

# GLogS: Interactive Graph Pattern Matching Query at Large Scale

Longbin Lai<sup>1\*</sup>, Yufan Yang<sup>2\*</sup>, Zhibin Wang<sup>3</sup>, Yuxuan Liu<sup>2</sup>, Haotian Ma<sup>2</sup>, SiJie Shen<sup>1</sup>,  
Bingqing Lyu<sup>1</sup>, Xiaoli Zhou<sup>1</sup>, Wenyuan Yu<sup>1</sup>, Zhengping Qian<sup>1</sup>, Chen Tian<sup>3</sup>, Sheng  
Zhong<sup>3</sup>, Yeh-Ching Chung<sup>2</sup> and Jingren Zhou<sup>1</sup>

<sup>1</sup>Alibaba Group, China

<sup>2</sup>The Chinese University of Hong Kong, Shenzhen

<sup>3</sup>Nanjing University

# Outline


---

- 1 Background and Motivation
- 2 System Overview
- 3 Feature Highlights
- 4 Evaluation
- 5 Conclusion


# Graph and Graph Pattern Matching



Social Network Graph





Knowledge Graph




Molecular Graph

Graph is Prevalent!



(a) Pattern p



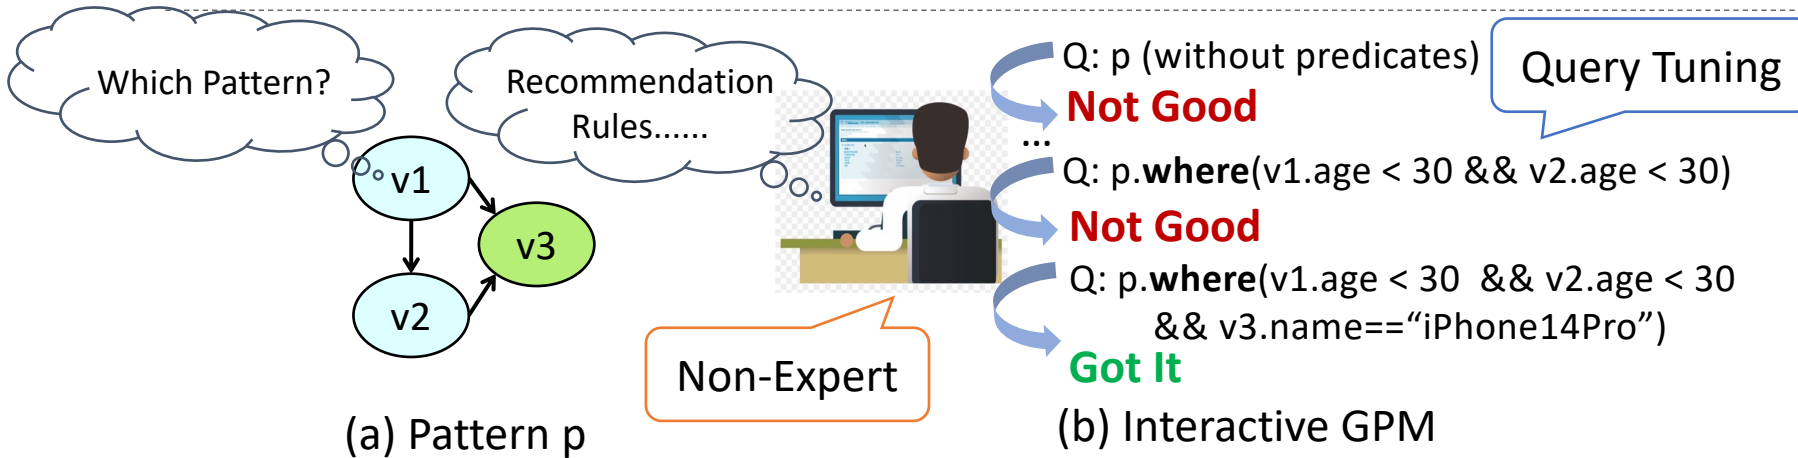

(b) Graph G

v1	v2	v3
u2	u1	u4
u2	u3	u5

(c) The mappings of p in G

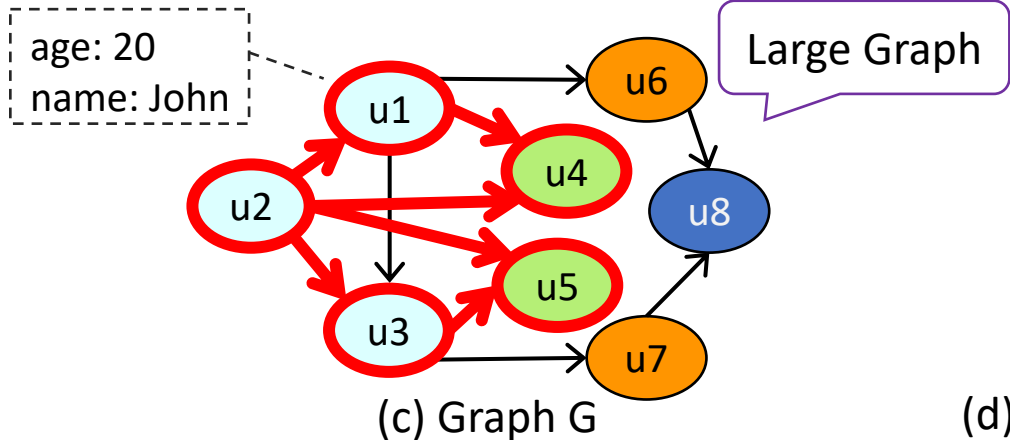
Graph Pattern Matching

# Interactive Graph Pattern Matching



Usability

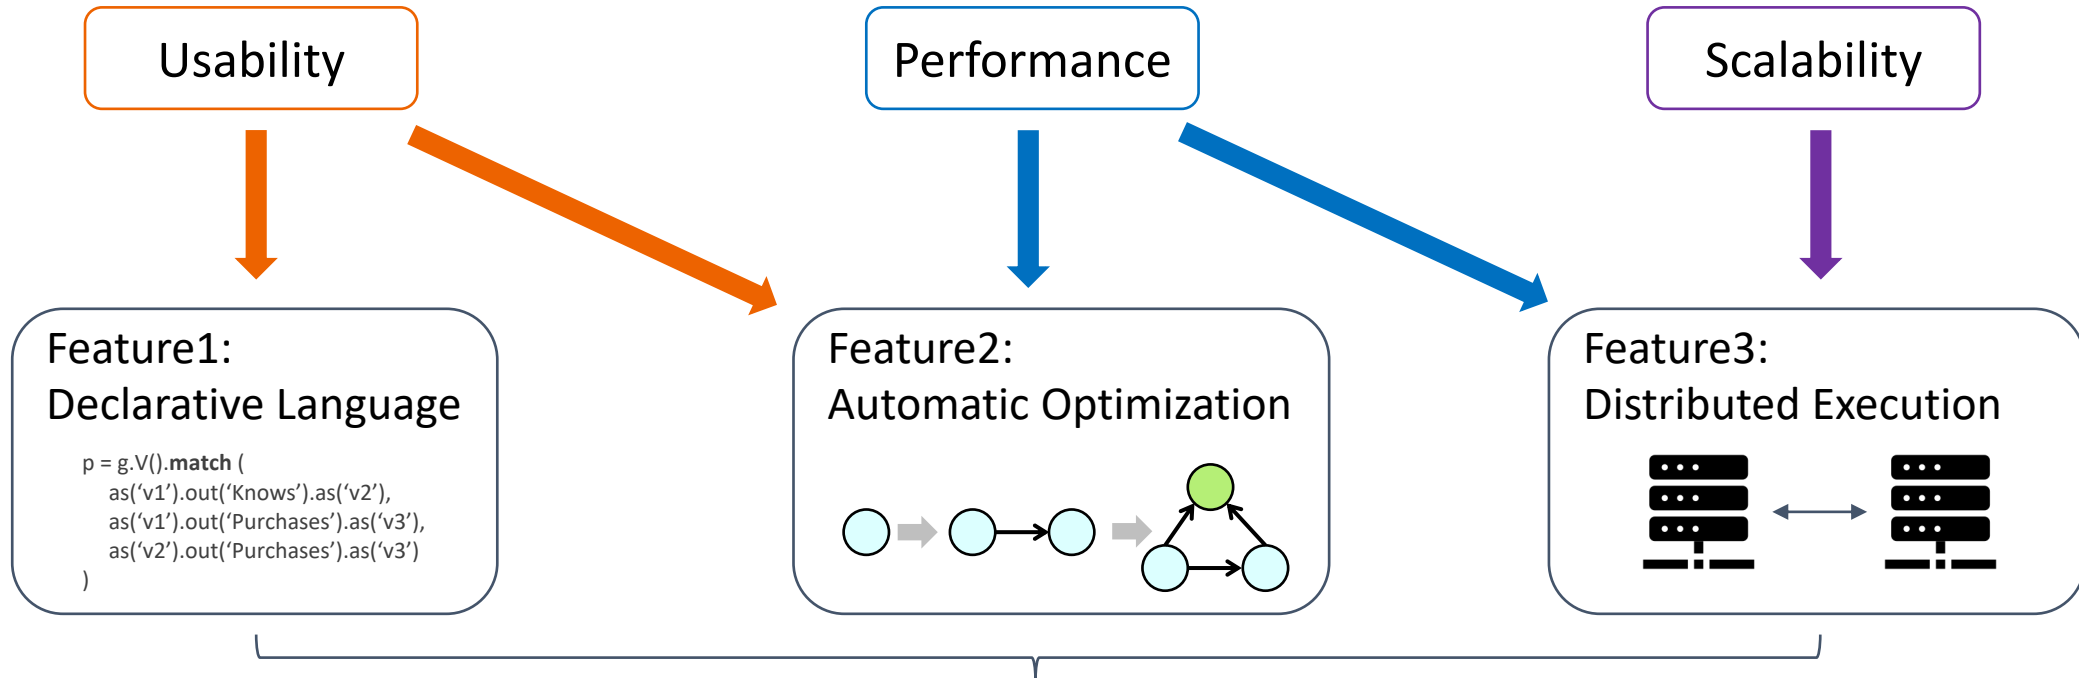
Performance



$v_1$	$v_2$	$v_3$
$u_2$	$u_1$	$u_4$
$u_2$	$u_3$	$u_5$

Scalability

# iGPM: Requirements and Features



Are there any existing solutions?



