query_metadata`, and `query_cell_history`, which capture transformations, dependencies, and result sets.

Once this metadata is collected, the pipeline transforms the information into a graph-compatible structure, ensuring that lineage relationships are accurately represented in Neo4j. The relationships between queries, data sources, transformations, and results are mapped as nodes and edges, mirroring the execution dependencies established within PostgreSQL. To prevent data redundancy, the system leverages Neo4j's MERGE operations, which check for the existence of nodes before insertion, ensuring that only new entries are added while maintaining lineage integrity.

To improve efficiency, the pipeline is designed to operate incrementally, processing only newly executed queries rather than scanning the entire database each time. This is achieved by tracking the last processed query ID, ensuring that subsequent runs only capture new executions without reprocessing historical data. Additionally, batch processing techniques are employed to group multiple updates into structured transactions, reducing the overhead associated with frequent insertions and updates in Neo4j. This approach minimizes performance bottlenecks, allowing the pipeline to scale efficiently even in high-volume analytical workloads.

Beyond its role in metadata transfer, the pipeline also enhances lineage usability by incorporating query execution timestamps, column dependencies, and transformation types into Neo4j. These metadata attributes enable users to perform advanced filtering and contextual analysis, making it easier to investigate query interactions, detect anomalies, and assess performance trends. By automating the synchronization between PostgreSQL and Neo4j, the pipeline provides a scalable, real-time solution for lineage tracking, ensuring that data flow remains transparent and auditable across complex analytical queries.

4.4 Performance Evaluation with TPC-H Benchmark

To assess the impact of lineage tracking on query performance, the system is evaluated using the TPC-H benchmark, a widely recognized industry standard designed to

measure the efficiency of database systems under analytical workloads. The TPC-H dataset and queries simulate real-world decision-support scenarios, making them an ideal framework for analyzing the computational overhead introduced by lineage tracking mechanisms. By systematically comparing execution time, resource utilization, and storage consumption, the evaluation provides a quantitative measure of how lineage tracking affects query performance and database scalability.

4.4.1 Evaluation Methodology

The benchmarking process is structured around two distinct execution modes to ensure a fair comparison:

1. **Baseline Execution:** Queries are executed without lineage tracking to establish a control group that represents normal query performance. These results serve as a reference point for understanding how queries perform without additional lineage-related operations.
2. **Tracked Execution:** The same queries are executed with full lineage tracking enabled, capturing metadata, transformations, and dependencies in both PostgreSQL and Neo4j. This mode allows for the quantification of the performance trade-offs associated with lineage logging.

By comparing the execution times and computational overhead between these two cases, the system's efficiency is assessed across several key dimensions.

4.4.2 Key Performance Metrics

The evaluation process focuses on three primary performance aspects:

1. Query Execution Overhead

One of the most critical factors in assessing lineage tracking is the additional time required to log metadata and update Neo4j. Since tracking involves extra insertions into lineage-related tables and graph updates, it is important to quantify the impact of these operations on overall query performance. Execution time for each query in both tracked and untracked modes is recorded, allowing for a direct comparison of how lineage tracking affects computational efficiency.

2. Storage Overhead and Growth Rate

Lineage tracking introduces additional metadata storage requirements in both PostgreSQL and Neo4j. This aspect of the evaluation measures the rate at which lineage metadata accumulates over time and how it affects overall system storage capacity. The analysis includes:

- Metadata growth in PostgreSQL, tracking the expansion of lineage tables such as `query_log`, `query_metadata`, and `query_results`.
- Graph expansion in Neo4j, analyzing how the number of nodes and relationships increases as more queries are executed and logged. By studying storage trends, the experiment identifies whether lineage tracking remains scalable as database usage increases.

3. Scalability and Dataset Size Variations

To ensure that lineage tracking remains efficient under different workload sizes, the system is tested with multiple dataset scales. The 100MB dataset (`tpch_small`) provides insights into how the system performs under moderate workloads, while the 1GB dataset (`tpch_large`) evaluates performance in more complex, high-volume environments. Running queries on datasets of increasing size allows for

an assessment of whether lineage tracking maintains efficiency as the system scales, or if noticeable performance degradation occurs.

4.4.3 Lineage Query Performance in Neo4j

Beyond evaluating execution and storage overhead, the study also examines the efficiency of lineage retrieval in Neo4j, a key component of the proposed architecture. Since one of the primary advantages of using a graph-based approach is the ability to rapidly trace data origins and analyze dependencies, specific tests are conducted to measure how quickly Neo4j can answer lineage queries.

The following types of lineage queries are evaluated:

- **Tracing Data Origins:** Determining where a specific data point originated within the system.
- **Impact Analysis:** Identifying how changes to a dataset propagate across dependent queries.
- **Historical Tracking:** Retrieving past transformations applied to a particular dataset over time.

These lineage queries simulate real-world auditing and debugging scenarios, where analysts must investigate how specific results were derived. Their performance is measured in terms of retrieval time, computational complexity, and scalability as the lineage graph expands over time.

4.4.4 Findings and Insights

By systematically analyzing these performance aspects, the benchmarking study provides a comprehensive understanding of the trade-offs introduced by lineage tracking. The key findings reveal:

1. Minimal overhead in execution time for most queries, though some complex queries experience noticeable slowdowns due to metadata logging and Neo4j updates.
2. A moderate increase in storage consumption, with PostgreSQL metadata tables growing at a predictable rate and Neo4j maintaining a manageable number of nodes and relationships.
3. Efficient lineage retrieval performance in Neo4j, where relationship-based queries outperform traditional relational approaches for complex dependency tracing.

These results help determine whether the proposed approach offers a balanced trade-off between lineage transparency and system performance, ensuring that lineage tracking remains viable for large-scale analytical environments without introducing prohibitive computational or storage costs.

4.5 End-to-End System Flow

The proposed system follows a structured and automated workflow that seamlessly integrates query execution, lineage tracking, and performance evaluation. By combining relational and graph-based components, it ensures efficient tracking of data provenance, transformations, and dependencies while maintaining system performance.

The process begins with query execution in PostgreSQL, where analytical queries run against the TPC-H dataset. In addition to producing results, queries automatically

trigger a metadata logging mechanism that records essential lineage details. This metadata—capturing table and column dependencies, transformations, and outputs—is systematically stored in PostgreSQL, ensuring a complete historical record for lineage analysis.

A Python-based synchronization pipeline continuously monitors PostgreSQL for newly logged queries. Upon detecting a new execution, it retrieves the relevant metadata and processes it into a structured format for graph-based representation in Neo4j. This transformation step involves mapping queries, tables, columns, and result cells into interconnected nodes and relationships, reflecting the logical flow of data transformations.

