ARTICLE IN PRESS

Computer Physics Communications ▮ (▮▮▮▮) ▮▮▮–▮▮▮

Contents lists available at ScienceDirect

Computer Physics Communications

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

Asynchronous replica exchange software for grid and heterogeneous computing

Emilio Gallicchio [a,*], Junchao Xia [b], William F. Flynn [b], Baofeng Zhang [a], Sade Samlalsingh [a], Ahmet Mentes [b], Ronald M. Levy [b,*]

[a] Department of Chemistry, Brooklyn College of the City University of New York, Brooklyn, NY, United States

[b] Center for Biophysics and Computational Biology, Institute of Computational Molecular Science and Department of Chemistry, Temple University, Philadelphia, PA, United States

ARTICLE INFO

Article history:
Received 3 April 2015
Received in revised form
6 June 2015
Accepted 14 June 2015
Available online xxxx

Keywords:
Replica exchange molecular dynamics
Grid computing
BOINC
Distributed computing
Protein–ligand binding
Peptide dimerization

ABSTRACT

Parallel replica exchange sampling is an extended ensemble technique often used to accelerate the exploration of the conformational ensemble of atomistic molecular simulations of chemical systems. Inter-process communication and coordination requirements have historically discouraged the deployment of replica exchange on distributed and heterogeneous resources. Here we describe the architecture of a software (named ASyncRE) for performing asynchronous replica exchange molecular simulations on volunteered computing grids and heterogeneous high performance clusters. The asynchronous replica exchange algorithm on which the software is based avoids centralized synchronization steps and the need for direct communication between remote processes. It allows molecular dynamics threads to progress at different rates and enables parameter exchanges among arbitrary sets of replicas independently from other replicas. ASyncRE is written in Python following a modular design conducive to extensions to various replica exchange schemes and molecular dynamics engines. Applications of the software for the modeling of association equilibria of supramolecular and macromolecular complexes on BOINC campus computational grids and on the CPU/MIC heterogeneous hardware of the XSEDE Stampede supercomputer are illustrated. They show the ability of ASyncRE to utilize large grids of desktop computers running the Windows, MacOS, and/or Linux operating systems as well as collections of high performance heterogeneous hardware devices.

© 2015 Elsevier B.V. All rights reserved.

1. Introduction

Many physiochemical processes, such as phase transitions [1] and the folding and binding of proteins, occur on time scales difficult to access even with the fastest supercomputers available [2–4]. Enhanced conformational sampling algorithms have been developed to accelerate the modeling of these processes so that they can be studied in reasonable time frames [5–8]. An important class of methods is based on the imposition of thermodynamic or mechanical biasing forces which can speed up, often by many orders of magnitude, conformational interconversions otherwise too rare to be observed in traditional simulations [9–20]. Results produced by biased sampling methods are typically analyzed using post-processing to recover true (unbiased) thermodynamic observables [21–26].

Replica Exchange (RE) algorithms [27,28] are recognized as among the most powerful enhanced conformational sampling tools available [29–34], yielding converged results orders of magnitude faster than conventional approaches [7]. RE methods are based on replicating the system across combinations of thermodynamic and potential energy parameters [35] (in this work primarily alchemical receptor–ligand coupling λ and temperature T), each evolving independently with the exception of occasional exchanges of thermodynamic parameters between replicas. Hence, each replica travels in conformational space as well as in thermodynamic space. Accelerated conformational sampling is achieved because conformational transitions can occur at the thermodynamic state where they are most likely, rather than only at the thermodynamic state of interest, where they may be rare [36].

Because RE is inherently a parallel sampling algorithm, it is particularly suited for large high performance computing (HPC) clus-

* Corresponding authors.
E-mail addresses: egallicchio@brooklyn.cuny.edu (E. Gallicchio), ronlevy@temple.edu (R.M. Levy).

http://dx.doi.org/10.1016/j.cpc.2015.06.010
0010-4655/© 2015 Elsevier B.V. All rights reserved.

ters. In conventional *synchronous* implementations of RE, parameter exchanges occur simultaneously for all replicas after, for example, the completion of a given number of MD steps. Synchronous RE is a suitable algorithm for dedicated High Performance Computing (HPC) clusters where MD threads can efficiently execute in parallel at equal speeds for extended periods of time without failures. In these cases exchanges among replicas are implemented using inter-process communication libraries, such as the Message Passing Interface (MPI). On HPC resources a high rate of exchanges, which is beneficial for sampling efficiency [37,38], can be achieved with minimal overhead.

However, unlike methods based on swarms of independent trajectories [39–41], traditional synchronous RE approaches are not well suited to distributed and dynamically allocated resources, such as those available on computational grids [42]. In these environments direct communication across compute nodes is typically not available, and the pool of compute nodes varies unpredictably. Synchronous RE is also not well suited to heterogeneous HPC environments, such as the XSEDE Stampede cluster, which are increasingly becoming common place. Typically these clusters provide large computational throughput by utilizing highly parallel co-processors (many-core integrated co-processors, MIC's, or general programmable graphic processing units, GPGPU's) attached to conventional CPU nodes. Effective use of these computing devices is challenging due to their wide range of performance profiles and modes of operation. Legacy MD packages for biomolecular modeling often lack the ability to exploit multiple heterogeneous computing devices concurrently; in part because of the large diversity of hardware configurations that would need to be supported.