TigerGraph

# iGPM: Existing Solutions and Challenges



Feature1:  
Declarative Language

✔ Support Cypher

$p_i$	$\mathcal{F}(p_i)$
	$f_1$
	$f_2$
	$f_3$
.....	.....

low-order

high-order

Feature2:  
Automatic Optimization

○ Limited Optimization

Cannot guarantee worst-case optimal

- Only support binary join operation
- large intermediate results size

Only use low-order statistics to estimate plan cost

- Inaccurate plan cost estimation

Feature3:  
Distributed Execution

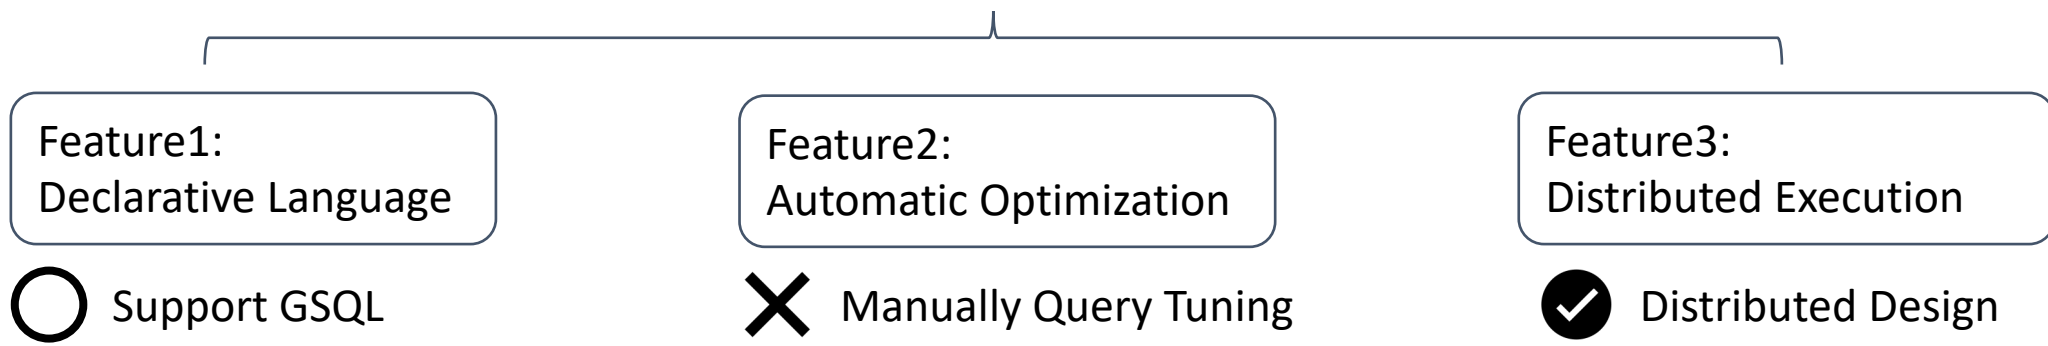
✗ Single Machine Design

Bad execution plan!

Challenge1: Optimization

- Guarantee worst-case optimal
- Adopt high-order statistics

# iGPM: Existing Solutions and Challenges



However, it requires query pre-installation!

- Pre-install the queries before they can be executed in the TigerGraph
- Involves native code generation and compilation
  - 1~3 minutes per query, too long!