Once processed, the metadata is inserted into Neo4j, utilizing MERGE operations to prevent duplication and maintain consistency. These results in an interactive lineage graph, where users can efficiently trace the origin of data points, analyze the downstream impact of changes, and assess dependencies across multiple queries. Compared to traditional relational databases that require complex JOIN operations for lineage tracking, Neo4j enables faster, more intuitive lineage exploration through its graph-based model and Cypher query language.

To evaluate the system's efficiency, TPC-H benchmark tests are conducted under two execution modes: baseline queries, which run without lineage tracking, and tracked queries, where full lineage logging is enabled. This comparison provides insights into execution overhead, storage footprint, and retrieval efficiency, ensuring that lineage tracking remains scalable without introducing prohibitive performance costs. Additionally, Neo4j's lineage query performance is assessed to validate its efficiency in real-world analytical workloads.

By integrating relational metadata storage with graph-based lineage analysis, this architecture delivers an automated, scalable, and non-intrusive solution for query lineage tracking. The combination of PostgreSQL for structured metadata storage and Neo4j for lineage exploration provides a comprehensive, performance-conscious approach to managing data provenance in large-scale analytical environments.

5. Data Description

5.1 Overview of the Dataset

The dataset used in this study is derived from the TPC-H benchmark, a widely adopted dataset for evaluating query performance and decision-support systems. TPC-H is designed to simulate real-world business transactions by providing a structured relational database that contains sales, orders, customers, and supplier-related information. It consists of multiple interrelated tables that collectively represent an enterprise managing orders, inventory, and financial operations.

This dataset is well-suited for query lineage tracking since it provides predefined queries that operate on complex relationships, making it ideal for testing the efficiency and accuracy of lineage tracking mechanisms. Additionally, the structured nature of TPC-H enables lineage metadata extraction at different granularity levels, from table-level tracking to column and cell-level dependencies.

For this study, two dataset scales are used: a 1GB dataset, which serves as the primary dataset for lineage tracking implementation, and a 100MB dataset, which is used for evaluation and benchmarking. These two dataset sizes allow for a performance comparison between different query workloads and help assess how lineage tracking scales with increasing data volume.

5.2 Structure of the TPC-H Database

The TPC-H dataset consists of eight relational tables, each representing different aspects of business transactions within an enterprise. The dataset is designed to support complex analytical queries, making it a suitable benchmark for evaluating query execution and lineage tracking. The database schema follows a structured relational model, where tables are interconnected through primary and foreign key relationships, ensuring data consistency and facilitating efficient querying. The core tables in the dataset include:

Query Execution Flow for Lineage Tracking in PostgreSQL

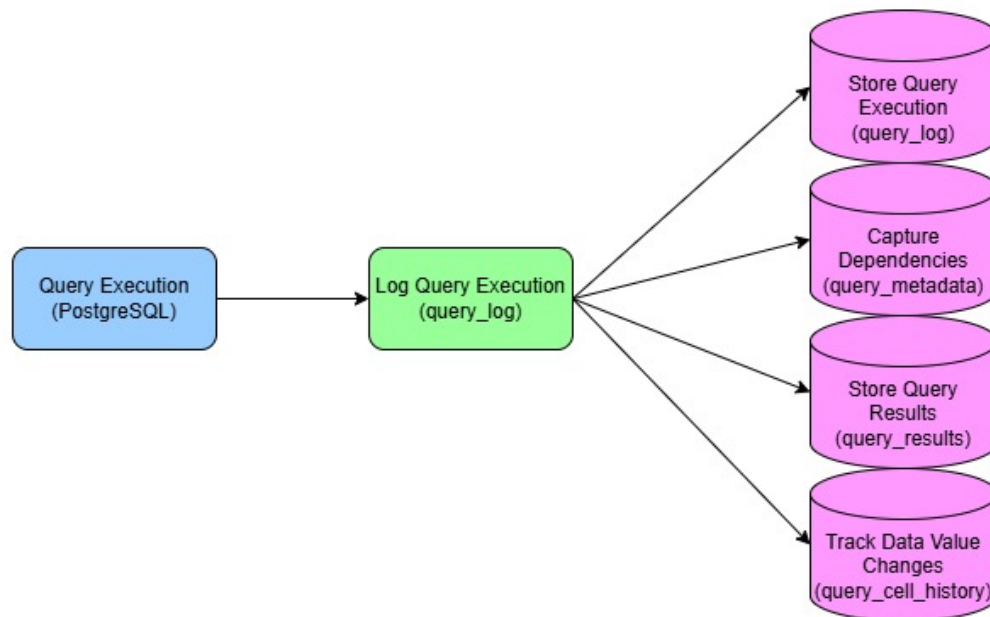


Figure 5.1: Query Execution Flow for Lineage Tracking in PostgreSQL

- **LINEITEM:** This table stores detailed information about individual items in customer orders, including attributes such as shipping details, pricing, discounts, and tax rates. Due to its granular nature, it is frequently used in analytical queries.
- **ORDERS:** This table contains records of customer orders, including order priority, order status, order date, and total price. It is linked to LINEITEM through a foreign key relationship.
- **CUSTOMER:** Holds customer-related data such as names, addresses, and associated market segments. It is commonly used for analyzing customer demographics and purchasing behavior.
- **SUPPLIER:** Represents suppliers providing products, containing details such as supplier names, locations, and contact information. It serves as a key component in supplier performance analysis.
- **PART:** Stores product-related information, including names, sizes, and manufacturing costs. This table is crucial for inventory management and pricing analysis.

- **PARTSUPP**: Acts as a bridge table between **PART** and **SUPPLIER**, linking products to their suppliers and storing details such as supply costs and stock availability.
- **NATION**: Defines geographical regions for customers and suppliers, supporting market segmentation and trade analysis.
- **REGION**: Represents broader geographic regions, grouping multiple nations into categories useful for macroeconomic and business trend evaluations.

Each table is interconnected through foreign key constraints, ensuring that queries can efficiently retrieve related data while maintaining referential integrity. The **LINEITEM** and **ORDERS** tables are particularly significant for query lineage tracking, as they are central to many analytical queries that drive business intelligence processes.

The structured nature of the TPC-H database facilitates tracking at different levels of granularity. Queries executed on this dataset generate interactions between tables that can be captured in metadata logs, making it possible to analyze dependencies between tables, columns, and even individual data values. This structured design enables efficient lineage tracking and allows for precise analysis of how data flows across multiple queries.

