EYWA: a Distributed Graph Engine in the Huawei MIND Platform

Yinglong Xia
Huawei Research America
02/09/2017

About Huawei

170,000
employees

170
countries

$75B
revenue

14
regional HQ

16
R&D centers

32%
YoY growth

HUAWEI

ICT Enterprise Business

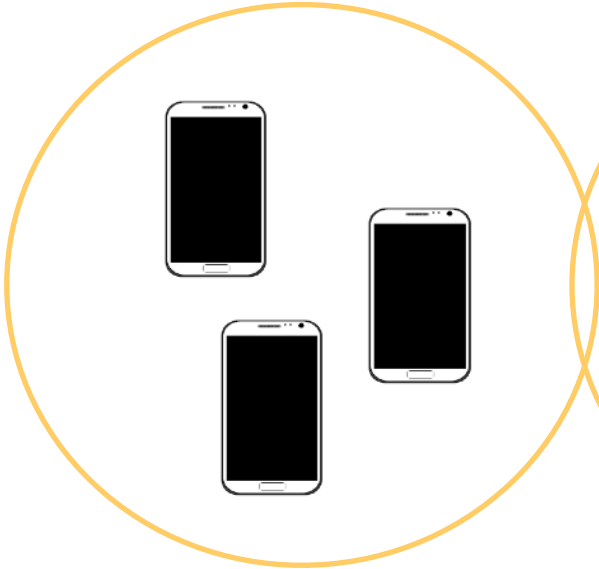

Increase user experience for HW's consumers

Maintain large scale systems and predict/ diagnose faults

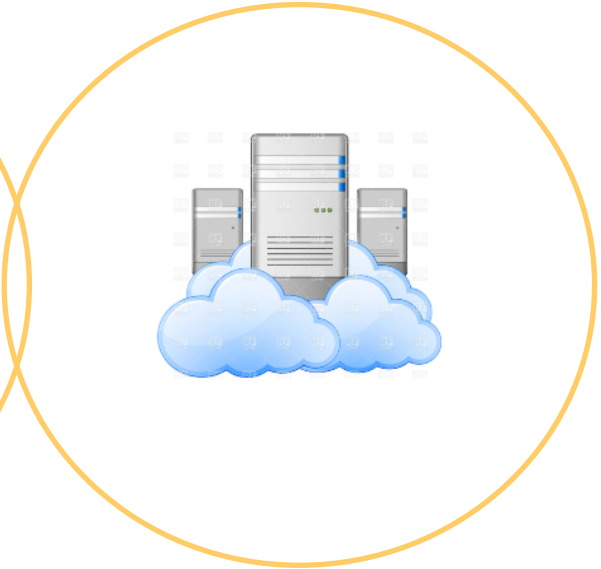

Powerful back-end datacenter for cloud services

HUAWEI

Glance at Huawei MIND Platform

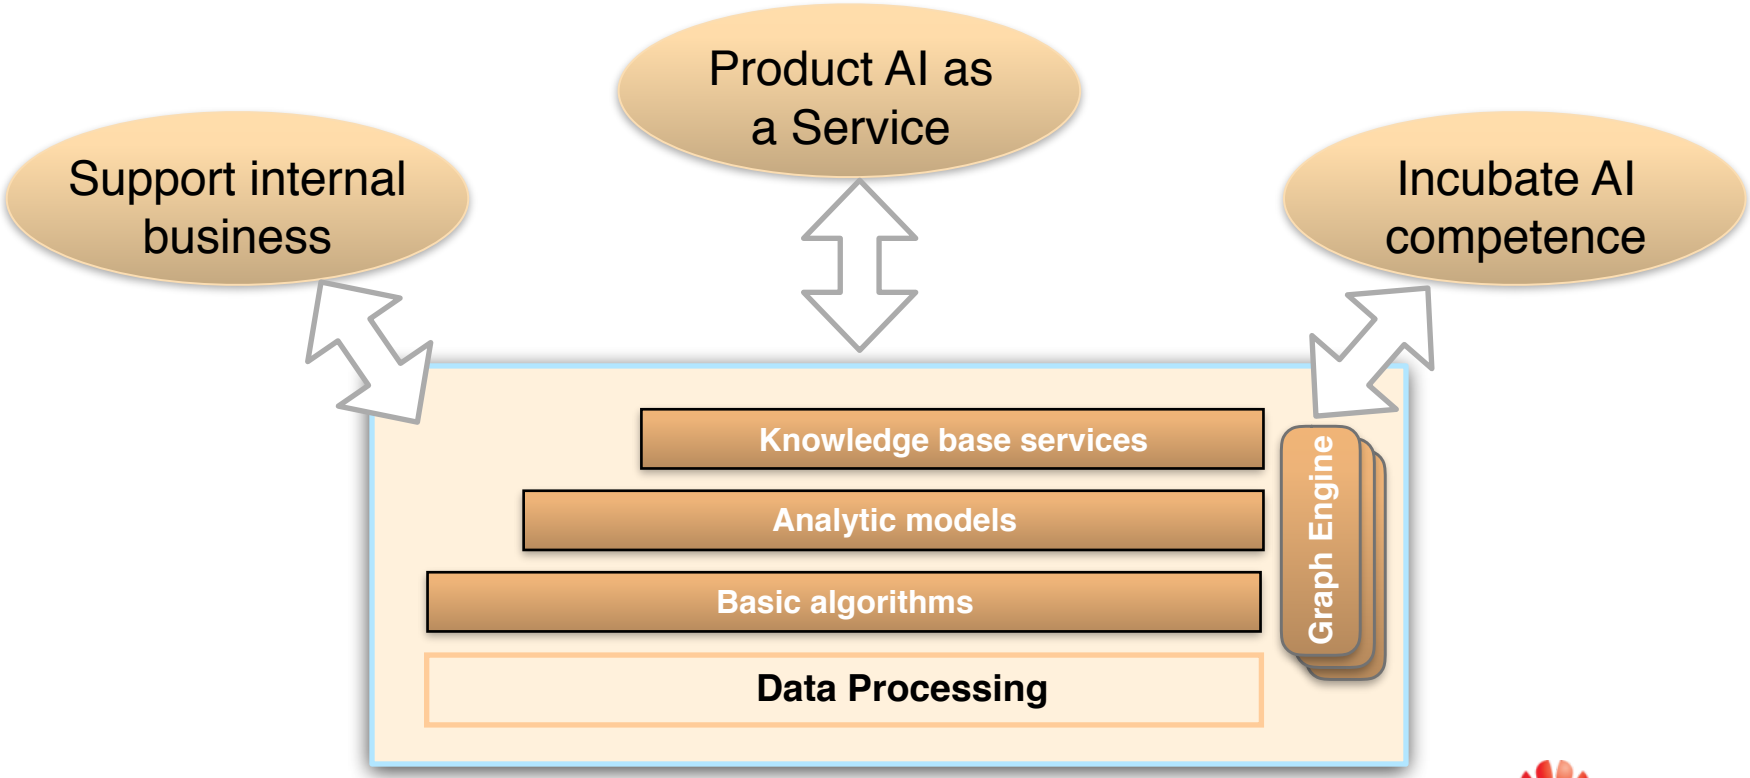

HUAWEI

Graphs in AI: Intelligent Service

Use Case:

