

ANKH: A Generalized $O(N)$ Interpolated Ewald Strategy for Molecular Dynamics Simulations

Igor Chollet,* Louis Lagardère,* and Jean-Philip Piquemal*

Cite This: *J. Chem. Theory Comput.* 2023, 19, 2887–2905


Read Online

ACCESS |

Metrics & More

Article Recommendations

ABSTRACT: To evaluate electrostatics interactions, molecular dynamics (MD) simulations rely on Particle Mesh Ewald (PME), an $O(N\log(N))$ algorithm that uses Fast Fourier Transforms (FFT's) or, alternatively, on $O(N)$ Fast Multipole Methods (FMM) approaches. However, the FFT's low scalability remains a strong bottleneck for large-scale PME simulations on supercomputers. On the opposite, FFT-free FMM techniques are able to deal efficiently with such systems but they fail to reach PME performances for small- to medium-size systems, limiting their real-life applicability. We propose ANKH, a strategy grounded on interpolated Ewald summations and designed to remain efficient/scalable for any size of systems. The method is generalized for distributed point multipoles, and so for induced dipoles, which makes it suitable for high performance simulations using new generation polarizable force fields toward exascale computing.



1. INTRODUCTION

Performing Molecular Dynamics (MD) simulations requires to numerically solve Newton's equations of motion for a system of interacting particles. Modern simulations rely on Molecular Mechanics to evaluate the different physical interactions acting on an ensemble of particles. These approaches are very diverse and range from classical force fields (FFs)^{9,10,44} to a more evolved polarizable approach embodying many-body effects (PFFs).^{24,29,36} In any case, one needs to compute electrostatic interactions that are associated with the Coulomb energy and are an essential contribution to the systems total potential energy. To efficiently compute these quantities, MD softwares mainly exploit grid based methods such as the Particle Mesh Ewald (PME) approach, which is an $O(N \log(N))$ algorithm that uses Fast Fourier Transforms (FFT's). While PME is extremely efficient for small to medium system sizes, for very large ensembles of particles, the FFT's limited scalability is a major bottleneck for large-scale simulations on supercomputers. Historically, Fast Multipole Methods (FMM) have been considered as good $O(N)$ candidates to overcome such limitations. The FMM strategy is FFT-free and capable of efficiently dealing with such very large systems. Still, PME performance is yet to be reached for small- to medium-size systems, limiting the real-life applicability of FMMs. In this paper, we introduce ANKH, a new general strategy for the fast evaluation of the electrostatic energy using periodic boundary conditions and a density of charge due to point multipoles:


$$\mathcal{E} := \frac{1}{2} \sum_{\mathbf{t} \in 2r_B \mathbb{Z}^3} \sum_{\mathbf{x} \in X} \sum_{\mathbf{y} \in Y} \mathcal{D}_{\mathbf{x}} \mathcal{D}_{\mathbf{y}} \left(\frac{1}{|\mathbf{x} - \mathbf{y} + \mathbf{t}|} \right) \quad (1)$$

where $X, Y \subset B \subset \mathbb{R}^3$ are two point clouds (of atoms) that can be equal, $\mathcal{D}_{\mathbf{u}} := q_{\mathbf{u}} + \mu_{\mathbf{u}} \cdot \nabla + \Theta_{\mathbf{u}} : \nabla^2$, $q_{\mathbf{u}}$ is the charge of the atom \mathbf{u} , $\mu_{\mathbf{u}}$ its dipole moment, $\Theta_{\mathbf{u}}$ is its quadrupole moment, and ∇, ∇^2 respectively denote the gradient and Hessian operators, B is the cubical simulation box of radius r_B . Each point cloud X and Y is composed of N atoms, named *particles* to fit with the usual notation used in hierarchical methods.²⁰ Notice that we restricted ourselves to charges, dipoles, and quadrupoles for the sake of clarity, but that all the theory presented in this paper can be straightforwardly extended to any multipole order and therefore to induced dipoles, providing solutions for polarizable force fields.