5.3 Metadata Storage for Query Lineage Tracking

To enable query lineage tracking, a dedicated set of metadata tables is created within PostgreSQL. These tables systematically store information about executed queries, the database elements they interact with, and how query results are derived. By maintaining structured metadata, the system ensures that each executed query is logged with sufficient detail to support comprehensive lineage analysis. The key metadata tables include `query_log`, `query_metadata`, `query_results`, and `query_cell_history`, each serving a distinct role in tracking the lineage of data transformations.

5.3.1 Query Execution Log (`query_log`)

The `query_log` table serves as the core repository for recording query executions. Each entry in this table represents an individual query execution, storing essential metadata that allows for reconstructing execution history and analyzing query behavior over time.

The table consists of the following key attributes:

- `log_id` (Primary Key) – A unique identifier assigned to each executed query.
- `query_name` – The name of the executed query function, allowing queries to be categorized and identified.
- `execution_time` – A timestamp indicating when the query was executed, which aids in temporal lineage tracking.
- `parameters` – A JSON object that stores the input parameters supplied to the query at execution time, allowing for reproducibility of query conditions.
- `result_count` – The number of rows returned by the query, providing an overview of the volume of retrieved data.

By storing these execution details, the `query_log` table forms the foundation of the lineage tracking system. It maintains a historical record of all executed queries, enabling users to trace when a query was run, with what parameters, and how many results it produced. This table is critical for monitoring query usage patterns, auditing database activity, and diagnosing performance issues.

5.3.2 Query Metadata (query_metadata)

The `query_metadata` table extends lineage tracking by recording dependencies between queries and the underlying database elements they interact with. This table provides structured insights into how different queries utilize tables and columns, making it possible to trace how data is transformed across multiple query executions.

The table includes the following attributes:

- `metadata_id` (Primary Key) – A unique identifier for each metadata entry.
- `log_id` (Foreign Key to `query_log`) – A reference to the specific query execution that generated this metadata.
- `table_name` – The name of the database table accessed by the query.
- `column_name` – The specific column in the table that was used in the query execution.
- `dependency_type` – The nature of the relationship between the query and the referenced table-column pair. The dependency types include:
 - `source` – The column provides input data for the query.
 - `derived` – The column contains values that are computed or transformed within the query.
 - `computed` – The column is used in aggregations or arithmetic calculations.
 - `foreign_key` – The column serves as a foreign key that links different tables.

By mapping out these dependencies, the `query_metadata` table enables users to conduct impact analysis, where they can assess how changes to a specific column or table might affect downstream queries. This level of tracking is particularly valuable for understanding how data flows through analytical workflows and ensuring schema modifications do not unintentionally disrupt critical queries.

5.3.3 Query Results (query_results)

The `query_results` table stores the actual outputs generated by executed queries, providing a way to trace how specific results were derived. This table facilitates result-driven lineage analysis, allowing users to track the transformation of input data into final outputs.

The table contains the following attributes:

- `result_id` (Primary Key) – A unique identifier assigned to each stored query result.
- `log_id` (Foreign Key to `query_log`) – A reference linking the result to its corresponding query execution.
- `result_data` – The full query output stored in JSONB format, preserving structured data for lineage tracking and future analysis.

By storing query results in a structured format, this table ensures that users can investigate how specific outputs were generated. The JSONB format allows for efficient storage and retrieval of structured query results, supporting advanced lineage queries where analysts may need to trace the origins of a particular result value.

5.3.4 Query Cell History (query_cell_history)

For fine-grained lineage tracking, the `query_cell_history` table records individual data values that were referenced or modified during a query execution. This table allows for cell-level lineage analysis, enabling precise tracking of data transformations down to individual values within a table.

The table includes the following attributes:

- `log_id` (Foreign Key to `query_log`) – Links each entry to a specific query execution, establishing a direct lineage connection.
- `query_name` – The name of the executed query function, allowing for function-level lineage analysis.
- `parameters` – The input parameters used in the query execution, stored in JSONB format for efficient reference.
- `table_name` – The name of the table that contains the tracked data value.
- `column_name` – The specific column associated with the recorded value.
- `cell_value` – The actual data value that was accessed or modified during the query execution.

This level of detail allows users to trace transformations at the data level, answering critical questions such as:

- Which queries have accessed or modified a specific value?
- How did an individual data point evolve over time across multiple queries?
- What was the original source of a derived value?

By maintaining cell-level tracking, the `query_cell_history` table plays a crucial role in auditing, debugging, and ensuring data integrity. It provides granular insights into how queries interact with individual values, which is essential for industries requiring detailed data traceability, such as finance, healthcare, and regulatory compliance.

5.4 Relationships Between Metadata Tables

The metadata tables in PostgreSQL are designed with strict referential integrity to ensure accurate query lineage tracking. These tables establish structured relationships between query executions, data dependencies, and results, enabling seamless tracing of how queries interact with different elements of the database. The relationships between metadata tables facilitate multi-level lineage tracking, from high-level query execution records to fine-grained cell-level transformations.

At the core of the lineage tracking system is the `query_log` table, which serves as a central repository for all executed queries. Each query recorded in `query_log` has associated metadata in other tables, creating structured relationships that define query dependencies.

One of the key relationships exists between `query_log` and `query_metadata`. Each query execution may involve multiple database tables and columns, and this information is captured in `query_metadata`. By linking query executions to the specific tables and columns they access, this relationship allows for structured tracking of how queries interact with database elements. This is particularly useful for impact analysis, where changes in table structures or column values need to be traced across queries.

Another critical relationship connects `query_log` and `query_results`. Every query execution generates an output, which is stored in `query_results` and linked back to the corresponding entry in `query_log`. This connection allows for result-driven lineage analysis, where users can trace the derivation of specific query results. Since results are stored in JSONB format, structured data can be efficiently retrieved for auditing and verification purposes.

At a more granular level, `query_cell_history` is linked to `query_log` to provide detailed lineage tracking at the individual data value level. While `query_metadata` tracks dependencies at the table and column level, `query_cell_history` extends this by capturing the specific data points referenced or modified during a query execution. This enables precise tracking of data transformations, ensuring that individual values can be traced back to their original sources. This relationship is particularly important for debugging and compliance scenarios where precise data lineage is required.

Together, these relationships form a hierarchical model for lineage tracking:

- At the highest level, `query_log` maintains a historical record of executed queries.
- At the intermediate level, `query_metadata` captures table and column dependencies, while `query_results` links queries to their outputs.
- At the most detailed level, `query_cell_history` tracks data at the individual value level.

This structured approach ensures that lineage tracking is both scalable and flexible, allowing analysts to trace query dependencies at different levels of granularity based on their analytical needs. By leveraging these relationships, the system provides a comprehensive lineage tracking framework that supports auditing, debugging, and performance optimization.