AI assists Technical Services

Private Knowledge Graph

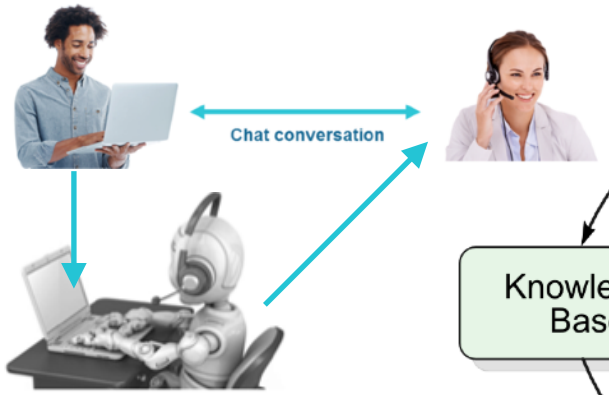

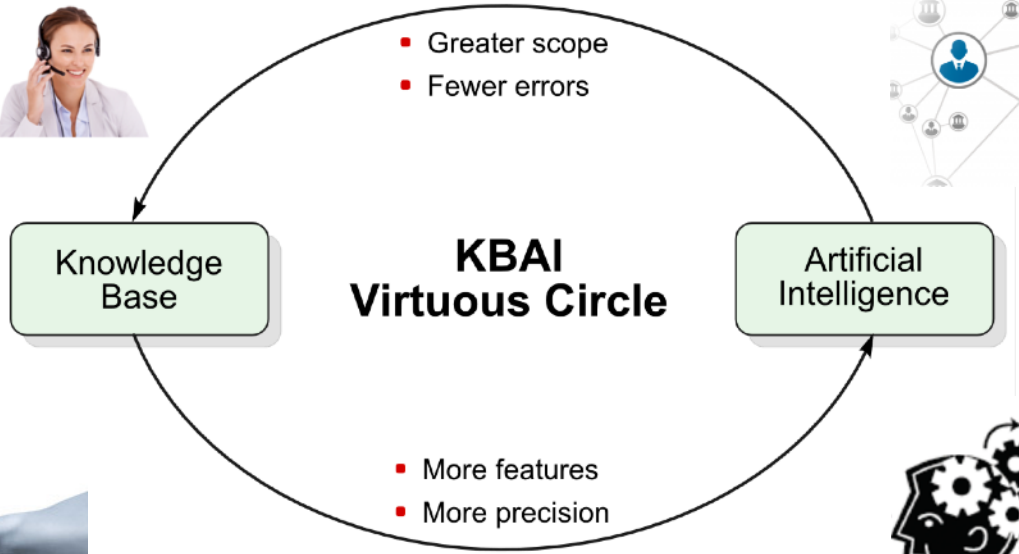

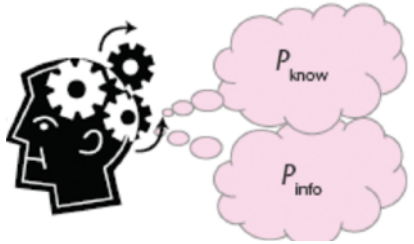

Human-Machine Collaboration

Inference & Reasoning

Graphs in AI: Knowledge based Search

Use Case:

AI helps stuff productivity

Domain specific questions

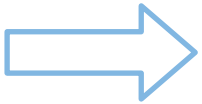

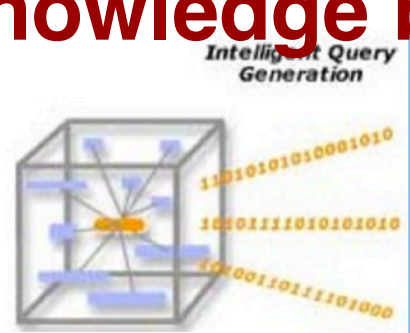

Private Knowledge Graph

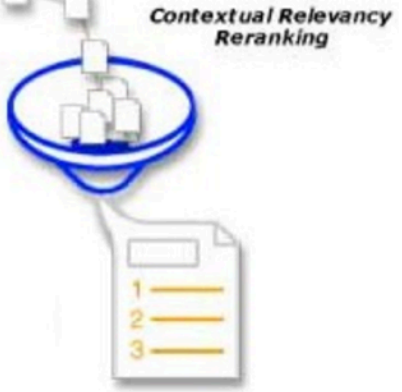

Search results

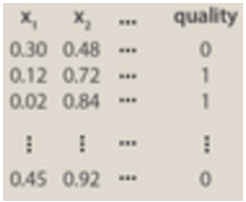

x_1	x_2	...	quality
0.30	0.48	...	0
0.12	0.72	...	1
0.02	0.84	...	1
⋮	⋮	...	⋮
0.45	0.92	...	0

Raw data & Evidence

Models & Rules

Business procedures

HUAWEI

Graphs in AI: Private Knowledge Graph

Use Case:

Domain knowledge persistency

Unstructured/ Semistructured content

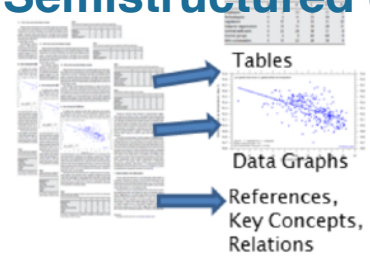

External domain data

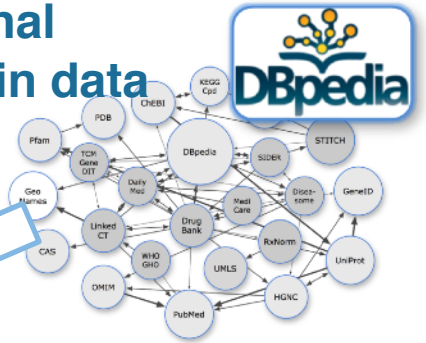

Domain ontology

Enterprise & User data

PMML
Predictive Model Markup Language

Models, Dataset, etc.

HUAWEI

Reasoning w/ Uncertainty: Graphical Model

Use Case:

AI for (network) automatic DevOps

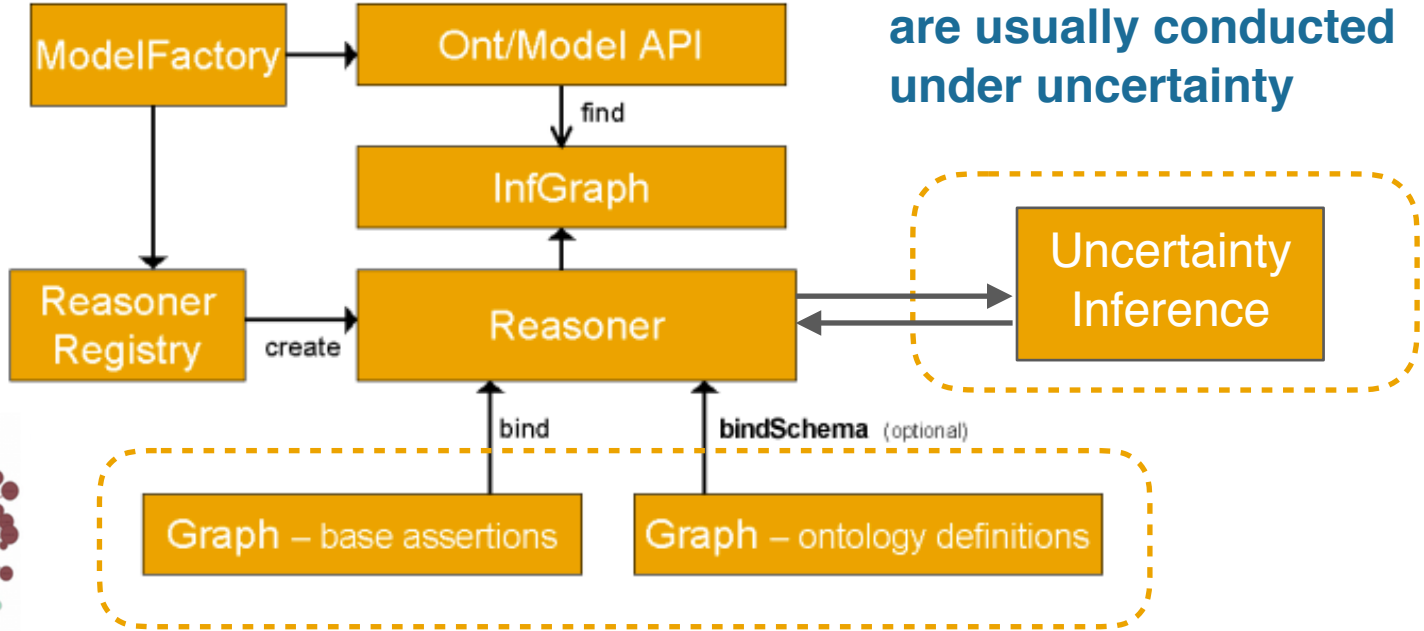

System diagnosis or root causality analysis are usually conducted under uncertainty

Extended Jena reasoner and rules engine

HUAWEI

Graph Acceleration

Use Case:

Real-time interactive exploration

Acceleration is critical for some analytics in streaming mode and/or interactive mode

- Intelligent Search & knowledge inference
- Machine Learning & Deep Learning
- Graph-based reasoning
- Large scale optimization

