Efficient Parallel All-Pairs Computation Framework: using Computation - Communication Overlap

Venkata Kasi Viswanath Yeleswarapu
Electrical and Computer Engineering
Iowa State University
E-mail: yvk@iastate.edu

April 1, 2017

Abstract

The advent of parallel computing systems enabled the users with huge computation power to efficiently process huge work loads. Most of the recent applications, which are data intensive, require parallel computing power to complete the job efficiently. To facilitate efficient computing there is a necessity for simplified abstraction of the parallel computing systems.

We propose one such parallel computation abstraction, designed to solve All-Pairs problems which fits the needs of several data intensive applications. All-Pairs problems require each data element to be paired with every other data element. This framework aims to address recurring problems of scalability, distributing equal work load to all nodes and reducing memory footprint. Our framework reduces memory footprint of All-Pairs problems, by reducing memory requirement from $N/\sqrt{P}$ to $3N/P$. A bio-informatics application is implemented to demonstrate the scalability (ranging up to 512 cores), redundancy management and speed up performance of the framework(super-linear speed up).