9th European Congress on Computational Methods in Applied Sciences and Engineering (ECCOMA2024)
June 3-7, 2024, Lisbon, Portugal

ADVANCED PARALLEL ALGORITHMS FOR EXTREME-SCALE SIMULATIONS

TRACK NUMBER 1700

XAVIER ÁLVAREZ-FARRÉ[*], ÀDEL ALSALTI-BALDELLOU[†], MANUEL RODRIGUES[*], VALERIU CODREANU[*], TAKAYUKI AOKI[‡], F. XAVIER TRIAS[†]

[*] SURF, High-Performance Computing and Visualization Team,
Science Park 140, 1098 XG Amsterdam, The Netherlands
{xavier.alvarezfarre, manuel.torresrodrigues, valeriu.codreanu}@surf.nl

[†] Technical University of Catalonia, Heat and Mass Transfer Technological Center,
Carrer Colom 11, 08222 Terrassa (Barcelona), Spain
{adel.alsalti, francesc.xavier.trias}@upc.edu

[‡] Tokyo Institute of Technology, Global Scientific Information and Computing Center
2-12-1, O-okayama, Meguro-ku, Tokyo 152-8550, Japan
taoki@gsic.titech.ac.jp

Key words: Heterogeneous systems, distributed memory, parallel computing, extreme-scale simulations, sparse algebra, linear solvers.

ABSTRACT

The evolution in hardware technologies enables scientific computing to reach further aims. Nowadays, the use of massively parallel supercomputers is rather common in the solution of both industrial and academic scale problems. Unfortunately, the literature reveals that the performance of numerous scientific computing applications is often limited to tiny fractions of the system's peak performance. While the peak performance has been growing 9 and 4.5 times faster than network and main memory bandwidth per decade over the past 30 years, many algorithms employed in scientific computing have very low operational intensity. Numerical simulation codes are increasingly memory-bounded, making processors suffer from serious data starvation. Moreover, distributed memory systems pose an additional bottleneck: network bandwidth and latency. The relative weight of data exchanges grows as parallel applications are distributed among a greater number of processors, making it advantageous to increase the problem size per node. However, hardware accelerators with high-bandwidth memory, which are major contributors to the increased computational power of individual nodes, have smaller memory capacities than traditional processors. Application bottlenecks becoming more severe exacerbates the already difficult job of architects, programmers, and researchers. In this context, this mini-symposium aims to bring together people working on advanced, cutting-edge parallel methods for solving extreme-scale simulations with a special focus on efficiency, portability, and sustainability. Works on domain-specific languages, parallel linear solvers, sparse algebra kernels, eigenvalue problems, parallel-in-time methods, load balancing, adaptive mesh refinement, and memory footprint reduction techniques, among others, are welcomed.