When Graph Meets Big Data: Opportunities and Challenges

Yinglong Xia
Huawei Research America
11/13/2016

Introduction

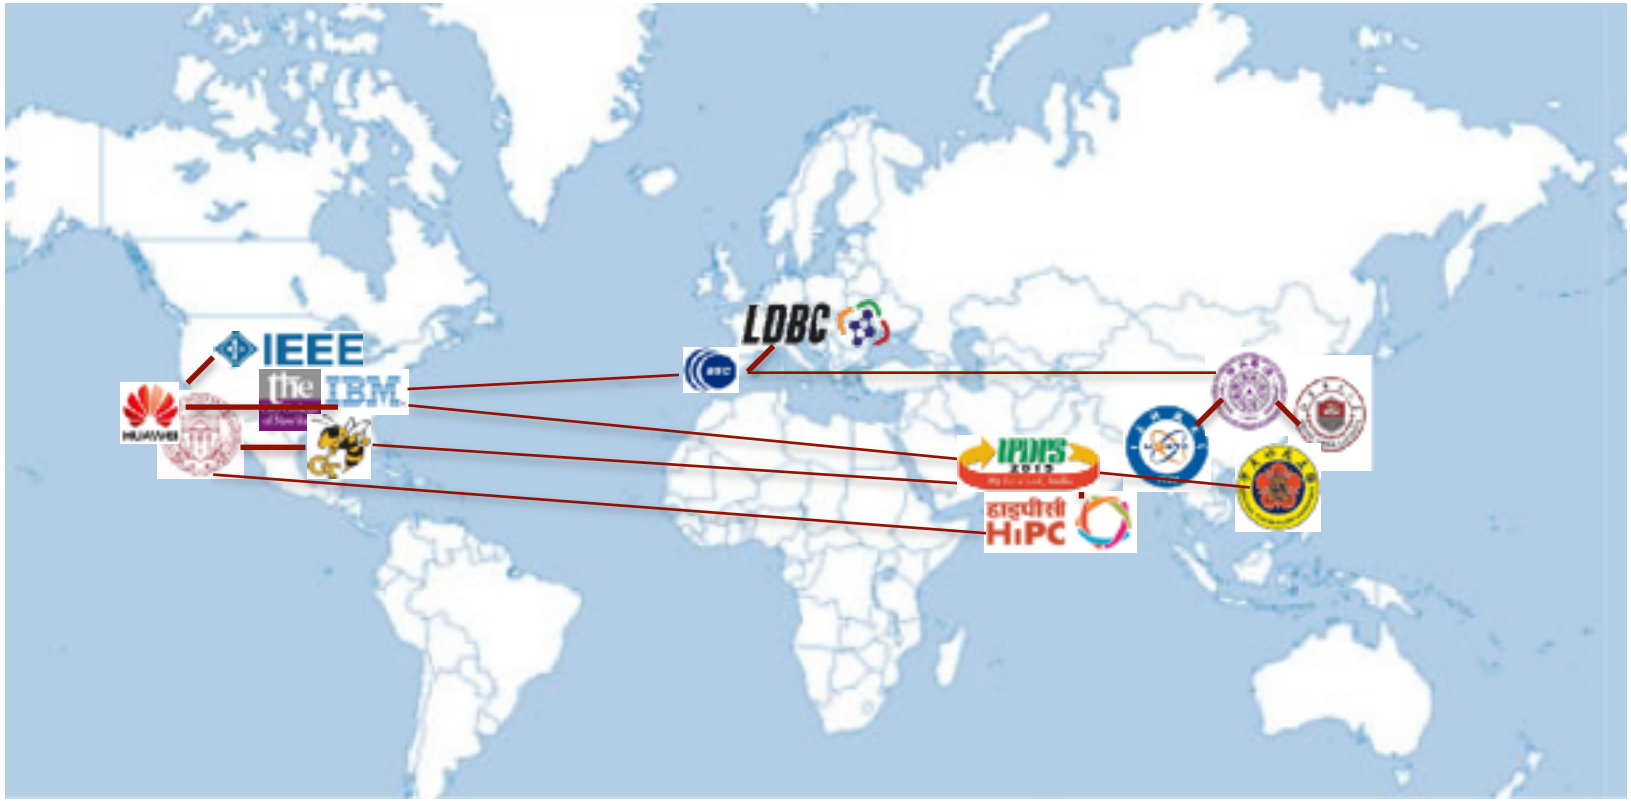

HUAWEI

Recent Growth

Huawei has business in over **170 countries**, with **150,000 employees**, approximately **70,000** of which are engaged in Research & development. Huawei operates a global network of **14 regional headquarters, 16 R&D Centers, 28 Innovation Centers** jointly operated with customers, and 45 Training Centers.

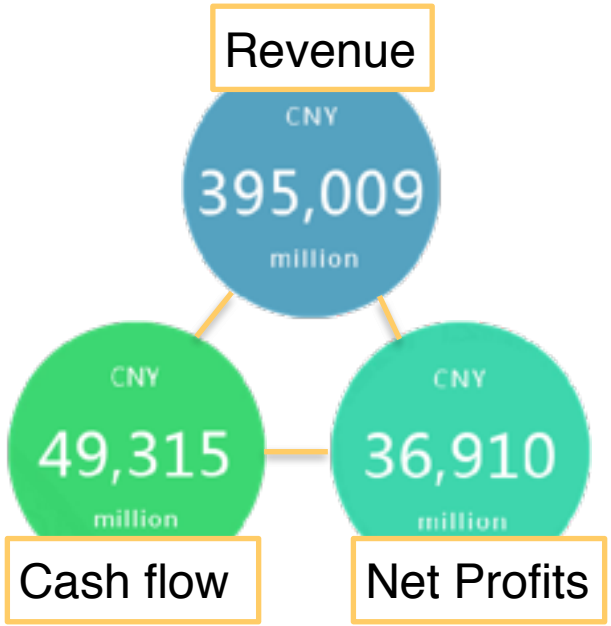

CNY Million	2015	2014	YoY
Carrier Business	232,307	191,381	21.4%
Enterprise Business	27,609	19,201	43.8%
Consumer Business	129,128	74,688	72.9%
Others	5,965	2,927	103.8%
Total	395,009	288,197	37.1%

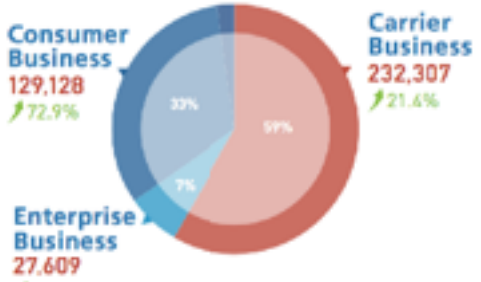

CNY Million	2015	2014	YoY
China	167,690	108,674	54.3%
EMEA	128,016	100,674	27.2%
Asia Pacific	50,527	42,409	19.1%
Americas	38,976	30,844	26.4%
Others	9,800	5,596	75.1%
Total	395,009	288,197	37.1%