5.5 Integration with Query Lineage Tracking

To seamlessly integrate query lineage tracking into the execution process, predefined TPC-H queries are modified into PL/pgSQL functions that log metadata during execution. This ensures that each query execution is automatically recorded in the lineage tracking system without requiring external logging mechanisms. By embedding metadata capture directly into the execution workflow, the system provides a real-time, automated approach to tracking query dependencies, transformations, and outputs.

The integration follows a structured approach, ensuring that each executed query generates a complete set of metadata entries across multiple tables. When a query function is executed in PostgreSQL, it triggers a sequence of lineage tracking events that capture essential details at different levels of granularity:

1. **Recording Query Execution in `query_log`.** Every time a query function runs, it is logged in the `query_log` table. This entry includes the query name, execution timestamp, input parameters, and the number of rows returned. This structured execution history provides a chronological record of all queries, enabling users to reconstruct past executions, analyze query usage patterns, and audit database activities.
2. **Capturing Table and Column Dependencies in `query_metadata`.** The system then logs detailed table and column dependencies in the `query_metadata` table. Each query execution is analyzed to determine which tables and columns were accessed or modified. These dependencies are classified into source, derived, computed, and foreign_key relationships, allowing users to track how data is

transformed across queries. By explicitly documenting these relationships, the system enables impact analysis and schema change tracking.

3. Logging Query Results in `query_results`. To support result-driven lineage analysis, the outputs of each executed query are stored in `query_results`. The result data is logged in JSONB format, allowing for efficient retrieval and structured storage. This linkage between query executions and their generated results ensures that users can trace how a particular output was derived and validate the correctness of query processing.

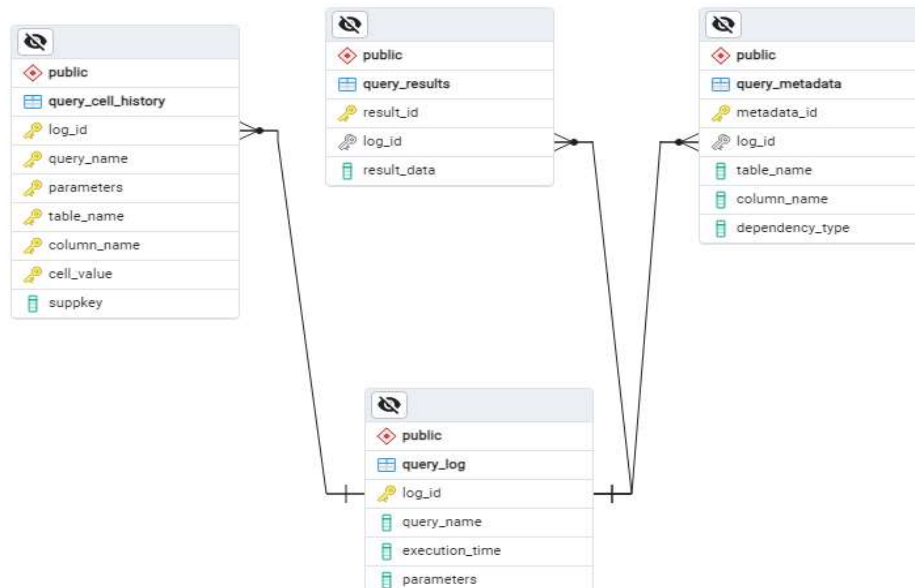


Figure 5.2: Metadata Storage Schema for Query Lineage Tracking

4. Tracking Fine-Grained Data Lineage in `query_cell_history`. In addition to table and column dependencies, the system records individual data values accessed or modified during query execution. These values are logged in `query_cell_history`, capturing the specific table, column, and cell values involved. This level of detail allows users to:
 - Identify which queries referenced a specific value in past executions.
 - Trace how individual data points were transformed through multiple queries.
 - Perform granular impact analysis, ensuring that modifications to a dataset can be traced with precision.

By incorporating lineage tracking directly into query execution, this integration ensures that all relevant metadata is captured automatically and in real time. The use of PL/pgSQL functions eliminates the need for external logging tools while maintaining high efficiency and minimal performance overhead.

This structured logging mechanism enhances data provenance, debugging capabilities, and compliance auditing, ensuring that users can analyze query dependencies, validate results, and trace data transformations at multiple levels of granularity. By maintaining a fully automated tracking system, this approach provides a scalable, robust solution for query lineage analysis within analytical environments.

6. Implementation

6.1 System Setup and Environment

The implementation of the query lineage tracking system is built on PostgreSQL and Neo4j, with a Python-based synchronization pipeline ensuring seamless data consistency between them. PostgreSQL serves as the primary query execution engine, handling analytical workloads while simultaneously logging metadata related to query execution. Neo4j, on the other hand, provides an efficient graph-based model for storing, retrieving, and analyzing query lineage. The system is designed to capture metadata at various levels of granularity, ranging from query execution records to detailed data dependencies, ensuring comprehensive lineage tracking.

The development environment consists of PostgreSQL 16.2, Neo4j 5.24.0 (Enterprise Edition), and Python 3.13.1, running on a local Windows 10 machine. The implementation leverages the `psycopg2` library for PostgreSQL connectivity and the official `neo4j` Python driver for executing queries in the graph database. PostgreSQL's built-in JSONB support is utilized to store query results in a structured format, enabling efficient retrieval and analysis. The system operates without modifications to the default configurations of PostgreSQL and Neo4j. No additional indexing, memory tuning, or query optimizations were introduced, as the default query planner in PostgreSQL and Cypher optimizations in Neo4j were deemed sufficient for ensuring efficient execution.

6.2 Query Execution and Metadata Logging

The query lineage tracking mechanism is embedded directly within the execution process by modifying predefined TPC-H queries into PL/pgSQL functions that handle metadata logging. This ensures that every executed query automatically generates a corresponding record in the lineage tracking system without requiring external monitoring mechanisms. The metadata logging process is seamlessly integrated into the query execution workflow, capturing essential details regarding each execution.

When a query function is executed, it stores its metadata in the `query_log` table, which acts as the primary execution log. Each entry records critical execution details, including the query name, the execution timestamp, the query parameters, and the number of result rows retrieved. These details help reconstruct query execution history and facilitate impact analysis. To further extend lineage tracking, table and column dependencies are captured in the `query_metadata` table, where each query logs the database objects it interacts with. These dependencies are classified into four types: source, representing raw input data used in the query; derived, representing computed columns resulting from transformations; computed, referring to columns involved in aggregations or arithmetic operations; and foreign key, capturing relational links between tables.