**Challenge2: Compilation**

- Timeliness
- Support interactive manner

# Outline

---

1

Background and Motivation

2

System Overview

3

Feature Highlights

4

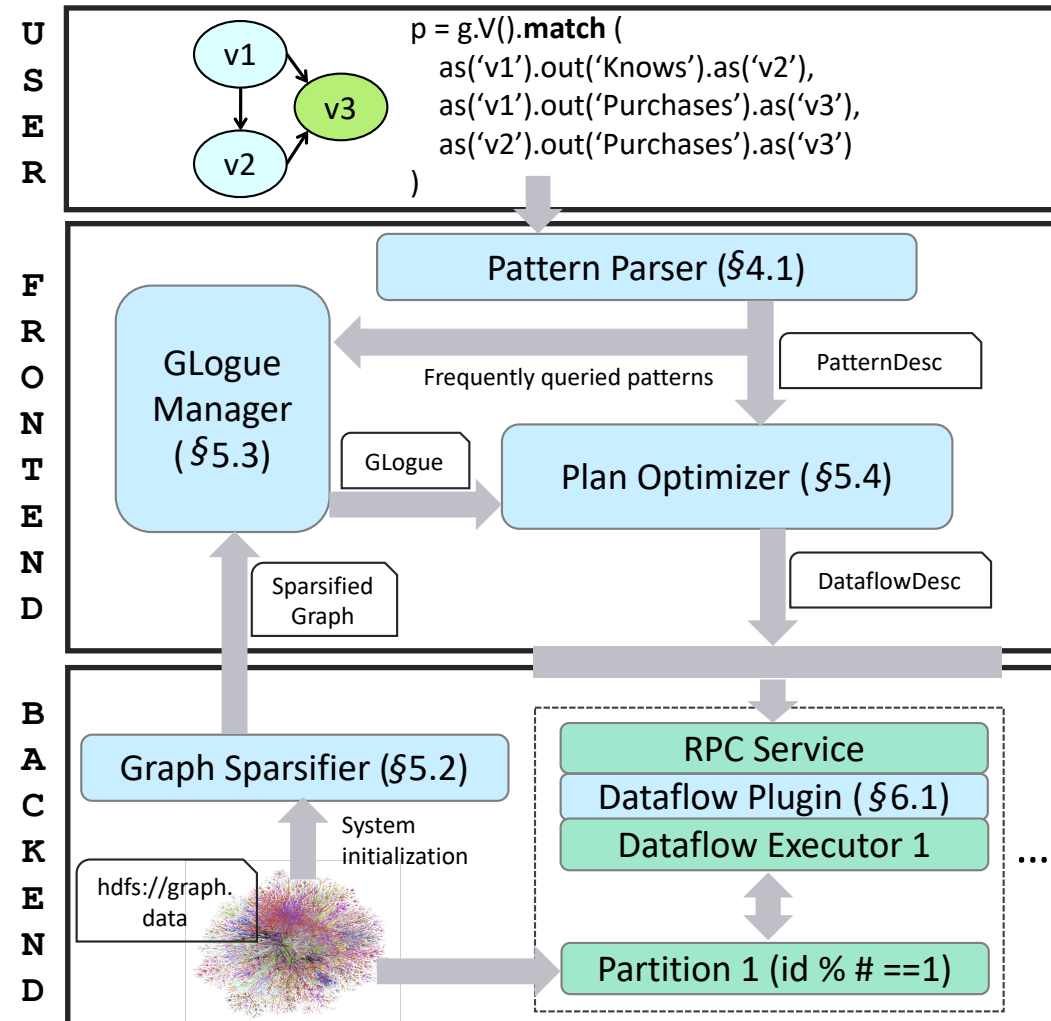
Evaluation

5

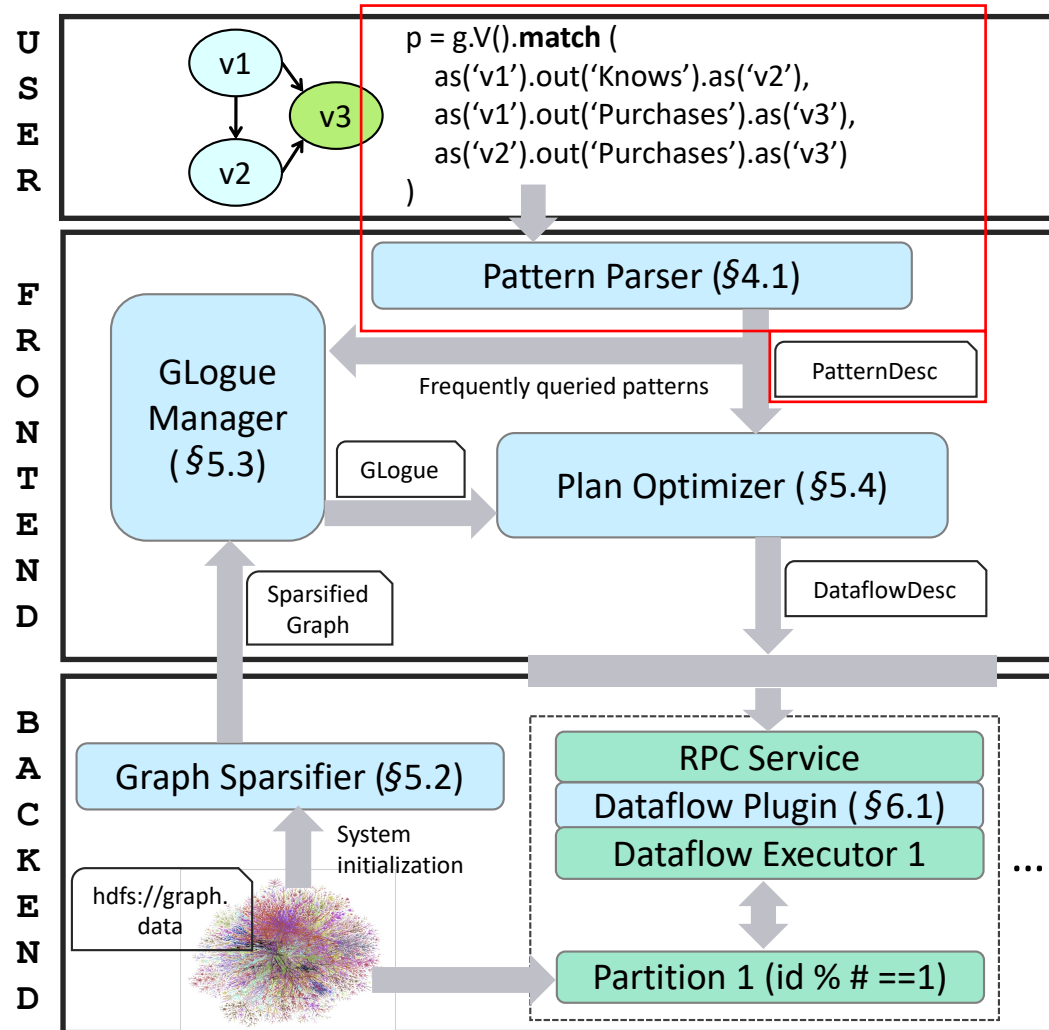
Conclusion



# System Overview



# Frontend: Pattern Parser

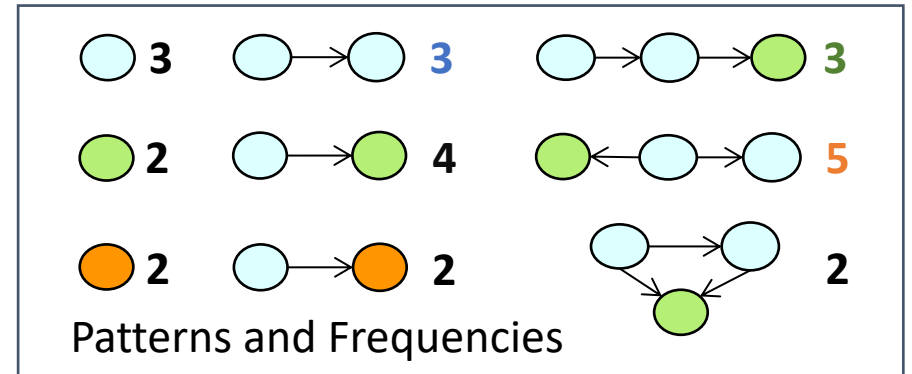
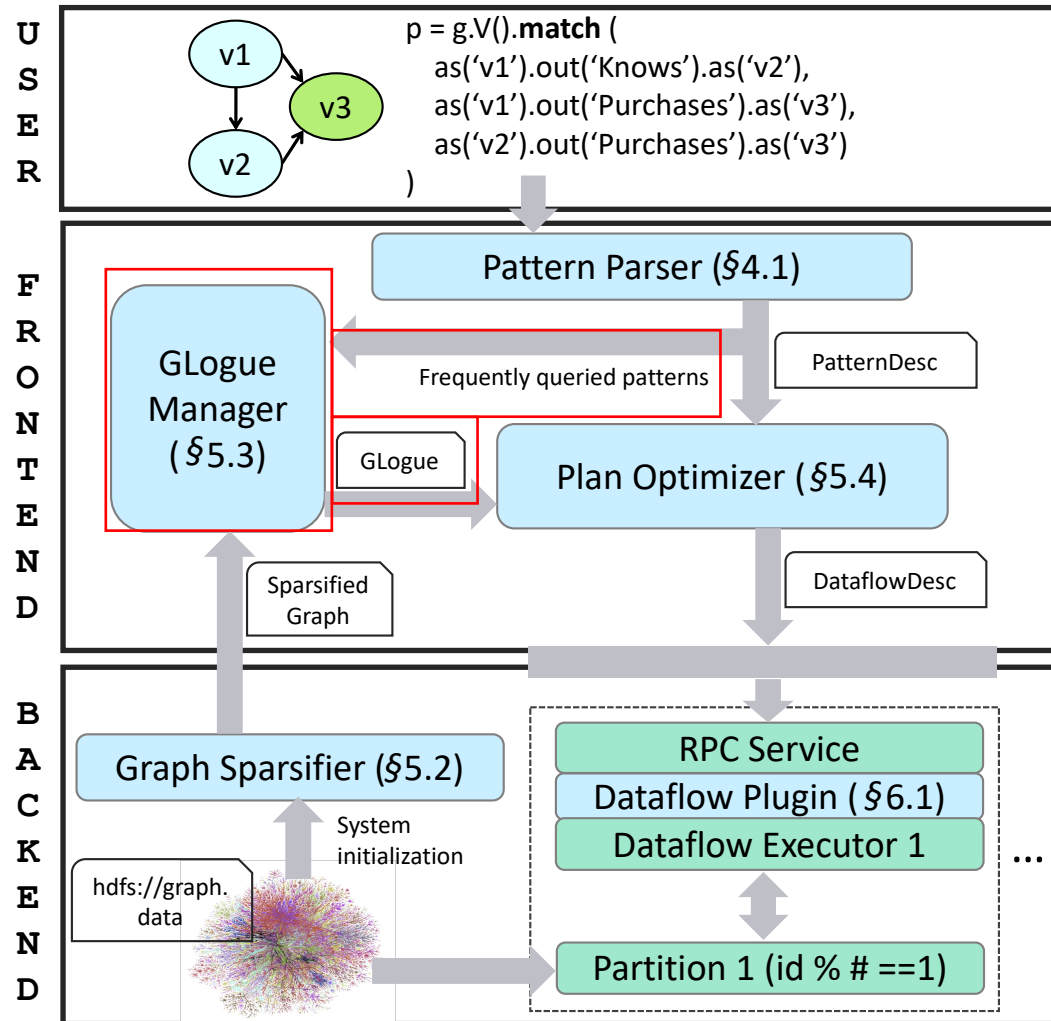


<b>Gremlin</b>	<code>g.V().hasLabel('Person').match(   as('v1').outE('Knows').as('e1').inV('Person').as('v2'),   as('v1').outE('LivesIn').inV('City').as('v3'),   as('v2').outE('LivesIn').inV('City').as('v3'))</code>
<b>Cypher</b>	<code><b>MATCH</b> (v1: Person) -[e1:Knows]-&gt; (v2: Person),   (v1: Person) -[LivesIn]-&gt; (v3: City),   (v2: Person) -[LivesIn]-&gt; (v3: City)</code>

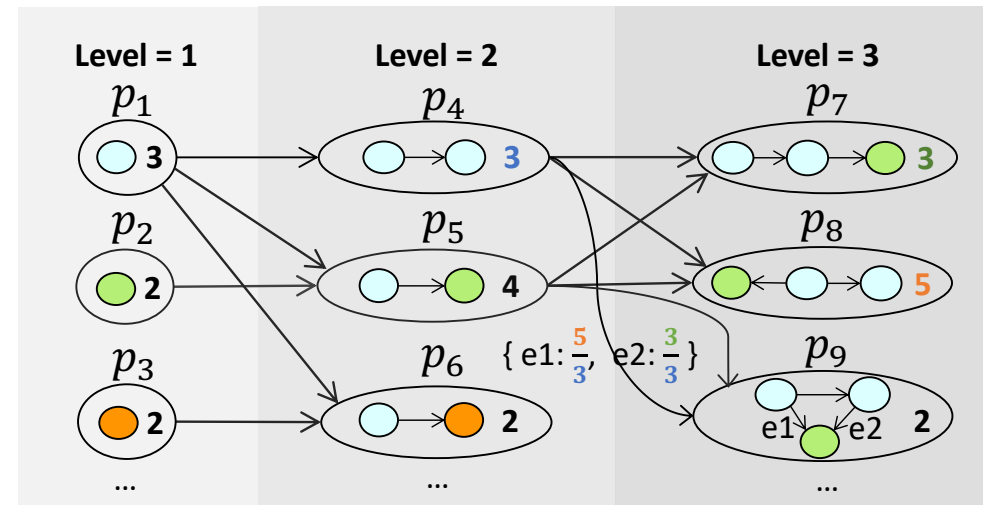
Parse

```
PatternDesc { sentences: [ [
  GetV("", 'v1', Person, None)],
  GetE('v1', 'e1', Knows, Out)],
  GetV('e1', 'v2', Person, Target) ], ... ] }
```

# Frontend: GLogue Manager

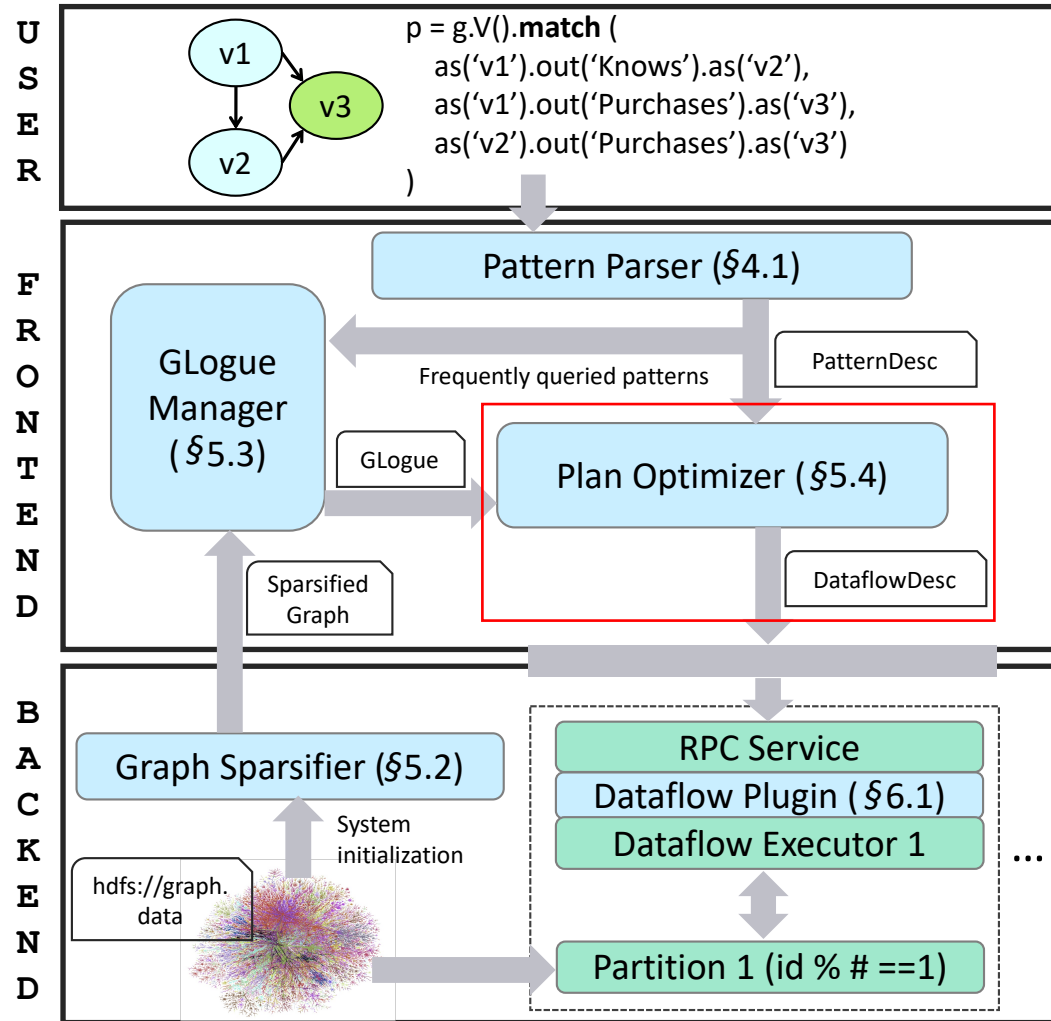


Build, Update, Maintain




GLogue

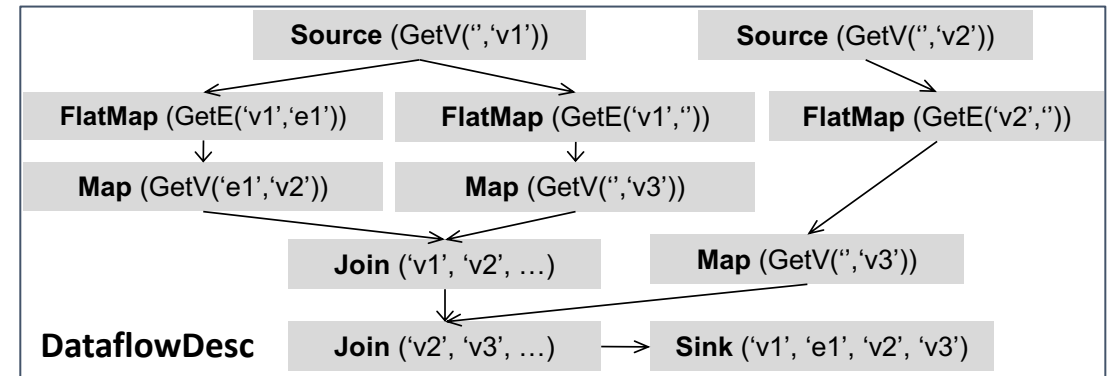
# Frontend: Plan Optimizer



**PatternDesc** { sentences: [ [ GetV("", 'v1', Person, None)], GetE('v1', 'e1', Knows, Out)], GetV('e1', 'v2', Person, Target) ], ... ] }



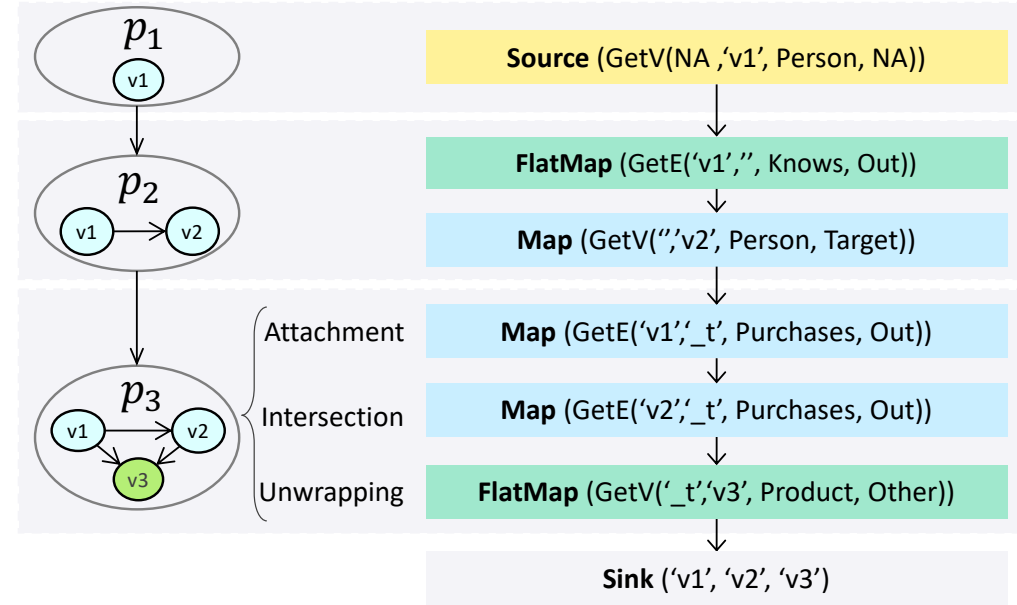
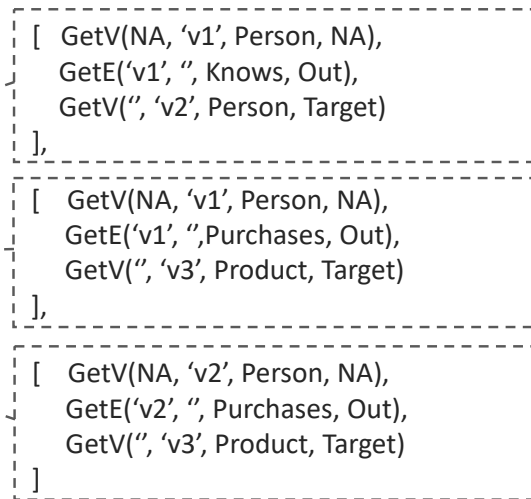
## Plan Generation