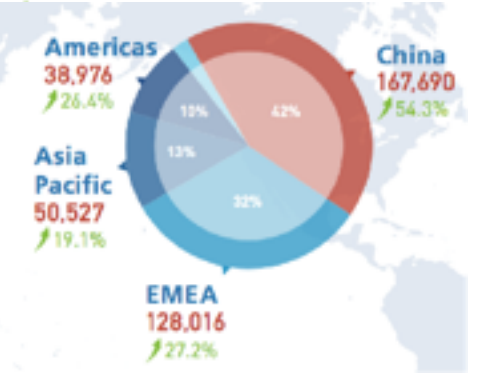

http://www.huawei.com/en/about-huawei

Graph Analytics Basics

Brief History

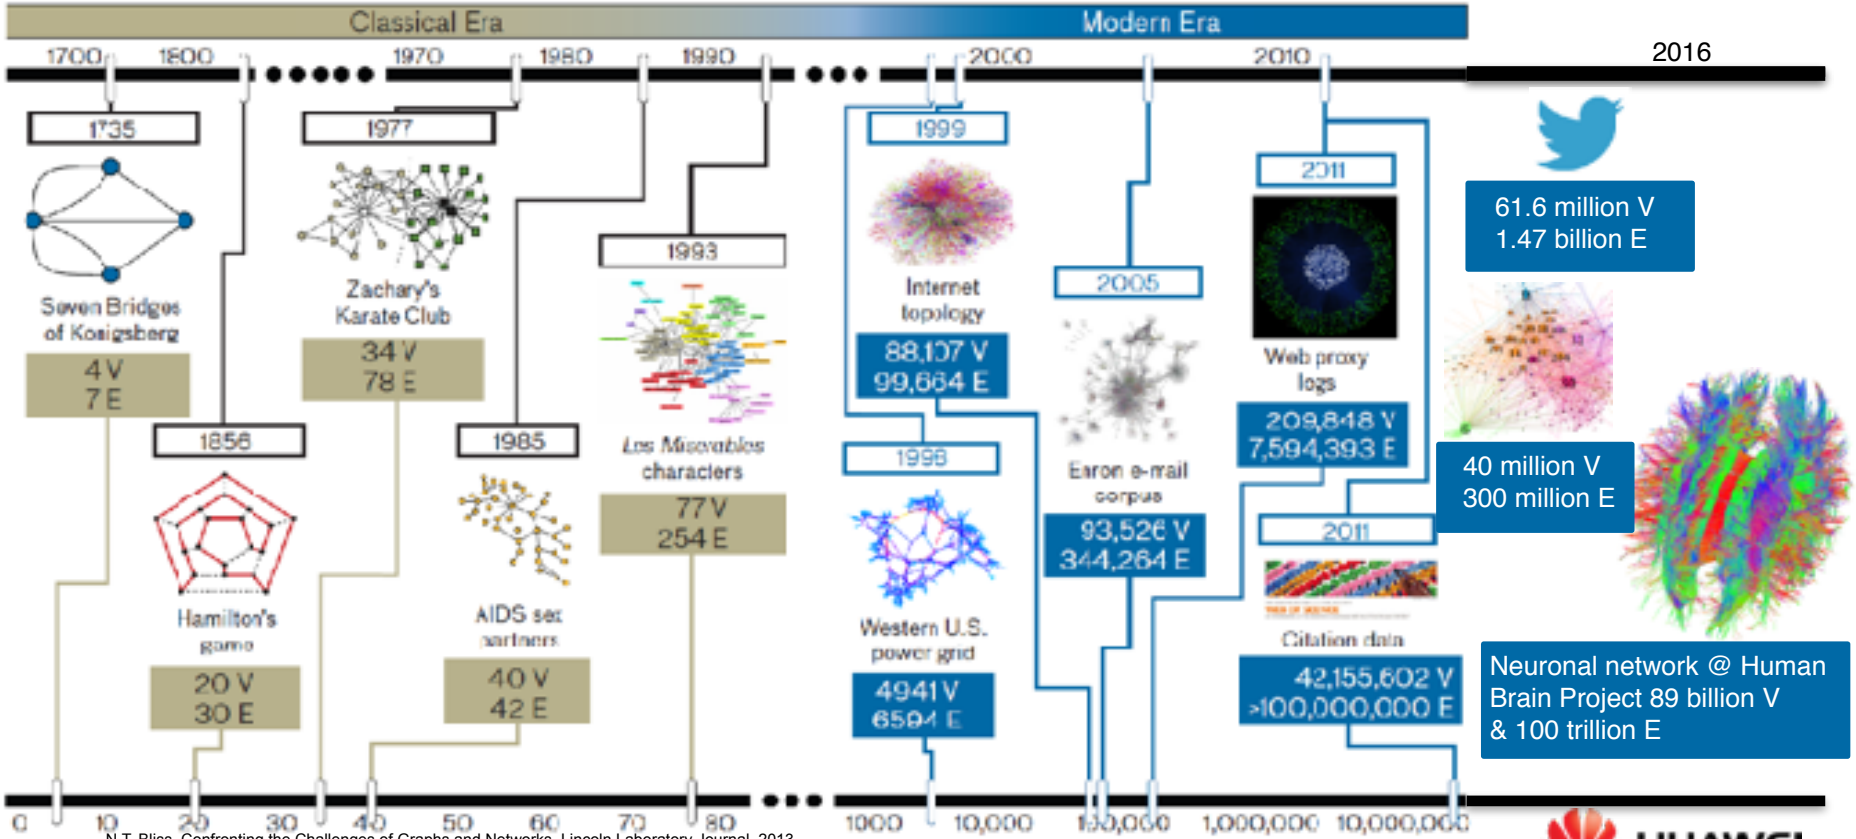

N.T. Bliss, Confronting the Challenges of Graphs and Networks, Lincoln Laboratory Journal, 2013

HUAWEI

Complex Network Analysis

Real world **complex networks** include WWW, Social Network, Biological network, Citation Network, Power Grid, Food Web, Metabolic network, etc.

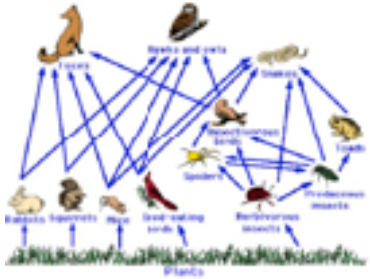

Import properties/metrics:

- Small-world effect
- Betweenness
- Eccentricity/**Centrality**
- Transitivity
- Resilience
- Community structure
- Clustering coefficient
- Matching index

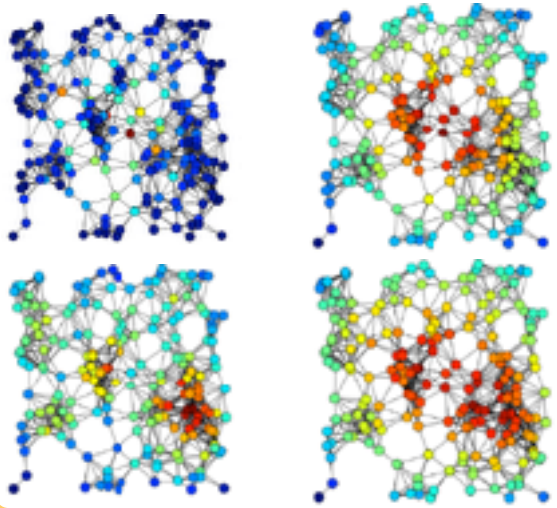

Complex network models:

- Poisson random graph
  - degree~Poisson
  - Small world effect
- Watts and Strogatz graph
  - Transitivity
  - Small world effect
- Barabasi and Albert graph
  - Small world
  - Power law

Diversity in Graph Technology

Dynamic graph helps analyze the spatial and temporal influence over the entities in the network

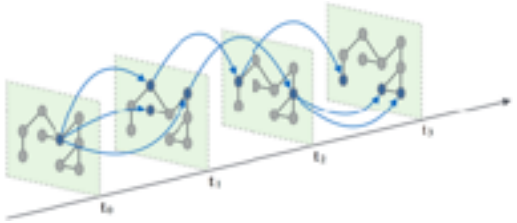

RDF graph enables knowledge inference over linked data

Streaming graph monitors sentiment propagation over time and how the graph structure can impact

Graph technology leads to rich analytic abilities

Property graph is widely used as a data storage model to manage the properties of entities as well as the interconnections

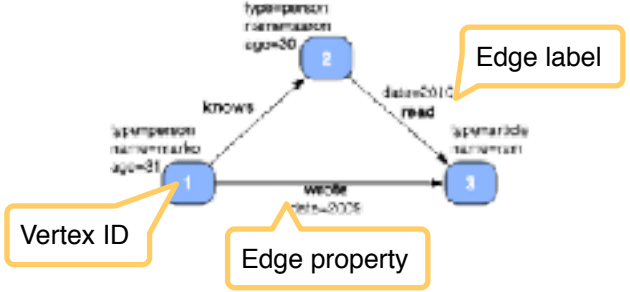

Graphical models leverages statistics to inference latent factors in a complex system

