UNIVERSITY OF AMSTERDAM

Acquiring Negative Polarity Items

J. Lin

In such an exciting age of information explosion, huge amount of visual data are produced continuously 24 hours, 7 days in both daily life and scientific research. Processing and storage of such a huge amount of data forms big challenges.Use of supercomputers tackles the need-for-speed challenge partially, but is blocked by its high cost of ownership and slow capacity growth. Distributed computing provides an attractive alternative which scale on demand, ranging from academic clusters to commercial clouds. In recent years, research on computer architecture also advanced significantly, resulting in wide deployment of multi-core and many core processors in both industries and daily life.To this end, this thesis explores how to apply the massively heterogeneous parallel distributed computing architectures for accelerating large-scale image processing applications effectively and efficiently. One step further, considering the actual gap between computing and imaging, the focus of this thesis invests efforts in exploring whether we can efficiently use high performance distributed computing systems to accelerate modern image processing algorithms.Combining the representative computing platforms with the selected image processing applications, four interesting aspects of fast image processing cases are studied in this thesis, namely: large-scale hyperspectral image analysis on multi-clusters platform, complex matrix manipulation on a cluster platform, matrix bidiagonalization on many-core platform, and fundamental matrix multiplication on a many-core accelerated cluster platform. Based on findings in each aspect, we conclude that parallel and distributed systems can give a lot of performance improvements to large-scale image processing applications, but are still hard to be deployed by image processing researchers. New research directions opened by this thesis include large-scale compute intensive application optimization with many-core clusters, complex time-consuming iterative matrix algorithms parallelization on grids, and identification of a generic hybrid multi-layer algorithmic parallelization methodology.