Query outputs are preserved in the `query_results` table, stored in JSONB format, allowing structured storage and retrieval of query results for result-driven lineage analysis. This enables users to trace the transformations applied to input data to derive the final output. Additionally, fine-grained lineage tracking is implemented using the `query_cell_history` table, which records individual data values accessed or modified during query execution. By tracking dependencies at multiple levels, this structured metadata storage provides a multi-layered lineage model, ensuring traceability and transparency in query execution.

6.3 Python Synchronization Pipeline

To ensure that the metadata stored in PostgreSQL is accurately reflected in Neo4j, a Python-based synchronization pipeline is implemented. This pipeline continuously monitors newly executed queries in PostgreSQL and transfers the relevant metadata to Neo4j, where it is restructured into a graph-based representation. The synchronization process follows an incremental approach, ensuring that only newly executed queries are processed, preventing redundant updates and minimizing computational overhead.

The synchronization process begins by retrieving unprocessed logs from the `query_log` table in PostgreSQL. Once a new execution is detected, the pipeline extracts execution details, such as query parameters, execution time, and result sets. Using this metadata, it constructs relationships between query executions, stored results, and individual data values. The extracted metadata is then inserted into Neo4j using MERGE operations, ensuring that redundant entries are avoided and referential integrity is maintained. The graph-based lineage model in Neo4j represents executed queries as `QueryLog` nodes, query outputs as `QueryResultEntry` nodes, and individual data values as `Cell` nodes. Relationships such as `HAS_RESULT`, `HAS_CELL`, and `CONTRIBUTED_FROM` are used to establish query dependencies.

Graph Schema for Query Lineage in Neo4j

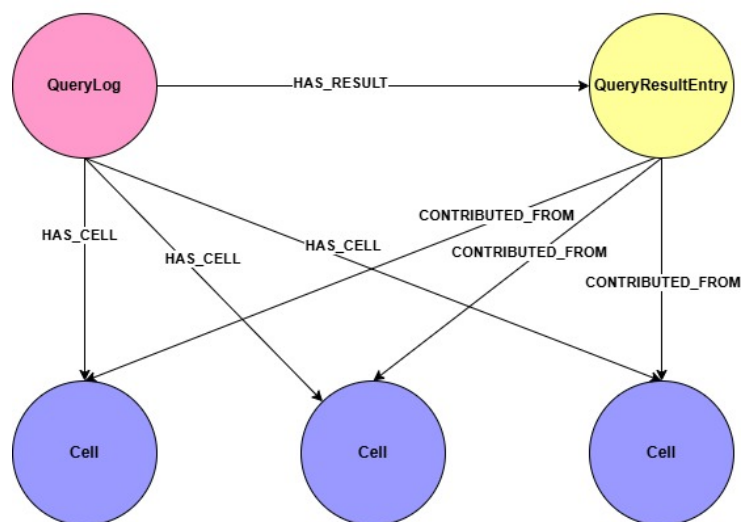


Figure 6.1: Graph Schema for Query Lineage in Neo4j

This automated synchronization process ensures that Neo4j remains up to date with the latest lineage information, allowing for real-time data dependency tracking. By maintaining a consistent lineage graph, the system facilitates interactive exploration of query dependencies, enabling users to trace how query results were derived, analyze changes in data, and evaluate the impact of modifications.

6.4 Query Performance with Lineage Tracking

To assess the impact of lineage tracking on query execution, a benchmarking experiment was conducted by comparing query execution times with and without metadata logging. A complex analytical query, designed to simulate a real-world workload, was selected for evaluation. The results demonstrate that lineage tracking introduces only a minimal increase in execution time, primarily due to the additional metadata insertions in PostgreSQL.

Without lineage tracking enabled, the query execution time remains within the baseline performance. When lineage tracking is activated, the execution time increases slightly, reflecting the additional computational overhead of logging metadata. However, the impact remains within an acceptable range, ensuring that the system remains efficient even under lineage-tracked workloads. The use of PL/pgSQL functions for metadata logging helps minimize this overhead by integrating lineage capture directly into the query execution process.

To evaluate storage overhead, database size measurements were recorded before and after multiple query executions. The total database size before lineage tracking was 1478 MB, and after multiple executions with metadata logging, it increased slightly to 1479 MB. This minor increase in storage consumption indicates that metadata storage overhead is well-managed, allowing the system to scale efficiently without excessive resource requirements.

In Neo4j, the lineage graph expands dynamically as queries are executed and tracked. The resulting dataset consists of 2692 nodes and 2872 relationships, with the majority of nodes representing query results and data values. Despite the expansion, Neo4j's query retrieval performance remains stable, as graph traversal algorithms optimize lineage queries. The most common operations—such as tracing data provenance and identifying impacted queries—execute in sub-second response times, confirming that graph-based lineage representation does not impose significant performance bottlenecks.

6.5 Summary of Implementation

The implementation of the query lineage tracking system successfully integrates PostgreSQL and Neo4j, providing a robust solution for structured metadata storage and real-time lineage visualization. The metadata logging framework ensures that query executions, dependencies, and results are automatically captured, while the Python synchronization pipeline keeps Neo4j updated with the latest lineage information.

Through performance evaluation, it is confirmed that lineage tracking introduces minimal execution overhead, ensuring that analytical queries remain efficient. The storage footprint remains manageable, demonstrating that the system scales effectively without excessive resource consumption. The integration of relational metadata storage in PostgreSQL and graph-based lineage analysis in Neo4j allows for an automated, scalable, and interactive approach to query lineage tracking.

This structured framework supports real-time monitoring of query dependencies, efficient debugging, and historical auditing of data transformations. The combination of relational and graph-based components ensures that query lineage tracking remains accurate, scalable, and suitable for large-scale analytical environments. By providing detailed insights into query execution, efficient storage management, and interactive lineage analysis, the system lays a solid foundation for data governance, impact analysis, and regulatory compliance in modern data-driven applications.

7. Evaluation & Benchmarks

7.1 Overview

Evaluating the performance and efficiency of the query lineage tracking system is crucial to understanding its impact on query execution and metadata storage. This chapter presents a comprehensive analysis of execution time overhead, storage requirements, and system resource usage, providing insights into how lineage tracking scales with increasing data volume. The evaluation is conducted by benchmarking

query execution with and without lineage tracking, analyzing storage consumption in both PostgreSQL and Neo4j, and assessing system resource utilization during query execution.

The experiments aim to quantify the additional computational and storage costs introduced by lineage tracking while ensuring that query performance remains within an acceptable range. By comparing different configurations, the feasibility of this approach for large-scale analytical environments is validated.