Industrial Use Cases

Financial Risk Management

Motivation

- 56 million small business owners in China
- 1/3 want to finance their ventures
- Weight to 24.3% of the GDP
- 16.3% received loads from banks
- 3/4 citizens in China has no credit history

Identifying Fraud by Graph

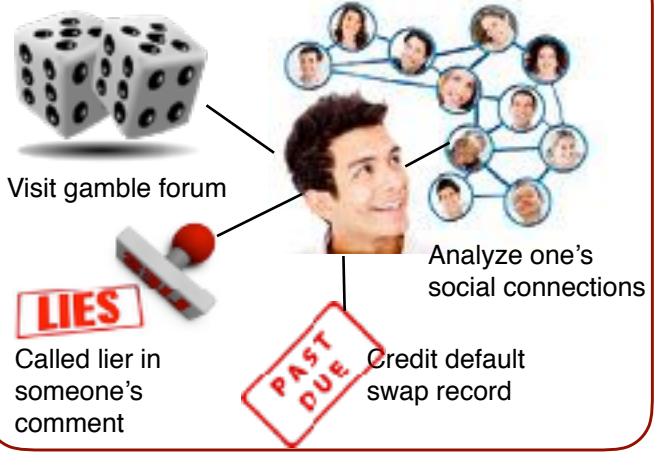

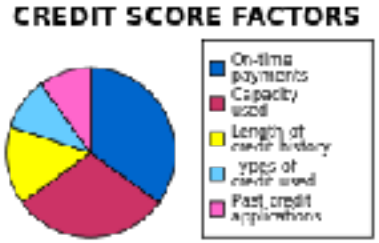

Traditional approach

- Strong variables
- Weak variables
- Individual regression
- collective regression

Emerging approach

- Customer profiling
- Precise engagement
- Anti-fraud control
- Credit scoring in realtime
- Post-load management
- Lost-customer engage

Credit Scoring

Traditional credit scoring organizations

Emerging credit scoring organizations

LendingClub vs PROSPER

Auditing

Samsung Group cross-ownership structure in Gephi network visualization

Auditing such labyrinthian beasts is enough to bring a grown auditor to tears. This is where network 'graph analytics' can be invaluable

Public Security

Use Case -1

IBM Insider Thread Solution

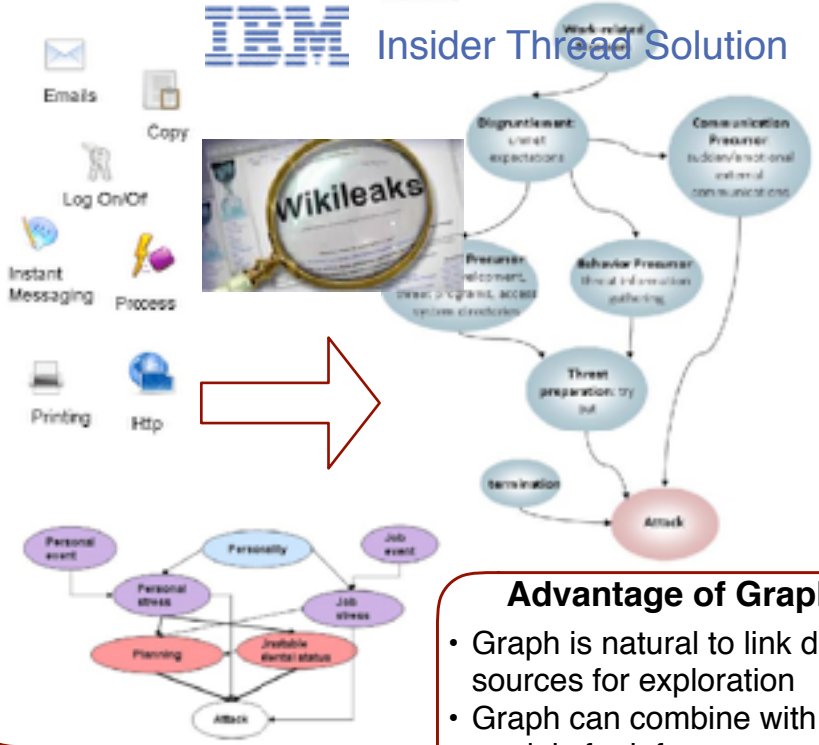

Use Case -2

Palantir

- Knowledge graph
- Heterogeneous data
- Information integration
- Complex search
- Security analysis

Advantage of Graph for Security

- Graph is natural to link data from different sources for exploration
- Graph can combine with probabilistic models for inference
- Graph helps find anomalous behaviors

HUAWEI

Telecom Fraud

In May 2016, Communications Fraud Control Association (**CFCA**) and the Forum for International Irregular Network Access (**FIINA**), operators ranging from AT&T, Vodafone, Korea Telecom to Orange and Deutsche Telekom shed light on how old and newer forms of fraud are detected and combatted

Step	Call Flow	Money Flow
1	Fraudster steals credentials of the service provider's subscribers and registers with the service provider's softswitch using those stolen credentials.	
2	Fraudster's subscribers place a call	Fraudster's subscribers pay for service
3	(A) Fraudster sends INVITE to service provider's softswitch	(B) Service provider's subscribers are billed for the call that was placed using their stolen credentials, but are unlikely to pay for the fraudulent calls.
4	Service provider routes calls to their destination.	Service provider must pay to complete the stolen call.

