Published in final edited form as:

Comput Phys Commun. 2018 June ; 227: 99–108. doi:10.1016/j.cpc.2018.02.005.

RPYFMM: Parallel Adaptive Fast Multipole Method for Rotne-Prager-Yamakawa Tensor in Biomolecular Hydrodynamics Simulations

W. Guan^a, X. Cheng^b, J. Huang^a, G. Huber^d, W. Li^c, J. A. McCammon^d, and B. Zhang^{e,*}

^aDepartment of Mathematics, University of North Carolina, Chapel Hill, NC 27599-3250, USA

^bMedicinal Chemistry and Pharmacognosy, College of Pharmacy, Ohio State University, Columbus, OH 43210, USA

^cSchool of Transportation and Vehicle Engineering, Shandong University of Technology, Zibo, Shandong, China

^dDepartment of Chemistry and Biochemistry and Department of Pharmacology, University of California at San Diego, La Jolla, CA 92093-0365

^eCenter for Research in Extreme Scale Technologies, School of Informatics, Computing, and Engineering, Indiana University, Bloomington, IN, 47404, USA

Abstract

RPYFMM is a software package for the efficient evaluation of the potential field governed by the Rotne-Prager-Yamakawa (RPY) tensor interactions in biomolecular hydrodynamics simulations. In our algorithm, the RPY tensor is decomposed as a linear combination of four Laplace interactions, each of which is evaluated using the adaptive fast multipole method (FMM) [1] where the exponential expansions are applied to diagonalize the multipole-to-local translation operators. RPYFMM offers a unified execution on both shared and distributed memory computers by leveraging the DASHMM library [2, 3]. Preliminary numerical results show that the interactions for a molecular system of 15 million particles (beads) can be computed within one second on a Cray XC30 cluster using 12, 288 cores, while achieving approximately 54% strong-scaling efficiency.

Keywords

Brownian dynamics; Rotne-Prager-Yamakawa tensor; Hydrodynamics interactions; Fast multipole method; DASHMM

*Corresponding author. zhang416@indiana.edu.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

1. Introduction

The dynamics of macromolecules inside a living cell usually takes place at very low Reynolds numbers, where the viscous forces dominate over inertial effects, and the *in vivo* macromolecule diffusion is strongly influenced by the hydrodynamic interactions (HIs). Due to its long-range nature and many-body character, HI is responsible for a wide variety of fascinating collective phenomena. Depending on whether HI is present or not, existing studies have shown qualitative differences in the aggregation or microphase separation of colloids [4]. In particular, HI was shown to facilitate barrier crossing during the microphase separation of block copolymers, whereas without HI, the system appeared to be trapped in a metastable state. HI was also found to greatly accelerate the kinetics of lipid membrane self-assembly [5].

There have been many research efforts to develop accurate mathematical models and efficient simulation tools for understanding the HI effects on biomolecular dynamics. In this paper, we focus on the bead model [6], a particular realization of the commonly used Ermak-McCammon model [7]. In this model, the biomolecular system is treated as a system of N beads, each representing a Brownian molecule, the translational and rotational displacement $\mathbf{X}_i$ of the i th bead during time step Δt , due to external force $\mathbf{F}_i$ acting on bead i , $i = 1, \dots, N$, and the random displacement $\mathbf{R}_i$ is given by

$$\Delta \mathbf{X}_i(\Delta t) = \sum_j \frac{D_{ij} \mathbf{F}_j}{k_B T} \Delta t + \sum_j \frac{\partial D_{ij}}{\partial \mathbf{r}_j} \Delta t + \mathbf{R}_i(\Delta t), \quad (1)$$

where k_B is the Boltzmann constant and T is the absolute temperature. The external forces $\mathbf{F}$ include the electrostatic and van der Waals interactions. The hydrodynamic forces are accounted for by the $6N \times 6N$ diffusion tensor $\mathbf{D}$ that describes the hydrodynamic coupling between the N beads with their three translational and three rotational degrees of freedom. The random displacements in $\mathbf{R}$ have mean zero and variance matrix $2\mathbf{D} \Delta t$. When the rotational motions of the N beads of radii a are neglected, the dimension of $\mathbf{D}$ is reduced to $3N \times 3N$ and the most common form is given by the Rotne-Prager-Yamakawa (RPY) tensor [7, 8] as follows:

$$D_{ii}(\mathbf{r}_i, \mathbf{r}_i) = \frac{k_B T}{6\pi\eta a} \mathbf{I}, \quad (2)$$

$$D_{ij}(\mathbf{r}_i, \mathbf{r}_j) = \frac{k_B T}{8\pi\eta r_{ij}} \left[\left(\mathbf{I} + \frac{\mathbf{r}_{ij} \otimes \mathbf{r}_{ij}}{r_{ij}^2} \right) + \frac{2a^2}{3r_{ij}^2} \left(\mathbf{I} - \frac{3\mathbf{r}_{ij} \otimes \mathbf{r}_{ij}}{r_{ij}^2} \right) \right], \quad i \neq j, r_{ij} \geq 2a \quad (3)$$

$$D_{ij}(\mathbf{r}_i, \mathbf{r}_j) = \frac{k_B T}{6\pi\eta r_{ij}} \left[\left(1 - \frac{9r_{ij}}{32a}\right)I + \frac{3}{32a} \frac{\mathbf{r}_{ij} \otimes \mathbf{r}_{ij}}{r_{ij}^2} \right], \quad i \neq j, r_{ij} < 2a \quad (4)$$

Here η is the solvent viscosity, i and j label bead indices, I is the 3×3 identity matrix, $\mathbf{r}_i = [x_i, y_i, z_i]^T$ is the 3×1 position vector of bead i , $\mathbf{r}_{ij} = \mathbf{r}_j - \mathbf{r}_i$, and $r_{ij} = \|\mathbf{r}_{ij}\|_2$ is the distance between beads i and j . In this paper, we focus on the positive definite RPY tensor. The generalized RPY tensor with both translational and rotational motions was presented in [9] and its efficient evaluation is still an active research topic.