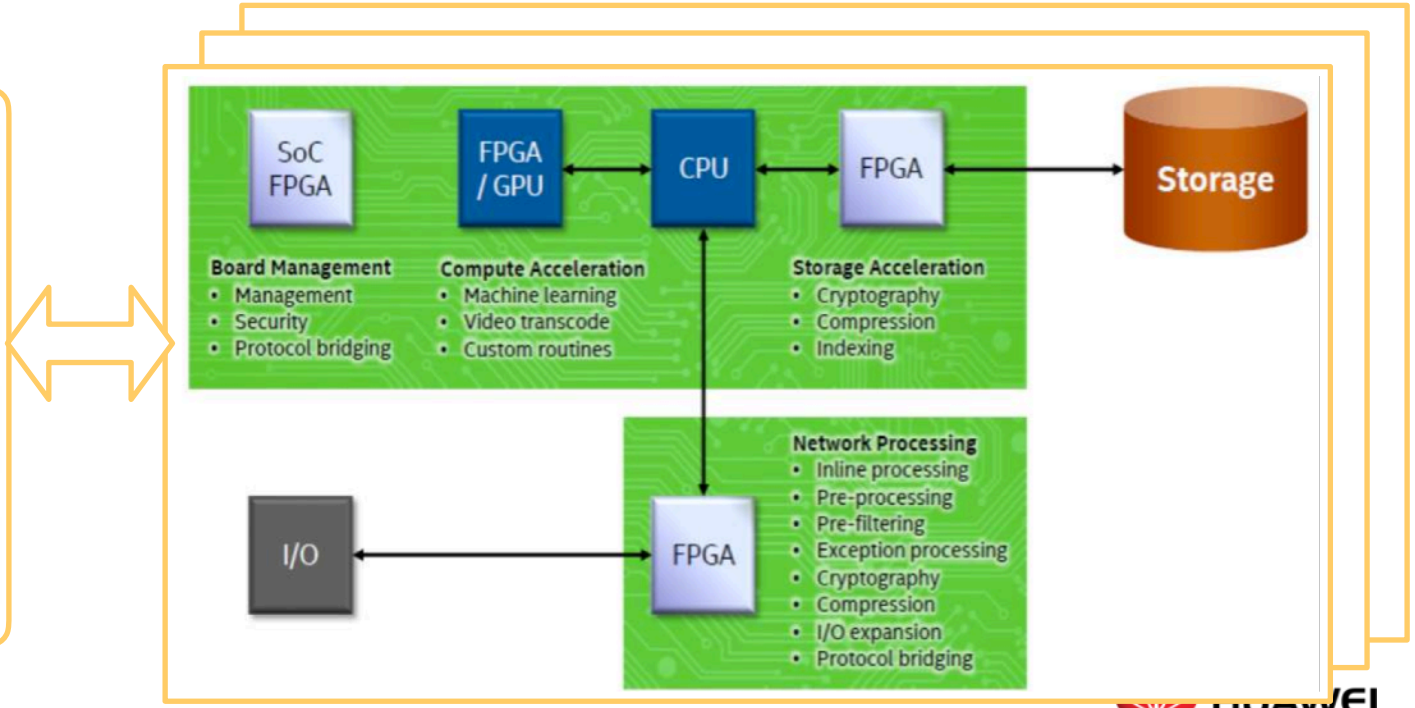

EYWA: Graph Engine for MIND

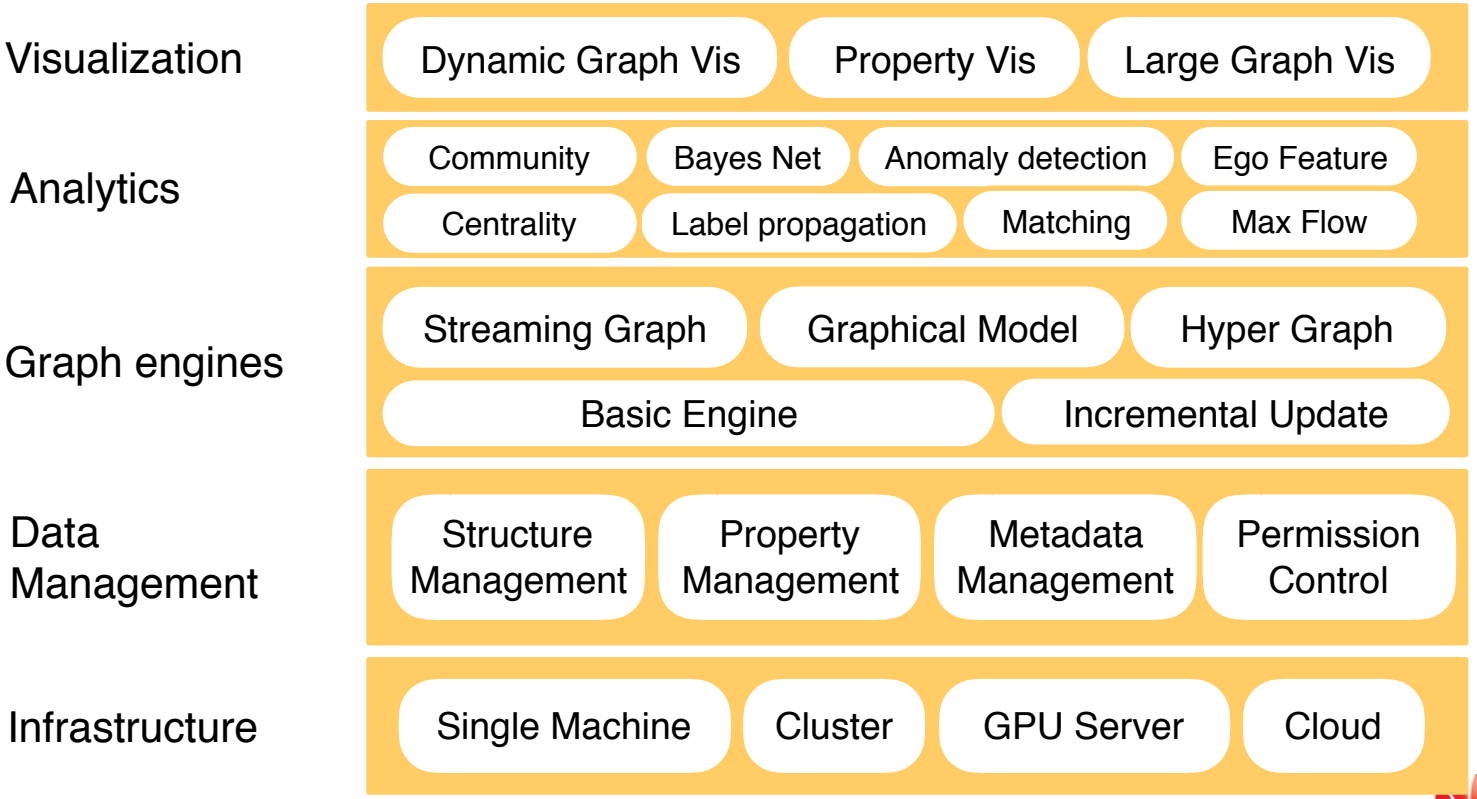

HUAWEI