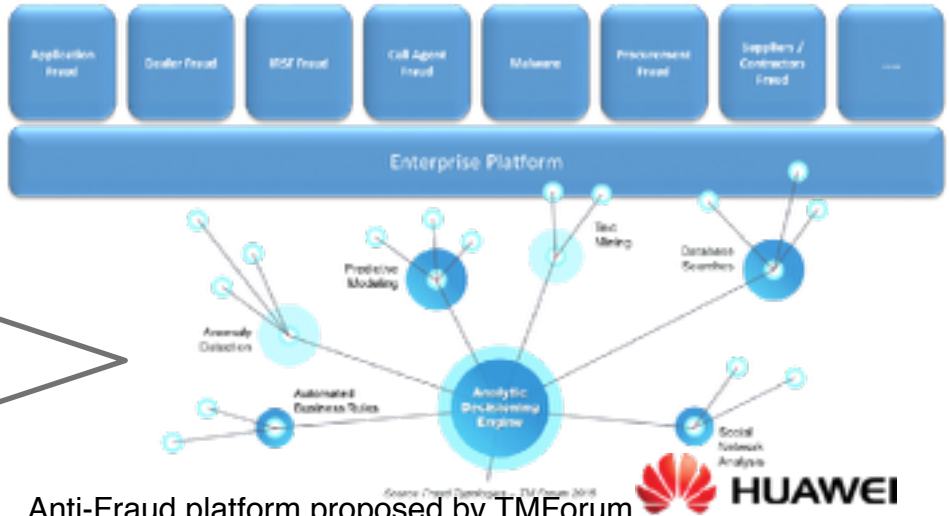

Anti-Fraud platform proposed by TMForum

HUAWEI

Existing Graph Systems

Some Existing Products

Visualization

Analytics

Frameworks

Storage

ScaleGraph

Flink/Gelly

HUAWEI

Neo4j System Architecture and Storage Format

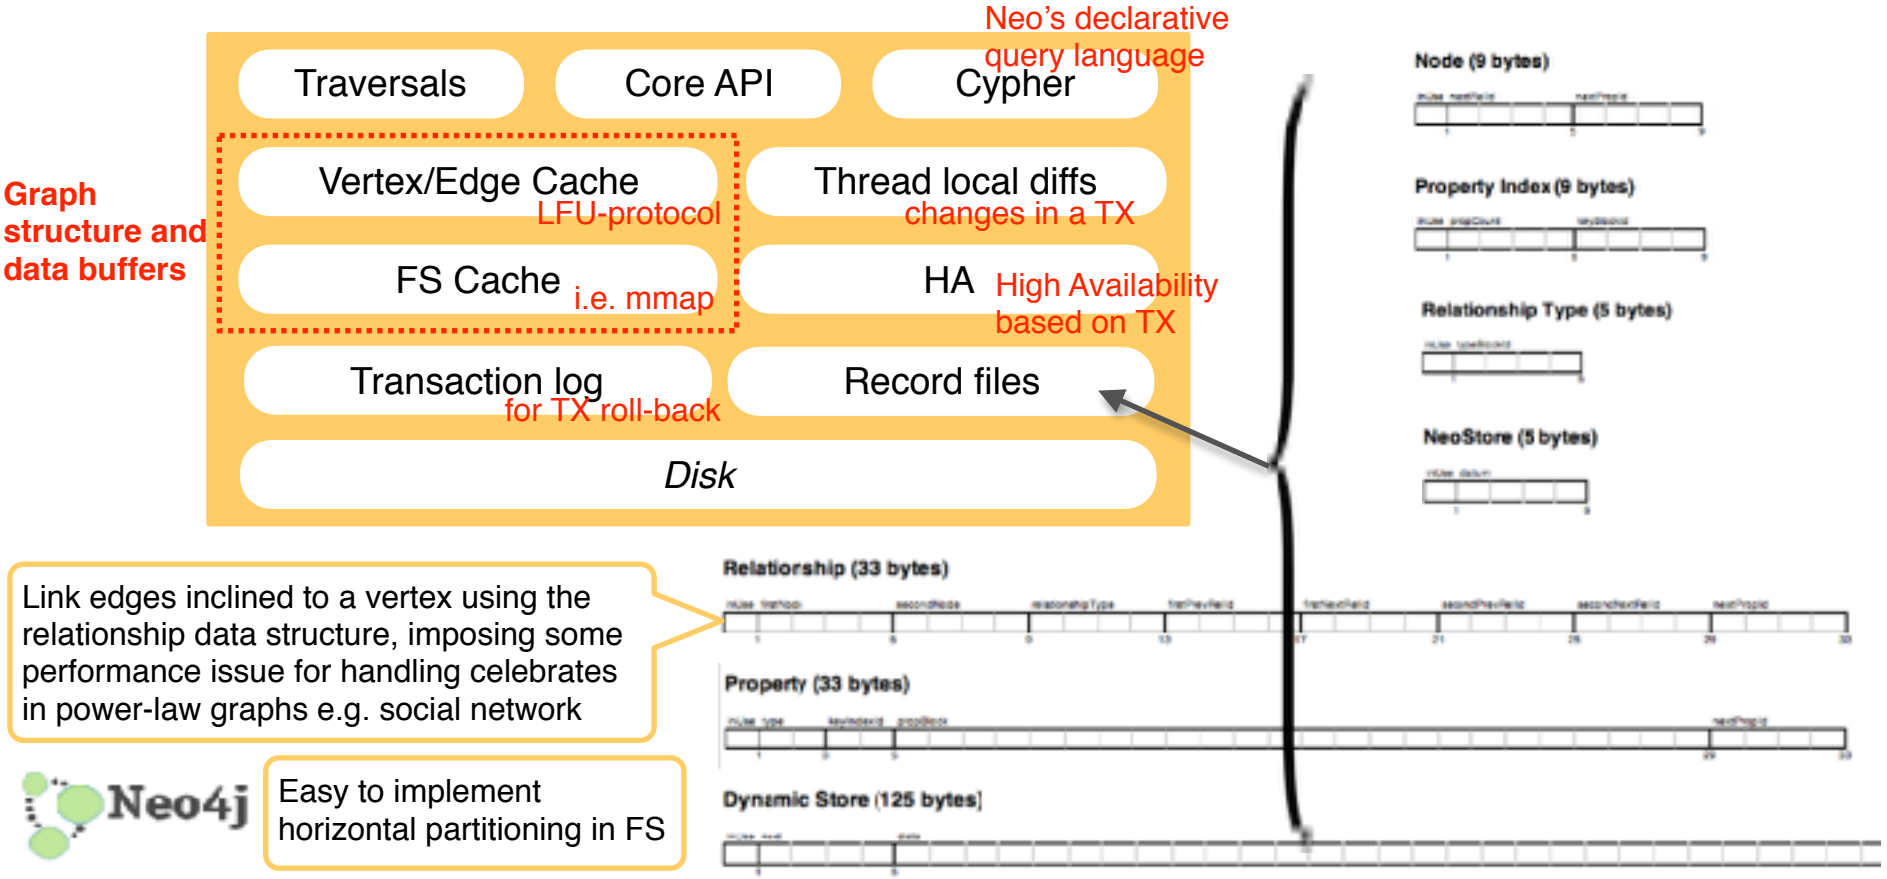

Link edges inclined to a vertex using the relationship data structure, imposing some performance issue for handling celebrates in power-law graphs e.g. social network