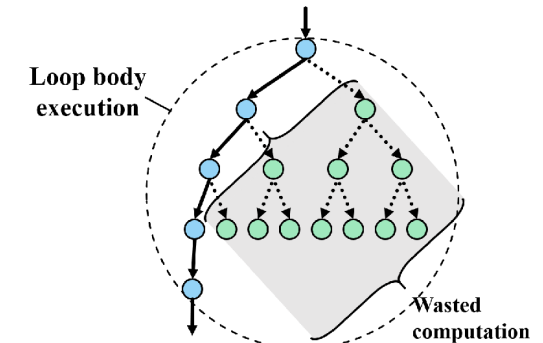

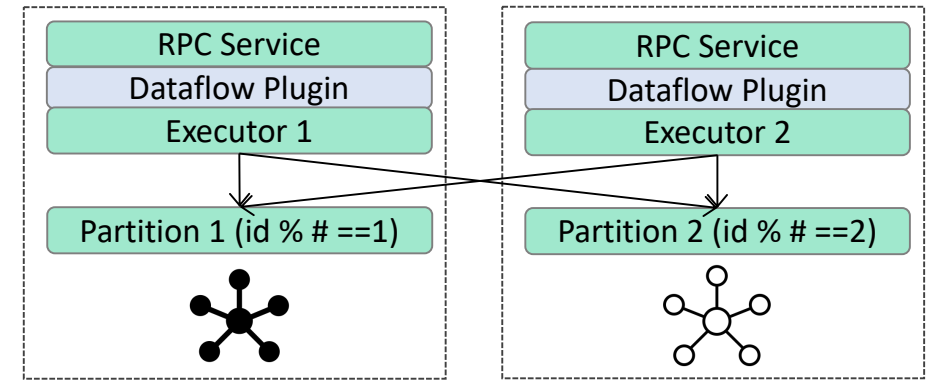
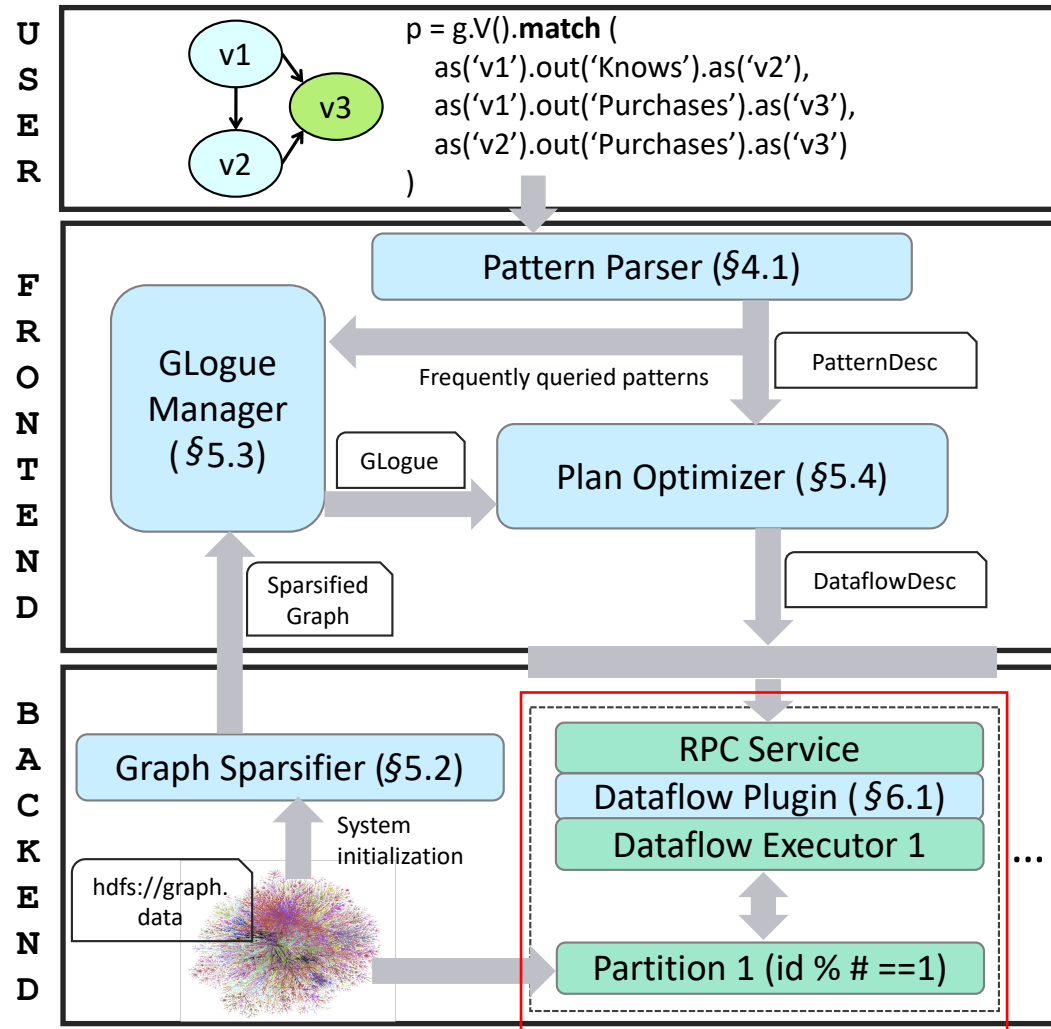
# Frontend

```

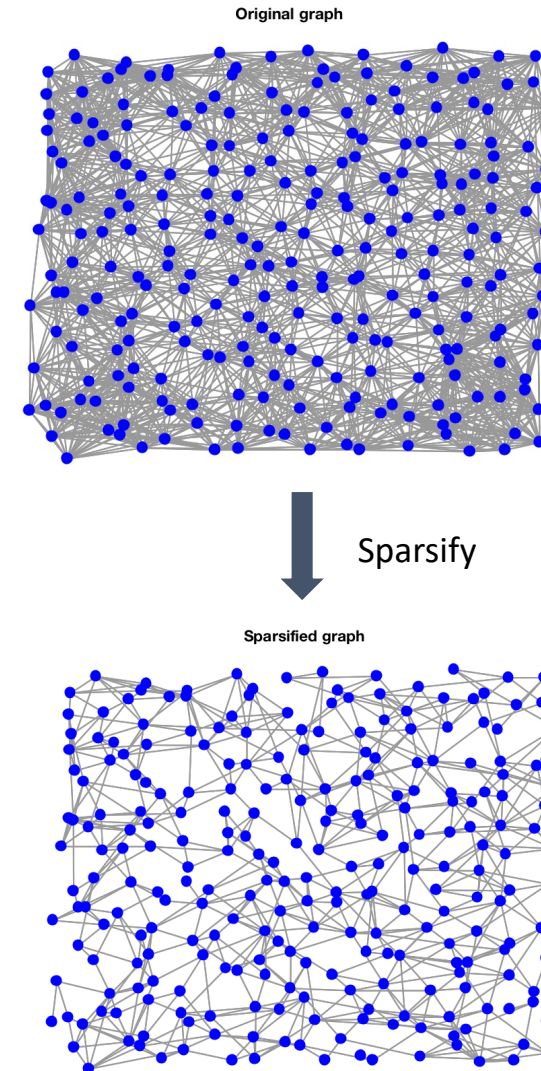
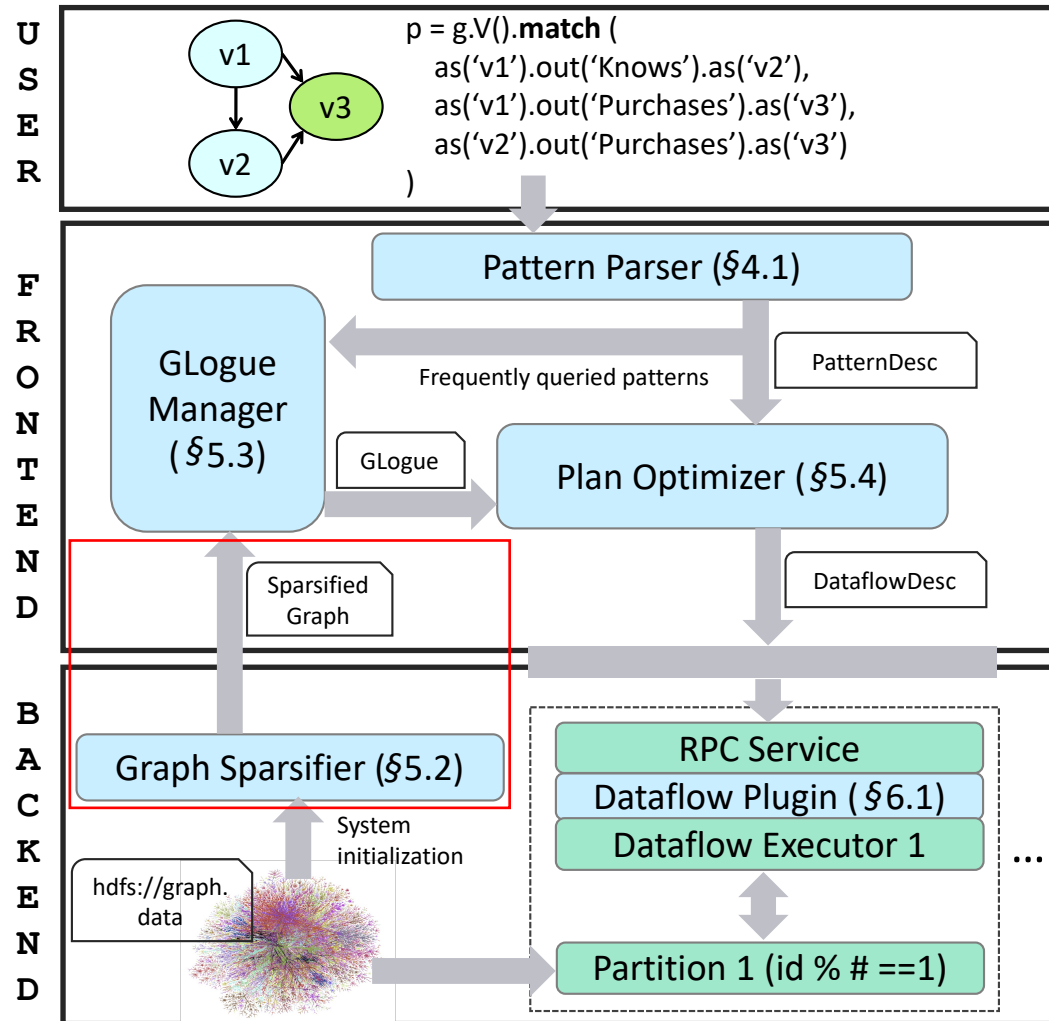
p = g.V().match (
  as('v1').out('Knows').as('v2'),
  as('v1').out('Purchases').as('v3'),
  as('v2').out('Purchases').as('v3')
)
  
