Adaptive Asynchronous Parallelization of Graph Algorithms

WENFEI FAN, University of Edinburgh and Beihang University and SICS, Shenzhen University
PING LU, BDBC, Beihang University
WENYUAN YU, JINGBO XU, QIANG YIN, XIAOJIAN LUO, and JINGREN ZHOU, Alibaba Group
RUOCHUN JIN, University of Edinburgh

This article 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. We show that BSP, AP, and Stale Synchronous Parallel model (SSP) are special cases of AAP. Better yet, AAP optimizes parallel processing by adaptively switching among these models at different stages of a single execution. Moreover, employing the programming model of GRAPE, AAP aims to parallelize existing sequential algorithms based on simultaneous fixpoint computation with partial and incremental evaluation. Under a monotone condition, AAP guarantees to converge at correct answers if the sequential algorithms are correct. Furthermore, we show that AAP can optimally simulate MapReduce, PRAM, BSP, AP, and SSP. Using real-life and synthetic graphs, we experimentally verify that AAP outperforms BSP, AP, and SSP for a variety of graph computations.

CCS Concepts: • **Information systems → Database management system engines**; **Parallel and distributed DBMSs**;

Additional Key Words and Phrases: Graph computations, parallel graph query engines, parallelizing sequential algorithms, convergence, simulation

ACM Reference format:
Wenfei Fan, Ping Lu, Wenyuan Yu, Jingbo Xu, Qiang Yin, Xiaojian Luo, Jingren Zhou, and Ruochun Jin. 2020. Adaptive Asynchronous Parallelization of Graph Algorithms. *ACM Trans. Database Syst.* 45, 2, Article 6 (July 2020), 45 pages.
https://doi.org/10.1145/3397491

W. Fan is supported in part by ERC 652976, Royal Society Wolfson Research Merit Award WRM/R1/180014, EPSRC EP/M025268/1, Shenzhen Institute of Computing Sciences, and Beijing Advanced Innovation Center for Big Data and Brain Computing. P. Lu is supported in part by NSFC 61602023.

Authors' addresses: W. Fan and R. Jin, University of Edinburgh and Beihang University, 10 Crichton Street, Edinburgh, EH8 9AB, UK and SICS, Shenzhen University, 9~10F, Building 26, Hongshan 6979, Longhua, Shenzhen, 518000, China; emails: wenfei@inf.ed.ac.uk, ruochun.jin@ed.ac.uk; P. Lu, BDBC, Beihang University, 37 Xue Yuan Road, Haidian District, Beijing, 100191, China; email: luping@buaa.edu.cn; W. Yu, J. Xu, Q. Yin (corresponding author), X. Luo, and J. Zhou, Alibaba Group, 969 West Wen Yi Road, Yu Hang District, Hangzhou, 311121, China; emails: {wenyuan.ywy, xujingbo.xjb, qiang.yq, lxj193371, jingren.zhou}@alibaba-inc.com.

0362-5915/2020/07-ART6 $15.00
https://doi.org/10.1145/3397491