

UC Berkeley

UC Berkeley Previously Published Works

Title

A Communication-Optimal N-Body Algorithm for Direct Interactions

Permalink

<https://escholarship.org/uc/item/4m17p8bw>

ISBN

978-1-4673-6066-1

Authors

Driscoll, Michael
Georganas, Evangelos
Koanantakool, Penporn
et al.

Publication Date

2013-05-01

DOI

10.1109/ipdps.2013.108

Peer reviewed

A Communication-Optimal N-Body Algorithm for Short-Range Interactions

Michael Driscoll*, Evangelos Georganas*, Penporn Koanantakool*, Edgar Solomonik*, and Katherine Yelick*†

**Computer Science Division, University of California, Berkeley*

†*Lawrence Berkeley National Laboratory, Berkeley, CA*

{*driscoll,egeor,penpornk,solomon,yelick*}@cs.berkeley.edu

Abstract

We consider the problem of communication avoidance in computing interactions between a set of particles in scenarios with and without cutoff radius for interaction. Our strategy, which we show to be optimal in communication, divides the work in the iteration space rather than simply dividing the particles over processors, so more than one processor may be responsible for computing updates to a single particle. Similar to a *force decomposition* in molecular dynamics, this approach requires \sqrt{p} times more memory than a *particle decomposition*, but reduces communication costs by a factor of \sqrt{p} and is often faster in practice than a particle decomposition [1]. We examine a generalized force decomposition algorithm that tolerates the memory limited case, i.e. when memory can only hold c copies of the particles for $c = 1, 2, \dots, \sqrt{p}$. When $c = 1$, the algorithm degenerates into a particle decomposition; similarly when $c = \sqrt{p}$, the algorithm uses a force decomposition. We present a proof that the algorithm is communication-optimal and reduces critical path latency and bandwidth costs by factors of c^2 and c , respectively. Performance results from experiments on up to 24K cores of Cray XE-6 and IBM BlueGene/P machines indicate that the algorithm reduces communication in practice. In some cases, it even outperforms the original force decomposition approach because the right choice of c strikes a balance between collective and point-to-point communication cost. Finally, we extend the analysis to include a cutoff radius for direct evaluation of force interactions. We show that with a cutoff, communication optimality still holds. We describe a generalized algorithm for multi-dimensional space and assess its performance for 1D and 2D simulations on the same systems.

REFERENCES

- [1] S. Plimpton, "Fast parallel algorithms for short-range molecular dynamics," *J. Comput. Phys.*, vol. 117, no. 1, pp. 1–19, Mar. 1995. [Online]. Available: <http://dx.doi.org/10.1006/jcph.1995.1039>