Trade-off Between Analytics and Query

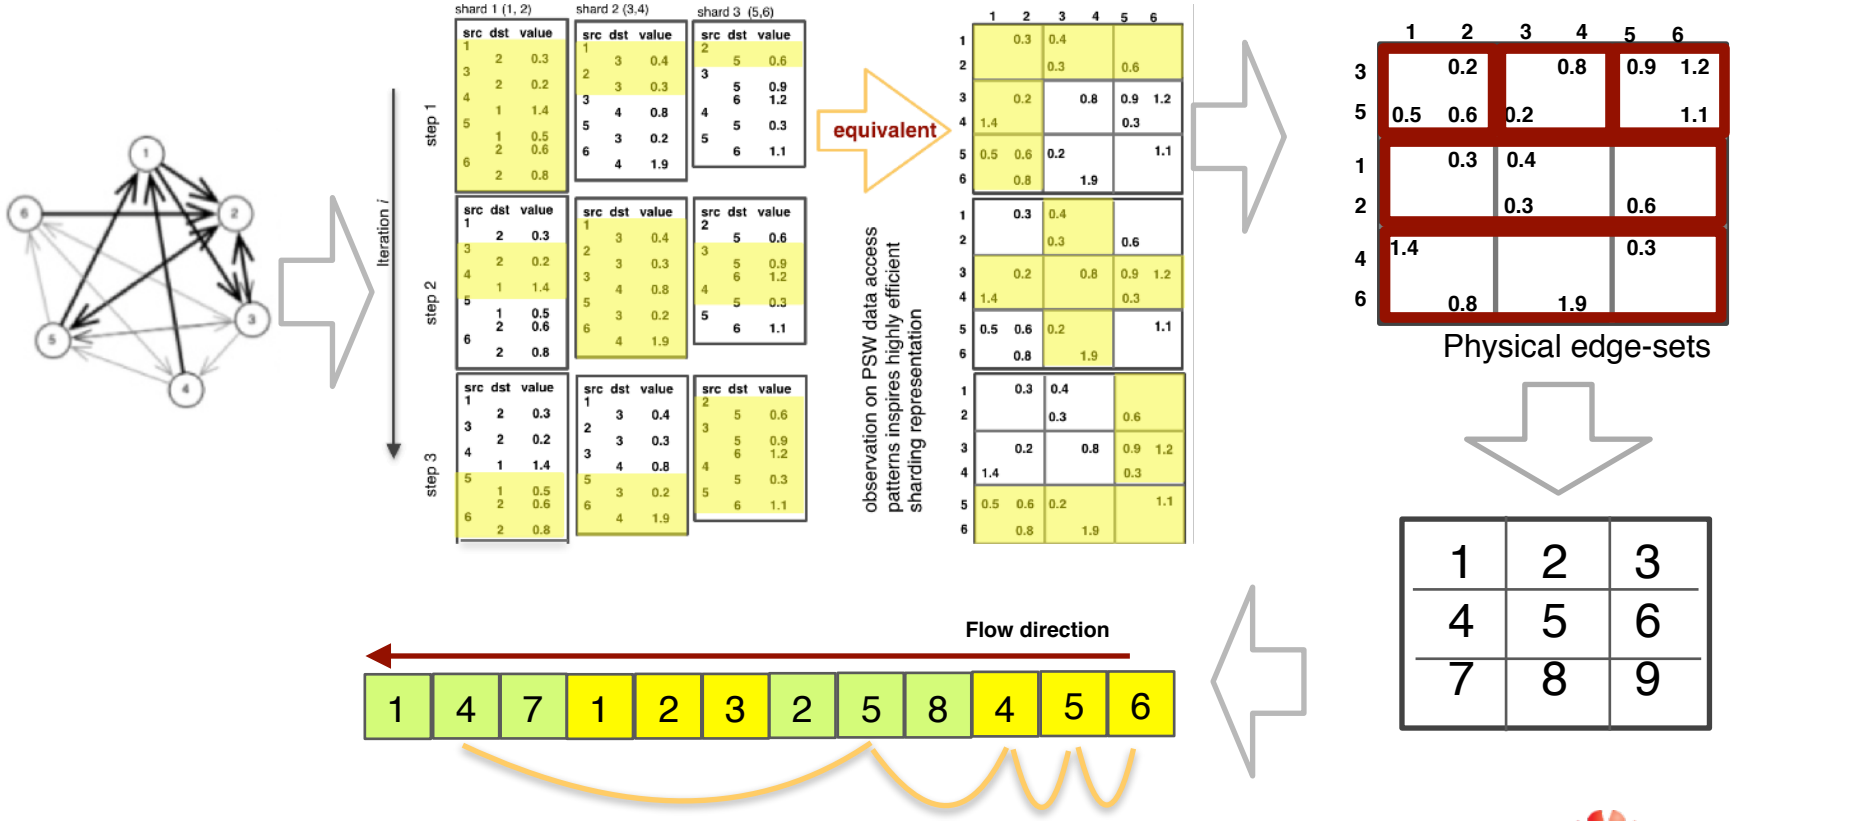

HUAWEI

Scheduler/Prefetch of Relaxed BSP

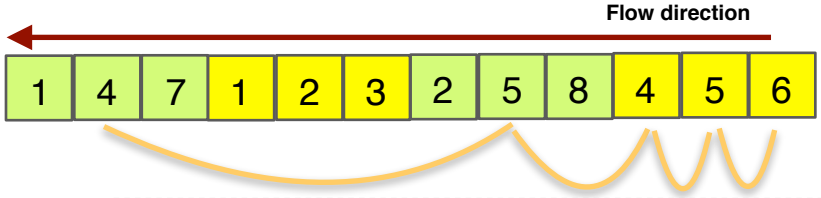

load shard | load PSW | local computation | queue updates | send updates | recv updates | update/storePSW

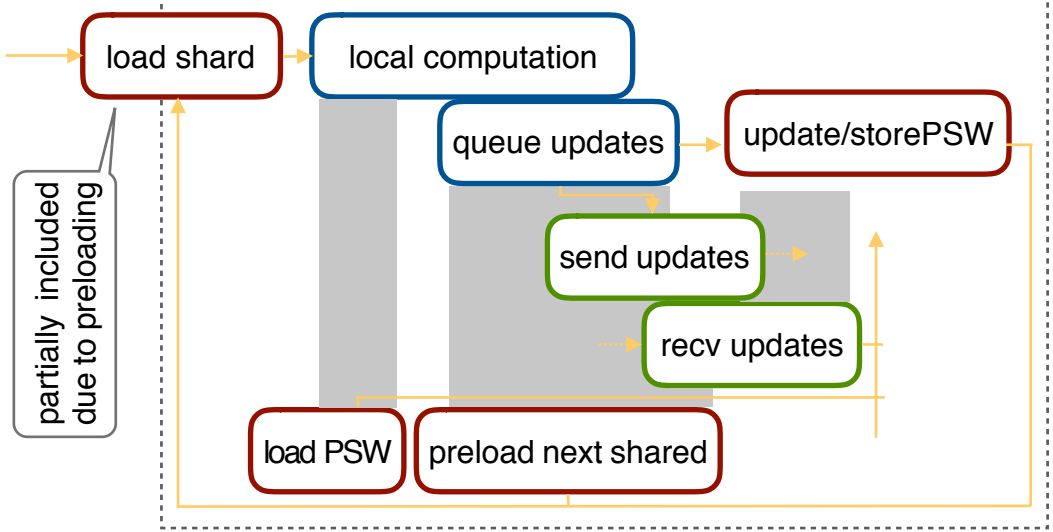

speedup due to overlapping

Dash box includes all tasks within an iteration

task type: disk IO | CPU | network IO

Support Very Large Scale Graphs

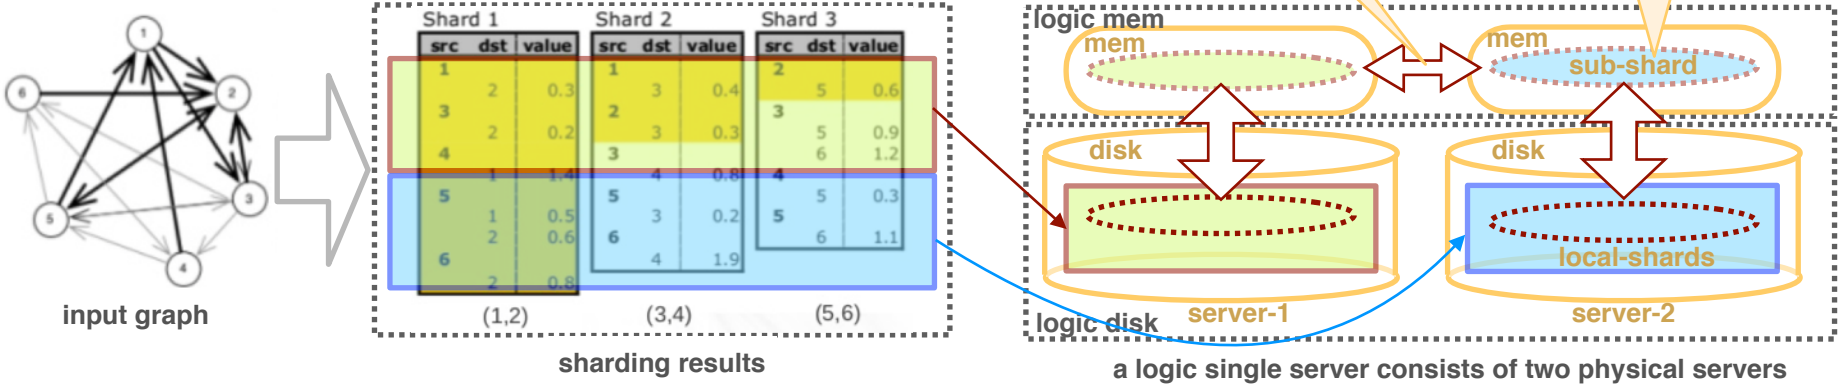

HUAWEI

How to Query & Analyze the Graphs

- **Business users**
  - **Web-based explorer**
  - **Interactive visualization front-end**
- **(Graph) technical users —> Bi-lingual**
  - **Cypher (w/ additions for active query)**
  - **Gremlin**
  - **(LDBC's effort on query language)**
- **(Non-graph) technical users**
  - **RESTful APIs**
  - **Service on Cloud**

Experiments

Name	Vertices	Edges
Kronecker (22)	4.1M	34.1M
Kronecker (24)	16.7M	165.2M
Kronecker (26)	67.1M	799.8M
LDBC-1000K	1M	28.8M
Twitter-2010	41.7M	1.4B