There are several numerical difficulties in solving (1) efficiently, including the evaluation of the electrostatic force field contributing to the external forces F , and the algebraic operations on the dense diffusion matrix D . For instance, direct calculation of all the two-body HI interactions requires $O(N^2)$ operations, and computing hydrodynamically correlated random displacement vectors R requires prohibitive $O(N^3)$ operations via the Cholesky factorization. For these reasons, in most previous studies, HI has been either completely neglected or considered only for a much smaller equivalent sphere system, where the detailed molecular shape was ignored, and each protein—a large set of beads—was modeled by an equivalent sphere of the same hydrodynamic radius.

There exist many research efforts aimed at reducing the computational complexity of solving (1), either for the steady state or dynamic settings. A few representative work include the parallel adaptive fast multipole method for evaluating the electrostatic potential modeled by the linearized Poisson-Boltzmann equation on distributed memory computers [10], the Particle-Mesh Ewald (PME) summation method for the matrix-vector multiplication (DF) which scales as $O(N \log N)$ [11], and different techniques for the efficient generation of the random vector R in [12–15].

The purpose of this paper is to present a parallel software package for the efficient evaluation of the HI interactions ($D \cdot F$) modeled by the RPY tensor. The package is also essential to the efficient generation of the random vectors R by approximating $\sqrt{D}v$ with $\{D^k v\}$, $k = 0, 1, 2, \dots, k$ via the preconditioned Krylov subspace iterations [16]. Our numerical algorithm uses the technique introduced in [17] for Stokeslet and [18] for RPY tensor to decompose the RPY tensor as a linear combination of four Laplace potentials and their derivatives. Each Laplace potential is evaluated using the adaptive new version fast multipole method [1] where the exponential expansion is used to diagonalize the multipole-to-local translation operators. The package is built on top of the open-source DASHMM library [2, 3] developed by some of the authors, and the Asynchronous Multi-Tasking HPX-5 runtime, providing a unified execution on both shared and distributed memory computers. Preliminary numerical results show that for a molecular system with 15 million beads, the package is able to compute the HI within one second on a Cray XC30 cluster using 12, 288 cores and achieves 54% strong-scaling efficiency.

This paper is organized as follows. Section 2 reviews the mathematical foundations of RPYFMM, including the decomposition of the RPY tensor and how to compute the gradient

and Hessian of the Laplace potentials. Section 3 describes the main components of HPX-5 runtime system and parallelization strategy of DASHMM library. Section 4 provides the installation guide and job examples, and demonstrates the strong-scaling efficiency for different accuracy requirements. Section 5 concludes the paper by discussing several related research topics in order to build the next generation of Brownian dynamics simulation package.

2. Mathematical foundations of RPYFMM

Similar to the electrostatic interaction, the hydrodynamics interaction (HI) modeled by the RPY tensor decays like $O(1/|\mathbf{r}_i - \mathbf{r}_j|)$ as $\|\mathbf{r}_i - \mathbf{r}_j\|_2 \rightarrow \infty$ and is therefore considered long range. To efficiently evaluate these long-range hydrodynamics interactions, we apply the technique first proposed for the Stokeslet [17] and later generalized to the RPY tensor [18] to decompose the far-field RPY tensor (3) as a linear combination of four scalar Laplace potentials and their derivatives. For a target bead i located at $\mathbf{r}_i = [x_i, y_i, z_i]^T$, denote the set of well-separated beads by Ω_i^{far} , where each bead $j \in \Omega_i^{far}$ is located at $\mathbf{r}_j = [x_j, y_j, z_j]^T$ exerting force $\mathbf{F}_j = [F_j^1, F_j^2, F_j^3]^T$. The four Laplace potentials are defined as follows:

$$L_1(\mathbf{r}_i) = \sum_{j \in \Omega_i^{far}} \frac{F_j^1}{|\mathbf{r}_i - \mathbf{r}_j|} \quad (5)$$

$$L_2(\mathbf{r}_i) = \sum_{j \in \Omega_i^{far}} \frac{F_j^2}{|\mathbf{r}_i - \mathbf{r}_j|} \quad (6)$$

$$L_3(\mathbf{r}_i) = \sum_{j \in \Omega_i^{far}} \frac{F_j^3}{|\mathbf{r}_i - \mathbf{r}_j|} \quad (7)$$

$$L_4(\mathbf{r}_i) = \sum_{j \in \Omega_i^{far}} \frac{F_j^1 x_j + F_j^2 y_j + F_j^3 z_j}{|\mathbf{r}_i - \mathbf{r}_j|} = \sum_{j \in \Omega_i^{far}} \frac{\mathbf{F}_j \cdot \mathbf{r}_j}{|\mathbf{r}_i - \mathbf{r}_j|} \quad (8)$$

To simplify the notations later, we further define

$$L_C(\mathbf{r}_i) = \frac{\partial L_1(\mathbf{r}_i)}{\partial x_i} + \frac{\partial L_2(\mathbf{r}_i)}{\partial y_i} + \frac{\partial L_3(\mathbf{r}_i)}{\partial z_i}. \quad (9)$$