Naive computation of the energy through direct implementation of eq 1 faces numerical convergence issues, requiring reformulation. The literature widely exploits Ewald summation techniques^{14,17,33,41,43} to provide numerically convergent expressions to compute eq 1. The previously mentioned limitations of 3D-FFT based techniques motivate the develop-

Received: January 4, 2023

Published: May 3, 2023



ment of alternatives,^{23,38,39} also based on Ewald summation, such as fast summation schemes for eq 1 directly^{25,27,31} and applied to molecular dynamics, mainly exploiting Fast Multipole Method²⁰ (FMM) and cutoff approaches.⁷ The latter has many advantages since it provides linear complexity and high scalability, despite the error that may occur when applying cutoff on images of B . Convergent alternatives^{11,35} also extend FMM for Coulomb potential to a periodic case when \mathcal{D}_x and \mathcal{D}_y are restricted to charges. Other kernels also benefit from periodic extensions.³²

Our ANKH approach aims at providing a theoretical framework well-suited for linear-complexity and scalable FMM-based energy computations, built on Ewald summation. The methodology presented here introduces various novelties, declined in two variants, both exploiting different ways of solving N -body problems appearing in Ewald summation. Among them, we introduce

- a new interpolated Ewald summation, leading to numerical schemes to handle differential operators of multipolar (polarizable) molecular dynamics, accelerated through two different numerical methods,
- alternative techniques to account for periodicity,
- explicit formulas to handle the mutual interactions.

Due to the positioning of our work, this article deals with mathematical and computer science topics in the precise context of molecular dynamics. For the sake of clarity, we provide both the theory and the algorithms that should allow the reader to minimize effort to implement our method. However, efficient implementations rely on various optimizations, and some details are beyond the scope of this paper. This is why we also provide our code, used to generate the results in the following.

The paper is organized as follows. We first briefly review in Sect. II the interpolation-based Fast Multipole Method (referred to as IBFMM in the remainder), then we introduce in Sect. III ANKH-FMM which is our new method for the fast Ewald Summation. In Sect. IV we introduce an entirely new approach named ANKH-FFT that overcomes some limitations of ANKH-FMM regarding modern High Performance Computing (HPC) architectures (at the cost of a slightly higher complexity than ANKH-FMM), then we present in Sect. V numerical results to emphasize the performance of our method as well as a comparison with a production code for molecular dynamics (namely Tinker-HP²⁶).

II. INTERPOLATION-BASED FMM

This section is dedicated to the presentation of interpolation-based FMM. We first recall in Sect. II.A the reasons of the efficiency of the original 3D FMM (for Coulomb potential) as well as the algorithm in the case of quasi-uniform particle distributions. Then, in Sect. II.B, we shortly describe one of the main kernel-independent FMM, namely the interpolation-based FMM, providing the corresponding operator formulas.

II.A. Generalities of FMMs. Fast Multipole Methods^{18,20,21} (FMM) is a family of fast divide and conquer strategies to compute N -body problems in linear or linearithmic complexity. Important efforts have been made in the literature to incorporate 3D FMM (i.e., for the Coulomb potential) in molecular dynamics simulation codes.^{25,27,37} Such a FMM scheme is based on a separable spherical harmonic expansion of the Coulomb potential $\frac{1}{|\mathbf{x} - \mathbf{y}|}$ writing as

$$\sum_{l=0}^{+\infty} \frac{4\pi r_{<}^l}{(2l+1)r_{>}^{l+1}} \sum_{m=-l}^l Y_l^m(\theta_x, \phi_x) Y_l^{-m}(\theta_y, \phi_y) \quad (2)$$

where $\mathbf{x} = (r_x, \theta_x, \phi_x)$, $\mathbf{y} = (r_y, \theta_y, \phi_y)$ in spherical coordinates, $r_{<} := \min\{r_x, r_y\}$, $r_{>} := \max\{r_x, r_y\}$ and Y_l^m refers to the spherical harmonic of order (l, m) . Restricting to subset of particles $\mathbf{x} \in t \subset \mathbb{R}^3$, $\mathbf{y} \in s \subset \mathbb{R}^3$ such that $\frac{\text{radius}(t) + \text{radius}(s)}{\text{distance}(t, s)} \leq \nu$, $\nu > 0$ sufficiently small, one may approximate eq 2 by truncating this series to a finite order P . This allows a derivation of a fast summation scheme for any $\mathbf{x} \in X$, for problems consisting in computing ϕ such that

