

Performance Characterization of a Molecular Dynamics Code on PC Clusters

Is there any easy parallelism in CHARMM?

Michela Taufer¹, Egon Perathoner^{1,3}, Andrea Cavalli², Amedeo Caflisch², Thomas Stricker¹.

¹ Department of Computer Science
ETH Zurich, Zentrum
CH-8092 Zurich, Switzerland
stricker,taufer@inf.ethz.ch

² Department of Biochemistry
University of Zurich
CH-8057 Zurich, Switzerland
caflisch@bioc.unizh.ch

³ Università degli Studi di Padova,
Dipartimento di Elettronica ed Informatica,
Via Gradenigo 6/a, Padova, Italy

The molecular dynamics code CHARMM is a popular research tool for computational biology. An increasing number of researchers are currently looking for affordable and adequate platforms to execute CHARMM or similar codes.

To address this need, we analyze the resource requirements of a CHARMM molecular dynamics simulation on PC clusters with a particle mesh Ewald (PME) treatment of long-range electrostatics, and investigate the scalability of the short-range interactions and PME separately. We look at the workload characterization and the performance gain of CHARMM with different network technologies and different software infrastructures and show that the performance depends more on the software infrastructures than on the hardware components. In the present study, powerful communication systems like Myrinet deliver performance that comes close to the MPP supercomputers of the past decade (e.g. Cray T3D), but improved scalability can also be achieved with better communication system software like SCORE without the additional hardware cost.

The experimental method of workload characterization presented can be easily applied to other codes. The detailed performance figures of the breakdown of the calculation into computation, communication and synchronization allow to derive good estimates about the benefits of moving applications to novel computing platforms such as widely distributed computers (grid).

Keywords: middleware, distributed computation, performance tuning, performance optimization, cluster of PCs, message passing, distributed molecular dynamics.