Using these notations, the far-field HI at target i due to contributions from Ω_i^{far} can be collected as

$$\begin{aligned} \phi_{far}(\mathbf{r}_i) = [u_i, v_i, w_i]^T &= \sum_{j \in \Omega_i^{far}} D_{ij}(\mathbf{r}_i, \mathbf{r}_j) F_j = \\ &= \begin{bmatrix} C_1 L_1 - C_1 \left(x_i \frac{\partial L_1}{\partial x_i} + y_i \frac{\partial L_2}{\partial x_i} + z_i \frac{\partial L_3}{\partial x_i} \right) + C_1 \frac{\partial L_4}{\partial x_i} + C_2 \frac{\partial L_C}{\partial x_i} \\ C_1 L_2 - C_1 \left(x_i \frac{\partial L_1}{\partial y_i} + y_i \frac{\partial L_2}{\partial y_i} + z_i \frac{\partial L_3}{\partial y_i} \right) + C_1 \frac{\partial L_4}{\partial y_i} + C_2 \frac{\partial L_C}{\partial y_i} \\ C_1 L_3 - C_1 \left(x_i \frac{\partial L_1}{\partial z_i} + y_i \frac{\partial L_2}{\partial z_i} + z_i \frac{\partial L_3}{\partial z_i} \right) + C_1 \frac{\partial L_4}{\partial z_i} + C_2 \frac{\partial L_C}{\partial z_i} \end{bmatrix} = C_1 \begin{bmatrix} L_1 \\ L_2 \\ L_3 \end{bmatrix} - C_1 (x_i \nabla L_1 + y_i \nabla L_2 \\ &+ z_i \nabla L_3) + C_1 \nabla L_4 + C_2 \nabla L_C \end{aligned} \quad (10)$$

where $C_1 = \frac{k_B T}{8\pi\eta}$ and $C_2 = \frac{k_B T a^2}{12\pi\eta}$. Notice that to compute the HI in (10), one has to compute the three Laplace potentials L_1, L_2, L_3 , the gradients $\nabla L_1, \nabla L_2, \nabla L_3, \nabla L_4$, and the Hessians of L_1, L_2 , and L_3 implicitly expressed in ∇L_C .

In the FMM, far-field potentials are collected in the form of *multipole* or *local* expansions. In RPYFMM, the multipole (M) and local (L) expansions for the Laplace potential are of the form

$$M = \sum_{n=0}^p \sum_{m=-n}^n \frac{M_n^m Y_n^m}{r^{n+1}} \quad \text{and} \quad L = \sum_{n=0}^p \sum_{m=-n}^n L_n^m Y_n^m r^n,$$

where the spherical harmonics Y_n^m is defined as

$$Y_n^m(\theta, \phi) = \sqrt{\frac{(n-|m|)!}{(n+|m|)!}} P_n^{|m|}(\cos\theta) e^{im\phi} = C_n^m P_n^{|m|} e^{im\phi}$$

Under this definition, it is easy to verify that $M_n^{-m} = \overline{M_n^m}$ and $L_n^{-m} = \overline{L_n^m}$ and one only saves the coefficients for $0 \leq m \leq n$ in the implementation.

To compute the gradient and Hessian of the above multipole (local) expansion, one defines operators $\nabla_0 = \frac{\partial}{\partial z}$, $\nabla_p = \frac{\partial}{\partial x} + i \frac{\partial}{\partial y}$, and $\nabla_m = \frac{\partial}{\partial x} - i \frac{\partial}{\partial y}$. If ϕ is a harmonic function, then

$$\nabla_m \nabla_p \phi = \nabla_p \nabla_m \phi = -\nabla_0 \nabla_0 \phi.$$

When the spherical harmonics follow the conventional definition

$$Y_n^m(\theta, \phi) = \sqrt{\frac{2n+1}{4\pi}} \sqrt{\frac{(n-m)!}{(n+m)!}} P_n^m(\cos\theta) e^{im\phi}, \quad Y_n^{-m} = (-1)^m \overline{Y_n^m}, \quad 0 \leq m \leq n,$$

the following relations hold

$$\nabla_0 \left(\frac{Y_n^m}{r^{n+1}} \right) = -\sqrt{\frac{(n+m+1)(n-m+1)}{(2n+1)(2n+3)}} (2n+1) \frac{Y_{n+1}^m}{r^{n+2}} \quad (11)$$

$$\nabla_p \left(\frac{Y_n^m}{r^{n+1}} \right) = \sqrt{\frac{(n+m+2)(n+m+1)}{(2n+1)(2n+3)}} (2n+1) \frac{Y_{n+1}^{m+1}}{r^{n+2}} \quad (12)$$

$$\nabla_m \left(\frac{Y_n^m}{r^{n+1}} \right) = -\sqrt{\frac{(n-m+2)(n-m+1)}{(2n+1)(2n+3)}} (2n+1) \frac{Y_{n+1}^{m-1}}{r^{n+2}} \quad (13)$$

$$\nabla_0 (Y_n^m r^n) = \sqrt{\frac{(n+m)(n-m)}{(2n-1)(2n+1)}} (2n+1) r^{n-1} Y_{n-1}^m \quad (14)$$

$$\nabla_p (Y_n^m r^n) = \sqrt{\frac{(n-m-1)(n-m)}{(2n-1)(2n+1)}} (2n+1) r^{n-1} Y_{n-1}^{m+1} \quad (15)$$

$$\nabla_m(Y_n^m r^n) = \sqrt{\frac{(n+m-1)(n+m)}{(2n-1)(2n+1)}}(2n+1)r^{n-1}Y_{n-1}^{m-1}. \quad (16)$$

As the Laplace potential is real, the gradient and Hessian of the multipole (local) expansion can be obtained

$$\partial_z \leftarrow \nabla_0, \quad \partial_x \leftarrow \Re \nabla_p, \quad \partial_y \leftarrow \Im \nabla_p$$

$$\partial_{xz}^2 \leftarrow \Re \nabla_p \nabla_0, \quad \partial_{yz}^2 \leftarrow \Im \nabla_p \nabla_0, \quad \partial_{zz}^2 \leftarrow \nabla_0 \nabla_0, \quad \partial_{xy}^2 \leftarrow \Im \nabla_p \nabla_p / 2$$

$$\partial_{xx}^2 \leftarrow \Re(\nabla_p \nabla_p - \nabla_0 \nabla_0) / 2, \quad \partial_{yy}^2 \leftarrow \Re(\nabla_p \nabla_p - \nabla_0 \nabla_0) / 2 \quad (17)$$