7.2 Query Execution Performance

To assess the impact of lineage tracking on query execution time, PostgreSQL's built-in timing functions were used to measure the execution duration of a predefined analytical query from the TPC-H benchmark. The function `tpch_query_aggregation` was selected for benchmarking due to its complexity, as it involves multiple tables, aggregations, and calculations—representing a realistic workload in analytical database environments. The total execution time included multiple phases: the query execution in PostgreSQL, metadata logging to capture dependencies, the extraction and transformation of metadata in Python, and the subsequent insertion of lineage data into Neo4j. To evaluate the performance overhead introduced by lineage tracking, the query was executed five times under two conditions: a baseline execution without lineage tracking, serving as a control, and a tracked execution, where the same query was executed with full lineage tracking enabled, including logging at multiple levels of granularity.

The benchmarking results indicated that lineage tracking introduces a measurable increase in execution time, with a 4.8× overhead for the 1GB dataset and a 13.5× overhead for the 100MB dataset. The disproportionate impact on the smaller dataset highlights the fixed cost of metadata logging, which remains relatively constant regardless of dataset size. Since the operations involved in recording query metadata—such as inserting records into `query_log`, `query_metadata`, `query_results`, and `query_cell_history`—occur regardless of how much data is processed, they represent a larger proportion of total execution time in small queries. Additionally, fine-grained lineage tracking amplifies this effect, as cell-level tracking requires logging individual data values that were referenced or modified, significantly increasing the number of metadata insertions. The Python-based synchronization pipeline also introduces a fixed delay in transferring metadata from PostgreSQL to Neo4j, which has a greater proportional impact on shorter queries. In larger datasets, this delay is amortized over a longer execution period, making the relative overhead less pronounced. Furthermore, the insertion of lineage metadata into Neo4j scales with the complexity of the query rather than the volume of data, which further explains the greater impact on small queries.

When compared to alternative lineage tracking approaches, the proposed graph-based method offers advantages in scalability despite the initial overhead. Traditional SQL-based lineage tracking, which relies on recursive joins and self-referencing queries, tends to increase computational complexity as data scales, making it impractical for large analytical workloads. In contrast, event-based lineage tracking systems, such as Apache Atlas or OpenLineage, capture lineage metadata passively without modifying query execution paths, reducing the direct overhead but often sacrificing fine-grained tracking capabilities. The Neo4j-based approach used in this study allows for efficient lineage queries and dependency analysis, which, while initially incurring additional logging costs, provides long-term benefits in terms of rapid traversal of data relationships and efficient impact analysis.




Figure 7.1: Query Execution Time Comparison (With and Without Lineage Tracking)

The results of this evaluation suggest that while fine-grained lineage tracking incurs execution overhead, particularly in small datasets where fixed metadata costs are more pronounced, the impact remains within practical limits for large analytical workloads. Despite the overhead, the system provides crucial benefits in data transparency, debugging, and compliance. This makes the trade-off worthwhile for enterprise-scale analytical environments, where accurate query tracking is critical. Future optimizations, such as batch processing for metadata logging and asynchronous synchronization with Neo4j, could further mitigate the impact on query performance while preserving the advantages of real-time lineage tracking. Overall, the findings demonstrate that integrating relational and graph-based lineage tracking is a viable approach for large-scale data analytics, balancing the need for transparency and traceability with computational efficiency.

7.3 Storage Overhead Analysis

To assess the storage impact of lineage tracking, the database size was recorded before and after executing queries. This analysis considers both PostgreSQL and Neo4j storage requirements.

7.3.1 PostgreSQL Storage Analysis

The size of PostgreSQL metadata tables was measured before and after query executions for both dataset scales.

For the tpch dataset (1GB):

- query_log: 32 kB
- query_metadata: 32 kB
- query_results: 32 kB
- query_cell_history: 688 kB
- Total Database Size: Increased from 1478 MB to 1479 MB after multiple executions.

For the tpch_small dataset (100MB):

- query_log: 32 kB
- query_metadata: 16 kB

- query_results: 16 kB
- query_cell_history: 66 MB

The query_cell_history table accounted for the largest metadata growth, particularly in smaller datasets, where cell-level lineage tracking significantly increased storage needs. This suggests that fine-grained lineage tracking—particularly at the cell level—substantially contributes to metadata storage. Despite this, the overall PostgreSQL storage overhead remains manageable, demonstrating that the system can scale efficiently even with detailed lineage tracking.

7.3.2 Neo4j Storage and Graph Growth

In Neo4j, the lineage graph expands as queries execute and metadata is transferred. The number of nodes and relationships in the lineage graph was measured for both dataset sizes.

For the tpch dataset (1GB):

- Nodes: 2,441 total
 - 1 QueryLog node
 - 2,436 Cell nodes
 - 4 QueryResultEntry nodes
- Relationships: 2,440 total
 - 2,436 HAS_CELL
 - 4 HAS_RESULT
 - 0 CONTRIBUTED_FROM
- Storage Impact: The total size of these nodes and relationships in Neo4j is approximately 0.34 MB, confirming that despite the detailed metadata storage, the overall impact on database size remains minimal.

For the tpch_small dataset (100MB):

- Nodes: 234 total
 - 1 QueryLog node
 - 230 Cell nodes
 - 3 QueryResultEntry nodes
- Relationships: 233
 - 230 HAS_CELL
 - 3HAS_RESULT
 - 0 CONTRIBUTED_FROM

The lineage tracking metadata increases significantly with dataset size, particularly due to fine-grained cell-level tracking, with the most significant metadata contribution coming from Cell nodes. However, despite the increase in metadata volume, the actual storage overhead in Neo4j remains relatively low. The 0.34 MB footprint for the 1GB dataset demonstrates that the system can effectively store and manage lineage data without excessive resource consumption.

Moreover, the impact of lineage metadata on Neo4j performance remains controlled. The storage growth, while linear with the number of recorded queries, does not

introduce significant bottlenecks. The efficient representation of relationships ensures that query execution for lineage retrieval remains performant even as the dataset scales. This further confirms the feasibility of integrating Neo4j for large-scale analytical lineage tracking.

7.3.3 Threshold-Based Lineage Granularity

To address the challenges posed by fine-grained lineage tracking, a threshold-based lineage granularity mechanism has been implemented within the Python synchronization pipeline. This feature dynamically adjusts the level of lineage detail based on query complexity to optimize performance and storage usage. When a query is executed, the system typically records every individual cell that contributes to the query result, generating a detailed lineage representation. However, for large queries that reference or modify an overwhelming number of data points, this level of granularity can lead to excessive metadata growth, making it difficult to analyze meaningful dependencies.