Neo4j

Easy to implement horizontal partitioning in FS

Titan System Architecture and Storage Format

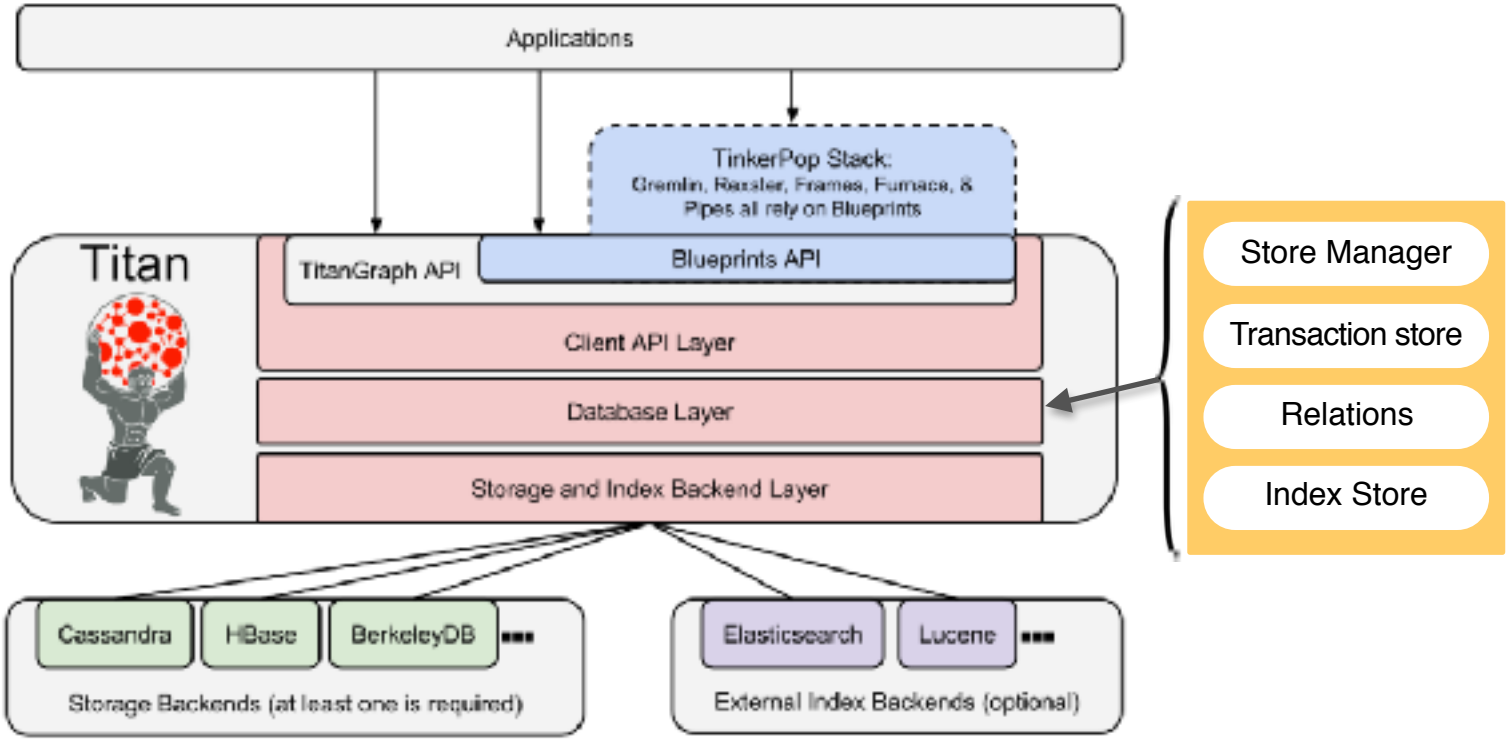

HUAWEI

OrientDB System Architecture

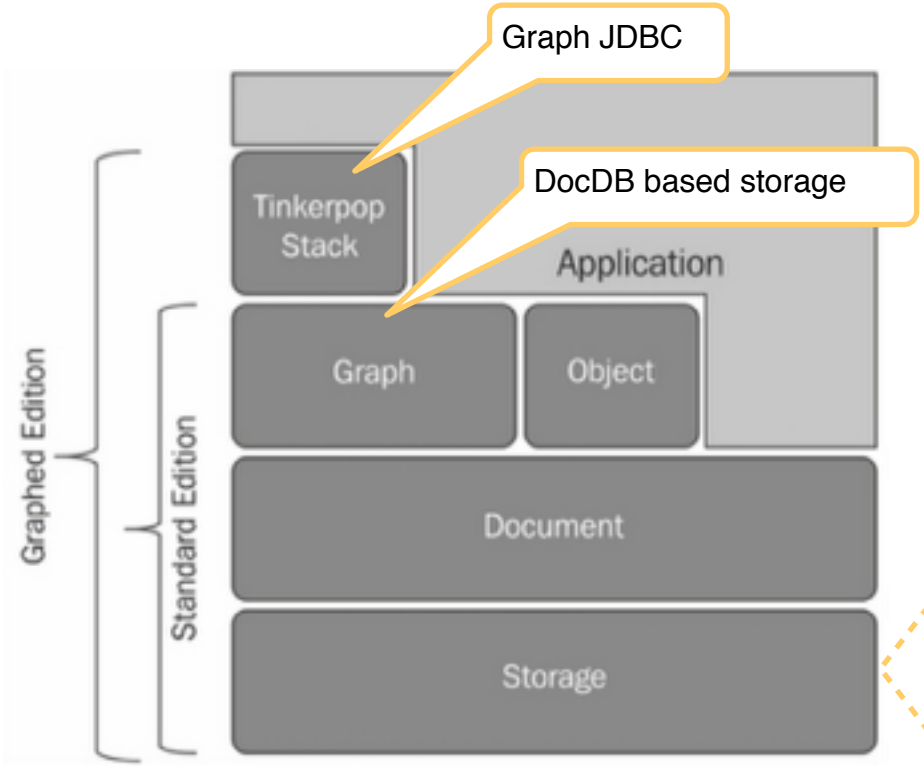

Support distributed platforms, offering key-value store, docDB, and graphDB in one system