Theorem 1

Let $M = \sum_{n=0}^p \sum_{m=-n}^n \frac{M_n^m C_n^m P_n^m e^{im\phi}}{r^{n+1}}$ be the multipole expansion. Then

$$\nabla_0 M = s_1 + 2\Re s_2, \quad \nabla_p M = s_3 + \overline{s_4},$$

$$\nabla_p \nabla_0 M = s_5 + \overline{s_6}, \quad \nabla_p \nabla_p M = s_7 + \overline{s_8}, \quad \nabla_0 \nabla_0 M = s_9 + 2\Re s_{10},$$

where

$$s_1 = - \sum_{n=0}^p \frac{M_n^0 P_{n+1}^0}{r^{n+2}}$$

$$s_2 = - \sum_{n=1}^p \sum_{m=1}^n \frac{M_n^m P_{n+1}^m e^{im\phi}}{r^{n+2}} \sqrt{\frac{(n-m+1)!(n-m+1)!}{(n-m)!(n+m)!}}$$

$$s_3 = \sum_{n=0}^p \sum_{m=0}^n \frac{M_n^m P_{n+1}^{m+1} e^{i(m+1)\phi}}{r^{n+2}} \sqrt{\frac{(n-m)!}{(n+m)!}}$$

$$s_4 = - \sum_{n=1}^p \sum_{m=1}^n \frac{M_n^m P_{n+1}^{m-1} e^{i(m-1)\phi}}{r^{n+2}} \sqrt{\frac{(n-m+2)!(n-m+2)!}{(n-m)!(n+m)!}}$$

$$s_5 = - \sum_{n=0}^p \sum_{m=0}^n \frac{M_n^m P_{n+2}^{m+1} e^{i(m+1)\phi}}{r^{n+3}} \sqrt{\frac{(n-m+1)!(n-m+1)!}{(n-m)!(n+m)!}}$$

$$s_6 = \sum_{n=1}^p \sum_{m=1}^n \frac{M_n^m P_{n+2}^{m-1} e^{i(m-1)\phi}}{r^{n+3}} \sqrt{\frac{(n-m+3)!(n-m+3)!}{(n-m)!(n+m)!}}$$

$$s_7 = \sum_{n=0}^p \sum_{m=0}^n \frac{M_n^m P_{n+2}^{m+2} e^{i(m+2)\phi}}{r^{n+3}} \sqrt{\frac{(n-m)!}{(n+m)!}}$$

$$s_8 = \sum_{n=2}^p \sum_{m=2}^n \frac{M_n^m P_{n+2}^{m-2} e^{i(m-2)\phi}}{r^{n+3}} \sqrt{\frac{(n-m+4)!(n-m+4)!}{(n-m)!(n+m)!}} - \sum_{n=1}^p \frac{M_n^1 P_{n+2}^1 e^{-i\phi}}{r^{n+3}} \sqrt{\frac{(n+1)!}{(n-1)!}}$$

$$s_9 = \sum_{n=0}^p \frac{M_n^0 P_{n+2}^0 e^{i(n+1)(n+2)\phi}}{r^{n+3}}$$

$$s_{10} = \sum_{n=1}^p \sum_{m=1}^n \frac{M_n^m P_{n+2}^m e^{im\phi}}{r^{n+3}} \sqrt{\frac{(n-m+2)!(n-m+2)!}{(n-m)!(n+m)!}}.$$

Proof 1

We show the result for $\nabla_P \nabla_0$. Apply $\nabla_P \nabla_0$ on each term of the multipole expansion. When $m \geq 0$,

$$\begin{aligned} \nabla_P \nabla_0 \left(\frac{M_n^m C_n^m P_n^m e^{im\phi}}{r^{n+1}} \right) &= - \nabla_P \left(M_n^m \sqrt{\frac{4\pi}{2n+1}} \sqrt{\frac{(n+m+1)(n-m+1)}{(2n+1)(2n+3)}} (2n+1) \frac{Y_{n+1}^m}{r^{n+2}} \right) \\ &= - \frac{M_n^m P_{n+2}^{m+1} e^{i(m+1)\phi}}{r^{n+3}} \sqrt{\frac{(n-m+1)!(n-m+1)!}{(n-m)!(n+m)!}} \end{aligned}$$

When $m = -1$,

$$\nabla_p \nabla_0 \left(\frac{\overline{M_n^m C_n^m P_n^m e^{-im\phi}}}{r^{n+1}} \right) = \nabla_p \nabla_0 \frac{\overline{M_n^m C_n^m P_n^m e^{-im\phi}}}{r^{n+1}} = \nabla_m \nabla_0 \frac{\overline{M_n^m C_n^m P_n^m e^{-im\phi}}}{r^{n+1}}$$

Similar algebraic work gives

$$\nabla_m \nabla_0 \left(\frac{M_n^m C_n^m P_n^m e^{-im\phi}}{r^{n+1}} \right) = \frac{M_n^m P_{n+2}^{m-1} e^{i(m-1)\phi}}{r^{n+3}} \sqrt{\frac{(n-m+3)!(n-m+3)!}{(n-m)!(n+m)!}}$$

The results for the other operators can be obtained with similar algebraic work.

One can similarly obtain the following result for the local expansion.