$$\phi(\mathbf{x}) := \sum_{\mathbf{y} \in Y} \frac{q(\mathbf{y})}{|\mathbf{x} - \mathbf{y}|} \quad (3)$$

where $q(\mathbf{y})$ is a charge associated with the particle \mathbf{y} . Such a problem in eq 3 is referred to as a N -body problem and such t and s are named *well-separated target (t) and source (s) cells*, where a *cell* denotes in this paper a cubical subset of \mathbb{R}^3 . These cells are obtained in practice through a recursive splitting of the computational box B into smaller boxes. Assuming that B is a cube, one may divide it into 8 other cubes and repeat this procedure until the induced cubes (the cells) are sufficiently small. A tree whose nodes are cells and such that the daughters of a given cell are the nonempty cubes coming from its decomposition is named an *octree*. The root cell is considered to belong at octree level 0 and any non-root cell belongs at level $E + 1$, where E denotes its mother level. In the case of perfect octrees (i.e., with maximal amount of daughter per non-leaf node, that is 8 daughters), we consider that two cells at the same tree level with strictly positive distance (i.e., non-neighbor cells) are well-separated.

The sum over $\mathbf{y} \in Y_s$ (the restriction of Y to the particles in s) can be switched, for any $\mathbf{x} \in X_t$ to the terms depending only on \mathbf{y} , due to the separability in the variables \mathbf{x} and \mathbf{y} of eq 2. Expression in eq 2 results²¹ in a three-step summation scheme expressing eq 3 in the form




Figure 1. 2D schematic representation of FMM operators. A source cell s at level 2 (orange) interacts with well-separated target cells t 's at level 2 (purple) through M2L operators (green arrows). At level 3, daughters of s aggregate their multipole expansions through M2M operators (orange arrows) to form multipole expansion in s . Each t generates 4 L2L operators to add influence of s to its daughters through L2L operators (purple arrows). Red and blue arrows in source and target cells respectively represent the action of the P2M and L2P operators respectively in the corresponding cell. Gray cells at level 3 correspond to the daughter of non-well-separated cells from s at level 2.

$$\sum_k U_k[t]^*(\mathbf{x}) \sum_l C_{k,l}[t, s] \sum_{\mathbf{y} \in Y_s} V_l[s](\mathbf{y}) q(\mathbf{y}) \quad (4)$$

or equivalently under matrix form $U^*[t]C[t, s]V[s]q_s$ where q_s refers to the vector of charges of particles in s . The result M_s of the product $V[s]q_s$ is named the *multipole expansion* of source cell s . The matrix $V[s]$ involved in this product is named Particle-To-Multipole operator (shortly denoted by P2M) since it maps information associated with particles (i.e., here the charges associated with atoms) to the multipole expansion of the cell in which they belong. This multipole expansion is transformed into a *local* one \mathcal{L}_t by means of product by the Multipole-To-Local operator (shortly denoted by M2L) $C[t, s]$. Similarly, the product by U^* is referred to as the Local-To-Particle operator (L2P). Nevertheless, the efficiency of the FMM algorithm is also guaranteed by two additional operators based on the idea that multipole expansions of a given non-leaf-cell can be approximating through combination of its daughter's multipole expansions by means of a Multipole-To-Multipole operator (M2M) and local expansions in a given non-root-cell can be approximated from its mother local expansion thanks to a Local-To-Local operator (L2L). Finally, two leaf-cells that are not well-separated still need to interact, but without expansion. This is done by applying the Particle-To-Particle operator (P2P), i.e., nothing else than a direct computation involving only the particles of these two cells. The relation between all these operators according to the tree structure is depicted in Figure 1.

In the case of quasi-uniform particle distributions (i.e., with approximately the same density of particles in all the simulation box), octrees can be considered as perfect. In this situation, the FMM algorithm is easily provided (and usually referred to as the *uniform FMM algorithm* in the literature), according to the definition of the different FMM operators. This is summarized in Alg. 1, where we used the notation $V[s, s']$ to denote the M2M operator between the source cell s and one of its daughters s' ; $U^*[t', t]$ to denote the L2L operator between the target cell t and one of its daughters t' ; for any target cell t , $\Lambda(t)$ is the interaction list of t , that is, the set of source cells s that are well-separated from t and whose ancestors in the octree are not well-separated from any ancestor of t ; \mathcal{T} refers to the octree; $\mathcal{L}(\mathcal{T})$ refers to the set of leaves of \mathcal{T} , and $K(\mathbf{x}, \mathbf{y}) = \frac{1}{|\mathbf{x} - \mathbf{y}|}$.

Algorithm 1 Uniform FMM algorithm

```
// Step 1: compute multipole expansions
for each leaf cell  $s \in \mathcal{L}(\mathcal{T})$  do
  // Compute multipole expansion  $M_s$  in  $s$  using P2M
   $M_s = V[s]q_s$ 
end for
for each non leaf level  $e$  starting from the deepest one do
  for each cell  $s$  at level  $e$  do
    // Aggregate multipole expansions through M2M
     $M_s = \mathbf{0}$ 
    for each daughter  $s'$  of  $s$  do
       $M_s = M_s + V[s, s']M_{s'}$ 
    end for
  end for
end for
// Step 2: Compute local expansions
for each target cell  $t$  do
  // Add far contribution using M2L
   $\mathcal{L}_t = \mathbf{0}$ 
  for each source cell  $s \in \Lambda(t)$  do
     $\mathcal{L}_t = \mathcal{L}_t + C[t, s]M_s$ 
  end for
end for
for each level  $e$  starting from the root do
  // Scatter local expansions using L2L
  for each target cell  $t$  at level  $e$  do
    for each daughter  $t'$  of  $t$  do
       $\mathcal{L}_t = \mathcal{L}_t + U^*[t', t]\mathcal{L}_{t'}$ 
    end for
  end for
end for
// Step 3: convert local expansions
for each leaf cell  $t \in \mathcal{L}(\mathcal{T})$  do
  // Apply L2P
   $\phi_t = U^*[t]\mathcal{L}_t$ 
end for
// Step 4: add near contribution
for each leaf cell  $t \in \mathcal{L}(\mathcal{T})$  do
  // Apply all P2P involving  $t$ 
  for each source cell  $s$  adjacent to  $t$  do
    for  $\mathbf{x} \in X_t$  do
       $\phi(\mathbf{x}) = \phi_t(\mathbf{x})$ 
      for  $\mathbf{y} \in Y_s$  do
         $\phi(\mathbf{x}) = \phi(\mathbf{x}) + K(\mathbf{x}, \mathbf{y})q(\mathbf{y})$ 
      end for
    end for
  end for
end for
```

Alg. 1 results in $O(N)$ complexity²¹ with the prefactor depending on the user required precision. Notations of Alg. 1 directly provide a way of approximating the influence $\phi_{t',s'}$ of s' a daughter of s on t' and daughter of t , provided that t and s are well-separated, under matrix form

$$\phi_{t',s'} \approx \underbrace{U^*[t']}_{L2P} \underbrace{U^*[t', t]}_{L2L} \underbrace{C[t, s]}_{M2L} \underbrace{V[s, s']}_{M2M} \underbrace{V[s']}_{P2M} q_{s'} \quad (5)$$

This last can be interpreted as a low-rank matrix approximation when seeing eq 3 as a matrix-vector product with matrix entries corresponding to an evaluation of the Coulomb potential on pairs of target and source cells. Notice that the interaction between equally located particles is not well-defined, and we consider in this article that they are just equal to 0.

In terms of parallel implementation, 3D FMM has shown its ability to scale in distributed memory context. Combined with its linear complexity, this makes this method particularly attractive for MD simulations. However, few limitations still remain: the presentation we made only takes into account charges and is considered a nonperiodic case. This last can be solved by means of periodic FMM.¹¹