To mitigate this issue, the system introduces an adaptive threshold that determines when to switch from fine-grained lineage tracking to a more aggregated approach. If the number of generated cell nodes surpasses a predefined threshold, the system automatically shifts to a higher-level lineage representation by logging only table- and column-level dependencies instead of individual data values. This ensures that lineage tracking remains informative without becoming overly complex or unmanageable.

The primary advantage of this approach is improved performance, as reducing cell-level tracking for large queries minimizes the execution overhead associated with metadata logging. By restricting the number of nodes created in Neo4j, the system also reduces storage costs and prevents excessive database growth, making lineage tracking more scalable for high-volume analytical workloads. Additionally, in cases where queries generate millions of contributing data points, tracking every individual cell becomes impractical, as the number of dependencies may be too large to effectively analyze. The threshold-based mechanism simplifies query lineage representation, allowing users to trace dependencies at a more meaningful level without being overwhelmed by excessive detail.

However, this approach also presents trade-offs. By switching to table- and column-level tracking, the system sacrifices fine-grained detail, meaning that individual cell dependencies are no longer available when the threshold is exceeded. This could be a limitation for debugging specific data transformations, particularly in applications where understanding precise data modifications is crucial. Additionally, determining an optimal threshold requires careful tuning, as different queries may have varying complexity levels. Setting the threshold too low could result in premature loss of detailed tracking, while a threshold that is too high might still allow excessive metadata accumulation.

The implementation of this threshold-based approach provides a balance between maintaining detailed lineage tracking and ensuring system scalability. It is particularly beneficial for large-scale analytical environments where excessive lineage granularity does not necessarily provide additional value. By allowing users to configure the threshold based on specific analytical needs, the system offers a flexible and adaptive approach to lineage tracking that can be tailored to different use cases.

7.4 System Resource Usage

To evaluate system performance overhead, CPU and memory usage were monitored during query execution with lineage tracking enabled. The system maintained a stable

processing load, with CPU usage averaging 34% across PostgreSQL, Neo4j, and synchronization tasks. While metadata transfers increased processing demand, memory consumption remained stable due to PostgreSQL's efficient logging mechanisms. The additional CPU consumption was primarily attributed to the continuous monitoring and metadata transfer operations performed by the synchronization script.

Regarding memory usage, no significant increase was observed in PostgreSQL, as the metadata inserts were lightweight and did not impose substantial memory demands. Similarly, Neo4j maintained stable memory consumption, leveraging its internal indexing mechanisms to optimize lineage queries. Despite the expansion of the lineage graph, the system efficiently handled node creation and relationship updates without excessive memory allocation.

However, some potential bottlenecks were identified. The rapid expansion of `query_cell_history` poses a long-term storage challenge, suggesting the need for pruning strategies in large-scale deployments. As cell-level tracking records individual data transformations, excessive metadata accumulation could lead to storage inefficiencies over time. Additionally, the speed of node creation in Neo4j depended on the number of lineage entries per query. While graph traversal queries remained efficient, the increase in node insertions could impact performance in environments with heavy query loads. These observations emphasize the need for optimized lineage tracking strategies to ensure long-term scalability.

7.5 Scalability Considerations

The evaluation results indicate key factors that must be addressed to optimize the scalability of the lineage tracking system. Optimizing lineage tracking requires balancing granularity and efficiency. One approach is selective logging, where only critical columns are tracked instead of recording all referenced values. This approach would significantly reduce the storage footprint and minimize query execution overhead without compromising the essential traceability of data transformations.

Another important consideration is the efficient pruning of metadata to manage storage growth over extended periods. Implementing periodic cleanup mechanisms for outdated lineage records would prevent unnecessary accumulation of historical metadata while ensuring that relevant query lineage information remains accessible for auditing and impact analysis.

Additionally, indexing and query optimization strategies in both PostgreSQL and Neo4j could further enhance system performance. While the current implementation relies on default indexing mechanisms, introducing additional indexes on frequently accessed metadata fields could improve query retrieval speeds and reduce computational overhead. Optimizing Cypher queries in Neo4j to leverage graph traversal efficiencies would also contribute to maintaining high query performance even as the dataset scales.

By incorporating these improvements, the system can maintain accurate lineage tracking while reducing execution time overhead and ensuring sustainable resource utilization.

7.6 Summary of Evaluation

The evaluation confirms that query lineage tracking introduces measurable execution time overhead, primarily due to metadata insertions and synchronization between PostgreSQL and Neo4j. However, the performance trade-off is justified by the

substantial benefits of enhanced traceability, auditing capabilities, and the ability to conduct advanced data provenance analysis.

The system incurs additional execution time but delivers strong gains in transparency, compliance, and debugging, making it suitable for large-scale data analytics. The primary source of execution overhead is the structured storage of metadata and the subsequent synchronization of lineage information to Neo4j. The impact is particularly noticeable in fine-grained tracking scenarios, where a significant number of cell-level lineage entries are generated. However, the introduction of the threshold-based lineage granularity mechanism provides a way to dynamically adjust the level of detail, ensuring that excessive metadata does not negatively impact performance or storage efficiency.

Storage utilization analysis highlights that the `query_cell_history` table is the largest contributor to metadata growth in PostgreSQL. While the storage overhead remains manageable for moderate workloads, large-scale deployments may require strategies to optimize lineage storage. In contrast, Neo4j scales efficiently, as graph-based lineage queries execute in sub-second response times, leveraging optimized relationship traversal mechanisms. The total size of nodes and relationships in Neo4j for the 1GB dataset is approximately 0.34 MB, confirming that graph-based storage remains lightweight even with detailed dependency tracking. Furthermore, the introduction of the threshold mechanism ensures that the number of nodes and relationships in Neo4j remains controlled, as switching to table- and column-level tracking significantly reduces the number of `HAS_CELL` relationships without compromising essential lineage information.

System resource monitoring reveals that CPU usage stabilizes at approximately 34%, reflecting the additional processing required for lineage metadata logging and synchronization. Memory utilization remains stable, with PostgreSQL and Neo4j efficiently handling metadata operations without excessive memory consumption. These findings confirm that the system maintains a balance between computational efficiency and detailed lineage tracking, ensuring its viability for large-scale analytical environments.

Overall, the proposed query lineage tracking system effectively integrates structured metadata storage in PostgreSQL with an interactive lineage graph in Neo4j. The system offers a scalable, automated, and performance-conscious approach to tracking data provenance, enabling users to analyze query dependencies, trace data transformations, and ensure regulatory compliance. The combination of fine-grained and adaptive lineage tracking mechanisms ensures that performance remains within acceptable limits while maintaining a high level of detail where necessary.