Theorem 2

Let $\mathbf{L} = \sum_{n=0}^p \sum_{m=-n}^n L_n^m C_n^m P_n^m e^{im\phi} r^n$ be the local expansion. Then

$$\nabla_0 \mathbf{L} = s_1 + 2\Re s_2, \quad \nabla_p \mathbf{L} = s_3 + \overline{s_4}$$

$$\nabla_p \nabla_0 \mathbf{L} = s_5 + \overline{s_6}, \quad \nabla_p \nabla_p \mathbf{L} = s_7 + \overline{s_8}, \quad \nabla_0 \nabla_0 \mathbf{L} = s_9 + 2\Re s_{10},$$

where

$$s_1 = \sum_{m=1}^p L_n^0 r^{n-1} P_{n-1}^0$$

$$s_2 = \sum_{n=1}^p \sum_{m=1}^{n-1} L_n^m r^{n-1} P_{n-1}^m e^{im\phi} \sqrt{\frac{(n+m)!(n-m)!}{(n+m-1)!(n+m-1)!}}$$

$$s_3 = \sum_{n=2}^p \sum_{m=0}^{n-2} L_n^m r^{n-1} P_{n-1}^{m+1} e^{i(m+1)\phi} \sqrt{\frac{(n-m)!}{(n+m)!}}$$

$$s_4 = - \sum_{n=1}^p \sum_{m=1}^n L_n^m r^{n-1} P_{n-1}^{m-1} e^{i(m-1)\phi} \sqrt{\frac{(n+m)!(n-m)!}{(n+m-2)!(n+m-2)!}}$$

$$s_5 = \sum_{n=3}^p \sum_{m=0}^{n-3} L_n^m P_{n-2}^{m+1} e^{i(m+1)\phi} r^{n-2} \sqrt{\frac{(n-m)!(n+m)!}{(n+m-1)!(n+m-1)!}}$$

$$s_6 = \sum_{n=2}^p \sum_{m=1}^{n-1} L_n^m P_{n-2}^{m-1} e^{i(m-1)\phi} r^{n-2} \sqrt{\frac{(n-m)!(n+m)!}{(n+m-3)!(n+m-3)!}}$$

$$s_7 = \sum_{n=4}^p \sum_{m=0}^{n-4} L_n^m P_{n-2}^{m+2} e^{i(m+2)\phi} r^{n-2} \sqrt{\frac{(n-m)!}{(n+m)!}}$$

$$s_8 = \sum_{n=2}^p \sum_{m=2}^n L_n^m P_{n-2}^{m-2} e^{i(m-2)\phi} r^{n-2} \sqrt{\frac{(n+m)!(n-m)!}{(n+m-4)!(n+m-4)!}} - \sum_{n=3}^p L_n^1 P_{n-2}^1 e^{-i\phi} r^{n-2} \sqrt{\frac{(n+1)!}{(n-1)!}}$$

$$s_9 = \sum_{n=2}^p L_n^0 P_{n-2}^{n(n-1)} r^{n-2}$$

$$s_{10} = \sum_{n=2}^p \sum_{m=1}^{n-2} L_n^m r^{n-2} P_{n-2}^m e^{im\phi} \sqrt{\frac{(n+m)!(n-m)!}{(n+m-2)!(n+m-2)!}}$$

We point out that there are at least two approaches to computing (10) at each target bead i . In the first approach, for each Laplace potential, one differentiates the multipole or local expansion, evaluates the potential, gradients, and Hessian at $\mathbf{r}_i$, and accumulates the result. In the second approach, one simplifies each component in (10) into a single expansion and then carries out the evaluation. Compared with the first approach, the second approach performs fewer expansion evaluation but consumes more storage to assemble the final expansion. In RPYFMM, the current implementation adopts the first approach.

3. HPX-5 runtime and DASHMM library

3.1. HPX-5 runtime

HPX-5 (High Performance ParalleX) is an experimental Asynchronous Multi-Tasking (AMT) programming model and runtime developed at Indiana University [19, 20]. Its design is governed by the ParalleX exascale execution model [21] and it aims to enable programs to run unmodified on systems from a single SMP to large clusters and supercomputers with thousands of nodes.

HPX-5 defines a broad API that covers most aspects of the system. Programs are organized as diffusive, message driven computation, consisting of a large number of lightweight

threads and active messages, executing within the context of a global address space, and synchronizing through the use of lightweight synchronization objects. The HPX-5 runtime is responsible for managing global allocation, address resolution, data and control dependence, and scheduling threads and the network.

The HPX-5 *global address space* provides a global shared memory space abstraction and serves as the basis for computation. Global allocation is performed through a set of dynamic allocators that provide individual, block cyclic, and user-defined allocation for blocks of memory. Access to data in the global address space is provided through an asynchronous memput/memget API. Explicit global address translation can be performed in order to operate on local machine virtual aliases. Finally, raw global addresses serve as the targets for HPX-5's active message *parcels*, described below. Localities (roughly equivalent to MPI processes) are mapped into the global address space and can be accessed through indices allowing messages to target localities as in other active message runtimes.

Parcels form the basis of parallel computation in HPX-5. They contain a description of the action to be performed, argument data, and continuation information and are sent to the global address on which the action is to be performed. The HPX-5 scheduler invokes parcels as lightweight threads once they reach their destination. This parcel-thread equivalence is key to abstracting the difference between shared and distributed execution in HPX-5. Sending a parcel is equivalent to, and the only means of, spawning a lightweight thread. In shared memory execution it just happens that all target addresses are on a single locality. Unlike many other AMT runtimes, HPX-5 is designed around cooperative threading and not simply run to completion tasks.

Program data and control dependencies are represented in memory by *local control objects* (LCOs). An LCO is an event-driven, lightweight, globally addressable synchronization object that co-locates data and control information. All LCOs have input slots, predicates that evaluate functions of the inputs and may determine that an LCO has been triggered, and continuations (i.e., dependent threads and parcels) that will be executed once the LCO is triggered. This allows the user to build fully dynamic dataflow networks managed by the runtime. A simple example of an LCO is a reduction that performs a sum across its inputs. HPX-5 is delivered with a number of classes of built-in LCOs, e.g., futures and reduction types, and permits user-defined LCO classes as well. A user-defined LCO encodes the data that it represents, the task performed when an input becomes available, and the predicate under which the LCO is considered to be triggered.

