GraphScope Flex: A Graph Computing Stack with LEGO-Like Modularity

GraphScope Team
DAMO Academy, Alibaba Group.

I. Introduction

Graph applications in real life have diverse workloads, programming interfaces, and storage formats. GraphScope is a one-stop solution that addresses these variations. However, challenges remain due to:

– Various deployment modes are needed, such as an offline analytical pipeline for low latency, online services requiring high throughput, or a learning task benefiting from heterogeneous hardware.

– An all-inclusive solution may not be the best option as it could lead to increased resource and cost requirements.

GraphScope Flex is the next iteration of GraphScope. It features a modular design that reduces resource and cost requirements while providing a seamless, user-friendly experience for flexible deployment.

II. Architecture & Design

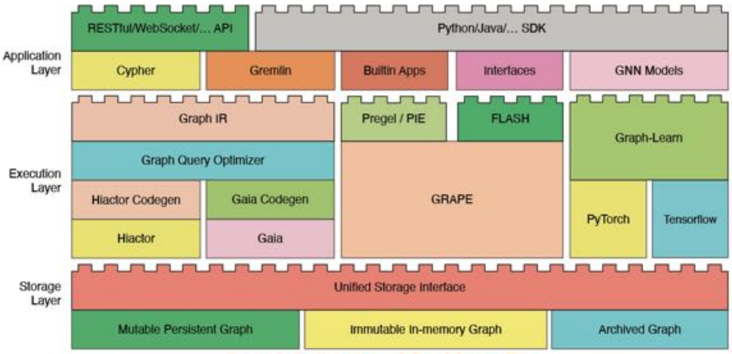

Figure 1. Architecture of GraphScope Flex

The GraphScope Flex stack (Fig. 1), consists of multiple components that users can combine like LEGO bricks to customize their graph computing deployments. The components are classified into three layers:

Application Layer, which includes pre-built libraries of algorithms and GNN models, as well as SDKs and APIs;

Execution Layer, which comprises multiple engines that are specialized for their respective domains;

Storage Layer, which establishes a uniform interface for managing graph data across various storage backends.

III. Case Study

Recap GraphScope, for one-stop graph computing

E-commerce platforms use graph computations, e.g., mining, LPA, and GNN techniques, to detect suspicious users and transactions.

Without GraphScope, it requires a workflow arcoss many systems to deal this case, while GraphScope provides a one-stop and efficient solution for a wide range of graph computations at scale.

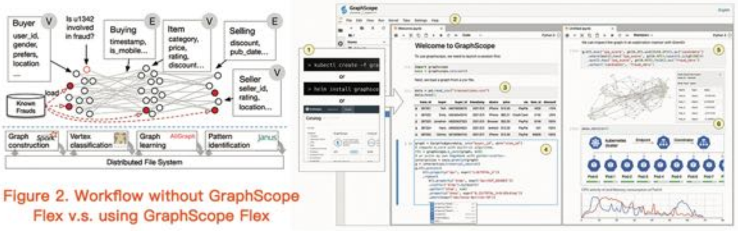

```
> kubectl create -f gra
> helm install graphsco
```

Figure 2. Workflow without GraphScope Flex v.s. using GraphScope Flex

Flexiable and Composable Deployments of GraphScope Flex

GraphScope Flex offers a no-code WebUI to configure deployment. It generates a customized .yaml file for selected components which can be parsed by helm to launch a K8s cluster.

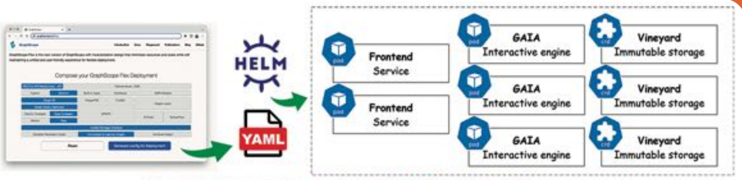
Compose your GraphScope Flex Deployment
HELM
YAML
Frontend Service
Frontend Service
GAIA Interactive engine
GAIA Interactive engine
GAIA Interactive engine
Vineyard Immutable storage
Vineyard Immutable storage
Vineyard Immutable storage
pod
crd

Figure 3. Composable deployment of GraphScope Flex

Case 1: For online BI analysis

BI analysis is for analysts who interactively analyze data in a WebUI. While high concurrency is unlikely, low latency for complex queries is crucial. GraphScope Flex compiles Cypher and Gremlin queries into a unified intermediate representation (IR) and optimizes it using a universal query optimizer and catalog module. The optimized IR is passed to Gaia Codegen and executed on Gaia, a distributed dataflow engine that reduces query latency through data parallelism. Graph data is accessed from a mutable persistent storage via a unified interface.

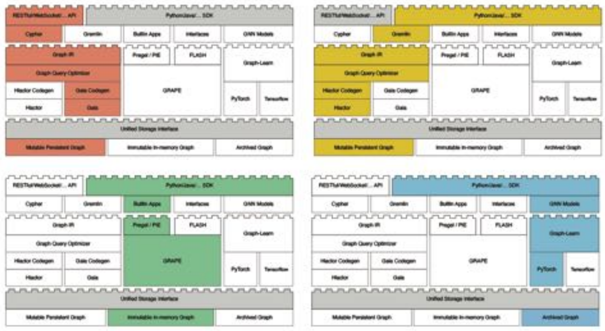

Figure 4. LEGO view of GraphScope

Case 2: For high QPS queries

For high QPS scenarios like recommendation or searching, GraphScope Flex can be deployed with a different component set. The compiler generates an optimized query plan and Hiactor Codegen produces a physical plan tailored for Hiactor, a high-performance and concurrent actor framework for OLTP-like queries.

Case 3: For offline graph analytics

GraphScope Flex is an efficient and user-friendly platform for performing graph analytics. It offers built-in algorithms, as well as interfaces for developing customized algorithms. The runtime, based on GRAPE, is fragment-centric and extensible, supporting multiple programming models like FLASH, PIE, and Pregel. Sequential algorithms can be easily parallelized or incrementalized using the Ingress component. To achieve high performance, an in-memory graph store is deployed in this stack.

Case 4: For graph learning tasks

GraphScope Flex's GNN framework supports billion-scale graphs in industrial scenarios. It provides GNN model development paradigms, example models, and the flexibility to choose between TensorFlow or PyTorch as the training backend. Furthermore, the framework employs decoupled sampling and training processes, which can be independently scaled for optimal end-to-end throughput, providing superior performance.

IV. Performance

GraphScope Flex deployments are flexible yet highly performant. It outperformed other systems from 2X to magnitudes in LDBC SNB and Graphalytics Benchmarks.

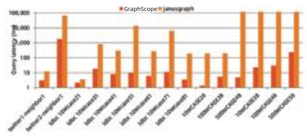
GraphScope
Janusgraph
Query Latency (ms)

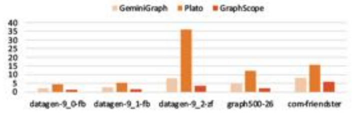
GeminiGraph
Plato
GraphScope
datagen-9_0-fb
datagen-9_1-fb
datagen-9_2-zf
graph500-26
com-friendster