Contribution ID: **24** Type: **Talk**

Effective Partitioning and Communication Strategies for Parallel Sparse Tensor Factorization

Wednesday 6 November 2019 14:15 (25 minutes)

Sparse tensor decompositions have become increasingly popular in the literature due to their capability to naturally model high dimensional sparse data with many features, and glean hidden relations owing to underlying low-rank structures within the data. They have been successfully employed in many application settings including recommender systems, graph analytics, healthcare data analytics, web search, cyber security, and many others. The aptitude of tensor methods for such big data analysis applications solicited the development of efficient parallel tensor factorization algorithms capable of handling datasets of billions of entries, which has been among the most trending areas of research in the HPC community in the recent past.

In this talk, we will discuss parallelization techniques for sparse tensor factorization together with various partitioning strategies for balancing computation/memory costs and reducing communication. We will compare advantages and limitations of various approaches and touch upon outstanding challenges for better parallel scalability. We will conclude the talk with an overview of the capabilities of the PACOS library (PArtitioning and COmmunication framework for Sparse irregular applications) that enable devising efficient and scalable parallel sparse tensor factorization kernels as well as partitioning routines, and facilitate reproducibility of scalability results.

Author: KAYA, Oguz (Université Paris-Saclay and Laboratoire de Recherche en Informatique (LRI))

Presenter: KAYA, Oguz (Université Paris-Saclay and Laboratoire de Recherche en Informatique (LRI))

Session Classification: Day I

Track Classification: Talks