```



# Backend: Distributed Dataflow Engine



# Backend: Graph Sparsifier



# Outline



---

- 1 Background and Motivation
- 2 System Overview
- 3 Feature Highlights
- 4 Evaluation
- 5 Conclusion




# GPM Execution Plan


- Graph Pattern Matching Plan Space:
  - Worst-Case-Optimal Join Plan, based on Vertex Expansion



- Binary Join Plan



- Hybrid Plan

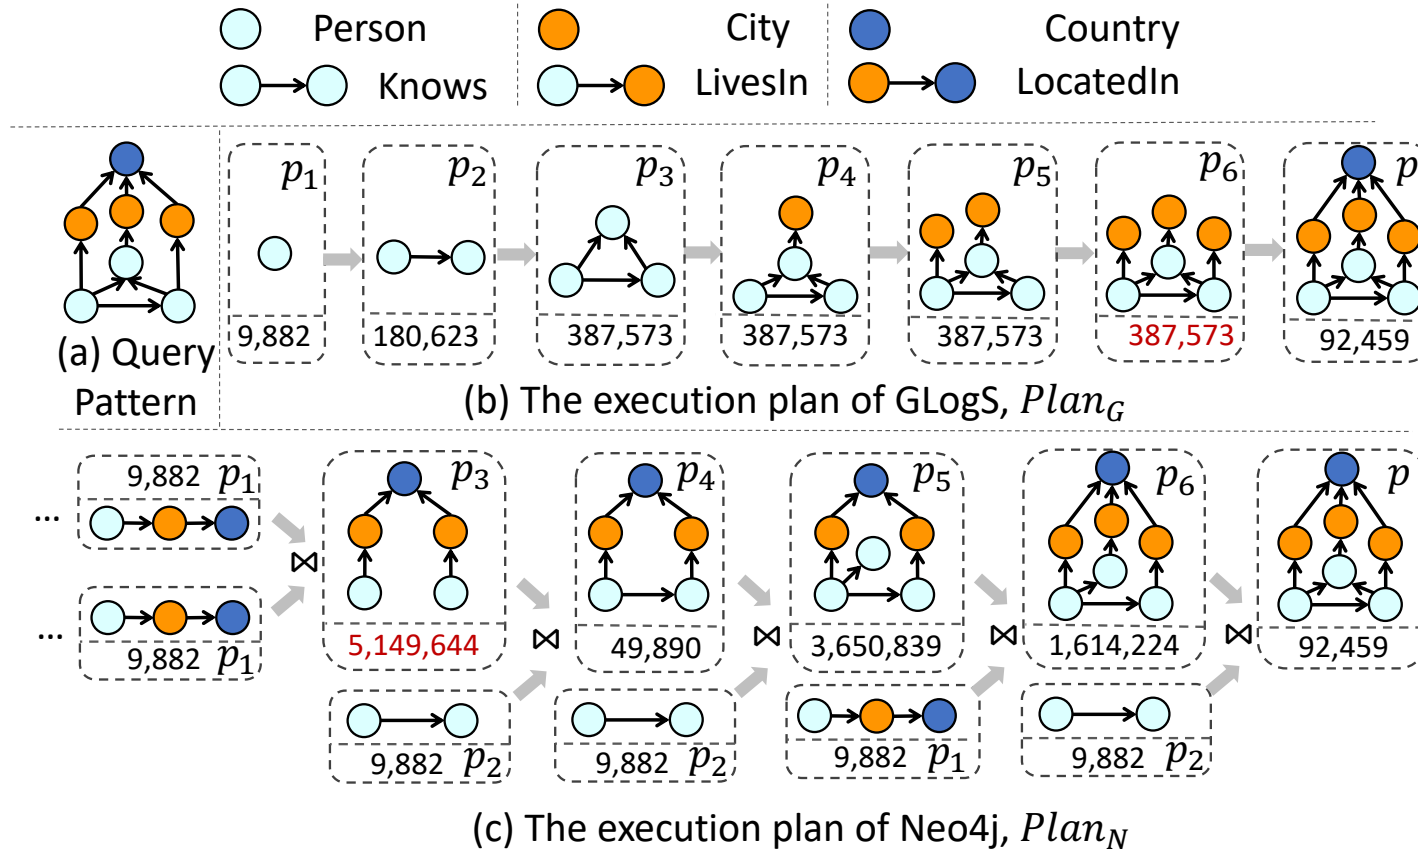


Best GPM execution plan lies in the hybrid plan space!

➔ Vertex Expansion      ⋈ Binary Join

# GPM Execution Plan

- Key to GPM Execution Plan Optimization:
  - Minimizing Intermediate Results
- Comparison between the plans generated by GLogS and Neo4j for LDBC BI11 query separately

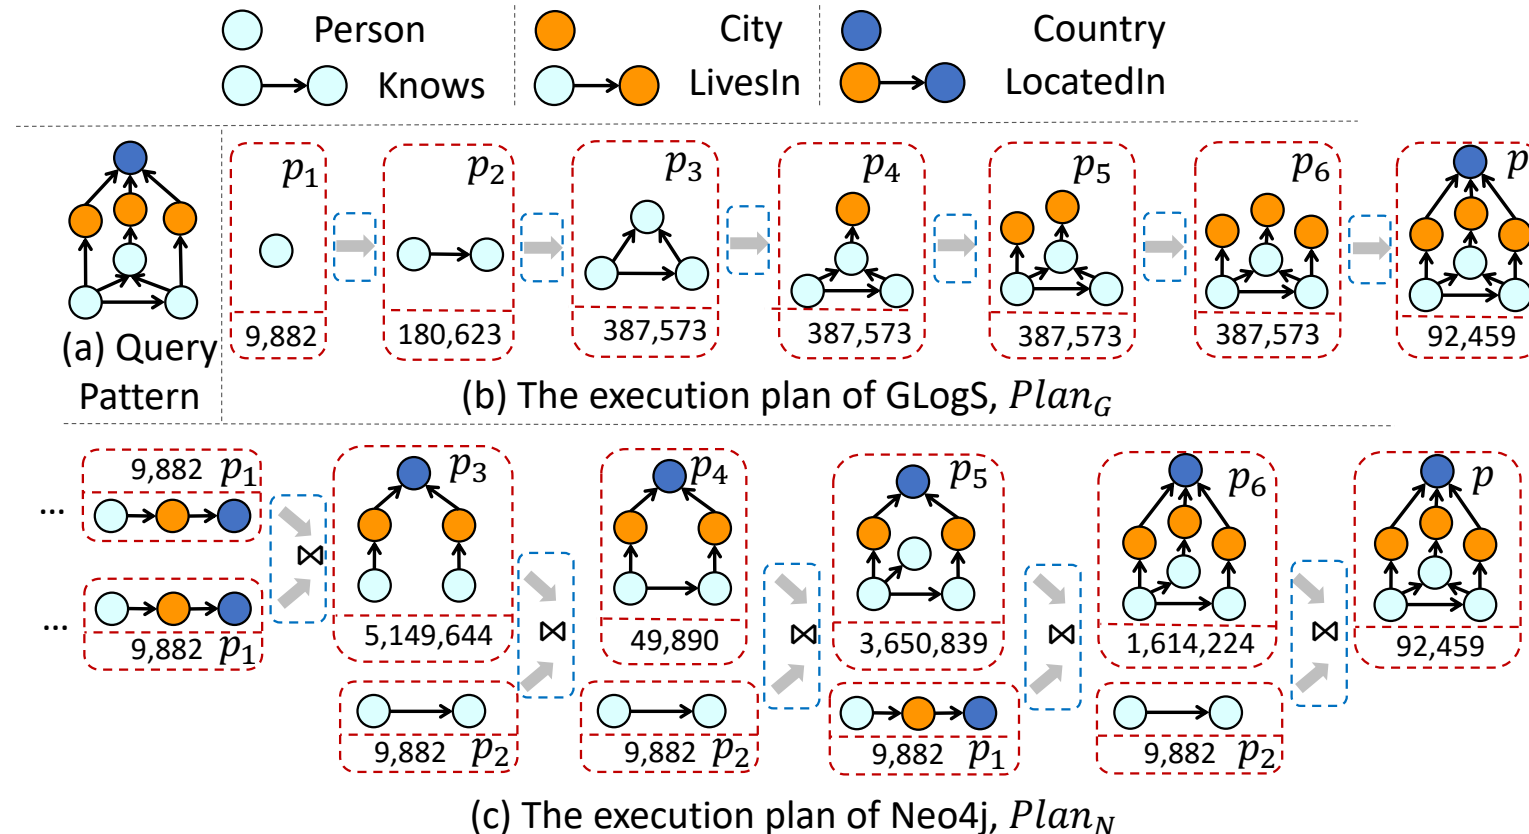


Why is GLogS's plan better?

- Hybrid Execution plan
- High-order statistics

# Cost Model for GPM Execution Plan

- Define a **Cost Model** to estimate the cost of every possible GPM execution plan
  - Cost of **accessing intermediate results** from the memory (either local or remote memory)
  - Cost of **operations** (vertex expansion, binary join)



# Cost Model for GPM Execution Plan

---

- $\text{Plan}(p) = (\Phi = \{p_1, p_2, \dots, p_n = p\}, \Gamma = [\tau_1, \tau_2, \dots, \tau_m])$ 
  - $p_i$ : intermediate pattern
  - $\tau_i$ : operation on intermediate patterns, Join or Vertex Expansion
  - $\mathcal{F}(p_i)$ : Frequency(count number) of the  $p_i$  in the graph.
  - $\text{Cost}(\text{Plan}(p)) = \sum_{p' \in \Phi} \mathcal{F}(p') + \sum_{\tau \in \Gamma} \text{Cost}(\tau)$
  - $\text{Cost}(\text{Join}(\{p_{s_1}, p_{s_2}\} \rightarrow p_t)) = \alpha_j (\mathcal{F}(p_{s_1}) + \mathcal{F}(p_{s_2}))$
  - $\text{Cost}(\text{Expand}(p_s \rightarrow p_t)) = \alpha_{ve} \sum_{f \in Q(p_s)} \sum_{i=1}^k \sigma_{e_i}(f) = \alpha_{ve} \mathcal{F}(p_s) \sum_{i=1}^k \overline{\sigma}_{e_i}$
- We find that  $\mathcal{F}(p_i)$  is the key to the entire cost model!

# Cost Model for GPM Execution Plan

---

How to extract various of  $\mathcal{F}(p_i)$  from graph efficiently?



Graph Sparsifier

What data structure should be used to organize all kinds of  $p_i$ ,  $\mathcal{F}(p_i)$  and  $\tau_i$ ?