3.2. DASHMM

The Dynamic Adaptive System for Hierarchical Multipole Methods is an open-source scientific software library built on top of HPX-5 runtime system that aims to provide an easy-to-use system that can provide scalable, efficient, and unified evaluations of general HMMs on both shared and distributed memory architectures. Unlike conventional practice in many existing MPI+X implementations, which use static partitioning of the global tree structure and bulk-synchronous communication of the locally essential tree [22–24], DASHMM considers the distribution of the directed acyclic execution graph, represented implicitly as a network of expansion objects. By leveraging the LCO construct in HPX-5,

the expansion object not only encloses data, but also encodes dependency and continuation. The execution of DASHMM is therefore completely data-driven and asynchronous.

DASHMM can be used for a wide variety of applications, and it does so with an easy-to-use and general interface. Further, DASHMM's interface is HPX-5 oblivious, meaning that no knowledge of HPX-5 is required to use DASHMM in either its basic or advanced forms. The basic interface allows users to rapidly apply the built in methods and kernels in end-science applications. DASHMM currently provides three built in multipole methods: Barnes-Hut [25], the classic fast multipole method [26, 27], and a variant of FMM that uses the exponential expansion to reduce arithmetic complexity [1]. The last method is referred to as FMM97 in DASHMM. DASHMM also provides three built in kernels: the scaling invariant Laplace potential, the scaling variant Yukawa potential, and the oscillatory Helmholtz potential in low-frequency regime. The advanced interface allows users to implement variants of the multipole method, new kernels, or new distribution policies that guide the placement of expansion data.

The simplicity and generality of DASHMM are achieved by adopting two strategies. First, DASHMM uses C++ templates to accept `SourceData`, `TargetData`, `Method`, and `Expansion` as parameters for automatic generalization over user-specified types. Second, DASHMM extends the `Expansion` abstraction from its mathematical counterpart and introduces three additional operators. Each expansion object in DASHMM is either a normal expansion that corresponds to the usual multipole/local expansion or an intermediate expansion that is needed for advanced multipole methods (see Figure 1). Additionally, each expansion object in DASHMM is a collection of views, each of which is a mathematical approximation stored in compact form. Together, they facilitate the implementation of various multipole and multipole-like methods and concurrent kernel evaluation on the same input data.

4. Software Installation and Numerical Examples

4.1. Installation

RPYFMM depends on two external libraries: the HPX-5 runtime system, and the DASHMM library. The current version of RPYFMM depends on version 4.1.0 of HPX-5 or later, which can be downloaded from <https://hpx.crest.iu.edu/download>. DASHMM is automatically downloaded in the RPYFMM when the library is built.

Users must first install HPX-5 on their systems. HPX-5 can be built without or with network transports. For the latter, HPX-5 currently specifies two network interfaces: the `ISend/IRecv` (ISIR) interface with the MPI transport, and `Put-with-Completion` (PWC) interface with the Photon transport. Assume that you have unpacked the HPX-5 source into the folder `/path/to/hpx` and want to install HPX-5 into `/path/to/install`. The following steps build and install HPX-5.

```
> cd /path/to/hpx
> % without network transport
```

```

> ./configure --prefix=/path/to/install
> % with ISIR interface
> % ./configure --prefix=/path/to/install --enable-mpi
> % with PWC interface
> % ./configure --prefix=/path/to/install --enable-mpi --enable-photon
> % ./configure --prefix=/path/to/install --enable-pmi --enable-photon
> make
> make install

```

The `--enable-mpi` or `--enable-pmi` option for the PWC network is used to build support for `mpirun` or `aprun` bootstrapping because HPX-5 does not provide its own distributed job launcher. Please see the official documentation for more detailed installation instructions for certain Cray machines. To finish the setup for HPX-5, one modifies the following environment variables.

```

export PATH=/path/to/install/bin:$PATH
export LD_LIBRARY_PATH=/path/to/install/lib:$LD_LIBRARY_PATH
export PKG_CONFIG_PATH=/path/to/install/lib/pkgconfig:$PKG_CONFIG_PATH

```

Once HPX-5 is installed, assume the RPYFMM is unpacked in directory `/path/to/rpyfmm`, will be built in directory `/path/to/rpyfmm/build`, and be installed into directory `/path/to/rpyfmm/install`, the library can be built in the following steps using CMake, version 3.4 or higher.

```

> cd /path/to/rpyfmm/build
> cmake ../ -DCMAKE_INSTALL_PREFIX=/path/to/rpyfmm/install
> make
> make install

```

4.2. Example

Included with RPYFMM is a test code that demonstrates a simple use of the library. This code is given in `/path/to/rpyfmm/demo/`. The demonstration code is not built/installed by `make install`. To build it, run `make demo` in `/path/to/rpyfmm/build/demo/`. User can request a summary of the options to the demo code by running the code with `--help` as a command line argument. In the following, we provide some walk through of the demo code.

The basic usage of RPYFMM is through an `Evaluator` object of the DASHMM library. The `Evaluator` object is a template over four types: the source type, the target type, the expansion type and the method type. For example,

```
dashmm::Evaluator<Source, Target,
    dashmm::RPY, dashmm::FMM97> rpy_eval{}
```

declares an `Evaluator` for the RPY kernel using the advanced FMM method for two user-defined types `Source` and `Target`. The minimum requirements for the `Source` and `Target` types are

```
struct Source {
    dashmm::Point position;
    double q[3];          // "charges" of the source
    //...
};
struct Target{
    dashmm::Point position;
    double value[3];      // store the results
    //...
};
```

Users can declare the object `Evaluator` object with a single type if that type satisfies both the minimum requirements for `Source` and `Target` (see type `Bead` of the demo code).

There are four parameters associated with the RPY kernel, specifying the radius of the bead, the Boltzmann constant, the absolute temperature, and the solvent viscosity. This information needs to be passed to the library by declaring.

```
std::vector<double> kparams(1, radius);
```

which uses the default values for the rest, or

```
std::vector<double> kparams(4);
kparams[0] = radius;
kparams[1] = boltzmann_constant;
kparams[2] = temperature;
kparams[3] = viscosity
```

Finally, an evaluation of the RPY kernel on a set of source and target points can be completed by instantiating the `FMM97` method and calling the `evaluate` member function of the `Evaluator` object.

```
dashmm::FMM97<Bead, Bead, dashmm::RPY> fmm97{ };
```

```
err = rpy_eval.evaluate(bead_handle, bead_handle, threshold,
                        fmm97, accuracy, &kparams);
```

where `threshold` specifies the refinement limit (max number of beads allowed in a childless leaf node).

