

A Modular Graph-Native Query Optimization Framework

Bingqing Lyu, Xiaoli Zhou, Longbin Lai, Yufan Yang, Yunkai Lou, Wenyuan Yu, *Ying Zhang, Jingren Zhou

Tongyi Lab, Alibaba Group, China *Zhejiang Gongshang University, China

> Presented by Longbin Lai 2025/06/25







CONTENTS

01



Background

O2 GOpt Overview 03





Experimental Results





CONTENTS

01



Background

O2 GOpt Overview 03



04

Experimental Results



Background





Declarative Query Languages

Query Optimizer is a KEY component in (G)DBMS!



Graph Database

*The image is downloaded from the internet. If there is any infringement, we will delete it



Relational Database



A Monolithic Pipeline for query optimization



[1]GraphScope: a unified engine for big graph processing, Fan et al, VLDB 2021

*The image is downloaded from the internet. If there is any infringement, we will delete it

A Monolithic Pipeline for query optimization (Cont.)

*The image is downloaded from the internet. If there is any infringement, we will delete it

A Relational Optimization Pipeline

Standard Relational Query Language

*The image is downloaded from the internet. If there is any infringement, we will delete it

Relational Algebra

Relational Optimizer

A Graph Optimization Pipeline?

*The image is downloaded from the internet. If there is any infringement, we will delete it

Optimizer?

A Non-graph-native Pipeline for graph optimization

[1] Modern Techniques for Querying Graph-Structured Relations: Foundations, System Implementations, and Open Challenges, Amine Mhedhbi, Semih Salihoğlu, VLDB 2022 [2] Implementing Graph queries in a Relational Database, https://blog.whiteprompt.com/implementing-graph-queries-in-a-relational-database-7842b8075ca8?gi=fb6853dc262e

*The image is downloaded from the internet. If there is any infringement, we will delete it

A Non-graph-native Pipeline for graph optimization (Cont.)

[1] Optimizing Subgraph Queries by Combining Binary and Worst-Case Optimal Joins, Amine Mhedhbi, Semih Salihoğlu, VLDB 2019 [2] Glogs: interactive graph pattern matching query at large scale, Longbin Lai et al, ATC 2023

*The image is downloaded from the internet. If there is any infringement, we will delete it

GOpt: A Modular Graph-native Optimization Framework

*The image is downloaded from the internet. If there is any infringement, we will delete it

A unified Graph Intermediate **Representation (GIR)**: uniform the logical expression of various graph query languages (Cypher, Gremlin)

GOpt: A Modular Graph-native Optimization Framework

*The image is downloaded from the internet. If there is any infringement, we will delete it

A unified Graph Intermediate Representation (GIR): uniform the logical expression of various graph query languages

Advanced graph query optimization techniques: hybrid join, high-order statistics, and type inference

GOpt: A Modular Graph-native Optimization Framework

*The image is downloaded from the internet. If there is any infringement, we will delete it

CONTENTS

01

Background

O2 GOpt Overview 03

04

Experimental Results

[-]阿里云

A Simple Query Example

Schema

MATCH a triangle pattern containing a Place vertex, and then FILTER with a specified place name, GROUP the matchings by the name and the TOP 10 results are returned?

MATCH $(v1) - [e1] \rightarrow (v2)$ $(v2) - [e2] \rightarrow (v3)$ MATCH $(v1) - [e3] \rightarrow (v3:Place)$ WHERE v3.name = "China" WITH v2, COUNT(v2) as cnt RETURN v2, cnt ORDER BY cnt LIMIT 10

User Query

Typical Optimization Pipeline in GOpt

NOTE: Case can be more complex in practice

Step 1: Compile (Graph Intermediate Representation, GIR)

Context for pattern matching, including graph operators

MATCH (v1) - [e1] -> (v2)(v2) - [e2] -> (v3)MATCH (v1) -[e3]-> (v3:Place) WHERE v3.name = "China" WITH v2, COUNT(v2) as cnt RETURN v2, cnt ORDER BY cnt LIMIT 10

Compile

(a) Cypher Query

(b) Logical Plan in GIR

Step 2: Rule-based Optimization (RBO)

Built-in Predefined Rules

(a) Logical Plan

(b) Optimized Plan After RBO

[-] 阿里云

Step2: RBO (Cont.)

Built-in Rule: JoinToPattern

[-] 阿里云

Optimized Plan After RBO

Step 3: Type Inference (Cont.)

Graph Schema

Step 3: Type Inference (Cont.)

Step 3: Type Inference (Cont.)

Step 3: Type Inference (Cont.)