	App.	Prep.	Load+Reconstr.	Comp.
GraphChi	PageRank	1184.64	2535.76	1030.76
	BFS	958.991	197.6692	82.5688
	SSSP	966.342	217.1342	82.9308
Edge-Set	PageRank	786.199	431.693	557.81
	BFS	703.41	46.1713	11.7145
	SSSP	712.216	41.5668	25.1802

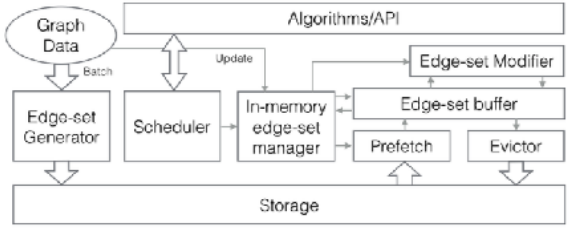

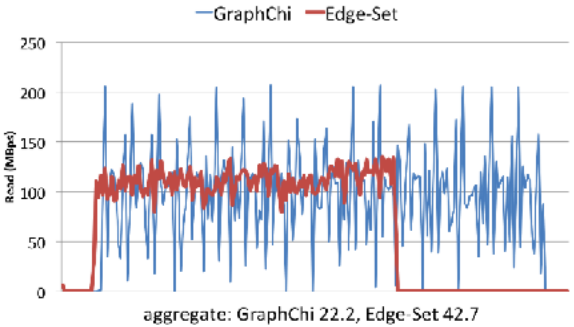

Disk read bandwidth over time by GraphChi and EdgesSet. Edge-Set showed up to 2x aggregate bandwidth and more constant IO usage.

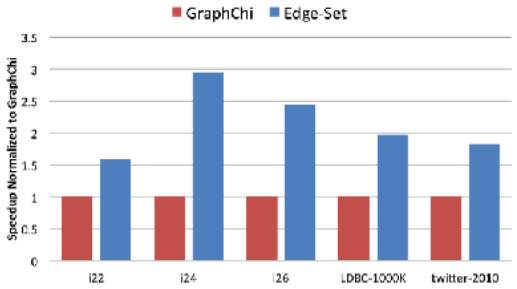

Performance improvement of SSSP against GraphChi including data ingestion

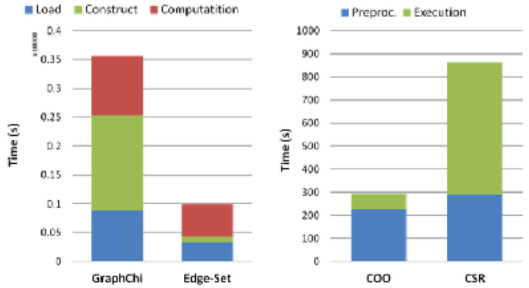

Execution time breakdown on Pagerank running Twitter-2010

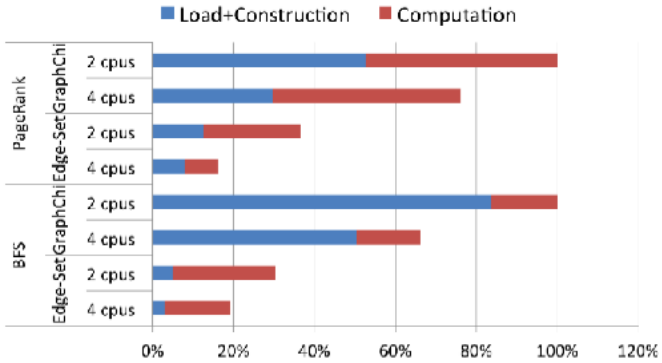

Execution time breakdown on Pagerank running Twitter-2010

HUAWEI