4.3. Numerical Results

We demonstrate the performance of the RPYFMM library using the demo code, particularly focusing on the resulting scalability. The tests were performed on a Cray XC30 cluster at Indiana University, running Linux kernel 3.0.101-0.47-102. Each compute node has two Intel Xeon E5-2650 v3 processors at 2.3 GHz clock rate and 64 GB of DDR3 RAM. All compute nodes are connected through the Cray Aries interconnect. The RPYFMM library and the demo code were compiled using GNU compiler 6.2.0 with `-O3` optimization flags. The configuration of the tests can be summarized as follows:

- Two data distributions are tested: (a) Uniform distribution inside a cube and (b) Uniform distribution on a spherical surface.
- The problem size is 15 million for a cube distribution and 8 million for a sphere distribution.
- Three, six, and nine digits of accuracy are tested. The refinement limit for three, six, and nine digits are 80, 100, and 120, respectively.
- Tests requiring three digits of accuracy start from one compute node. Tests with six and nine digits of accuracy start from two compute nodes. Tests use up to 512 compute nodes for strong scaling evaluation.
- All tests are repeated five times and the average is reported here.

The accuracy results of the tests are given in Table 1. They were computed according to formula [28, Eq.(57)] at 400 randomly selected points. The scaling results of the tests are summarized in Figure 2. At 12, 288 cores using 512 compute nodes, RPYFMM is able to compute both problem within one second at three-digit accuracy requirement.

5. Conclusion

In this paper, we present RPYFMM, a parallel adaptive fast multipole method (FMM) software package on shared and distributed memory computers for the Rotne-Prager-Yamakawa tensor in biomolecular simulations. RPYFMM decomposes the RPY tensor into a linear combination of Laplace potentials, which are evaluated using the advanced FMM [1]. RPYFMM is an essential building block to enable full cell dynamics simulation, which is within reach in the near term as demonstrated by our numerical results.

The performance of RPYFMM can be further improved depending on the application context and underlying architecture. First, future version of DASHMM library will support heterogeneous memory architecture. As a result, the near field and possibly other part of the computation can be offloaded to the accelerators. Second, the FMM for the Laplace potential in RPYFMM is based on spherical harmonics and exponential expansions.

Spherical harmonic expansions are the orthogonal basis functions for the Laplace operator for a unit sphere and the number of expansion terms is optimal for spherical geometry. However, compared with polynomial expansion based approach [29, 30], the benefits might not be obvious, especially at lower accuracy requirement, because polynomials can be evaluated much more efficiently than spherical harmonics and exponential functions. Future version of RPYFMM will either have a more tuned evaluations of harmonic and exponential expansions or will internally switch to a polynomial based approach when the accuracy requirement is low.

From a broader perspective, authors of this paper are also working on several closely related research projects in solving (1). Some examples include introducing more realistic HI models by adding the rotational motions of the beads; accelerating the simulations when rigid body assumptions are valid (the relative locations of the atoms in portions of the molecules are fixed in time); developing efficient mesh generation tools for dynamic simulations; developing parallel and more efficient time integration schemes; developing better preconditioners to further reduce the number of iterations in the iterative schemes when solving the dynamic equations. Research results along these directions will be discussed in the future.

Acknowledgments

The authors gratefully acknowledge the inspiring discussions with Profs. David Keyes and Rio Yokota on different parallelization strategies for our solver. BZ was supported in part by National Science Foundation grant number ACI-1440396. GH was supported in part by National Institute of Health grant number GM 31749. This research was supported in part by Lilly Endowment, Inc., through its support for the Indiana University Pervasive Technology Institute. Part of the work was finished when JH was a visiting professor at the King Abdullah University of Science and Technology.

References

1. Greengard L, Rokhlin V. A new version of the fast multipole method for the Laplace equation in three dimensions. *Acta Num.* 1997; 6:229–269.
2. DeBuhr J, Zhang B, Tsueda A, Tilstra-Smith V, Sterling T. DASHMM: Dynamic Adaptive System for Hierarchical Multipole Methods. *Comm Comput Phys.* 2016; 20:1106–1126.
3. DeBuhr J, Zhang B, Sterling T. Revision of DASHMM: Dynamic Adaptive System for Hierarchical Multipole Methods. *Comm Comput Phys.* 2018; 23:296–314.
4. Tanaka H, Araki T. Simulation method of colloidal suspensions with hydrodynamic interactions: Fluid particle dynamics. *Phys Rev Lett.* 2000; 85:1338. [PubMed: 10991546]
5. Ando T, Skolnick J. On the importance of hydrodynamic interactions in lipid membrane formation. *Biophys J.* 2013; 104:96–105. [PubMed: 23332062]
6. Wang N, Huber GA, McCammon JA. Assessing the two-body diffusion tensor calculated by the bead models. *J Chem Phys.* 2013; 138(20):204117. [PubMed: 23742464]
7. Ermak DL, McCammon JA. Brownian dynamics with hydrodynamic interactions. *J Chem Phys.* 1978; 69:1352.
8. Batchelor G. Brownian diffusion of particles with hydrodynamic interaction. *J Fluid Mech.* 1976; 74(01):1–29.
9. Wajnryb E, Mizerski KA, Zuk PJ, Szymczak P. Generalization of the Rotne-Prager-Yamakawa mobility and shear disturbance tensors. *Journal of Fluid Mechanics.* 2013; 731:R3.
10. Zhang B, DeBuhr J, Niedzielski D, Mayolo S, Lu B, Sterling T. DASHMM accelerated adaptive Fast Multipole Poisson-Boltzmann solver on distributed memory architecture. *arXiv preprint arXiv:1710.06316.*