Step 4: Cost-based Optimization (CBO)

(a) Optimized Plan After Type Infer

GLogS: Interactive Graph Pattern Matching Query At Large Scale, Lai et al, ATC 2023

(b) Optimized Plan After CBO

[-] 阿里云

Step 4: CBO (Cont.) : Cost Model

- $Plan(p) = (\Phi = \{p_1, p_2, ..., p_n = p\}, \Gamma = [\tau_1, \tau_2, ..., \tau_m])$
 - p_i : intermediate pattern
 - τ_i : operation on intermediate patterns, Join or Vertex Expansion
 - $\mathcal{F}(p_i)$: Frequency(count number) of the p_i in the graph.
 - $\operatorname{Cost}(\operatorname{Plan}(p)) = \sum_{p' \in \Phi} \mathcal{F}(p') + \sum_{\tau \in \Gamma} \operatorname{Cost}(\tau)$
- GLogueQuery provides $\mathcal{F}(p_i)$ based on high-order statistics
 - GLogue provides precomputed statistics of small patterns

• $Cost(\tau_i)$ is registrable according to backend implementations, e.g.,

- Neo4j (ExpandInto): $\operatorname{Cost}(\operatorname{Expand}(p_s \to p_t)) = \alpha_{ve} \mathcal{F}(p_s) \sum_{i=1}^k \overline{\sigma_{e_i}}$
- GraphScope (WCOJ): $Cost(Expand(p_{s 28} \rightarrow p_t)) = n\mathcal{F}(p_1)$

If pattern is not in GLogue, compute by $\mathcal{F}(p) = Avg_{p_1,p_2} \frac{\mathcal{F}(p_1) \times \mathcal{F}(p_2)}{\mathcal{F}(p_1 \cap p_2)}$, $p = p_1 \cup p_2$

Step 4: CBO (Cont.): Plan Optimizer

A top-down search framework utilizing registered operator costs and cardinality estimation from GLogueQuery to find optimal plans with minimum estimated costs

CONTENTS

01

Background

O2 GOpt Overview

03

GOpt Integration

- GOpt Architecture
- Integration with Neo4j
- Integration with GraphScope

04

System Architecture of GOpt

User-friendly interfaces in GOpt

Input: various clients to the unified GIR

Output: convert to alternative execution plans

Rules: heuristic rules for using in RBO

GOpt Integration – Neo4j

GOpt Integration – Neo4j (Cont.)

GOpt Integration – GraphScope

GOpt Integration – GraphScope (Cont.)

CONTENTS

01

Background

O2 GOpt Overview

03

04

Experimental Results

Experimental Settings

Dataset

Queries

LDBC Graph scaled at sf=30, 100, 300, 1000, up to 2.7 billion vertices and 17.8 billion edges.

LDBC Interactive (IC) and Business Intelligence (BI)

workloads, challenging the systems with real-world scenarios.

Compared Systems

- Neo4j with Cypher on a single machine,
- GraphScope with
 Gremlin on distributed
 machines

- Neo4j-Plan: Plans by Neo4j
 Optimizer
- **Gopt-plan**: Plans by GOpt

Experimental Results: Comprehensive Comparison

- Neo4j-plan v.s. GOpt-plan on Neo4j
 - GOpt-plan outperformed Neo4j-plan in 16/29 queries
 - Avg. speedup 9.2x
 - Max. speedup 48.6x (IC6)
- Neo4j-plan v.s. GOpt-plan on GraphScope
 - GOpt-plan outperformed Neo4j-plan in 17/29 queries
 - Avg. speedup 33.4x
 - Max. speedup 78.7x (IC6) since hybrid-join strategy and optimized wcoj implementation

Comparison of Neo4j-plan and GOpt-plan on Neo4j and GraphScope resp.

Case Study: Money Mule

Money Mule: Fraudsters transfer funds through multiple intermediaries before another fraudster withdraws the money

MATCH (p1: PERSON)-[p:*6]->(p2: PERSON) WHERE p1.id IN \$S1 AND p2.id IN \$S2 RETURN p

(a) Cypher Query Money Mule Detection

Optimal Execution Plan of ST3, ST4, ST5

(b) GOpt-Plan for ST1 to ST5 with Different Src/Dst Vertices

(c) Neo4j-Plan for ST1 to ST5 with Different Src/Dst Vertices (all plans are single direction expansion)

GOpt automatically determines the best join vertex

- based on CBO
- GOpt-plans outperform alternatives by 3x 20x
- All Neo4j-plans run OT

Case Study (cont.): LDBC SNB Bnechmark

The Case of IC6 in LDBC SNB

AUDITED RESULTS USING A DECLARATIVE QUERY LANGUAGE

Benchmark setup	SF	Hardware	Performance	Performance (price-adj.)	Documents
 O System: GraphScope Flex 0.31.3 O Test sponsor: Alibaba 	30	Alibaba Cloud ecs.g8a.48xlarge 96×AMD EPYC 9T24 @ 3.7GHz vCPUs, 768GiB RAM	37,650.00 ops/s 339.97 O Full disclosure report O Executive summary O Signatures	 Full disclosure report Executive summary Signatures 	
Cloud O Date: 2025-04-21 O Query language: Cypher	100 2025-04-21 language: Cypher	Alibaba Cloud ecs.g8a.48xlarge 96×AMD EPYC 9T24 @ 3.7GHz vCPUs, 768GiB RAM	79,244.18 ops/s	715.55	O Supplementary package
 System cost: 808,444.48 RMB (110.745,82 USD) 	300	Alibaba Cloud ecs.g8a.48xlarge 96×AMD EPYC 9T24 @ 3.7GHz vCPUs, 768GiB RAM	80,510.79 ops/s	726.99	

The full disclosure report from LDBC

More Details in GOpt Technique Report

THANKS

Please star us on GraphScope. We will soon announce a **new open-source product** for better graph data management

42