Future improvements will refine granularity control, metadata pruning, and indexing optimizations, ensuring that lineage tracking remains efficient as dataset complexity grows. The results confirm that Neo4j's storage efficiency and retrieval performance make it a viable long-term solution for scalable query lineage tracking, offering a balanced trade-off between granularity and system overhead.

8. Conclusions & Future Work

8.1 Conclusions

Query lineage tracking is a fundamental component of modern data analytics, allowing organizations to trace the origin, transformations, and dependencies of data. This dissertation proposed a graph-based lineage tracking approach that integrates

PostgreSQL for query execution and Neo4j for lineage analysis, delivering an efficient and scalable solution.

This work makes three key contributions: (1) automation, by enabling PostgreSQL to capture query lineage with minimal manual effort; (2) scalability, by leveraging a graph-based lineage representation in Neo4j that allows for efficient dependency traversal; and (3) integration, through a Python-based synchronization pipeline that bridges relational and graph-based storage systems in real time. Performance evaluation using the TPC-H benchmark confirms that lineage tracking remains viable in large-scale analytical environments, with execution overhead ranging from 4.8× for large datasets to 13.5× for smaller datasets. While metadata storage growth is an inherent trade-off, the benefits of enhanced query traceability and debugging outweigh these costs. Furthermore, the use of Neo4j significantly outperforms traditional relational JOIN-based approaches for dependency retrieval, reinforcing the effectiveness of graph-based models in complex data environments requiring real-time impact analysis.

8.2 Limitations

While effective, the system introduces several trade-offs that impact performance and scalability. Fine-grained lineage tracking—especially at the cell level—adds execution overhead, as metadata must be logged for every referenced or modified value. This results in an increased query execution time, particularly for analytical workloads that involve frequent transformations.

Another key limitation is the rapid growth of the `query_cell_history` table, which significantly expands with each query execution, posing a long-term scalability concern for metadata storage. Additionally, synchronization between PostgreSQL and Neo4j introduces periodic latency, as metadata transfers occur in batches rather than instantaneously. This can lead to slight delays in real-time lineage retrieval, especially in high-frequency workloads.

Furthermore, the current implementation does not leverage advanced indexing or query optimizations in either PostgreSQL or Neo4j. While the default query planners provide acceptable performance, implementing indexing strategies and query optimizations could improve retrieval speeds and reduce storage costs.

8.3 Future Work

Addressing these limitations will improve system efficiency and scalability. Future research should focus on optimizing execution time, reducing storage overhead, and enhancing multi-platform integration. The key areas for improvement include:

8.3.1 Performance Optimization

One area for improvement is refining the threshold mechanism introduced in Section 7.3.3, which dynamically switches between cell-level and table/column-level tracking based on query complexity. While this feature significantly reduces execution overhead, an important consideration is how it affects Neo4j's structure. Specifically, reducing cell-level tracking decreases the number of nodes and edges in Neo4j, which improves storage efficiency but may limit certain lineage queries that rely on fine-grained details. Future work could investigate strategies to balance lineage granularity while preserving query effectiveness, potentially through adjustable or hybrid tracking approaches

8.3.2 Scalable Storage Management

To mitigate storage concerns, intelligent metadata pruning can automate the removal of outdated lineage records, reducing storage overhead while preserving essential query dependencies. Additionally, implementing hybrid lineage tracking, where critical columns retain fine-grained metadata while less significant attributes use coarse-grained tracking, would balance traceability and storage efficiency.

8.3.3 Integration with Machine Learning

Another promising direction is the use of machine learning for threshold tuning. The current threshold-based system relies on a fixed cutoff, which may not be optimal for all queries. Instead, an adaptive model could learn from query patterns, predict query complexity, and automatically determine the appropriate level of lineage granularity. Heuristic-based or ML-driven approaches could analyze metadata volume, dependency complexity, and query execution time to dynamically adjust the threshold—ensuring the most efficient tracking configuration. This would prevent excessive metadata growth while still maintaining fine-grained lineage where needed.

8.3.4 Expanding to Multi-Platform Environments

Extending lineage tracking to cloud-based databases like AWS Redshift and Google BigQuery would make the system more adaptable to modern, distributed analytics environments. Additionally, enhancing compatibility with Apache Spark and Snowflake would expand lineage tracking to distributed data pipelines, where large-scale transformations require real-time dependency analysis.

8.3.5 Enhancing Regulatory Compliance Features

Automating regulatory reporting would enable enterprises to generate compliance-ready lineage documentation, streamlining audits for GDPR, CCPA, and HIPAA requirements. Implementing user access tracking would allow organizations to monitor interactions with specific data elements, improving security and accountability in lineage tracking.

8.4 Final Remarks

With data lineage tracking becoming essential for governance, security, and compliance, this research lays a foundation for future advancements. By combining relational and graph-based approaches, it offers a scalable model that supports transparency and performance—key requirements for modern data-driven organizations. Future enhancements, particularly in performance optimization, storage management, and integration with machine learning, will further improve the system's applicability in enterprise-scale environments. This research provides a strong basis for advancing lineage tracking methodologies in modern data ecosystems, ensuring that organizations can efficiently manage and analyze their data dependencies in increasingly complex environments.

REFERENCES

- [1] S. Robinson, "Graph Data Lineage for Financial Services: Avoiding Disaster," *Graphable Blog*, Sep. 2021. [Online]. Available: <https://www.graphable.ai/blog/graph-data-lineage-graph-database/>
- [2] "Data Lineage: Using Knowledge Graphs for Deeper Insights," *Neo4j Blog*, Aug. 2022. [Online]. Available: <https://neo4j.com/blog/knowledge-graph/data-lineage-using-knowledge-graphs-deeper-insights-data-pipelines/>..
- [3] W. Trappers, "Data Lineage Using a Graph Database," *Medium*, Aug. 2022. [Online]. Available: <https://medium.com/plumbersofdatascience/data-lineage-using-a-graph-database-8cbf0497d5a6..>
- [4] "IBM Manta - Graph Database & Analytics," *Neo4j Customer Stories*, [Online]. Available: <https://neo4j.com/customer-stories/ibm-manta-data-lineage/>..
- [5] A. Amin, "Enabling Data Lineage Using a Graph Database," *Medium*, Sep. 2021. [Online]. Available: <https://medium.com/if-tech/enabling-data-lineage-using-a-graph-database-fd40be4bf768..>
- [6] E. Roberts, "Benefits of Using a Graph Database to Map Data Lineage," *The Left Join*, Aug. 2024. [Online]. Available: <https://theleftjoin.com/benefits-of-using-a-graph-database-to-map-data-lineage/>.