11. Liu X, Chow E. Large-scale hydrodynamic Brownian simulations on multicore and manycore architectures. IPDPS. 2014
12. Fixman M. Construction of Langevin forces in the simulation of hydrodynamic interaction. *Macromolecules*. 1986; 19:1204–1207.
13. Banchio AJ, Brady JF. Accelerated Stokesian dynamics: Brownian motion. *J Chem Phys*. 2003; 118:10323.
14. Geyer T, Winter U. An $O(N^2)$ approximation for hydrodynamic interactions in Brownian dynamics simulations. *J Chem Phys*. 2009; 130:114905. [PubMed: 19317564]
15. Ando T, Chow E, Saad Y, Skolnick J. Krylov subspace methods for computing hydrodynamic interactions in Brownian dynamics simulations. *J Chem Phys*. 2012; 137:064106. [PubMed: 22897254]
16. Liang Z, Gimbutas Z, Greengard L, Huang J, Jiang S. A fast mul-tipole method for the Rotne-Prager-Yamakawa tensor and its applications. *Journal of Computational Physics*. 2013; 234:133–139.
17. Tornberg AK, Greengard L. A fast multipole method for the three-dimensional Stokes equations. *Journal of Computational Physics*. 2008; 227(3):1613–1619.
18. Jiang S, Liang Z, Huang J. A fast algorithm for Brownian dynamics simulation with hydrodynamic interactions. *Mathematics of Computation*. 2013; 82(283):1631–1645.
19. Kulkarni A, Dalessandro L, Kissel E, Lumsdaine A, Sterling T, Swamy M. Network-managed virtual global address space for message-driven runtimes. *Proceedings of the 25th International Symposium on High Performance Parallel and Distributed Computing (HPDC 2016)*. 2016
20. Kissel E, Swamy M. Photon: Remote memory access middleware for high-performance runtime systems. *Proceedings of the 1st Emerging Parallel and Distributed Runtime Systems and Middleware (IP-DRM) Workshop*. 2016
21. Cimini M, Siek JG, Sterling T. The Semantics of ParalleX, v1.0. School of Informatics and Computing, Indiana University Bloomington; May, 2016 Tech. Rep. TR726
22. Warren M, Salmon J. Astrophysical N-body simulation using hierarchical tree data structures. SC 92': *Proceedings of the 1992 ACM/IEEE Conference on Supercomputing*. 1992
23. Ying L, Biros G, Zorin D, Langston H. A new parallel kernel-independent fast multipole method. SC '03: *Proceedings of the 2003 ACM/IEEE Conference on Supercomputing*. 2003
24. Kurzak J, Pettitt BM. Communications overlapping in fast multipole particle dynamics methods. *J Comput Phys*. 2005; 203:731–743.
25. Barnes J, Hut P. A hierarchical $O(N \log N)$ force-calculation algorithm. *Nature*. 1986; 324:446–449. DOI: 10.1038/324446a0
26. Greengard L, Rokhlin V. A fast algorithm for particle simulations. *J Comput Phys*. 1987; 73(2): 325–348.
27. Carrier J, Greengard L, Rokhlin V. A fast adaptive multipole algorithm for particle simulations. *SIAM J Sci Stat Comp*. 1988; 9:669–686.
28. Cheng H, Greengard L, Rokhlin V. A fast adaptive multipole algorithm in three dimensions. *J Comput Phys*. 1999; 155:468–498.
29. Yokota R, Narumi T, Sakamaki R, Kameoka S, Obi S, Yasuoka K. Fast multipole methods on a cluster of GPUs for the meshless simulation of turbulence. *Computer Physics Communications*. 2009; 180(11):2066–2078.
30. Yokota R, Barba LA, Narumi T, Yasuoka K. Petascale turbulence simulation using a highly parallel fast multipole method on GPUs. *Computer Physics Communications*. 2013; 184(3):445–455.

PROGRAM SUMMARY

Program Title: RPYFMM: Parallel Adaptive FMM for RPY Tensor

Program Files doi: <http://dx.doi.org/10.17632/zpbjvy8whp.1>

Licensing provisions: GNU General Public License, version 3

Programming language: C++

Nature of problem: Evaluate the Rotne-Prager-Yamakawa tensor matrix-vector multiplications describing the hydrodynamics interaction in biomolecular systems.

Solution method: The Rotne-Prager-Yamakawa tensor is decomposed as a linear combination of four Laplace interactions, each of which is evaluated using the new version of adaptive fast multipole method [1].

Additional Comments: RPYFMM is built on top of the DASHMM library and the Asynchronous Multi-Tasking HPX-5 runtime system. DASHMM is automatically downloaded during installation and HPX-5 is available at <http://hpx.crest.iu.edu/>.

Figure 1.

Operator diagrams in DASHMM. Basic multipole methods use multipole (M) and local (L) expansions, and eight operators (shown in solid lines) that connect them to the sources (S) and targets (T). Advanced multipole methods use intermediate expansions (I) and three additional operators (shown in dashed lines). The M-to-L operation is decomposed into a chain of M-to-I, I-to-I, and I-to-L operations in advanced multipole methods.

Figure 2.

The time to completion t_n (top) and speed relative t_*/t_n (bottom) as a function of locality cores n . Each locality has 24 cores and 64 GB of RAM. For three digits of accuracy, $t_* = t_1$. For six and nine digits of accuracy, $t_* = t_2$. The problem size for cube distribution is 15 million and the problem size for sphere distribution is 8 million.

Table 1

Accuracy results of the RPYFMM.

	Cube	Sphere
3-digit	2.1410e-3	2.1420e-3
6-digit	1.4115e-7	1.3949e-7
9-digit	2.6447e-9	2.6347e-9

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript