Journal of Computational Physics 230 (2011) 3556–3569

Contents lists available at ScienceDirect

Journal of Computational Physics

journal homepage: www.elsevier.com/locate/jcp

Fast analysis of molecular dynamics trajectories with graphics processing units—Radial distribution function histogramming

Benjamin G. Levine [a,*,1], John E. Stone [b,1], Axel Kohlmeyer [a]

[a] Institute for Computational Molecular Science and Department of Chemistry, Temple University, Philadelphia, PA, United States
[b] Beckman Institute, University of Illinois at Urbana-Champaign, Urbana, IL, United States

ARTICLE INFO

Article history:
Received 15 September 2010
Received in revised form 22 January 2011
Accepted 31 January 2011
Available online 26 February 2011

Keywords:
Pair distribution function
Two-point correlation function
GPGPU

ABSTRACT

The calculation of radial distribution functions (RDFs) from molecular dynamics trajectory data is a common and computationally expensive analysis task. The rate limiting step in the calculation of the RDF is building a histogram of the distance between atom pairs in each trajectory frame. Here we present an implementation of this histogramming scheme for multiple graphics processing units (GPUs). The algorithm features a tiling scheme to maximize the reuse of data at the fastest levels of the GPU's memory hierarchy and dynamic load balancing to allow high performance on heterogeneous configurations of GPUs. Several versions of the RDF algorithm are presented, utilizing the specific hardware features found on different generations of GPUs. We take advantage of larger shared memory and atomic memory operations available on state-of-the-art GPUs to accelerate the code significantly. The use of atomic memory operations allows the fast, limited-capacity on-chip memory to be used much more efficiently, resulting in a fivefold increase in performance compared to the version of the algorithm without atomic operations. The ultimate version of the algorithm running in parallel on four NVIDIA GeForce GTX 480 (Fermi) GPUs was found to be 92 times faster than a multithreaded implementation running on an Intel Xeon 5550 CPU. On this multi-GPU hardware, the RDF between two selections of 1,000,000 atoms each can be calculated in 26.9 s per frame. The multi-GPU RDF algorithms described here are implemented in VMD, a widely used and freely available software package for molecular dynamics visualization and analysis.

1. Introduction

The increase in available computing power in recent years has been a boon for computational chemists wishing to simulate larger systems over longer timescales, but the ability to create massive quantities of molecular dynamics trajectory data also creates difficulties. Without advanced data analysis software, computationally expensive analysis tasks can become a bottleneck in the discovery process. One such task is the calculation of the radial distribution function (RDF).

The RDF is an important measure of the structure of condensed matter for several reasons. Radial distribution functions can be determined both experimentally and from simulation, allowing direct comparison. In addition, all thermodynamic quantities can be derived from an RDF under the assumption of a pair-wise additive potential energy function [1,2]. The RDF has long been applied as a descriptor of the structure of liquids such as water [3–6], and though they can be very computationally expensive to calculate, RDFs derived from large-scale molecular dynamics (MD) simulations have been useful in

* Corresponding author.
E-mail addresses: ben.levine@temple.edu (B.G. Levine), johns@ks.uiuc.edu (J.E. Stone), akohlmey@temple.edu (A. Kohlmeyer).
[1] B.G.L. and J.E.S. contributed equally to this work.

0021-9991/$ - see front matter © 2011 Elsevier Inc. All rights reserved.
doi:10.1016/j.jcp.2011.01.048

a wide range of applications. For example, Kadau and coworkers investigated shock wave induced phase transitions in metals using radial distribution functions calculated from simulations of systems with eight million atoms [7]. Radial distribution functions calculated from systems of several hundred thousand to one million atoms have also been useful in studies of radiation damage in nuclear waste [8] and long-range order in self-assembled alkanethiol monolayers [9]. The RDF is also widely used in astrophysics, where stars replace atoms and the function is typically known as the two-point correlation function [10].