Finally, conventional synchronous RE implementations require that computational resources be secured for all of the replicas before the simulation can begin execution, and that these are maintained until the simulation is completed. This effectively precludes the use of general computing resources, such as farms of office desktops, that may be too small to run large RE applications all at once. Conversely, zero fault tolerance complicates the deployment of multi-dimensional RE algorithms, employing hundreds to thousands of replicas, on HPC resources.

The replica exchange method itself does not impose the restriction that exchanges should necessarily occur synchronously across all processors [43]. In particular the RE method itself does not require that all of the replicas be running at the same time. There are therefore no obstacles in principle preventing the deployment of *asynchronous* RE algorithms over distributed and heterogeneous computing infrastructures. Previous efforts have showed that the asynchronous prescription has significant advantages over the conventional synchronous implementation in terms of scalability and flexibility in the choice of exchange schemes [44].

Towards the goal of filling the current gap in availability of RE software tools capable of harnessing the latent computational power of distributed and heterogeneous computing resources, in this work we present an asynchronous replica exchange software framework, named ASyncRE, based on the idea of avoiding centralized synchronization steps by allowing replicas to progress at different rates and enabling parameter exchanges among arbitrary sets of replicas independently from other replicas. As illustrated here, the proposed ASyncRE framework is sufficiently flexible to drive RE molecular simulations on high end heterogeneous supercomputers, such as the XSEDE Stampede cluster, as well as on distributed computing grids, such as those built around the BOINC software platform. Campus computational grids that utilize existing general computing resources, such as those in student PC labs, when they would be otherwise idle, are attractive research computing alternatives to HPC clusters. It would be very beneficial to extend the domain of applicability of RE to campus-wide distributed computing grids and much larger collections of computers such as the World Community Grid (worldcommunitygrid.org).

To achieve these goals, while addressing as generally as possible the computational challenges described above, a number of design choices were made. For optimal flexibility, ASyncRE is implemented as a managing agent of the underlying MD engine, rather than being integral part of it. To ensure resiliency, ASyncRE issues MD executions of replicas in discrete chunks, typically lasting from a few minutes to several hours, after which the replica is checkpointed on a coordination server. Executions can be launched on a diverse pool of computing devices ranging from parallel co-processors to user desktops. Parameter exchanges are performed by the coordination server by manipulating output and input files of the replicas currently checkpointed, thereby avoiding synchronization bottlenecks and the need for direct communication links between computing devices. The specific algorithms and techniques used to achieve these characteristics are described in detail below, followed by illustrative applications of the software to the study of supramolecular and macromolecular complexes. The ASyncRE software is freely available at https://github.com/ComputationalBiophysicsCollaborative/AsyncRE.

2. Methods

2.1. Replica exchange conformational sampling

The objective of equilibrium molecular simulations is to efficiently sample the canonical distribution, $\exp[-\beta U(x;\theta)]$, where x represents a molecular conformation, $\beta = 1/k_B T$ is the inverse temperature, and $U(x;\theta)$ is the potential energy function, dependent on a set of parameters collectively denoted by θ. Collecting sample sets of sufficient quality is often very challenging due to the high dimensionality of the sampling space and the ruggedness of the potential. General-purpose sampling algorithms such as Monte Carlo (MC) and Molecular Dynamics (MD) often remain trapped into metastables conformations isolated by high potential energy barriers which are crossed only rarely at standard conditions. Sampling can be accelerated by performing biased sampling at higher temperatures (small β) or at suitable values of potential energy parameters, at which the likelihood of trapping is decreased [45–48]. The problem of correcting the probabilities of samples so that they reflect the experimental conditions of interest is addressed using thermodynamic reweighting techniques [24,25].

In replica exchange the potential energy parameters and/or the temperature are sampled across a discrete set of values, which include those of interest, together with molecular conformations in such a way that the canonical distribution is sampled at each choice of parameters and temperature. Defining the reduced potential energy function as

$$u(x;s) = \beta U(x;\theta), \tag{1}$$

where in the most general case $s = (\beta, \theta)$ denotes a joint thermodynamic and mechanical state of the system, the Replica Exchange (RE) algorithm involves simulating M replicas of the system at states $s_i = (\beta_i, \theta_i)$. Conformational sampling of each replica at fixed s_i is performed using standard MD. A replica can change its thermodynamic/mechanical state (hereafter referred to simply as a "state") by exchanging its current state with that of another replica. The probability of exchange is regulated to ensure microscopic reversibility, that is equilibration towards the canonical distribution at every value of $s_i = (\beta_i, \theta_i)$.

Formally, the RE ensemble distribution

$$p_{\mathrm{RE}}(x_1, x_2, \ldots, x_M | \{s\}) = \exp\left[-\sum_{i=1}^{M} u(x_i; s_i)\right] \Big/ Z_{\mathrm{RE}} \tag{2}$$

is introduced, where x_i is the molecular configuration of replica i and s_i is the state assigned to it, which is taken from a discrete set

Please cite this article in press as: E. Gallicchio, et al., Asynchronous replica exchange software for grid and heterogeneous computing, Computer Physics Communications (2015), http://dx.doi.org/10.1016/j.cpc.2015.06.010

Download English Version:

https://daneshyari.com/en/article/6919840

Download Persian Version:

https://daneshyari.com/article/6919840

Daneshyari.com