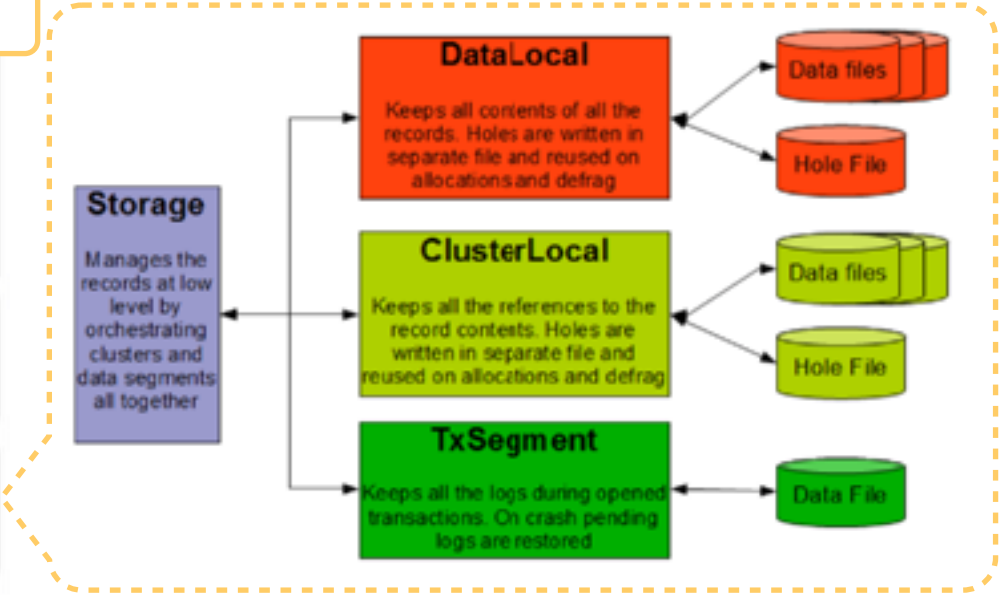

HUAWEI

Glance at Graph Computing Engines

Spark/GraphX

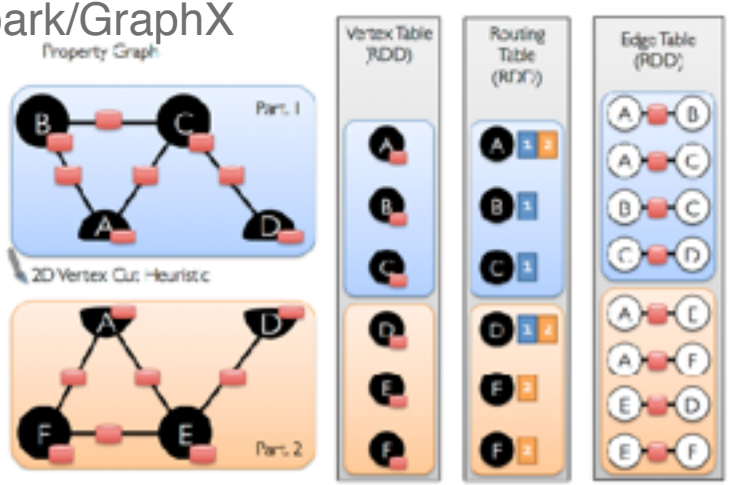

Dato Architecture

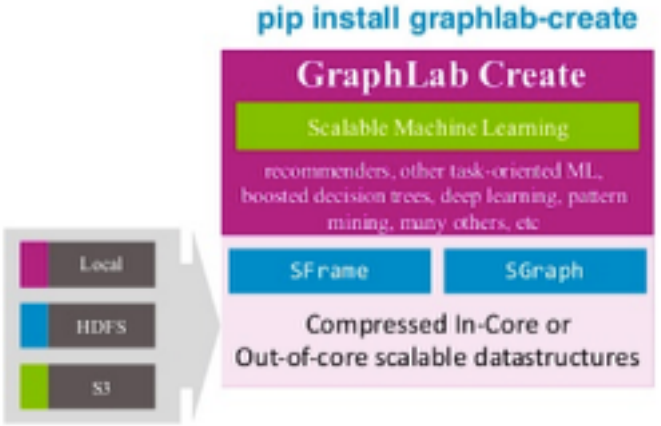

GraphChi

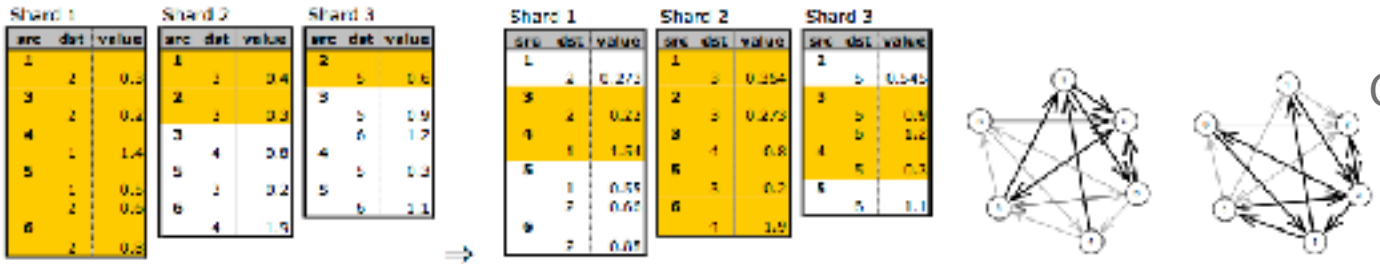

HUAWEI

Graph in ONOS

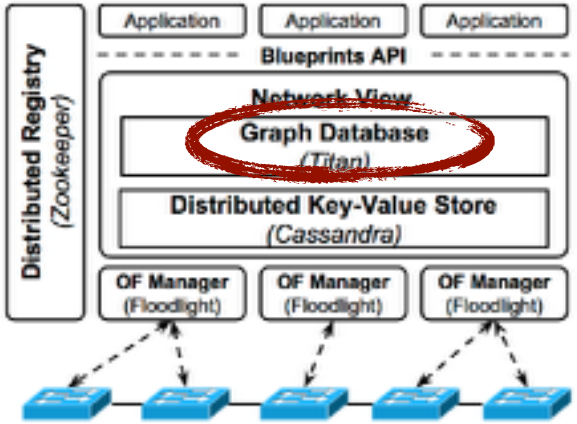

Figure 2: Prototype 1 Architecture

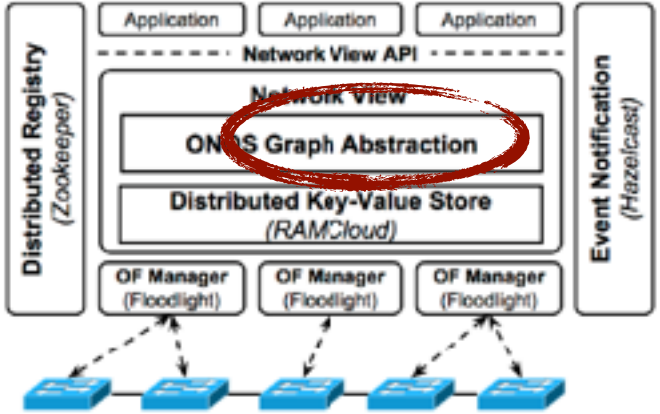

Figure 3: Prototype 2 Architecture

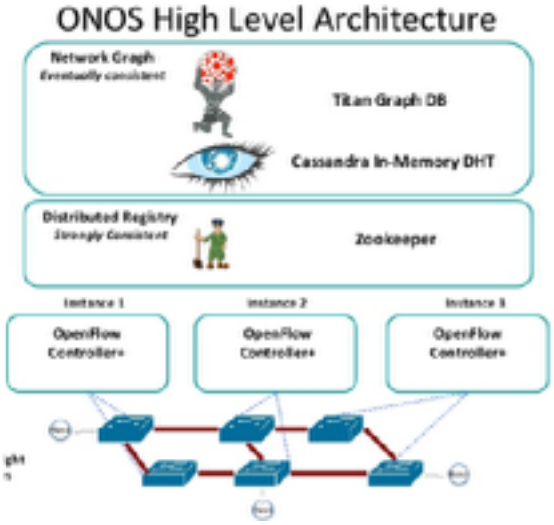

HotSDN'2014

HUAWEI

Challenges

Challenges - Performance

- **Understand performance bottleneck by breaking down the execution time**
  - **Bottleneck comes from the memory sub-system**
    - **DTLB is inefficient**
    - **Cache performs well**
    - **Cache MPKI rate is high**

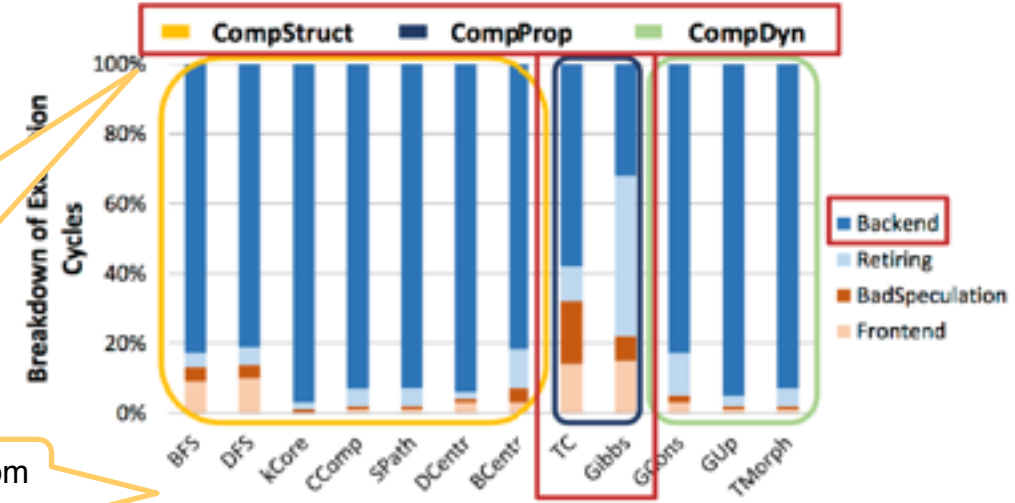

3 different types of graph computing, with focus on structural traversal, property processing, and graph editing, respectively

Core graph algorithms from 21 real-world use cases

HUAWEI

Challenges — Input Sensitivity

- **Impact from graph topology**
  - Power-law graph results in imbalanced workload due to dense vertices
    - Dense subgraph, sparse backbone
    - Dense subgraph can be converted into matrices
    - Iterative update in a subgraph
  - Road net is easy to decompose
- **Property type matters**
  - More time spent on property management
  - Computing performance can be negatively impacted

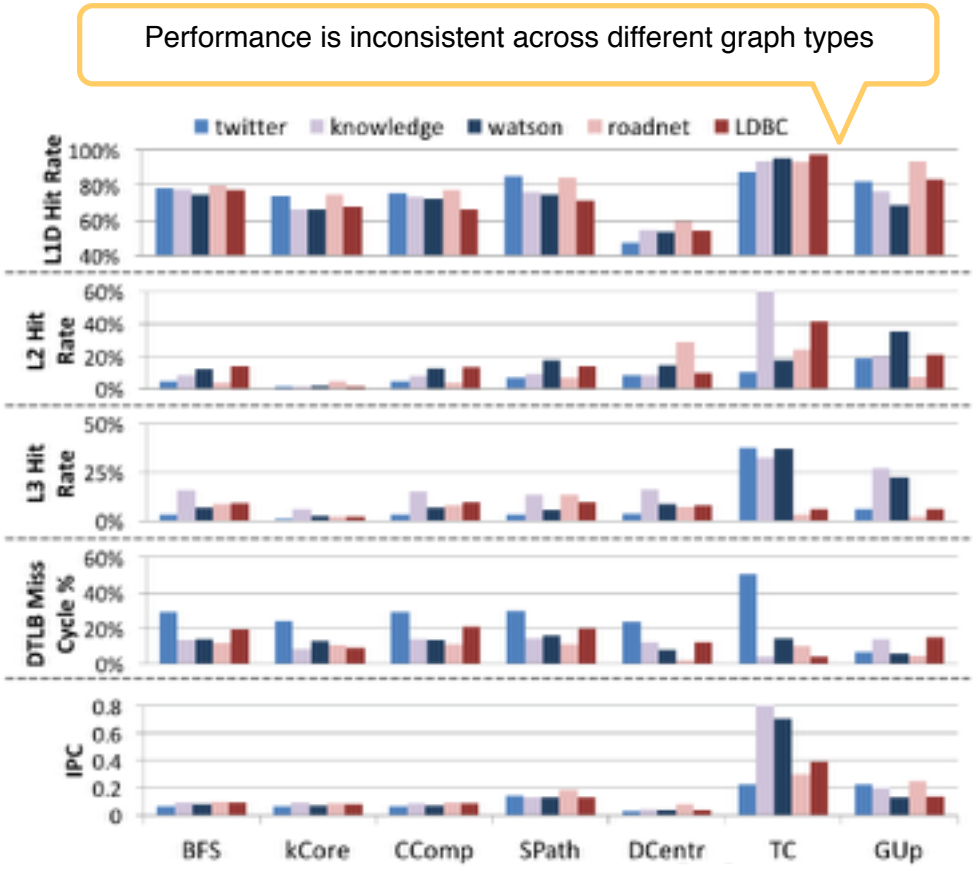

Challenges – Impact of H/W Accelerator

- **GPU can be helpful**
  - **Sufficient acceleration by GPU**
  - **Requires re-design of the algorithms**
- **Challenges**
  - **Data must be transferred to GPU**
  - **Cost of Host to Device data transfer**
  - **Difficulty in putting large graph into GPU (Double buffering)**
  - **Sensitive to input graph data**

