Published in final edited form as:

SIAM J Sci Comput. 2018 ; 40(3): A1345–A1361. doi:10.1137/17M1117744.

HIERARCHICAL ORTHOGONAL MATRIX GENERATION AND MATRIX-VECTOR MULTIPLICATIONS IN RIGID BODY SIMULATIONS*

FUHUI FANG[†], JINGFANG HUANG[‡], GARY HUBER[§], J. ANDREW MCCAMMON[¶], BO ZHANG^{||}

[‡]Department of Mathematics, University of North Carolina at Chapel Hill, Chapel Hill, NC 27599-3250.

[§]Howard Hughes Medical Institute, University of California at San Diego, La Jolla, CA 92093-0365.

[¶]Department of Chemistry and Biochemistry and Department of Pharmacology, University of California at San Diego, La Jolla, CA 92093-0365.

^{||}Center for Research in Extreme Scale Technologies, Indiana University, Bloomington, IN 47404.

Abstract

In this paper, we apply the hierarchical modeling technique and study some numerical linear algebra problems arising from the Brownian dynamics simulations of biomolecular systems where molecules are modeled as ensembles of rigid bodies. Given a rigid body p consisting of n beads, the $6 \times 3n$ transformation matrix Z that maps the force on each bead to p 's translational and rotational forces (a 6×1 vector), and V the row space of Z , we show how to explicitly construct the $(3n - 6) \times 3n$ matrix $\tilde{Q}$ consisting of $(3n - 6)$ orthonormal basis vectors of $V^\perp$ (orthogonal complement of V) using only $\mathcal{O}(n \log n)$ operations and storage. For applications where only the matrix-vector multiplications $\tilde{Q}V$ and $\tilde{Q}^T V$ are needed, we introduce asymptotically optimal $\mathcal{O}(n)$ hierarchical algorithms without explicitly forming $\tilde{Q}$. Preliminary numerical results are presented to demonstrate the performance and accuracy of the numerical algorithms.

Keywords

Brownian dynamics; hierarchical modeling; orthogonal linear algebra; fast algorithms; 15B10; 65F25; 65F50; 65Y20; 70E55

1. Background and problem statement.

In the Brownian dynamics simulations of biomolecules with hydrodynamic interactions, the complex molecular system is modeled as multiple (hundreds or thousands) rigid bodies to

*Submitted to the journal's Methods and Algorithms for Scientific Computing section February 21, 2017; accepted for publication (in revised form) January 16, 2018; published electronically May 10, 2018.

[†]Corresponding author. Department of Mathematics, University of North Carolina at Chapel Hill, Chapel Hill, NC 27599-3250 (fangf@live.unc.edu).

reduce the numerical “stiffness” due to the local chemical bond type interactions between atoms that cause very high frequency oscillations and subsequently require extremely small step size when marching in time. Instead of a thorough listing of existing literature on the Brownian dynamics models and hydrodynamic interactions, we focus on the “shell-bead” model (see, e.g., [5, 8, 21]) that describes the hydrodynamic forces exerted on a protein.

In the shell-bead model, the molecular system is represented by m rigid bodies, where rigid body j is modeled by n_j spherical beads (often of the same radius) placed on the molecular surface with complex geometry. The total number of beads is $n = \sum_{j=1}^m n_j$. Let $\mathbf{f}_k^j$ be the external force applied on bead k (located at $\mathbf{r}_k^j$) of rigid body j , then rigid body j 's *resultant force* $\mathbf{F}_j$ and *torque* $\boldsymbol{\tau}_j$ are given by

$$\mathbf{F}_j = \sum_{k=1}^{n_j} \mathbf{f}_k^j, \quad \boldsymbol{\tau}_j = \sum_{k=1}^{n_j} \mathbf{r}_k^j \times \mathbf{f}_k^j, \quad (1)$$

or in matrix form

$$\begin{bmatrix} \mathbf{F}_j \\ \boldsymbol{\tau}_j \end{bmatrix}_{6 \times 1} = \begin{bmatrix} I & I & \dots & I \\ A_1^j & A_2^j & \dots & A_{n_j}^j \end{bmatrix}_{6 \times 3n_j} \begin{bmatrix} \mathbf{f}_1^j \\ \mathbf{f}_2^j \\ \dots \\ \mathbf{f}_{n_j}^j \end{bmatrix}_{3n_j \times 1} = \mathbf{Z}^j \mathbf{f}^j, \quad (2)$$

where we refer to the $6 \times 3n_j$ matrix as the *Z-matrix*. In the formula, I is the 3×3 identity matrix and A_k^j is the transformation matrix of $A_k^j \cdot \mathbf{f}_k^j \triangleq \mathbf{r}_k^j \times \mathbf{f}_k^j$. Particularly, if $\mathbf{r}_k^j = (x_k^j, y_k^j, z_k^j)$, then

$$A_k^j = \begin{bmatrix} 0 & -z_k^j & y_k^j \\ z_k^j & 0 & -x_k^j \\ -y_k^j & x_k^j & 0 \end{bmatrix}.$$

The shell-bead model assumes that the hydrodynamic effects are related to the deterministic forces through $D\mathbf{f} = \mathbf{v}$, where the $3n \times 3n$ matrix D is the symmetric Rotne–Prager–Yamakawa (RPY) tensor whose entries are determined by the bead locations (see, e.g., [3, 9] for the explicit form of the RPY tensor), and $\mathbf{f}$ and $\mathbf{v}$ in MATLAB notation are given by

$$\mathbf{f} = [\mathbf{f}_1^1; \dots; \mathbf{f}_{n_1}^1; \dots; \mathbf{f}_1^m; \dots; \mathbf{f}_{n_m}^m], \quad \mathbf{v} = [\mathbf{v}_1^1; \dots; \mathbf{v}_{n_1}^1; \dots; \mathbf{v}_1^m; \dots; \mathbf{v}_{n_m}^m].$$

Let $[[F_1; \tau_1]; \dots; [F_m; \tau_m]]$ be the external deterministic force and torque vector acting on the m rigid bodies and $[[V_1; \omega_1]; \dots; [V_m; \omega_m]]$ be the translational and angular velocity vector of the m rigid bodies. Under the rigid body constraint, the corresponding deterministic velocity field of the rigid bodies can be obtained by solving

$$\underbrace{Z_{6m \times 3n} (D^{-1})_{3n \times 3n} (Z^T)_{3n \times 6m}}_{\text{forces on rigid bodies}} \underbrace{\begin{bmatrix} V_1 \\ \omega_1 \\ \vdots \\ V_m \\ \omega_m \end{bmatrix}}_{\text{velocities on beads}} = \underbrace{\begin{bmatrix} F_1 \\ \tau_1 \\ \vdots \\ F_m \\ \tau_m \end{bmatrix}}_{\text{forces on beads}}, \quad (3)$$

where $Z = \text{diag}(Z^1, \dots, Z^m)$ is block diagonal. Here, the force and translational velocity field are denoted by uppercase letters for the rigid bodies and lowercase letters for the beads. Equation (3) simply means that the unknown rigid body velocity field, when mapped onto individual beads, should yield the force field on the beads via $\mathbf{f} = D^{-1}\mathbf{v}$. Then, the force acting on the rigid body can be obtained by integrating all the bead forces using $\mathcal{A}$ and the result should match the given external forces on the right-hand side.

