Adaptive Asynchronous Parallelization of Graph Algorithms



Wenfei Fan^{1,2,3}, Ping Lu², Xiaojian Luo³, Jingbo Xu^{2,3}, Qiang Yin²,Wenyuan Yu^{2,3}, Ruiqi Xu¹

¹University of Edinburgh ²BDBC, Beihang University ³7 Bridges Ltd.



INTRODUCTION

We proposes an Adaptive Asynchronous Parallel (**AAP**) model for graph computations. As opposed to Bulk Synchronous Parallel (**BSP**) and Asynchronous Parallel (**AP**) models, AAP reduces both stragglers and stale computations by dynamically adjusting relative progress of workers.

- BSP, AP and Stale Synchronous Parallel model (SSP) are
- **Simulation Theorem:** MapReduce, BSP, AP, SSP and PRAM models are optimally simulated by AAP

DYNAMIC ADJUSTMENT

$$\mathsf{DS}_{i} = \begin{cases} +\infty & \neg S(r_{i}, r_{\min}, r_{\max}) \lor (\eta_{i} = 0) \\ T_{\mathsf{L}_{i}}^{i} - T_{\mathsf{idle}}^{i} & S(r_{i}, r_{\min}, r_{\max}) \land (1 \le \eta_{i} < \mathsf{L}_{i}) \\ 0 & S(r_{i}, r_{\min}, r_{\max}) \land (n_{i} \ge \mathsf{L}_{i}) \end{cases}$$

special cases of AAP.

- Better yet, **AAP** optimizes parallel processing by adaptively switching among these models at different stages.
- Employing the programming model of GRAPE, **AAP** parallelizes existing sequential algorithms based on fixpoint computation with partial and incremental evaluation.
- Under a monotone condition, **AAP** guarantees to converge at correct answers.
- It can optimally simulate MapReduce/PRAM/**BSP**/**AP**/**SSP**.



MOTIVATION

 $O \qquad O('_i, '_{\min}, '_{\max}) \land ('_i \leq \Box_i)$

 $S(r_i, r_{\min}, r_{\max})$: whether P_i should be suspended immediately. L_i "predicts" how many messages should be accumulated, $T_{L_i}^i$ estimates how longer P_i should wait to accumulate L_i messages. T_{idle}^i is the idle time of worker P_i after the last round.



IMPLEMENTATION

- **BSP**: stragglers keep other workers idle, thus wasting resources.
- **AP**: with redundant computation and communication
- SSP: ad-hoc staleness, incremental improvement over BSP/AP
- Is it possible to have a simple parallel model that inherits the benefits of BSP and AP, and reduces both stragglers and stale computations, without explicitly switching between the two?

PROGRAMMING MODEL

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).

A fixpoint computation Φ (R₁, ..., R_n)



PERFORMANCE

<u>Algorithms and their applications:</u> (1) SSSP (traffic analysis); (2) Connected Component (social analysis); (3) PageRank; (4) Collaborate Filtering (recommendation, machine learning).

> Outperform BSP, AP and SSP by 4.3X, 14.7X and 4.7X on average Outperform the state-of-the-art systems





- Convergence guarantee:
- T1: Update parameters take values from a finite domain
- T2: IncEval is contracting (the same run)
- T3: IncEval is monotonic (different runs)

T1 + T2: termination

T1 + T2 + T3: the Church-Rosser property (all asynchronous runs converge at the same correct result)



We have proposed AAP to remedy the limitations of BSP and AP by reducing both stragglers and redundant stale computations. More information about GRAPE: <u>https://7bridges.io</u>

RESEARCH POSTER PRESENTATION DESIGN © 2015 www.PosterPresentations.com