

Guest Editorial

Special Issue on the 9th International Symposium on Parallel and Distributed Computing

Alper Sen and Can Ozturan

Department of Computer Engineering, Bogazici University, Istanbul, Turkey
E-mails: {alper.sen, ozturaca}@boun.edu.tr

Welcome to this special issue on the ninth edition of International Symposium on Parallel and Distributed Computing (ISPDC) held in Istanbul in July 2010. ISPDC is dedicated to the advancement of state-of-the-art in the area of parallel and distributed computing. The symposium program featured 28 full papers accepted out of 71 submitted papers. The program covered a wide spectrum of topics including parallel models and algorithms, interconnection topologies, GPU programming, grids, P2P, multicores and distributed systems. The three papers in this special issue were selected among the highest rated accepted papers and involve scientific programming on a wide spectrum of computational platforms, namely, Graphics Processing Units (GPUs), multi-cluster systems and grids. These three papers were extended and revised after a new round of review process before being published in *Scientific Programming*.

The first paper titled “Improving accuracy for matrix multiplications on GPUs” by Badin et al. demonstrates the emergence of scientific programming research on Graphics Processing Units (GPU). The paper considers the problem of floating point arithmetic accuracy within the context of matrix multiplication on GPUs. On GPUs that lack double precision support, the problem of numerical accuracy becomes quite important. Hence, new algorithms that deliver high accuracy with little performance cost need to be developed. This paper addresses these issues and proposes doubly com-

pensated summation and Kahan’s compensated summation based algorithms with little performance costs.

The second paper titled “Performance models for the Spike banded linear system solver” by Manguoglu et al. investigates the problem of characterizing performance and scalability of banded preconditioners for linear system solvers. The authors develop parallel formulations of a banded solver, called the Truncated Spike solver. They present a highly accurate parallel performance model for the solver and establish the validity of the model on diverse heterogeneous multi-clusters. Their performance model supported by actual runs demonstrates high scalability of their solver.

The third paper “PetaShare: A reliable, efficient and transparent distributed storage management system” by Kosar et al. addresses the development of a distributed data management system for data-intensive e-science. The authors contribute a reliable and efficient distributed data storage system, called PetaShare, which spans multiple institutions over a wide geographic area. PetaShare provides a front-end with light-weight clients that enable easy, transparent and scalable access. A unified name space and efficient client data movement capability across geographically distributed storage sites are provided by PetaShare. The authors also present the details of their design and implementation and show performance results on reading/writing to remote PetaShare sites.