

UNIVERSITY OF SOUTH FLORIDA

Defense of a Doctoral Dissertation

Algorithms and Framework for Computing 2-body Statistics on Graphics Processing Units

by

Napath Pitaksirianan

For the Ph.D. degree in Computer Science and Engineering

Various types of two-body statistics (2-BS) are regarded as essential components of low-level data analysis in scientific database systems. In relational algebraic terms, a 2-BS is essentially a Cartesian product between two datasets (or two instances of the same dataset) followed by a user-defined aggregate. The quadratic complexity of these computations hinders timely processing of data. Use of modern parallel hardware has thus become an obvious solution to meet such challenges. This dissertation presents our recent work on designing and optimizing parallel algorithms for 2-BS computation on Graphics Processing Units (GPUs). While a typical 2-BS problem can be summarized into a straightforward parallel computing pattern, traditional knowledge from (general) parallel computing often falls short in delivering the best possible performance. Therefore, we present a suite of techniques to decompose 2-BS problems and methods for the effective use of computing resources on GPUs. We also develop analytical models that guide us towards finding the best parameters of our GPU programs. As a result, we achieve the design of highly-optimized 2-BS algorithms that significantly outperform the best-known GPU and CPU implementations.

Examining Committee

Lingling Fan, Ph.D., Chairperson
Yicheng Tu, Ph.D., Major Professor
Adriana Iamnitchi, Ph.D.
Yao Liu, Ph.D.
Charkgard Hadi, Ph.D.
Pandit Sagar, Ph.D.

Thursday, March 5, 2020
11:00 AM
ENB 313

THE PUBLIC IS INVITED

Publications

- 1) **Napath Pitaksirianan**, Zhila Nouri, and Yi-Cheng Tu. "Algorithms and Framework for Computing 2-body Statistics on GPUs", Distributed and Parallel Databases Journal (DAPD), December 2019, Volume 37, Issue 4
- 2) Jinghan Meng, **Napath Pitaksirianan**, and Yi-Cheng Tu, "Counting Frequent Patterns in Large Labeled Graphs: A Hypergraph-Based Approach", Data Mining and Knowledge Discovery Journal (DAMI) 2019, Accepted
- 3) Jinghan Meng, **Napath Pitaksirianan** and Yi-Cheng Tu "Generalizing Design of Support Measures for Counting Frequent Patterns in Graphs", 2019 IEEE International Conference on Big Data, Dec 2019,
- 4) Jinghan Meng, Yi-Cheng Tu and **Napath Pitaksirianan**. "Generalizing Design of Support Measures for Counting Frequent Patterns in Graphs", 31st International Conference on Scientific and Statistical Database Management (SSDBM), July 2019
- 5) **Napath Pitaksirianan**, Zhila Nouri, and Yi-Cheng Tu. "Efficient 2-Body Statistics Computation on GPUs: Parallelization and Beyond", In Procs. of 45th International Conference on Parallel Processing (ICPP), August, 2016.
- 6) "Computing 2-Body Statistics on Graphic Processing Units (GPUs)", Yi-Cheng Tu and **Napath Pitaksirianan**, US Patent Application No. 62/738,228

Robert Bishop, Ph.D.
Dean, College of Engineering

Dwayne Smith, Ph.D.
Dean, Office of Graduate Studies

Disability Accommodations:

If you require a reasonable accommodation to participate, please contact the Office of Diversity & Equal Opportunity at 813-974-4373 at least five (5) working days prior to the event.