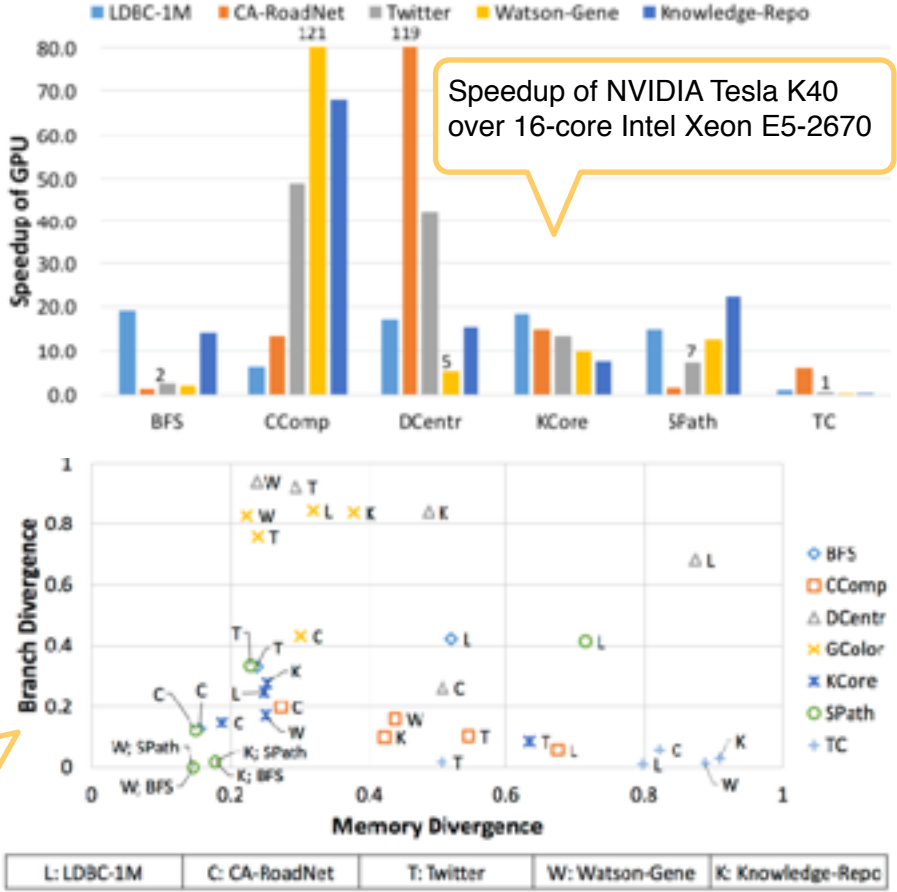

Memory divergency shows higher sensitivity for graph computing on GPU

Challenges – Scale-out Issue

- **Poor data locality and difficult partitioning result in challenges in scaling out the computing**
  - **Scale-out challenges can be seen in Graph500 analysis**
  - **Single machine with big memory can help**
  - **Must be cautious to use many computing nodes**

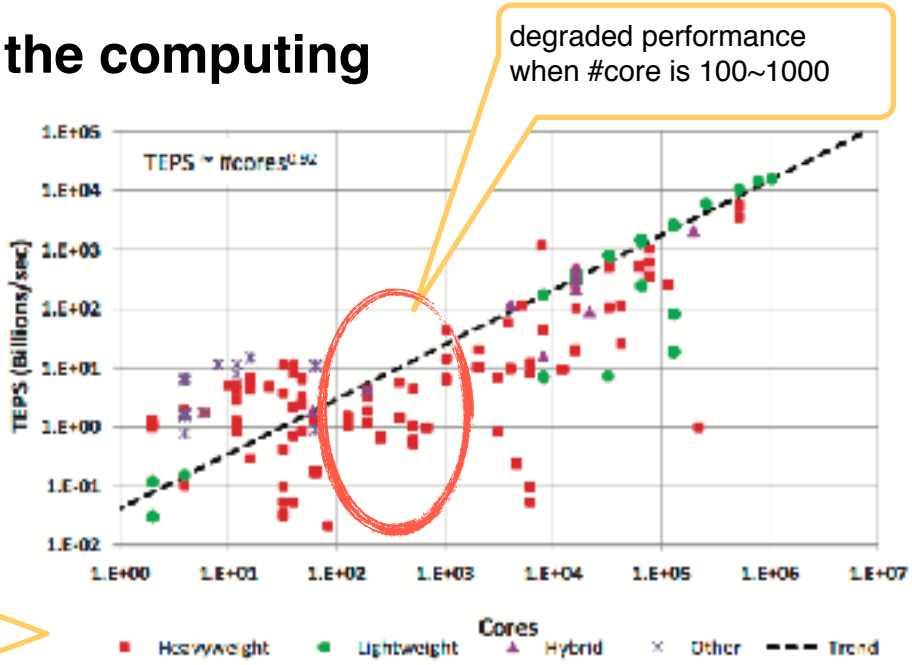

Analysis of data from Graph500

*from Peter Kogge

HUAWEI

Breakthrough

Graph Platform for Smart Big Data

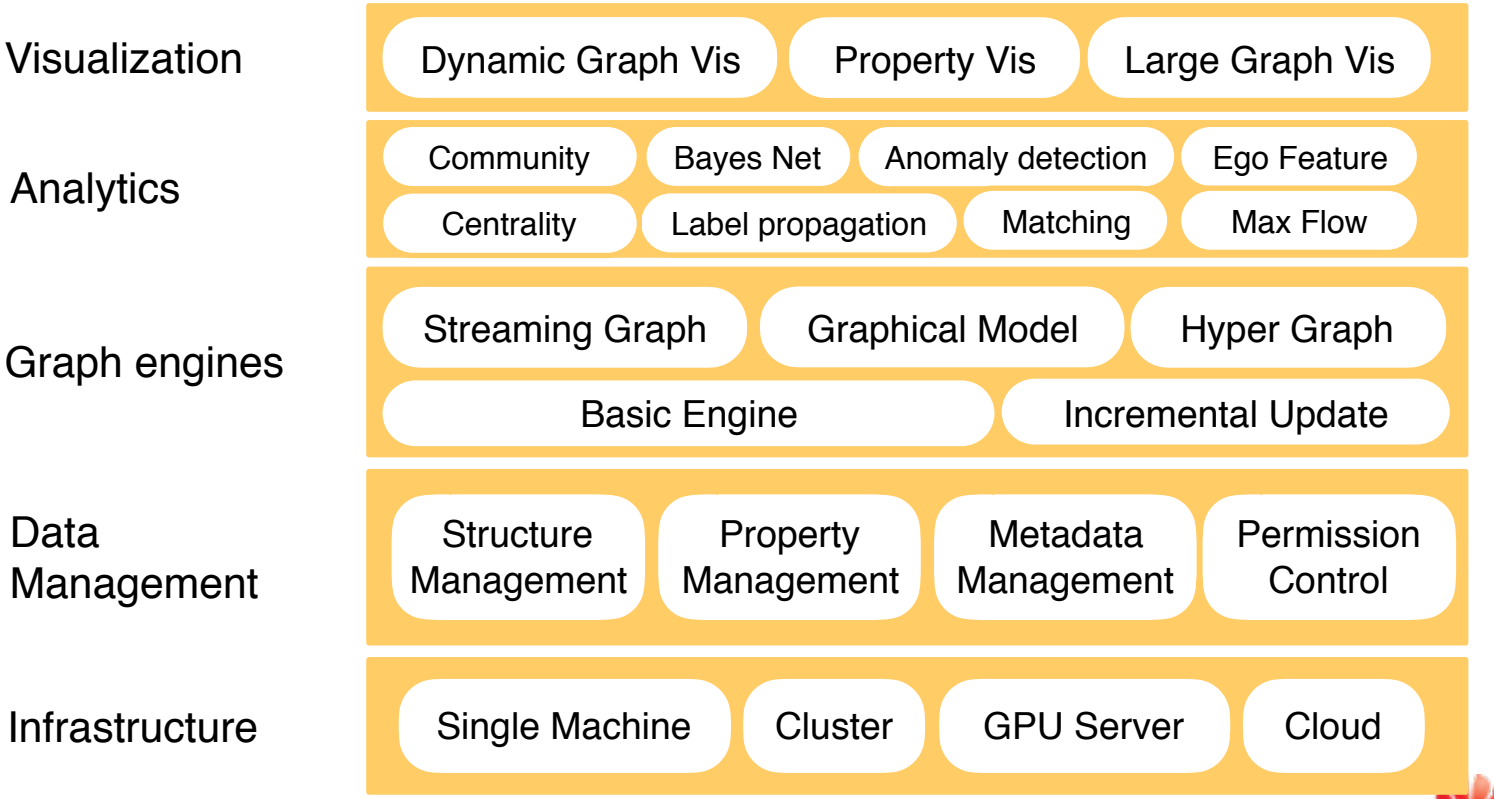

HUAWEI

Unified Graph Data Access Patterns

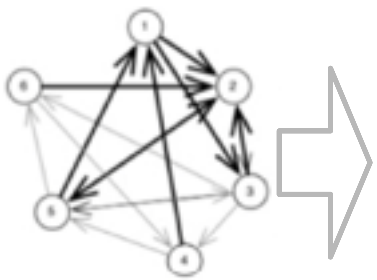

Iteration *i*

step 1

shard 1 (1, 2)

src	dst	value
1	2	0.3
3	2	0.2
4	1	1.4
5	1	0.5
	2	0.6
6	2	0.8

shard 2 (3,4)

src	dst	value
1	3	0.4
2	3	0.3
3	4	0.8
5	3	0.2
6	4	1.9

shard 3 (5,6)

src	dst	value
2	5	0.6
3	5	0.9
	6	1.2
4	5	0.3
5	6	1.1

step 2

src	dst	value
1	2	0.3
3	2	0.2
4	1	1.4
5	1	0.5
	2	0.6
6	2	0.8

src	dst	value
1	3	0.4
2	3	0.3
3	4	0.8
5	3	0.2
6	4	1.9

src	dst	value
2	5	0.6
3	5	0.9
	6	1.2
4	5	0.3
5	6	1.1

step 3

src	dst	value
1	2	0.3
3	2	0.2
4	1	1.4
5	1	0.5
	2	0.6
6	2	0.8

src	dst	value
1	3	0.4
2	3	0.3
3	4	0.8
5	3	0.2
6	4	1.9

src	dst	value
2	5	0.6
3	5	0.9
	6	1.2
4	5	0.3
5	6	1.1

equivalent

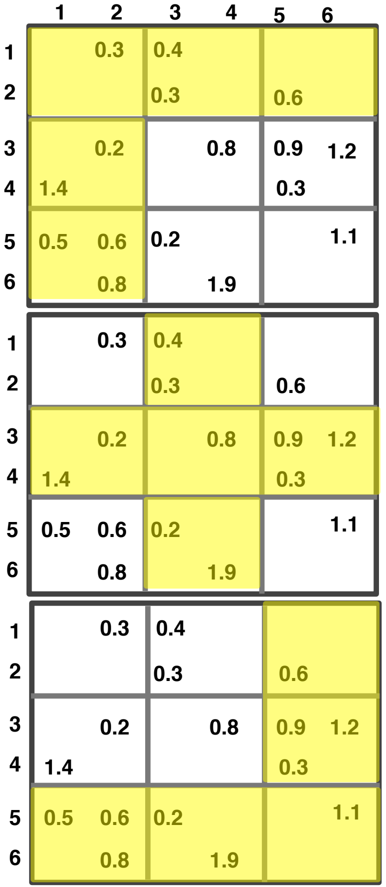

observation on PSW data access patterns inspires highly efficient sharding representation

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	712.236	41.5668	25.1802

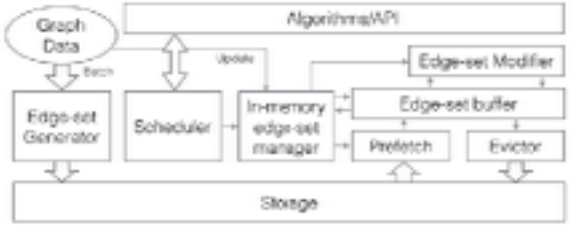

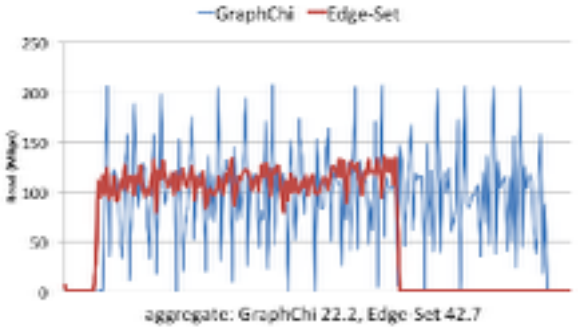

Disk read bandwidth over time by GraphChi and EdgesSet. Edge-Set showed up to 2x aggregate bandwidth and more constant IO usage.

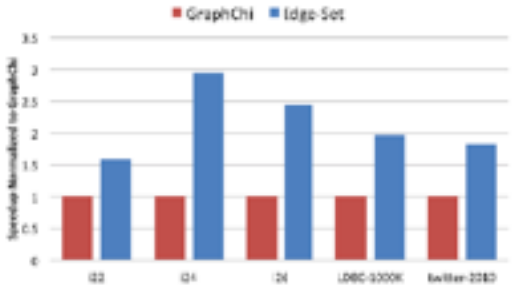

Performance improvement of SSSP against GraphChi including data ingestion

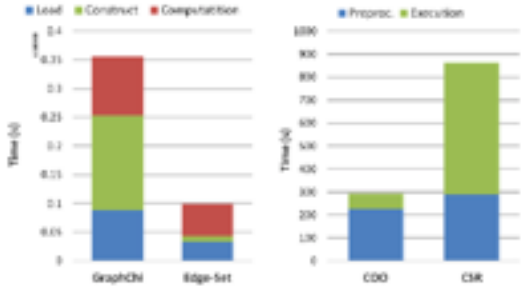

Execution time breakdown on Pagerank running Twitter-2010

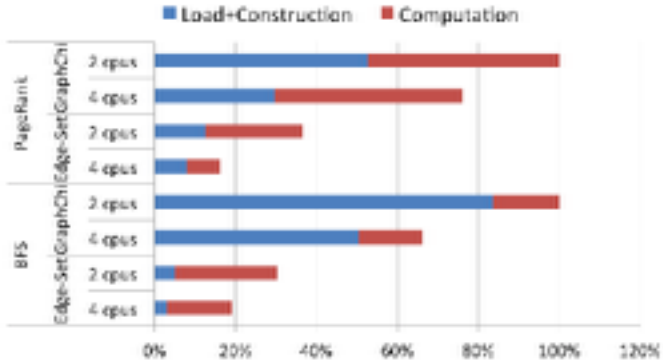

Execution time breakdown on Pagerank running Twitter-2010

HUAWEI

Opportunities in Graph Technology for Big Data

- Develop high performance graph computing kernels and primitives
  - Graph500 technique based architecture-awareness for graph computing
  - Heterogeneous computing and computing near-data technology
- Reinvent graph technology for supporting cognitive computing
  - One open platform with multiple graph and graph-related technologies
  - Integral consideration on graphical model, streaming graphs, etc. for AI/IoT
- Offer vertical solutions to break through separation among technique stacks
  - Holistic solution for rapidly building industry-level graph analytics solutions
  - Incorporating with market segmentation, such as security, finance, etc.
- Collaborations and Standardization
  - Foster collaboration with relevant professional communities to educate the market
  - Developing domain or cross-domain standardizations

THANK YOU

Yinglong Xia
yinglong.xia.2010@ieee.org