Massive molecular dynamics simulations like those cited above were once unusual, but now are becoming common. The extreme computational expense of data analysis of this type requires that we bring to bear computers as powerful as those used to run production simulations. Sometimes it is surprising which hardware offers the greatest performance to scientists, though. The introduction of the Beowulf cluster marked an important change in high performance computing [11]. Unlike previous high performance computers which were based on expensive, proprietary hardware, Beowulf clusters utilized inexpensive personal computers and commodity server hardware in large quantities to perform scientific tasks. Beowulf clusters soon became the standard in high performance computing because commodity hardware provided more computation per dollar spent than did the more expensive proprietary alternatives.

Recently the computer game market has driven the development of graphics processing units (GPUs) which provide much faster floating point performance than a typical CPU at a comparable price. As such they have been receiving a great deal of attention from scientists wishing to accelerate their applications [12]. Making use of massively parallel processors and high bandwidth memory systems, GPUs have already been applied to accelerate a wide variety of methods in computational chemistry and biomolecular simulation [13–34]. The first generation of large scale heterogeneous clusters based on highly parallel commodity processors are already online—e.g. Los Alamos National Laboratory's Roadrunner [35], the National Center for Supercomputer Applications' Lincoln [36], and Texas Advanced Computing Center's Longhorn [37]—and three GPU-based clusters are now among the ten fastest supercomputers in the world, with the top place currently held by a GPU-based cluster [38]. With additional large-scale GPU-based clusters planned [39], it appears that technology developed for the gaming market will increase the capability of available scientific computing resources dramatically.

One of the most attractive features of GPUs, however, is that they are already present in a typical desktop workstation where they accelerate visualization software. As such, it is natural to employ them not only to speed up large scale simulations, but also time consuming data analysis tasks which a scientist would typically perform on their local desktop machine. By executing such tasks on GPUs one accelerates the discovery process; data analysis that used to require a cluster can be run on a desktop, and time consuming tasks formerly run only in batch mode can be performed interactively.

One example of a visualization and analysis software package for molecular dynamics (MD) data which has begun to take advantage of GPU acceleration is VMD [40]. Specifically, a fast implementation of electrostatic and nonbonded force calculations is used to place ions and calculate time averaged potentials from MD trajectories [28,41].

In this work we have implemented the calculation of the RDF from molecular dynamics trajectory data on NVIDIA GPUs into VMD using the CUDA parallel programming architecture [42]. The computation time of the task, inherent data parallelism, and opportunity for data reuse make RDF calculation a perfect target for GPU acceleration. However, the calculation of an RDF requires histogramming, which can be difficult to parallelize. In Section 2 of this paper we define the RDF histogramming problem, describe the difficulties encountered in developing a parallel implementation, and present our GPU-accelerated solution. In Section 3 we present the results of our optimization and benchmarks that analyze the performance of our implementation on several generations of NVIDIA GPU hardware. In Section 4 we draw conclusions from our work.

2. Methods

The radial distribution function calculation contains several component algorithm steps. All of the steps can be formulated as data-parallel algorithms, but the histogramming operations are more difficult to adapt to the massively parallel architecture of GPUs, and are therefore the main focus of the discussion. Below we introduce the mathematical basis for computing radial distribution functions and describe how this relates to a naive serial implementation. We then describe high performance data-parallel algorithms for the histogram computation component of RDF calculation on multi-core CPUs and GPUs and the attributes that affect their performance.

2.1. RDF math and serial histogramming

The radial distribution function, $g(r)$, is defined,

$$g(r) = \lim_{dr \to 0} \frac{p(r)}{4\pi(N_{pairs}/V)r^2 dr} \tag{1}$$

where r is the distance between a pair of particles, $p(r)$ is the average number of atom pairs found at a distance between r and $r + dr$,V is the total volume of the system, and N_{pairs} is the number of unique pairs of atoms where one atom is from each of two sets (selections), sel_1 and sel_2. The definition of N_{pair} is given for two special cases by the following equations; the cases where $sel_1 = sel_2$ and where there are no atoms shared between sel_1 and sel_2 are given in (2) and (3), respectively.