

Parallelizing Sequential Graph Computations

Wenfei Fan^{1,2}, Jingbo Xu^{1,2}, Yinghui Wu³, Wenyuan Yu², Jiaxin Jiang⁴, Zeyu Zheng⁵, Bohan Zhang⁵, Yang Cao¹, Chao Tian^{1,2}

¹Univ. of Edinburgh ²Beihang Univ. ³Washington State Univ. ⁴Hong Kong Baptist Univ. ⁵Peking Univ.

30

25

(seconds)

Time

INTRODUCTION

GRAPE, parallelizing existing sequential algorithms as a whole, and guaranteeing convergence and correctness when the sequential algorithms provided are correct.

- Sequential graph algorithms can be "plugged into" GRAPE with minor additions, and get parallelized.
- Foundation: a simultaneous fixed point computation with



partial evaluation and incremental evaluation.



MOTIVATION



- **Ease of programming.** Only need to provide three (existing) sequential (incremental) algorithms for Q with minor additions.
- Semi-automated parallelization. Guarantee to converge at correct answers under a monotonic condition, if the three sequential algorithms provided are correct.
- Graph-level optimization. GRAPE inherits all optimization strategies available for sequential algorithms and graphs.
- **Scale-up.** GRAPE achieves comparable performance to the state-of- the-art graph systems.

PERFORMANCE

Sequential Algorithms and their applications: (1) **SSSP** (traffic analysis); (2) Connected Component (social analysis); (3) Graph Simulation (social recommendation); (4) Keyword Search/ **Subgraph Isomorphism**(Web/knowledge mining); (5)

- It is nontrivial for one to learn how to program in the new parallel models, e.g., "think like a vertex".
- Graph computations have been studied for decades, and a number of sequential graph algorithms are already in place. Can we use them without recasting?
- Do existing parallel graph engines guarantee termination and correctness?

GRAPE API

Data partitioned parallelism (shared-nothing architecture). Fragmented graph $G = (G_1, ..., G_n)$, distributed to workers. **GRAPE API:** three core functions for a graph query class **Q PEval**: a (existing) sequential algorithm for Q, for partial evaluation; **IncEval**: a (existing) sequential incremental algorithm for Q; **Assemble**: a sequential algorithm (taking a union of partial results).

FOUNDATION OF GRAPE

Collaborate Filtering (recommendation, machine learning).

2 times faster than Giraph, orders of magnitudes less data shipment GRAPE -E 250 GraphLab - 🗗 -Ġiraph ····×··· JraphLab --⊕--Blogel -O-200 Communication iraph ···×·· Blogel 150 100 50

(c) Varying n: SSSP (DBpedia)(c) Varying n: SSSP (DBpedia)

APPLICATION SCENARIO

Social media marketing with Graph Pattern Association Rules.

GRAPE Online 0.91 alpha	Plug Play	Results Analytics API Doc	Graph Pattern Assoc
QUERY RESULTS		QUERY PANEL	iation Rules (GPARs):
User uid=3180342 follow User	User uid=1923822 follow User	Dataset Weibo_Social_Network	$Q(x,y) \Rightarrow p(x,y),$
		Query Input	



- **R**^r: partial results at processor i in round r
- **M_i** : message to processor i



Assurance Theorem: termination and correctness guaranteed if the sequential algorithms provided are correct, and M_i is "monotonic"! **Simulation Theorem**: MapReduce, BSP (bulk synchronous parallel) and PRAM models are optimally simulated by GRAPE



Left figure shows a GPAR **φ**: if among the people followed by x_o (a) at least 80% of them recommend Huawei Mate 9, and (b) no one gives it a bad rating, then the chances are that x_o will buy this mobile phone.

12

CONCLUSION

GRAPE: make parallel graph computations accessible to users who know conventional graph algorithms from undergraduate textbooks. Welcome to play with GRAPE (alpha): <u>http://grapedb.io</u>