One major numerical difficulty in solving (3) accurately and efficiently is the calculation of D^{-1} as D is dense. For large n , even with the acceleration of the fast direct solvers [11, 17] or $\mathcal{H}$ -matrix techniques [14, 16] that utilize the hierarchically “low-rank” structures in the matrix D , computing D^{-1} at each time marching step is simply too expensive for dynamic simulations. It is possible to avoid the explicit computation of D^{-1} by a reformulation of (3). To simplify the discussion, we assume that with an easy QR procedure on both sides of (3) for each rigid body j , the six vectors in the diagonal block Z^j become orthogonal. Denote the orthogonal version of the matrix Z^j by Q_Z^j . Define $Q_Z = \text{diag}(Q_Z^1, \dots, Q_Z^m)$ and $\tilde{Q}$ the orthogonal vectors such that $Q = [Q_Z; \tilde{Q}]$ is a $3n \times 3n$ orthogonal matrix. Here, we mildly abuse the notation by using $\mathbf{V} = [V_1; \omega_1; \dots; V_m; \omega_m]$ and introducing $\mathbf{F}$ to represent the new right-hand side after the QR process. Padding $(3n - 6m)$ zeros to $\mathbf{V}$ and introducing $(3n - 6m)$ unknowns $\mathbf{g}$ to $\mathbf{F}$, we form

$$\begin{bmatrix} Q_Z \\ \tilde{Q} \end{bmatrix} D^{-1} \begin{bmatrix} Q_Z^T & \tilde{Q}^T \end{bmatrix} \begin{bmatrix} \mathbf{V} \\ \mathbf{0} \end{bmatrix} = \begin{bmatrix} \mathbf{F} \\ \mathbf{g} \end{bmatrix}. \quad (4)$$

As the inverse of the orthogonal matrix Q is simply its transpose, some algebraic manipulations show that for any given vector $\mathbf{F}$, one can first find the unknown $\mathbf{g}$ by solving

$$\mathbf{0} = \tilde{Q}DQ_Z^T\mathbf{F} + \tilde{Q}D\tilde{Q}^T\mathbf{g} \quad (5)$$

via a preconditioned Krylov subspace iterative method. Physically, the vector $\mathbf{g}$ can be envisioned as the constraint forces needed to keep the beads in each rigid body together during the imposed external forces. Then, the velocity vector $\mathbf{V}$ can be computed by

$$\mathbf{V} = Q_Z D(Q_Z^T\mathbf{F} + \tilde{Q}^T\mathbf{g}). \quad (6)$$

The fundamental building blocks required by this new formulation are the fast matrix-vector multiplications of $D\mathbf{v}$, $\tilde{Q}^T\mathbf{v}$, and $\tilde{Q}\mathbf{v}$ with any given vector $\mathbf{v}$. The $D\mathbf{v}$ operation can be carried out efficiently using the fast multipole methods (FMMs) as discussed in [19, 22]. This paper considers the $\tilde{Q}\mathbf{v}$ and $\tilde{Q}^T\mathbf{v}$ operations. Notice that the matrix Z is block diagonal and each block corresponds to a rigid body; this structure allows us to construct $\tilde{Q}^j$ for each Q_Z^j separately and then form the $\tilde{Q}$ matrix using $\tilde{Q} = \text{diag}(\tilde{Q}^1, \dots, \tilde{Q}^m)$. We therefore focus on one rigid body and drop the index j in the following problem statement.

Problem statement.

Given the Z -matrix of a rigid body p with n beads, how do we efficiently construct and store $\tilde{Q}$ explicitly (if needed)? And for any given vector $\mathbf{v}$ of proper size, how do we efficiently compute $\tilde{Q}\mathbf{v}$ and $\tilde{Q}^T\mathbf{v}$?

We want to mention that other reformulations of (3) are also possible. For instance, the Schur complement was used for a similar problem with rigid particle constraints [1, 2, 20]. These reformulations are equivalent mathematically; however, they have different stability and efficiency properties. As far as we know, our orthogonal matrix based reformulation in (5)–(6) has not been applied previously; one reason is the lack of stable and efficient numerical tools for computing $\tilde{Q}$, $\tilde{Q}\mathbf{v}$, and $\tilde{Q}^T\mathbf{v}$.

The main contributions of this paper are three novel algorithms for computing $\tilde{Q}$, $\tilde{Q}\mathbf{v}$, and $\tilde{Q}^T\mathbf{v}$ that are optimal in complexity, stability, and storage requirements. The discussions of these algorithms are organized as follows. In section 2, we introduce the hierarchical tree structure and present a hierarchical model for constructing $\tilde{Q}$. To better preserve orthogonality, in section 3, we apply tools from the orthogonal linear algebra and present the first algorithm to explicitly construct $\tilde{Q}$ using $\mathcal{O}(n \log n)$ operations and storage. In section 4, we present two asymptotically optimal algorithms to compute the matrix-vector multiplications $\tilde{Q}\mathbf{v}$ and $\tilde{Q}^T\mathbf{v}$ using only $\mathcal{O}(n)$ operations and storage, without explicitly forming $\tilde{Q}$. In section 5, numerical results are presented to demonstrate the algorithms' performance and orthogonality properties. Finally in section 6, we summarize our results and discuss several related research topics.

2. Hierarchical tree and hierarchical model.

We apply the hierarchical modeling technique to study the orthonormal basis vectors in $\tilde{Q}$. The hierarchical modeling technique identifies any low-rank, or low-dimensional, or other compact features in the system, and the compressed representations are then recursively collected from children to parents and transmitted between different nodes on a hierarchical tree structure using properly compressed translation operators. In its numerical implementation, the hierarchical models are often reexpressed as recursive algorithms, which can be easily interfaced with existing dynamical schedulers from the high-performance computing community for optimal parallel efficiency.

Different aspects of the hierarchical modeling technique have been known and addressed by different research communities previously. Examples include the classical fast Fourier transform (FFT) [6], where the halving lemma shows how data can be compressed and the odd-even term splitting of the polynomials creates a hierarchical tree to allow recursively processing the compressed information efficiently; the multigrid method (MG) [4, 15], where the hierarchical tree structure is formed via adaptively refining the computational domain, and data compression and transmission are performed using the relaxation (smoother) and projection (restriction) operators by analyzing the frequency domain behaviors of the error functions between different levels of the (adaptive) tree to effectively reduce the high frequency errors; and the FMM [12, 13], where the particle information inside a box is first compressed to the multipole expansion, and the compressed information is transmitted recursively to parent levels using the multipole-to-multipole translation in the upward pass on the hierarchical tree structures. The collected and compressed information is later transmitted to target boxes using the multipole-to-local translations and propagated to child levels using the local-to-local translations in the downward pass. When there are n terms (FFT) in the polynomial or n approximately uniformly distributed particles (MG or FMM), the depth of the hierarchical tree is normally $\mathcal{O}(\log n)$ and the number of tree nodes is approximately $\mathcal{O}(n)$. Therefore, if each level only requires $\mathcal{O}(n)$ operations (e.g., FFT), the algorithm complexity will be $\mathcal{O}(n \log n)$. If each tree node only requires a constant amount of operations (e.g., MG or FMM), the algorithm complexity will be asymptotically optimal $\mathcal{O}(n)$. In this section, we discuss how to use the hierarchical modeling technique to answer the questions in the problem statement in section 1.

2.1. Adaptive hierarchical tree structure.

We first consider generating a spatial adaptive hierarchical tree when simulating a molecular system modeled by multiple rigid bodies in the shell-bead model. We assume each rigid body is “discretized” into a number of beads to capture the hydrodynamic interactions between rigid bodies. A hierarchical partition is then performed to divide the beads domain into nested cubical boxes, where the root box is the smallest bounding box that contains all the beads. Without loss of generality, the root box is normalized to size 1 in each side. The root box is partitioned equally along each dimension. The partition continues recursively on the resulting box until the box contains no more than s beads, at which point it becomes a leaf node. Empty boxes encountered during partition are pruned off. Here, we set $s = 1$ to

simplify the discussions in the following sections. In our implementation, other values of s are allowed after modifying how the leaf nodes are processed.

Comment.

The octree can be modified to form a binary tree: consider a parent node p with eight child nodes. One can first add a “ghost” level by creating two “ghost” nodes of p each containing four of p ’s children separated along the z -direction. Then, the four children in each ghost node can be separated by the y -direction, creating four more ghost nodes at a finer ghost level. Finally, these four ghost nodes are partitioned along the x -direction, creating the actual eight child nodes of p in the original octree. With the two additional ghost levels, the original two-level parent-child structure in an *octree* is modified to a four-level structure in a *binary* tree. For a uniform octree with $L + 1$ levels, the total number of levels of the modified binary tree is $3L + 1$, where $L = O(\log n)$. Such modification is not fundamental but could significantly simplify both the notation and descriptions of the algorithms in the rest of the paper. For this reason, we focus on the binary tree and set $s = 1$.

2.2. Divide-and-conquer strategy and hierarchical model.

We consider a particular choice of the orthonormal vectors in $\tilde{Q}$ using the divide-and-conquer strategy on the *hierarchical* tree structure. We start from a two-level setting where the parent rigid body p consisting of n beads is partitioned into two child nodes, child x with n_x beads and child y with n_y beads, where $n_x + n_y = n$. Let

$$Z_x = \begin{bmatrix} I & I & \cdots & I \\ A_1 & A_2 & \cdots & A_{n_x} \end{bmatrix}, \quad Z_y = \begin{bmatrix} I & I & \cdots & I \\ B_1 & B_2 & \cdots & B_{n_y} \end{bmatrix}$$

be the Z -matrices of x and y , respectively. Assume both Z_x and Z_y are full rank, and the orthogonal matrices $\tilde{Q}_x$ of size $(3n_x - 6) \times 3n_x$ and $\tilde{Q}_y$ of size $(3n_y - 6) \times 3n_y$ that satisfy $\tilde{Q}_x \perp Z_x$ and $\tilde{Q}_y \perp Z_y$ are available in compact form. The key observation comes from the study of the matrix

$$H = \begin{bmatrix} Z_x & Z_y \\ 0 & Z_y \\ \tilde{Q}_x & 0 \\ 0 & \tilde{Q}_y \end{bmatrix}_{3n \times 3n}$$

and the fact that $Z_p = [Z_x, Z_y]$. It is straightforward to verify that the vectors in the lower $(3n - 12)$ rows of H are normalized, orthogonal to each other and to the first 12 rows of H . This means that $\tilde{Q}_p$ of parent p can readily “receive” the lower $(3n - 12)$ rows of vectors from its two children. For the remaining 6 row vectors in $\tilde{Q}_p$, a Gram–Schmidt procedure on the first 12 row vectors can be performed and the last six orthonormal vectors will be orthogonal to the vectors in Z_p and the $(3n - 12)$ vectors from the children. In a multilevel setting for a rigid body with n beads, there will be approximately $O(\log n)$ levels and for each level, the

Gram–Schmidt procedure requires approximately $\mathcal{O}(n)$ operations and storage to explicitly generate all the orthogonal vectors for that level. The total storage and operations required are therefore both $\mathcal{O}(n \log n)$.

Unfortunately, straightforward implementation of this divide-and-conquer idea will result in an algorithm with stability issues. In particular, the matrix $Q = [Q_Z; \tilde{Q}]$ will lose orthogonality, and the results computed via the new formulation in (5)–(6) will contain huge errors. To see this, we generate n beads that are randomly distributed inside a unit box and compute the infinity norm of $QQ^T - I$ and $Q^T Q - I$. The results are shown in Table 1. We see that the algorithm condition number increases when more beads are added to the box. One source for this instability is the ill-conditioning of the Z -matrices. For instance, when all the beads are located exactly on a straight line, the rank of the Z -matrix is only 5 instead of 6, the matrix H becomes singular, and the algorithm fails. In the following sections, we show how to resolve the instability issues using the orthogonal linear algebra techniques.

3. Stable $\mathcal{O}(n \log n)$ orthogonal matrix generation algorithm.

We start from Theorem 3.1.

THEOREM 3.1. *Assume the centroid of a rigid body p is located at the origin $\mathbf{0}$; then the first three rows of the Z -matrix are orthogonal to the last three rows.*

The proof follows from the identities

$$\sum x_i = \sum y_i = \sum z_i = 0$$

if the centroid of the rigid body is chosen as the origin. It suggests that in order to get Q_Z , the orthogonal version of the Z -matrix of p , the Gram–Schmidt procedure only needs to be applied to the last three rows of Z_p . We adopt this assumption in the following discussions and formulas and revisit the two-level setting in section 2.2. To preserve orthogonality, we assume Z_x and Z_y are already decomposed with respect to their respective centers as

$$Z_x = C_x \cdot Q_{Z_x} = \begin{bmatrix} \sqrt{n_x} I_{3 \times 3} & 0 \\ 0 & R_{22} \end{bmatrix} \cdot \begin{bmatrix} I_{n_x} \\ V_{3 \times 3n_x} \end{bmatrix}, \quad (7)$$

$$Z_y = C_y \cdot Q_{Z_y} = \begin{bmatrix} \sqrt{n_y} I_{3 \times 3} & 0 \\ 0 & S_{22} \end{bmatrix} \cdot \begin{bmatrix} I_{n_y} \\ W_{3 \times 3n_y} \end{bmatrix}, \quad (8)$$

where

$$I_{n_x} = \frac{1}{\sqrt{n_x}} \begin{bmatrix} I_{3 \times 3} & \cdots & I_{3 \times 3} \end{bmatrix}_{3 \times 3n_x} \text{ and } I_{n_y} = \frac{1}{\sqrt{n_y}} \begin{bmatrix} I_{3 \times 3} & \cdots & I_{3 \times 3} \end{bmatrix}_{3 \times 3n_y}.$$

In the formula, as each child uses its own centroid, the first three rows of the Z -matrix are orthogonal to the last three rows by Theorem 3.1 and matrix C is block diagonal.

Assuming orthogonality preserving results are available for x and y , i.e., the orthogonal submatrices

$$Q_{Z_x} = \begin{bmatrix} I_{n_x} \\ V \end{bmatrix}_{6 \times 3n_x}, \quad \tilde{Q}_x \perp Q_{Z_x}, \quad Q_{Z_y} = \begin{bmatrix} I_{n_y} \\ W \end{bmatrix}_{6 \times 3n_y}, \quad \tilde{Q}_y \perp Q_{Z_y}$$

are already constructed with excellent orthogonality properties, we study in the following how to stably find $\tilde{Q}_p$ as well as T_{22} and Q_{Z_p} in the decomposition of p 's Z -matrix

$$Z_p = C_p \cdot Q_{Z_p} = \begin{bmatrix} \sqrt{n} I_{3 \times 3} & 0 \\ 0 & T_{22} \end{bmatrix} \cdot \begin{bmatrix} I_{n_p} \\ U_{3 \times 3n_x} \end{bmatrix}, \quad (9)$$

$$I_{n_p} = \frac{1}{\sqrt{n}} \begin{bmatrix} I_{3 \times 3} & \cdots & I_{3 \times 3} \end{bmatrix}_{3 \times 3n},$$

where the new origin is located at the centroid of p that can be easily computed from the centroids of x and y . Following the ideas in section 2.2, we consider a particular orthogonal matrix $Q_p = \begin{bmatrix} Q_{Z_p} \\ \tilde{Q}_p \end{bmatrix}$, where

$$\tilde{Q}_p = \begin{bmatrix} \text{Residue Vectors} \\ \tilde{Q}_x & 0 \\ 0 & \tilde{Q}_y \end{bmatrix}$$

and the row vectors in $[Q_{Z_p}; \text{Residue Vectors}]$ form the same subspace as that spanned by the

12 row vectors in $\begin{bmatrix} Q_{Z_x} & 0 \\ 0 & Q_{Z_y} \end{bmatrix}$. Clearly, $\tilde{Q}_p$ contains two parts, the lower $3n - 12$ rows are

processed at child levels and require no additional operations or storage (only nonzero values are stored), and the first 6 row vectors are the “left-over” vectors (referred to as the *Residue Vectors*) after identifying and removing Q_{Z_p} components from the subspace of dimension 12. Therefore, in the hierarchical modeling technique, given compressed

information Q_{Z_x} , Q_{Z_y} , R_{22} , and S_{22} at the child level, we study how to compute parent's compressed information Q_{Z_p} , T_{22} as well as the one time output *Residue Vectors*.

We prefer the well-conditioned Q_{Z_x} and Q_{Z_y} to preserve orthogonality properties in the hierarchical model to the original Z_x and Z_y discussed in section 2.2 because the latter may become rank deficient if all the beads are located on the same line. Note that the subspace spanned by the Z -matrix is always a subset of the subspace spanned by the corresponding Q_Z . Also, to take advantage of Theorem 3.1, instead of considering the row vectors in $\begin{bmatrix} Q_{Z_x} & 0 \\ 0 & Q_{Z_y} \end{bmatrix}$, we consider the orthogonal vectors in the matrix C presented in Theorem 3.2,

where the first three row vectors are simply I_{n_p} . This strategy saves operations when performing the QR decomposition.

THEOREM 3.2. *For a parent node p with child nodes x and y , the 12 row vectors of the matrix*

$$C = \begin{bmatrix} \sqrt{n_x/n} I_{n_x} & \sqrt{n_y/n} I_{n_y} \\ V & 0 \\ 0 & W \\ \sqrt{n_y/n} I_{n_x} & -\sqrt{n_x/n} I_{n_y} \end{bmatrix}$$

are orthonormal and span the same subspace as the row vectors in $\begin{bmatrix} Q_{Z_x} & 0 \\ 0 & Q_{Z_y} \end{bmatrix}$, and the first three rows satisfy $[\sqrt{n_x/n} I_{n_x}, \sqrt{n_y/n} I_{n_y}] = I_{n_p}$.

This theorem is simply the result of $C \cdot C^T = I_{12 \times 12}$. As the orthonormal basis vectors in C are constructed analytically, their orthogonality properties are well-preserved.

To compute Q_{Z_p} , T_{22} , and the *Residue Vectors*, we first represent Z_p using Z_x and Z_y by shifting the centroid using the 3×3 matrices R_{21} and S_{21} as in

$$Z_p = \begin{bmatrix} I & 0 \\ R_{21} & I \end{bmatrix} \cdot Z_x \cdot \begin{bmatrix} I & 0 \\ S_{21} & I \end{bmatrix} \cdot Z_y.$$

Substituting the orthogonal decompositions of Z_x and Z_y available in (7) and (8), we have

$$Z_p = \left[\begin{array}{c|c} \sqrt{n_x} I_{n_x} & \sqrt{n_y} I_{n_y} \\ \hline \sqrt{n_x} R_{21} I_{n_x} + R_{22} V & \sqrt{n_y} S_{21} I_{n_y} + S_{22} W \end{array} \right]. \quad (10)$$

Notice that $I_{n_x} \perp V, I_{n_y} \perp W$, and

$$\left[\sqrt{n_x} I_{n_x} \sqrt{n_y} I_{n_y} \right] \perp \left[\sqrt{n_x} R_{21} I_{n_x} + R_{22} V \sqrt{n_y} S_{21} I_{n_y} + S_{22} W \right].$$

One can easily derive

$$n_x R_{21} I_{n_x} + n_y S_{21} I_{n_y} = 0 \Rightarrow R_{21} = -\frac{n_y}{n_x} S_{21}.$$

As the first three rows of $Z_p(1:3, :)$ are simply $\sqrt{n} I_n$, by Theorem 3.1, we only need to consider the last three rows of Z_p in (10) reformulated as

$$Z_p(4:6, :) = \begin{bmatrix} R_{22} & S_{22} \sqrt{\frac{n_x n}{n_y}} R_{21} \end{bmatrix} \begin{bmatrix} V & 0 \\ 0 & W \\ \sqrt{\frac{n_y}{n}} I_{n_x} & -\sqrt{\frac{n_x}{n}} I_{n_y} \end{bmatrix}$$

using the orthonormal vectors in the C matrix in Theorem 3.2. Applying existing orthogonality preserving QR algorithms from the orthogonal linear algebra packages (e.g., MATLAB qr command), we can derive the QR decomposition of the 3×9 matrix

$$\begin{bmatrix} R_{22} & S_{22} \sqrt{\frac{n_x n}{n_y}} R_{21} \end{bmatrix} = \begin{bmatrix} r_{11} & r_{21} & r_{31} \\ r_{12} & r_{22} & r_{32} \\ r_{13} & r_{23} & r_{33} \end{bmatrix} \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}_{3 \times 9} \begin{bmatrix} Q_{11}^T & Q_{12}^T & Q_{13}^T \\ Q_{21}^T & Q_{22}^T & Q_{23}^T \\ Q_{31}^T & Q_{32}^T & Q_{33}^T \end{bmatrix}_{9 \times 9}, \quad (11)$$

where each Q_{ij} is a 3×3 matrix. Then, the last three rows of Z_p become

$$\begin{aligned} Z_p(4:6, :) &= \left[\sqrt{n_x} R_{21} I_{n_x} + R_{22} V \sqrt{n_y} S_{21} I_{n_y} + S_{22} W \right] \\ &= \begin{bmatrix} r_{11} & r_{21} & r_{31} \\ r_{12} & r_{22} & r_{32} \\ r_{13} & r_{23} & r_{33} \end{bmatrix} \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} \begin{bmatrix} Q_{11}^T V + \sqrt{\frac{n_y}{n}} Q_{13}^T I_{n_x} & Q_{12}^T W - \sqrt{\frac{n_x}{n}} Q_{13}^T I_{n_y} \\ Q_{21}^T V + \sqrt{\frac{n_y}{n}} Q_{23}^T I_{n_x} & Q_{22}^T W - \sqrt{\frac{n_x}{n}} Q_{23}^T I_{n_y} \\ Q_{31}^T V + \sqrt{\frac{n_y}{n}} Q_{33}^T I_{n_x} & Q_{32}^T W - \sqrt{\frac{n_x}{n}} Q_{33}^T I_{n_y} \end{bmatrix}. \end{aligned}$$

Notice that the row vectors in the second matrix are orthonormal and comparing with (9), we derive

$$T_{22} = \begin{bmatrix} r_{11} & r_{21} & r_{31} \\ r_{12} & r_{22} & r_{32} \\ r_{13} & r_{23} & r_{33} \end{bmatrix}, \quad (12)$$

$$U = \left[Q_{11}^T V + \sqrt{\frac{n_y}{n}} Q_{13}^T I_{n_x} \quad Q_{12}^T W - \sqrt{\frac{n_x}{n}} Q_{13}^T I_{n_y} \right], \quad (13)$$

$$\text{Residue Vectors} = \begin{bmatrix} Q_{21}^T V + \sqrt{\frac{n_y}{n}} Q_{23}^T I_{n_x} & Q_{22}^T W - \sqrt{\frac{n_x}{n}} Q_{23}^T I_{n_y} \\ Q_{31}^T V + \sqrt{\frac{n_y}{n}} Q_{33}^T I_{n_x} & Q_{32}^T W - \sqrt{\frac{n_x}{n}} Q_{33}^T I_{n_y} \end{bmatrix}. \quad (14)$$

There are two cases that require special treatments in the algorithm. The first case is when both x and y are leaf nodes, containing one bead each. The second case is when x is a leaf node with a single bead and y is a multibead rigid body, or vice versa. The special treatments are presented as follows.

Case I: Bead and bead.

Consider a parent node with two child leaf nodes x and y , containing a single bead each. The locations of the beads are given by (a, b, c) and $(-a, -b, -c)$ such that the parent's centroid is located at the origin. Assume further that $c \neq 0$; then the parent's U and T_{22} matrices are given by

$$T_{22} = \begin{bmatrix} \sqrt{2(b^2 + c^2)} & 0 & 0 \\ -\frac{\sqrt{2}ab}{\sqrt{b^2 + c^2}} & \sqrt{\frac{2c^2(a^2 + b^2 + c^2)}{b^2 + c^2}} & 0 \\ -\frac{\sqrt{2}ac}{\sqrt{b^2 + c^2}} & -\sqrt{\frac{2b^2(a^2 + b^2 + c^2)}{b^2 + c^2}} & 0 \end{bmatrix}, \quad (15)$$

$$U = \begin{bmatrix} \frac{1}{\sqrt{2(b^2+c^2)}}[0 \ -c \ b \ 0 \ c \ -b] \\ \frac{1}{\sqrt{2(a^2+b^2+c^2)(b^2+c^2)}}[b^2+c^2 \ -ab \ -ac \ -b^2-c^2 \ ab \ ac] \\ \frac{1}{\sqrt{2(a^2+b^2+c^2)}}[-a \ -b \ -c \ a \ b \ c] \end{bmatrix}, \quad (16)$$

and there are no *Residue Vectors* generated. Note that in this case the rank of T_{22} is only two, but U is orthogonal with full rank.

Case II: Rigid body and bead.

Assume child x has $n_x > 1$ beads and child y has a single bead; Z_p can be reformulated as

$$\left[\begin{array}{c|c} \sqrt{n_x} I_{n_x} & I \\ \hline \sqrt{n_x} R_{21} I_{n_x} + R_{22} V & S_{21} \end{array} \right].$$

Applying the QR algorithm from the orthogonal linear algebra package, the last three rows of Z_p could be reformulated as

$$\begin{aligned} Z_p(4:6, :) &= \left[\sqrt{n_x} R_{21} I_{n_x} + R_{22} V \middle| S_{21} \right] \\ &= \begin{bmatrix} R_{22} & \sqrt{n_x} R_{21} \end{bmatrix} \left[\begin{array}{c|c} V & 0 \\ \hline \sqrt{\frac{1}{n}} I_{n_x} & -\sqrt{\frac{n_x}{n}} I \end{array} \right] \\ &= \begin{bmatrix} r_{11} & r_{21} & r_{31} \\ r_{12} & r_{22} & r_{32} \\ r_{13} & r_{23} & r_{33} \end{bmatrix} \begin{bmatrix} Q_{11}^T & Q_{12}^T \\ Q_{21}^T & Q_{22}^T \end{bmatrix}_{6 \times 6} \left[\begin{array}{c|c} V & 0 \\ \hline \sqrt{\frac{1}{n}} I_{n_x} & -\sqrt{\frac{n_x}{n}} I \end{array} \right]_{6 \times 6n} \\ &= \begin{bmatrix} r_{11} & r_{21} & r_{31} \\ r_{12} & r_{22} & r_{32} \\ r_{13} & r_{23} & r_{33} \end{bmatrix} \begin{bmatrix} Q_{11}^T V + \sqrt{\frac{1}{n}} Q_{12}^T I_{n_x} & -\sqrt{\frac{n_x}{n}} Q_{12}^T \\ Q_{21}^T V + \sqrt{\frac{1}{n}} Q_{22}^T I_{n_x} & -\sqrt{\frac{n_x}{n}} Q_{22}^T \end{bmatrix}. \end{aligned}$$

The orthogonal decomposition and *Residue Vectors* of parent p are given by

$$T_{22} = \begin{bmatrix} r_{11} & r_{21} & r_{31} \\ r_{12} & r_{22} & r_{32} \\ r_{13} & r_{23} & r_{33} \end{bmatrix}, \quad (17)$$

$$U = \left[Q_{11}^T V + \sqrt{\frac{1}{n}} Q_{12}^T I_{n_x} - \sqrt{\frac{n_x}{n}} Q_{12}^T \right]_{3 \times 3n}, \text{ and} \quad (18)$$

$$\text{Residue Vectors} = \left[Q_{21}^T V + \sqrt{\frac{1}{n}} Q_{22}^T I_{n_x} - \sqrt{\frac{n_x}{n}} Q_{22}^T \right]_{3 \times 3n}. \quad (19)$$

Note that the dimension of the *Residue Vectors* in this case is 3.

Given an adaptive binary tree structure, Algorithm 3.1 shows the pseudocode of the recursive function Q_gen that explicitly generates the matrix $Q = [Q_Z; \tilde{Q}]$. The orthogonal matrix is generated by calling Q_gen on the root node.

Algorithm 3.1 Recursive algorithm for explicit generation of Q

```

1: function  $Q\_gen(p)$ 
2:   if  $p$  is a leaf node then
3:     centroid = bead location,  $U = \mathbf{0}$ ,  $T_{22} = \mathbf{0}_{3 \times 3}$ 
4:   else
5:     Find child nodes  $x$  and  $y$  of node  $p$ 
6:      $Q\_gen(x)$ 
7:      $Q\_gen(y)$ 
8:     Compute centroid of  $p$  and form  $R_{21}$  and  $S_{21}$ 
9:     if both  $x$  and  $y$  are leaf nodes then
10:      Compute  $T_{22}$  and  $U$  using equations (15) and (16)
11:     else if only one of  $x$  and  $y$  is a leaf node then
12:      Compute  $T_{22}$ ,  $U$ , and Residue Vectors using (17), (18) and (19)
13:      Output Residue Vectors of size 3
14:     else
15:      Compute  $T_{22}$ ,  $U$ , and Residue Vectors using (12), (13) and (14)
16:      Output Residue Vectors of size 6
17:     if  $p$  is the root then
18:       Output  $U$ 

```

Algorithm complexity.

To estimate the algorithm complexity and storage requirement, we consider a system with n beads and a tree with $\mathcal{O}(\log n)$ levels. For each node in the tree, the number of operations to compute U and *Residue Vectors* and the storage required for these vectors are both constant times the number of beads in the node. Therefore, approximately $\mathcal{O}(n)$ operations and storage are required for each level of the tree structure and the overall complexity and memory requirement for the algorithm are $\mathcal{O}(n \log n)$.

4. Hierarchical $\mathcal{O}(n)$ algorithms for $Q\mathbf{v}$ and $Q^T\mathbf{v}$.

In the Brownian dynamics applications, one only needs the results of $Q\mathbf{v}$ and $Q^T\mathbf{v}$ instead of generating Q and Q^T explicitly. In this section, we show how to apply the hierarchical modeling technique to further compress the information and reduce the operations and storage for each tree node to a constant, independent of the number of beads contained in the node. As a result, the overall algorithm complexity and storage both become asymptotically optimal $\mathcal{O}(n)$.

4.1. Upward pass for computing $Q\mathbf{v}$.

We first consider $Q\mathbf{v}$ and discuss how to compress the information in the vectors U and *Residue Vectors*. We introduce the following definitions.

DEFINITION 4.1. For a nonleaf node p containing $n > 1$ beads in the tree structure, its Info-set $\mathbf{M}_p$ and Residue ϵ_p are respectively defined as

$$\mathbf{M}_p = Q_{Z_p} \cdot \begin{bmatrix} \mathbf{f}_1 \\ \vdots \\ \mathbf{f}_n \end{bmatrix}, \quad \epsilon_p = \text{Residue Vectors} \cdot \begin{bmatrix} \mathbf{f}_1 \\ \vdots \\ \mathbf{f}_n \end{bmatrix}, \quad (20)$$

where $\mathbf{f}_j$ is the force acting on bead j , Q_{Z_p} is the orthogonal matrix in the decomposition of p 's Z -matrix, and we assume the origin is located at the centroid of p . For a leaf node with only one bead, we define $\mathbf{M}_p = I \cdot \mathbf{f}$.

Ordering $\mathbf{f}$ for parent p as $[\mathbf{f}_x; \mathbf{f}_y]$, from (13) and (14), we have

$$\begin{aligned} \begin{bmatrix} Q_{Z_p} \\ \text{Residue Vectors} \end{bmatrix} \cdot \mathbf{f} &= \begin{bmatrix} I_{n_p} \mathbf{f} \\ U\mathbf{f} \\ \text{Residue Vectors} \cdot \mathbf{f} \end{bmatrix} \\ &= \begin{bmatrix} \sqrt{\frac{n_x}{n}} I_{n_x} \mathbf{f}_x + \sqrt{\frac{n_y}{n}} I_{n_y} \mathbf{f}_y \\ \begin{bmatrix} Q_{11}^T & Q_{12}^T & Q_{13}^T \\ Q_{21}^T & Q_{22}^T & Q_{23}^T \\ Q_{31}^T & Q_{32}^T & Q_{33}^T \end{bmatrix} \begin{bmatrix} V & 0 \\ 0 & W \\ \sqrt{\frac{n_y}{n}} I_{n_x} & -\sqrt{\frac{n_x}{n}} I_{n_y} \end{bmatrix} \begin{bmatrix} \mathbf{f}_x \\ \mathbf{f}_y \end{bmatrix} \end{bmatrix} \\ &= \begin{bmatrix} \sqrt{\frac{n_x}{n}} \mathbf{M}_x(1:3) + \sqrt{\frac{n_y}{n}} \mathbf{M}_y(1:3) \\ \begin{bmatrix} Q_{11}^T & Q_{12}^T & Q_{13}^T \\ Q_{21}^T & Q_{22}^T & Q_{23}^T \\ Q_{31}^T & Q_{32}^T & Q_{33}^T \end{bmatrix} \begin{bmatrix} \mathbf{M}_x(4:6) \\ \mathbf{M}_y(4:6) \\ \sqrt{\frac{n_y}{n}} \mathbf{M}_x(1:3) - \sqrt{\frac{n_x}{n}} \mathbf{M}_y(1:3) \end{bmatrix} \end{bmatrix}. \end{aligned}$$

Combined with (11) that states both T_{22} and Q_{ij} can be computed from children's R_{21} , R_{22} , S_{21} , and S_{22} , and all are 3×3 matrices, we see that both $\mathbf{M}_p$ and $\boldsymbol{\epsilon}_p$ can be computed using only a constant number of operations by

$$\mathbf{M}_p = \begin{bmatrix} \sqrt{\frac{n_x}{n}}\mathbf{M}_x(1:3) + \sqrt{\frac{n_y}{n}}\mathbf{M}_y(1:3) \\ \left[\begin{array}{c} \mathbf{M}_x(4:6) \\ \mathbf{M}_y(4:6) \\ \sqrt{\frac{n_y}{n}}\mathbf{M}_x(1:3) - \sqrt{\frac{n_x}{n}}\mathbf{M}_y(1:3) \end{array} \right] \\ \left[\begin{array}{ccc} Q_{11}^T & Q_{12}^T & Q_{13}^T \end{array} \right] \end{bmatrix}, \quad (21)$$

$$\boldsymbol{\epsilon}_p = \begin{bmatrix} \left[\begin{array}{ccc} Q_{21}^T & Q_{22}^T & Q_{23}^T \\ Q_{31}^T & Q_{32}^T & Q_{33}^T \end{array} \right] \left[\begin{array}{c} \mathbf{M}_x(4:6) \\ \mathbf{M}_y(4:6) \\ \sqrt{\frac{n_y}{n}}\mathbf{M}_x(1:3) - \sqrt{\frac{n_x}{n}}\mathbf{M}_y(1:3) \end{array} \right] \\ \left[\begin{array}{ccc} Q_{21}^T & Q_{22}^T & Q_{23}^T \\ Q_{31}^T & Q_{32}^T & Q_{33}^T \end{array} \right] \left[\begin{array}{c} \mathbf{M}_x(4:6) \\ \mathbf{M}_y(4:6) \\ \sqrt{\frac{n_y}{n}}\mathbf{M}_x(1:3) - \sqrt{\frac{n_x}{n}}\mathbf{M}_y(1:3) \end{array} \right] \end{bmatrix}. \quad (22)$$

Clearly, the *Info-set* of the child node contains the compressed information for generating parent node p 's *Info-set* and *Residue*. It is interesting to compare the *Info-set* with the “multipole expansion” in the FMM. Both provide effective ways to compress data contained in the node which will be sent to the interacting nodes in the tree structure.

Bead and bead and rigid body and bead cases.—When one or both of the child nodes are childless, the *Info-set* and *Residue* of the parent can be constructed using (16), (18), and (19). An alternative approach is to store a size 6 vector for both the leaf and nonleaf nodes such that a unified formula can be used for all cases. For a leaf node, the *Info-set* is simply $\mathbf{M} = [\mathbf{f}; \mathbf{0}]$, where the last three numbers are due to the fact that we choose the centroid as the origin.

Instead of the levelwise for-loop based execution in traditional FMM implementation, we present in Algorithm 4.1 a recursive implementation of the function `Compute_residue` for computing $Q\mathbf{v}$.

Algorithm complexity.—The complexity and storage of the algorithm can be estimated by checking the operations and memory requirement for each node in the tree structure. Notice that T_{22} is of size 3×3 , the matrix for storing Q_{ij} is of size

Algorithm 4.1 Recursive algorithm for computing Qv

```

1: function Compute_residue( $p$ )
2:   if node  $p$  is a leaf node then
3:     Construct Info-set  $\mathbf{M}_p$  directly
4:   else
5:     Find child nodes  $x$  and  $y$  of node  $p$ 
6:     Compute_residue( $x$ )
7:     Compute_residue( $y$ )
8:     Compute  $T_{22}$  and  $Q_{ij}$  using (11) and (12)
9:     Construct  $p$ 's Info-set  $\mathbf{M}_p$  and Residue  $\epsilon$  using (21) and (22)
10:    Output the Residue  $\epsilon$ 

```

9×9 , *Info-set* is a vector of size 6, and the size of the *Residue* is no more than 6. Therefore, both the number of operations and required storage are constants for each node. The overall algorithm complexity and storage are therefore proportional to the total number of nodes in the hierarchical tree structure. For most practical bead distributions, as the number of nodes in the tree structure is proportional to the number of beads n , the algorithm complexity and storage are both $\mathcal{O}(n)$.

4.2. Downward pass for computing $Q^T v$.

In rigid body dynamics, given the translational velocity $\mathbf{V}$ and angular velocity $\boldsymbol{\omega}$ of a rigid body p consisting of n beads, located at $\{\mathbf{r}_i\}$, and assuming that the reference point for the angular velocity is located at the centroid that is chosen as the origin, the velocity of bead i in the rigid body can be computed by

$$\mathbf{v}_i = \mathbf{V} + \boldsymbol{\omega} \times \mathbf{r}_i.$$

In matrix form, using the *Z-matrix*, the velocities of all the beads are given by

$$\mathbf{v} = Z_p^T \begin{bmatrix} \mathbf{V} \\ \boldsymbol{\omega} \end{bmatrix},$$

i.e., the velocity vector $\mathbf{v}$ is a linear combination of the column vectors in Z_p^T and the coupling coefficients are given by $\begin{bmatrix} \mathbf{V} \\ \boldsymbol{\omega} \end{bmatrix}$. To preserve orthogonality, we only consider the orthogonal basis vectors in the decomposition of Z_p and introduce the “generalized” velocities

$$\begin{bmatrix} I_{n_p}^T \vdots U^T \end{bmatrix} \begin{bmatrix} \mathbf{V} \\ \boldsymbol{\omega} \end{bmatrix} = Q_{Z_p}^T \begin{bmatrix} \mathbf{V} \\ \boldsymbol{\omega} \end{bmatrix}. \quad (23)$$

Notice that only six numbers in $\mathbf{V}$ and $\boldsymbol{\omega}$ are needed to construct the velocities of all the beads using (23). We define the vector containing these numbers as the *Velocity-set* of p as follows.

DEFINITION 4.2. *The Velocity-set of a rigid body p consisting of $n > 1$ beads is defined as a 6×1 vector $\mathbf{L} = [\mathbf{V}; \boldsymbol{\omega}]$ that describes the translational and angular velocities of p . The velocities of each individual bead in the rigid body p are given by (23).*

In the hierarchical modeling technique, consider a parent node p with child nodes x and y . The velocities of all the beads of p can be computed using p 's *Velocity-set* as $Q_{Z_p}^T \mathbf{L}_p$. Notice that both Q_{Z_p} and the *Residue Vectors* computed using (11)–(14) are combinations of

$\begin{bmatrix} Q_{Z_x}, \mathbf{0} \end{bmatrix}$ and $\begin{bmatrix} \mathbf{0}, Q_{Z_y} \end{bmatrix}$, which suggests that if we break the rigidity of p and assume both x and y are rigid bodies but permit their relative location to change, then both $Q_{Z_p}^T \mathbf{L}_p$ and the linear combination of the column vectors in the transpose of the *Residue Vectors* can be stored in $\begin{bmatrix} Q_{Z_x}^T \mathbf{L}_x \\ Q_{Z_y}^T \mathbf{L}_y \end{bmatrix}$, where $\mathbf{L}_x$ and $\mathbf{L}_y$ are the *Velocity-set* of children x and y , respectively. This is

summarized in the following theorem.

THEOREM 4.3. *Assume a rigid body p is partitioned into children x and y*

$$Q_{Z_p} = \begin{bmatrix} I_{n_p} \\ U \end{bmatrix}, \quad Q_{Z_x} = \begin{bmatrix} I_{n_x} \\ V \end{bmatrix}, \quad Q_{Z_y} = \begin{bmatrix} I_{n_y} \\ W \end{bmatrix};$$

then for any vector $\mathbf{v}$ containing the coupling coefficients of p 's Residue Vectors, there exist Velocity-set vectors $\mathbf{L}_x$ for child x and $\mathbf{L}_y$ for child y such that

$$Q_{Z_p}^T \mathbf{L}_p + (\text{Residue Vectors})^T \mathbf{v} = \begin{bmatrix} Q_{Z_x}^T \mathbf{L}_x \\ Q_{Z_y}^T \mathbf{L}_y \end{bmatrix}, \quad (24)$$

where

$$\mathbf{L}_x = \begin{bmatrix} \sqrt{\frac{n_x}{n}} I & \sqrt{\frac{n_y}{n}} Q_{13} \\ 0 & Q_{11} \end{bmatrix} \mathbf{L}_p + \begin{bmatrix} \sqrt{\frac{n_y}{n}} Q_{23} & \sqrt{\frac{n_y}{n}} Q_{33} \\ Q_{21} & Q_{31} \end{bmatrix} \mathbf{v}, \quad (25)$$

$$\mathbf{L}_y = \begin{bmatrix} \sqrt{\frac{n_y}{n}} I & -\sqrt{\frac{n_x}{n}} Q_{13} \\ \mathbf{0} & Q_{12} \end{bmatrix} \mathbf{L}_p + \begin{bmatrix} -\sqrt{\frac{n_x}{n}} Q_{23} & -\sqrt{\frac{n_x}{n}} Q_{33} \\ Q_{22} & Q_{32} \end{bmatrix} \mathbf{v}. \quad (26)$$

In other words, $Q_{Z_p} \mathbf{L}_p + (\text{Residue Vectors})^T \mathbf{v}$ can be compressed and stored in the *Velocity-set* $\mathbf{L}_x$ and $\mathbf{L}_y$ of child nodes x and y , respectively.

Theorem 4.3 is proved by applying $\begin{bmatrix} Q_{Z_x} \\ Q_{Z_y} \end{bmatrix}$ to both sides of (24) and we skip the details. It

suggests that $Q^T \mathbf{v}$ can be computed efficiently in a downward pass where the result of the matrix-vector product is considered as the velocity vectors of the beads. Starting from the root node's *Velocity-set*, the child nodes compute their *Velocity-set* using (25), (26) and the collection of *Velocity-set* at leaf nodes gives the result of $Q^T \mathbf{v}$. When only $\tilde{Q}^T \tilde{\mathbf{v}}$ is required, one only needs to replace the *Velocity-set* of the root node with a zero vector. For the special case when p is a leaf node containing a single bead, the velocity of the bead is $\mathbf{L}_p$ as we choose the location of the bead as the origin. It is also interesting to compare the *Velocity-set* with the “local expansion” in the FMMs. Both are used to store information inherited from parent levels.

We present the algorithm `Compute_velocity` in a recursive fashion in Algorithm 4.2. Here, we assume that Q_{ij} matrices for each node are already computed and stored in the upward pass discussed in section 4.1.

Algorithm complexity.—Similar to the upward pass to compute $Q\mathbf{v}$, as a constant amount of operations (and storage) is required for each node, the complexity of the recursive algorithm is $\mathcal{O}(n)$.

5. Preliminary numerical results.

We present some preliminary numerical results to show the complexity and orthogonality properties of the new algorithms introduced in this paper. The prototype implementations of the algorithms were done in MATLAB and the numerical tests were carried out on a personal laptop with an Intel Core i5 CPU at 2.6 GHz clock rate and 8 GB of RAM.

Algorithm 4.2 Recursive algorithm for computing $Q^T \mathbf{v}$

```

1: function Compute_velocity( $p$ )
2:   if  $p$  is a leaf node then
3:     Output the Velocity-set  $\mathbf{L}_p$ 
4:   else
5:     Find child nodes  $x$  and  $y$ 
6:     Compute  $x$  and  $y$ 's Velocity-set using (25) and (26)
7:     Compute_velocity( $x$ )
8:     Compute_velocity( $y$ )

```

5.1. Algorithm complexity.

We consider the memory requirement and CPU time of each algorithm for a variety number of beads n that ranges from 500 to 4000 in increments of 500. The memory requirement for each run was extracted from the MATLAB profiler's report with the memory option enabled and is summarized in Figure 1. Particularly, the fitted curves are

$$\begin{aligned}
 Q_gen: & \quad 8.87n \log n - 1.08 \cdot 10^4, \\
 \text{Compute_residue:} & \quad 5.31n + 1.69 \cdot 10^3, \\
 \text{Compute_velocity:} & \quad 3.17n + 9.08 \cdot 10^2,
 \end{aligned}$$

which agree with the analytic results. Clearly, the explicit generation of the matrix using Q_gen requires more memory than the implicit algorithms.

In Figure 2, we present the CPU time in seconds for the three algorithms. In the experiments, each algorithm was executed 10 times for each value of n and we present $\min_{1 \leq i \leq 10}$ (CPU time in i th run). The fitted curves are

$$\begin{aligned}
 Q_gen: & \quad 2.97 \cdot 10^{-5} n \log n + 1.51 \cdot 10^{-2}, \\
 \text{Compute_residue:} & \quad 1.56 \cdot 10^{-4} n + 3.57 \cdot 10^{-4}, \\
 \text{Compute_velocity:} & \quad 1.33 \cdot 10^{-4} n + 2.57 \cdot 10^{-3}
 \end{aligned}$$

and match the analytic results. The implicit methods are more efficient than the explicit function Q_gen . Furthermore, the downward pass Compute_velocity is more efficient than the upward pass Compute_residue because the upward pass needs to compute the matrices T_{22} and Q_{ij} that are also used in the downward pass.

5.2. Orthogonality properties.

By introducing the orthogonal numerical linear algebra technique, the hierarchical models and recursive algorithms better preserve the orthogonality properties of the matrix Q and Q^T . Table 2 shows various orthogonality properties of the algorithms, using the same beads generated in section 2.2. All the errors are measured in infinity norm. In columns two and three, we consider $QQ^T - I$ and $Q^T Q - I$ of matrix Q explicitly generated using Q_gen . In column four, we apply the explicitly generated Q to a given vector $\mathbf{v}$ and compare the output to that computed from Compute_residue . In column five, we apply the transpose of the

explicitly generated matrix Q to a given vector and compare the output to that computed from Compute_velocity. In column six, we first apply Compute_residue to compute $Q\mathbf{v}$ and then apply function Compute_velocity and compare it against the input vector. In column seven, we first apply Compute_velocity to compute $Q^T\mathbf{v}$ and then apply function Compute_residue and compare it against the input vector. Compared with the results in Table 1, all the errors in Table 2 are close to machine precision, providing strong evidence of the orthogonality preserving properties of the three algorithms.

6. Summary and related work.

In this paper, we apply the hierarchical modeling technique and present three recursive algorithms for generating special hierarchical structured orthogonal matrices with applications in Brownian dynamics simulations of biomolecular systems. By combining the orthogonal linear algebra techniques with the hierarchical models, our preliminary numerical experiments show that the implicit algorithms for computing $Q\mathbf{v}$ and $Q^T\mathbf{v}$ are both asymptotically optimal in complexity and have good orthogonality properties.

We are currently implementing the parallel versions of Algorithms 4.1 and 4.2 and developing toolboxes that will be integrated with our Brownian dynamics simulation package. As the algorithmic structure of these models is very close to the FMM, we plan to adapt and extend parallelization techniques in existing packages such as REC-FMM [23] and DASHMM [7]. The developed software will be released to the research community under open-source license agreement.

Our algorithms apply the divide-and-conquer strategy hierarchically to take advantage of the special structures in the Z -matrix. These ideas can be generalized to efficiently construct the “remaining” $n - k$ orthogonal vectors of more general matrices of size $n \times k$ with $n \gg k$ and may provide new formulations to some “linearly constrained” systems or problems with “low-rank” properties. Finally, we want to emphasize that in addition to the orthogonal matrix generation and matrix vector multiplication algorithms presented in this paper, other building blocks are also required to solve (5)–(6) and to build an optimal new Brownian dynamics simulation package. In particular, FMM-based parallel solvers have been developed for evaluating the electrostatic forces (contributing to the right-hand side of (3)) and for evaluating the RPY tensor matrix-vector multiplications $D\mathbf{v}$ [18, 19, 24]. The authors are also investigating effective preconditioning strategies for solving (5)–(6) where the fast direct solver or $\mathcal{H}$ -matrix techniques [11, 14, 16, 17] can be applied to each rigid body to generate a diagonal blocked preconditioner to further accelerate the convergence of the Krylov subspace iterative methods. Some of these topics are briefly discussed in [10]. Detailed results along these directions will be discussed in future papers.

Acknowledgments

Funding: The work of the first and second authors was supported by the National Science Foundation under grant DMS-1217080. The third and fourth authors’ work was supported by the Howard Hughes Medical Institute and National Institute of Health under grants GM31749 and GM103426. The fifth author’s work was supported by the NSF under grant ACI-1440396.

REFERENCES

- [1]. BALBOA USABIAGA F, KALLEMOV B, DELMOTTE B, BHALLA A, GRIFFITH B, AND DONEV A, Hydrodynamics of suspensions of passive and active rigid particles: A rigid multiblob approach, *Commun. Appl. Math. Comput. Sci*, 11 (2017), pp. 217–296.
- [2]. BAO Y, RACHH M, KEAVENY E, GREENGARD L, AND DONEV A, A Fluctuating Boundary Integral Method for Brownian Suspensions, preprint, arXiv:1709.01480, 2017.
- [3]. BATCHELOR G, Brownian diffusion of particles with hydrodynamic interaction, *J. Fluid Mech*, 74 (1976), pp. 1–29.
- [4]. BRANDT A, Multi-level adaptive solutions to boundary-value problems, *Math. Comp*, 31 (1977), pp. 333–390.
- [5]. CARRASCO B AND DE LA TORRE JG, Hydrodynamic properties of rigid particles: Comparison of different modeling and computational procedures, *Biophys. J*, 76 (1999), pp. 3044–3057. [PubMed: 10354430]
- [6]. COOLEY JW AND TUKEY JW, An algorithm for the machine calculation of complex fourier series, *Math. Comp*, 19 (1965), pp. 297–301.
- [7]. DEBUHR J, ZHANG B, TSUEDA A, TILSTRA-SMITH V, AND STERLING T, DASHMM: Dynamic Adaptive System for Hierarchical Multipole Methods, *Commun. Comput. Phys*, 20 (2016), pp. 1106–1126.
- [8]. DŁUGOSZ M AND ANTOSIEWICZ JM, Toward an accurate modeling of hydrodynamic effects on the translational and rotational dynamics of biomolecules in many-body systems, *J. Phys. Chem. B*, 119 (2015), pp. 8425–8439. [PubMed: 26068580]
- [9]. ERMAK DL AND MCCAMMON J, Brownian dynamics with hydrodynamic interactions, *J. Chem. Phys*, 69 (1978), pp. 1352–1360.
- [10]. Fang F, Recursive Tree Algorithms for Orthogonal Matrix Generation and Matrix-Vector Multiplications in Rigid Body Simulations, Undergraduate honors thesis, University of North Carolina at Chapel Hill, 2016.
- [11]. GREENGARD L, GUEYFFIER D, MARTINSSON PG, AND ROKHLIN V, Fast direct solvers for integral equations in complex three-dimensional domains, *Acta Numer*, 18 (2009), pp. 243–275.
- [12]. GREENGARD L AND ROKHLIN V, A fast algorithm for particle simulations, *J. Comput. Phys*, 73 (1987), pp. 325–348.
- [13]. GREENGARD L AND ROKHLIN V, A new version of the fast multipole method for the Laplace equation in three dimensions, *Acta. Numer*, 6 (1997), pp. 229–269.
- [14]. HACKBUSCH W, A sparse matrix arithmetic based on $\mathcal{H}$ -matrices. part I: Introduction to $\mathcal{H}$ -matrices, *Computing*, 62 (1999), pp. 89–108.
- [15]. HACKBUSCH W, *Multi-Grid Methods and Applications*, Vol. 4, Springer Science & Business Media, New York, 2013.
- [16]. HACKBUSCH W AND Khoromskij BN, A sparse $\mathcal{H}$ -matrix arithmetic, *Computing*, 64 (2000), pp. 21–47.
- [17]. HO KL AND GREENGARD L, A fast direct solver for structured linear systems by recursive skeletonization, *SIAM J. Sci. Comput*, 34 (2012), pp. A2507–A2532.
- [18]. JIANG S, LIANG Z, AND HUANG J, A fast algorithm for Brownian dynamics simulation with hydrodynamic interactions, *Math. Comp*, 82 (2013), pp. 1631–1645.
- [19]. LIANG Z, GIMBUTAS Z, GREENGARD L, HUANG J, AND JIANG S, A fast multipole method for the Rotne–Prager–Yamakawa tensor and its applications, *J. Comput. Phys*, 234 (2013), pp. 133–139.
- [20]. SPRINKLE B, USABIAGA FB, PATANKAR NA, AND DONEV A, Large Scale Brownian Dynamics of Confined Suspensions of Rigid Particles, preprint, arXiv:1709.02410, 2017.
- [21]. WANG N, HUBER GA, AND MCCAMMON JA, Assessing the two-body diffusion tensor calculated by the bead models, *J. Chem. Phys*, 138 (2013), p. 204117. [PubMed: 23742464]
- [22]. YING L, BIROS G, AND ZORIN D, A kernel-independent adaptive fast multipole algorithm in two and three dimensions, *J. Comput. Phys*, 196 (2004), pp. 591–626.

- [23]. ZHANG B, HUANG J, PITSIANIS NP, AND SUN X, recFMM: Recursive parallelization of the adaptive fast multipole method for coulomb and screened coulomb interactions, Commun. Comput. Phys, 20 (2016), pp. 534–550.
- [24]. ZHANG B, PENG B, HUANG J, PITSIANIS NP, SUN X, AND LU B, Parallel AFMPB solver with automatic surface meshing for calculation of molecular solvation free energy, Comput. Phys. Commun, 190 (2015), pp. 173–181.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Fig. 1. Memory usage in kilobytes versus the number of beads n for `Q_gen` (blue squares), `Compute_velocity` (green triangles), and `Compute_residue` (red diamonds).

Fig. 2.

CPU time in seconds (t) versus the number of beads n for Q_gen (blue squares), Compute_velocity (green triangles), and Compute_residue (red diamonds).

Table 1

Loss of orthogonality in the straightforward implementation.

n	500	1000	1500	2000	2500
$QQ^T - I$	2.6e-13	3.7e-13	1.5e-11	1.5e-11	1.5e-11
$Q^T Q - I$	1.9e-13	2.1e-13	6.5e-12	6.5e-12	6.5e-12
n	3000	3500	4000	8000	10000
$QQ^T - I$	4.2e-12	4.2e-12	4.2e-12	1.1e-11	1.8e-11
$Q^T Q - I$	1.6e-12	1.6e-12	1.6e-12	1.3e-11	1.7e-11

Table 2

Orthogonality preserving quality of the three algorithms.

n	$QQ^T - I$	$Q^T Q - I$	Qv	$Q^T v$	$Q^T(Qv) - v$	$Q(Q^T v) - v$
500	3.1e-15	6.6e-16	2.8e-15	2.3e-15	1.1e-15	4.8e-15
1000	1.3e-15	8.8e-16	8.4e-15	2.4e-15	2.2e-15	9.7e-15
1500	2.2e-14	1.1e-15	1.4e-14	1.8e-15	3.0e-15	7.3e-15
2000	1.9e-15	1.9e-15	4.0e-14	3.1e-15	1.8e-15	1.3e-14
2500	6.2e-15	8.8e-16	2.0e-14	1.8e-15	2.0e-15	1.7e-14
3000	7.1e-15	1.1e-15	1.9e-14	4.5e-15	3.2e-15	1.9e-14
3500	7.4e-14	1.1e-15	3.0e-14	5.6e-15	1.7e-15	2.8e-14
4000	2.8e-15	1.1e-15	3.4e-14	7.2e-15	1.6e-15	2.1e-14
8000	1.2e-13	8.8e-16	9.4e-14	5.3e-15	3.6e-15	2.3e-14
10000	7.1e-15	8.8e-16	7.5e-14	1.1e-14	2.9e-15	3.5e-14