GLogue

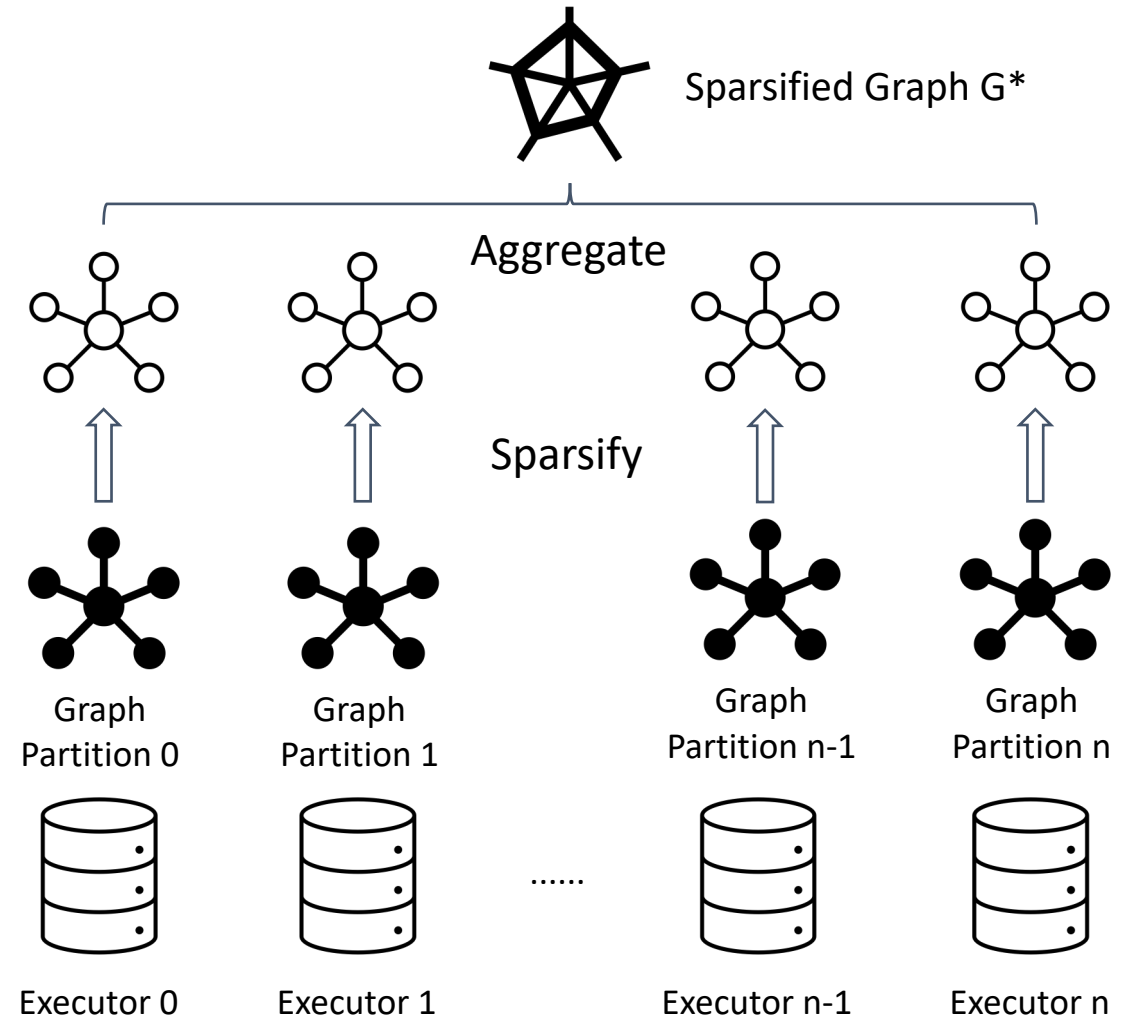
Based on  $p_i$ ,  $\mathcal{F}(p_i)$  and  $\tau_i$ , how to generate optimized plan for a given query?



Plan Optimizer

# Graph Sparsifier

- It's **cost-prohibitive** to compute  $\mathcal{F}(p_i)$  **directly** from the original large graph datasets.
- Conduct sparsification on each partition of the graph, and then aggregate them to form the sparsified graph  $G^*$
- Use  $\mathcal{F}_{G^*}(p_i)$  (with normalization) as an estimation of  $\mathcal{F}_G(p_i)$  for cost evaluation



# Graph Sparsifier: Stratified Sparsification

$e$	Knows	Purchases	LivesIn	LocatedIn
$\mathcal{F}(e)$	10000	50000	3000	100



Uniform Sparsification, rate 1%

$e$	Knows	Purchases	LivesIn	LocatedIn
$\mathcal{F}^*(e)$	$\approx 100$	$\approx 500$	$\approx 30$	$\approx 1$

Very likely to be 0!




Pattern Vanishing!

**Stratified Sparsification:** rate  $\min(1, \frac{M}{\sum \mathcal{F}(e)} * \frac{1}{\mathcal{F}(e)})$

$e$	Knows	Purchases	LivesIn	LocatedIn
rate	1%	0.2%	3%	100%
$\mathcal{F}^*(e)$	$\approx 100$	$\approx 100$	$\approx 100$	100

# GLogue

- How to organize  $p_i$ ,  $\mathcal{F}(p_i)$  and  $\tau_i$ ?
  - Table-based catalog isn't good enough!



$p_s$	$\tau$	$p_t$
	-[e: Knows]->	
	-[e: Knows]->	
...		...
	<-[e:LivesIn]-	
	<-[e1:LivesIn]-, [e2:LivesIn]->	
	-[e:LivesIn]->	
...		...


$p_i$	$\mathcal{F}(p_i)$
	$f_1$
	$f_2$
.....	.....
	$f_3$
.....	.....
	$f_4$
.....	.....

Many Redundancies!

Difficult for Plan Generation



# GLogue



edge

$p_s$	$\tau$	$p_t$
	-[e: Knows]->	
	-[e: Knows]->	
...		...
	<-[e:LivesIn]-	
	<-[e1:LivesIn]-, -[e2:LivesIn]->,	
	-[e:LivesIn]->	
...		...

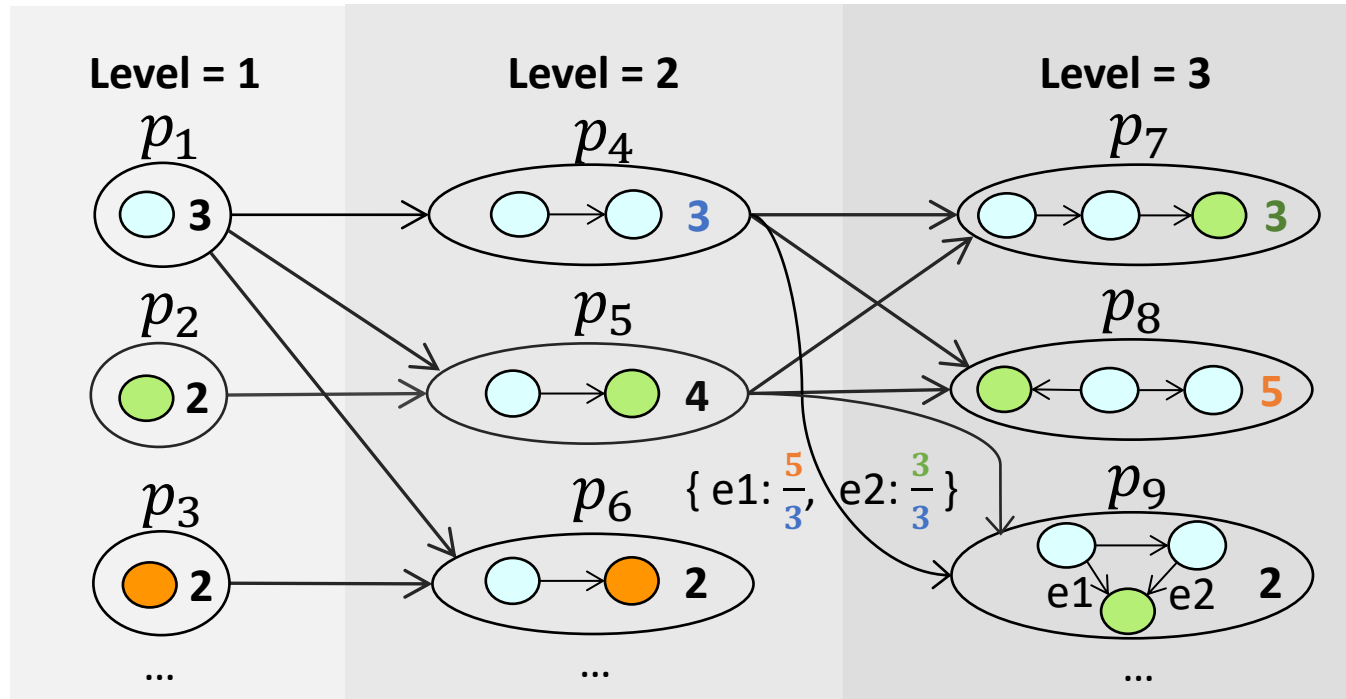

edge property

vertex

$p_i$	$\mathcal{F}(p_i)$
	$f_1$
	$f_2$
.....	.....
	$f_3$
.....	.....
	$f_4$
.....	.....

vertex property

# GLogue

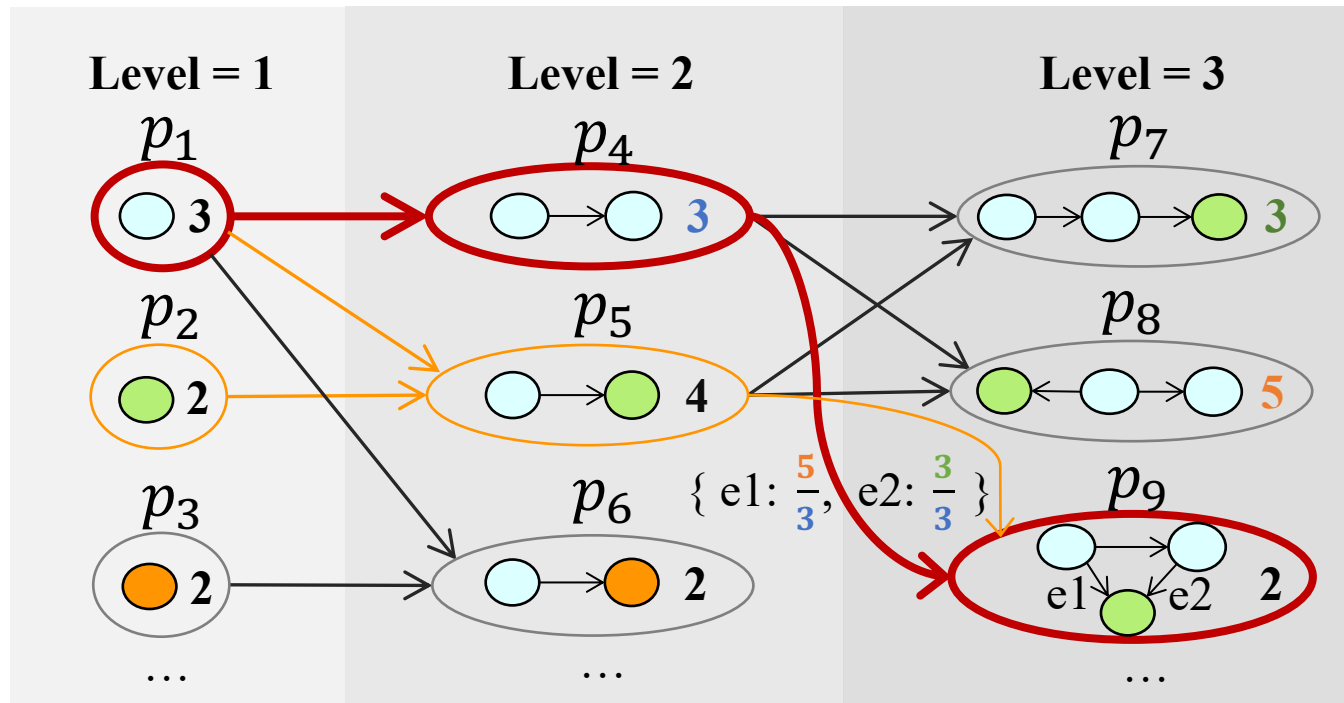


## Algorithm 1: Constructing GLogue: $(G^*, Level)$ .

```

1 Function constructGLogue ( $Level$ ): GLogue
2   Initialize GLogue with  $level = 1, 2$  precomputed ;
3   for  $3 \leq i \leq Level$  do
4     Let  $QSet$  maintain all  $p$  of  $i$  vertices in ascending
5     order via  $|E_p|$ ;
6     for  $p \in QSet$  do
7       updateFromPattern ( $p$ , GLogue);
8   return GLogue;
9
10 Function updateFromPattern ( $p$ , GLogue)
11   GLogue.addVertex ( $p$ ,  $\mathcal{F}(p)$ );
12   for  $p_s \subset p$  &  $|V_{p_s}| = |V_p| - 1$  do
13     if not GLogue.contains( $p_s$ ) then
14       updateFromPattern ( $p_s$ , GLogue);
15     for  $e = (s, t) \in E_p \setminus E_{p_s}$  do
16       let  $p'_s$  be the pattern with  $e$  adding to  $p_s$ ;
17       eMap.insert( $e$ ,  $\frac{\mathcal{F}(p'_s)}{\mathcal{F}(p_s)}$ );
18     GLogue.addEdge ( $(p_s, p)$ , eMap);
19
20 for  $p_{s_1}, p_{s_2} \subset p$  &  $E_{p_{s_1}} \cup E_{p_{s_2}} = E_p$  do
21   if not GLogue.contains( $p_{s_1}$ ) then
22     updateFromPattern ( $p_{s_1}$ , GLogue);
23   if not GLogue.contains( $p_{s_2}$ ) then
24     updateFromPattern ( $p_{s_2}$ , GLogue);
25   GLogue.addEdge ( $(p_{s_1}, p), (p_{s_2}, \mathcal{F}(p_{s_2}))$ );
  
```

# Plan Optimizer



—————> Best Execution Plan

—————> Potential Execution Plan

## Algorithm 1: The Plan Optimizer.

```


1 Function PlanOptimizer (GLogue, PatternDesc)
2   Construct a pattern  $p$  from the PatternDesc;
3   Let  $QSet$  organize all induced subgraphs of  $p$  by level;
4   Initialize a PlanMap to record  $\{p : (plan, cost)\}$  with
   patterns in level 1 and 2 pre-computed;
5   for  $3 \leq level \leq |V_p|$  do
6     for  $p \in QSet[level]$  do
7       searchPlan ( $p$ , PlanMap, GLogue);
8   return PlanMap.get( $p$ );

9 Function searchPlan ( $p$ , PlanMap, GLogue)
10  Initialize Plan( $p$ ) and Cost(Plan( $p$ ))  $\leftarrow \infty$ ;
11  for edge = ( $p_{s_1}, p$ )  $\in$  GLogue.getEdges( $p$ ) do
12    ( $plan1, cost1$ )  $\leftarrow$  PlanMap.get( $p_{s_1}$ );
13    if edge is a vertex extension then
14      Compute a new  $plan'$  by merging  $plan1$  and
      Expand( $p_{s_1} \rightarrow p$ );
15    else if edge  $\{(p_{s_2}, \mathcal{F}(p_{s_2}))\}$  is binary join then
16      ( $plan2, cost2$ )  $\leftarrow$  PlanMap.get( $p_{s_2}$ );
17      Compute a new  $plan'$  by merging  $plan1$ ,  $plan2$ 
      and Join( $\{p_{s_1}, p_{s_2}\} \rightarrow p$ );
18    Compute a new  $cost'$  of  $plan'$  by Equation 1;
19    if  $cost' < Cost(Plan(p))$  then
20      Update Plan( $p$ ) as  $plan'$  and the cost as  $cost'$ ;
21  PlanMap.insert( $p$ , (Plan( $p$ ), Cost(Plan( $p$ ))));

```

# Plan Optimizer

- If the pattern is not contained within the GLogue
  - Decompose the pattern step by step
  - Use this formula to estimate the frequencies of patterns not present in the GLogue.
  - $$\mathcal{F}(p) = \text{Avg}_{p_1, p_2} \frac{\mathcal{F}(p_1) \times \mathcal{F}(p_2)}{\mathcal{F}(p_1 \cap p_2)}, p = p_1 \cup p_2$$



# Outline

---

- 1 Background and Motivation
- 2 System Overview
- 3 Feature Highlights
- 4 Evaluation
- 5 Conclusion

# Environment and Datasets

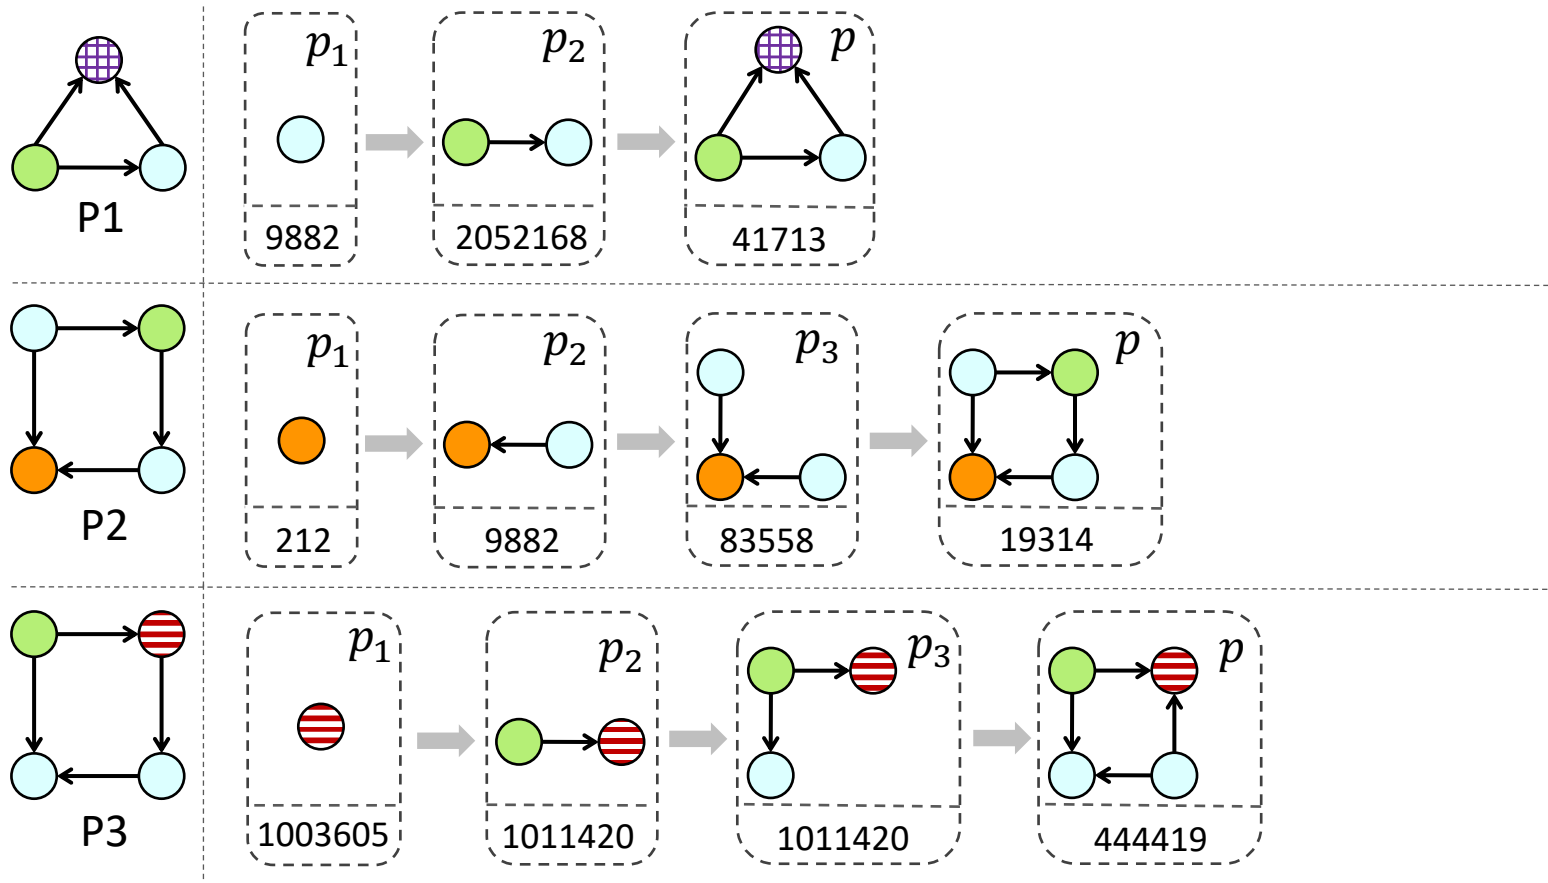
CPU	Memory	Disk	Network	Numbers
2*24-core Intel(R) Xeon(R) Platinum 8163 CPUs at 2.50GHz	512GB	2TB PCIE SSD	EDR 25Gbps InfiniBand network Full Bisection Bandwidth	1 frontend server Up to 16 backend servers





Table: Configurations for servers used in the experiments

Graph	$ V $	$ E $	Size	$\gamma$
$G_1$	3M	17M	1.5GB	100%
$G_{30}$	89M	541M	40GB	1%
$G_{100}$	283M	1,754M	156GB	0.1%
$G_{300}$	817M	5,269M	597GB	0.1%
$G_{1000}$	2,687M	17,789M	1,960GB	0.03%

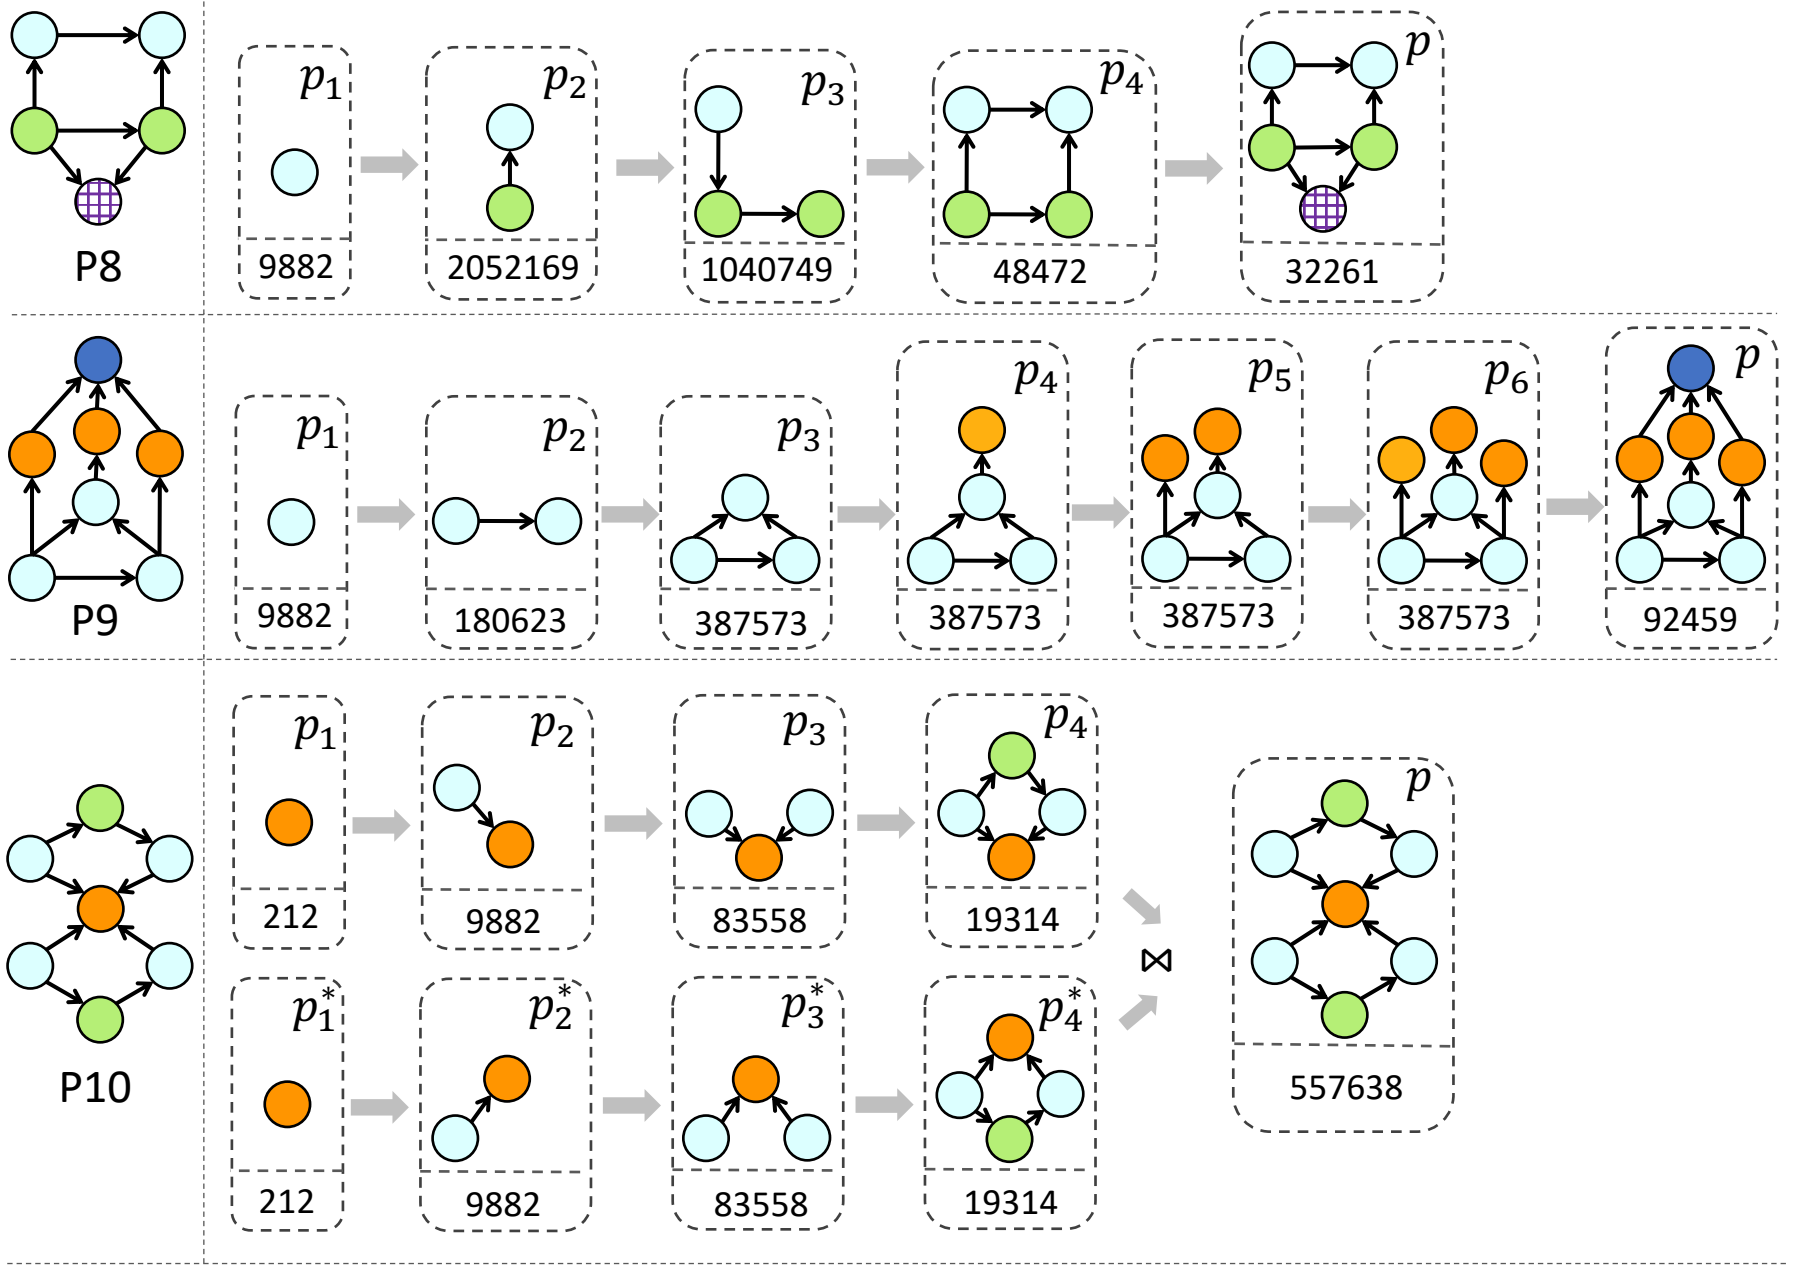
Table: LDBC graphs used in the experiments

# Pattern and Plan









# Experiment1: V.S. Neo4j and TigerGraph




Figure: GLogS V.S. Neo4j




Figure: GLogS V.S. TigerGraph

GLogs achieves 51× and 57% speedup compared with Neo4j and TigerGraph, respectively.

# Experiment2: High-order Statistics

---

	Level = 2	Level = 3	Level = 4
Slow-down(%)	966	245	243
Generation Time(s)	6	55	1664
Memory Usage(GB)	2	3	105
# Patterns	34	248	4164

Table: The effectiveness of high-order statistics

# Experiment3: Sparsification




Figure: Stratified V.S. Uniform Sparsification

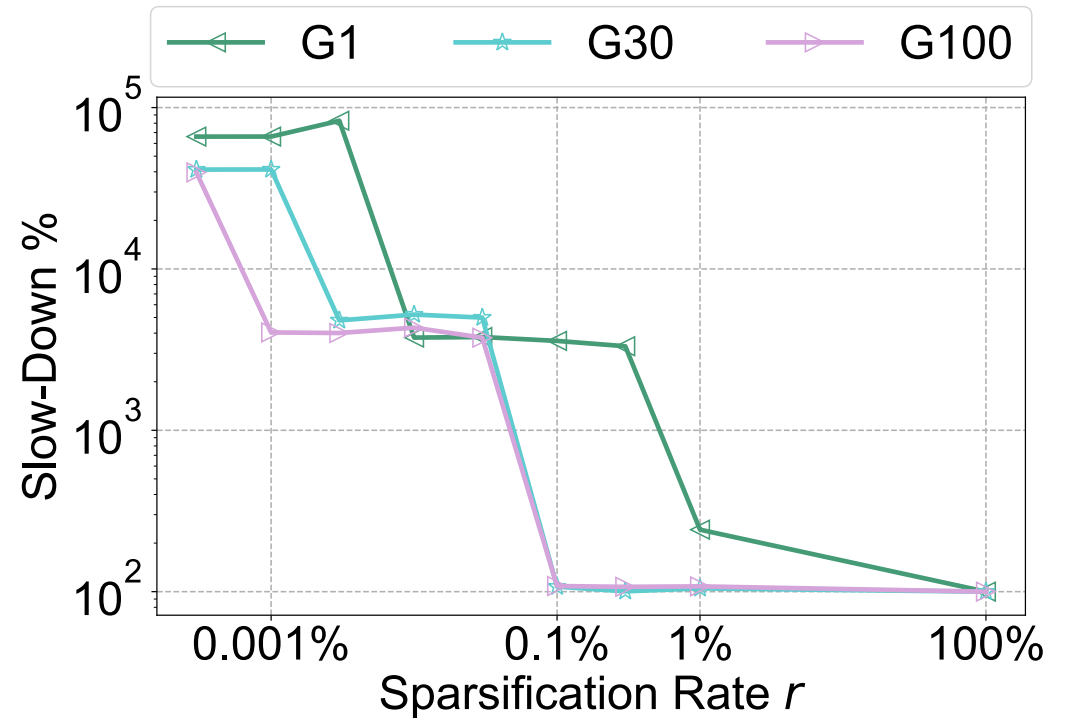


Figure: Sparsification of graphs with different scale

Uniform sparsification requires a higher rate to achieve same performance as stratified sparsification. Small graph requires a higher rate to have a good plan.

# Experiment4: Scalability




Figure: Scale-out experiments for group1 queries




Figure: Scale-out experiments for group2 queries

Most queries scale well, with up to  $15\times$  (average  $6\times$ ) performance gain from one machine to 16.

# Experiment4: Scalability




Figure: Scale-up experiments for group1 queries




Figure: Scale-up experiments for group2 queries

Up to  $23\times$  (average of  $10\times$ ) when increasing the number of threads from 1 to 32  
The scalability of  $p_2$  and  $p_{10}$  is insignificant due to a skewed workload.

# Experiment4: Scalability




Figure: Data-scale experiments for group1 queries




Figure: Data-scale experiments for group2 queries

As the graphs become larger, most queries demonstrate an almost linear trend in performance degradation. For  $p_6$ , its performance degrades by  $3\times$  from  $G_{30}$  to  $G_{1000}$ , due to its short duration and limited graph exploration. For  $p_{10}$ , its performance degrades by  $100\times$  from  $G_{30}$  to  $G_{1000}$ , due to the large cost of join operation.

# Outline

---

1

Background and Motivation

2

System Overview

3

Feature Highlights

4

Evaluation

5

Conclusion



# Conclusion

---

- GLogS system solves iGPM, meeting the requirements of:
  - usability
  - performance
  - scalability
- Allows user interactively submit GPM queries
- Supports automatic optimization for arbitrary GPM queries:
  - adopt high-order statistics
  - guarantee worst-case optimal
- Proposes a novel graph-based structure GLogue:
  - maintain the high-order statistics of the graph
- Capable of deployed in a large cluster to handle real-life large graphs

# Related Work

---

- GPC: A Pattern Calculus for Property Graphs
  - <https://dl.acm.org/doi/10.1145/3584372.3588662>
- GQL Standard
  - <https://www.gqlstandards.org/>

GPM Queries are gradually being standardized and becoming increasingly important!

GLogS is continuously evolving with the graph community and standards!



Scan to visit GraphScope Github Repo:  
<https://github.com/alibaba/GraphScope>

# THANK YOU