II.B. Interpolation-Based Kernel-Independent Scheme. Actually, a similar scheme can be derived for any kernel $K(\mathbf{x}, \mathbf{y})$ (not only the Coulomb one) that satisfies a

particular assumption, namely the *asymptotically smooth* behavior.⁸ There exist different approaches to obtain a fast separable formula in a *kernel-independent* way.^{18,46} Among them, we are especially interested into the interpolation-based FMM (IBFMM) because of its flexibility and its well documented error bounds, as well as the particular form of the approximated expansion of K it provides. Indeed, denoting by $S_k[t]$ (respectively $S_l[s]$) a multivariate Lagrange polynomial associated with the interpolation node \mathbf{x}_k^t (respectively \mathbf{y}_l^s) in t (respectively s), a Lagrange interpolation of K in $t \times s$ can be written as

$$K(\mathbf{x}, \mathbf{y}) \approx \sum_k S_k[t](\mathbf{x}) \sum_l K(\mathbf{x}_k^t, \mathbf{y}_l^s) S_l[s](\mathbf{y}) \quad (6)$$

for any $(\mathbf{x}, \mathbf{y}) \in t \times s$. Equation 6 clearly exhibits a separable expansion of K in the \mathbf{x} and \mathbf{y} variables. Hence, the restriction of eq 3 to particles in $t \times s$ (i.e., to X_t and Y_s and replacing the Coulomb kernel by a generic one K) can be approximated by

$$\underbrace{\sum_k S_k[t](\mathbf{x})}_{\text{L2P on } t} \underbrace{\sum_l K(\mathbf{x}_k^t, \mathbf{y}_l^s)}_{\text{M2L between } t \text{ and } s} \underbrace{\sum_{\mathbf{y} \in Y_s} S_l[s](\mathbf{y}) q(\mathbf{y})}_{\text{P2M on } s} \quad (7)$$

by analogy with eq 4. This identifies most of the FMM operators. Moreover, by interpreting multipole expansions generated this way as charges associated with interpolation nodes (now seen as particles), one can obtain a new N -body problem with the same kernel K . Hence, this procedure can be repeated all along the tree structure in order to approximate multipole expansions in the non-leaf cells (which is the analogous of the M2M operator). In the same way, evaluating interpolant at a non-leaf cell on interpolation nodes in its daughters allows to derive an analog of the L2L operator. To be more precise, we assume that the Lagrange interpolation applied on cell c uses the interpolation grid $\Xi_c = \{\mathbf{y}_0^s, \dots, \mathbf{y}_{p-1}^s\}$ with P interpolation nodes over c . In practice, Ξ_c is chosen as a product of one-dimensional interpolation nodes (such as Chebyshev nodes¹⁸). Hence, for any source cell s , the IBFMM M2M operator between s and s' , with s' a daughter of s , is defined as

$$V[s, s'] := \begin{bmatrix} S_0[s](\mathbf{y}_0^{s'}) & \dots & S_0[s](\mathbf{y}_{p-1}^{s'}) \\ \vdots & \ddots & \vdots \\ S_{p-1}[s](\mathbf{y}_0^{s'}) & \dots & S_{p-1}[s](\mathbf{y}_{p-1}^{s'}) \end{bmatrix} \quad (8)$$

Using eq 6, we can identify that the IBFMM L2P operator on leaf cell c is the adjoint of the IBFMM P2M operator on this c . Hence, for any non-leaf cell c , the IBFMM L2L operator between c and a daughter c' of c is the adjoint of the IBFMM M2M operator between c and c' , that is

$$U^*[c] = V[c]^*, \quad U^*[c', c] = V[c, c']^* \quad (9)$$

As for the classical FMM algorithm, the IBFMM reaches the linear complexity with respect to the number of particles in the system. In the remainder of this article, we consider that K is translationally and rotationally invariant, which is always verified in our applications. A schematic view of IBFMM is provided in Figure 2.

III. ANKH-FMM

Numerical issues arise when dealing with eq 1 since the series only conditionally converges, possibly preventing practical




Figure 2. Schematic representation of low-rank approximations obtained through interpolation over 1D quasi-uniform particle distributions. Perfect trees (binary trees here) are represented above (sources) and on the left (targets) of the matrix representation of the far field compression. There are admissible cell pairs (t, s) only at levels 2 (big dark gray blocks) and 3 (small dark gray blocks). In each of the corresponding blocks, we drew its low-rank approximation using red color for source polynomials (i.e., $V[s]$ matrix), blue color for target polynomials (i.e., $U^*[t]$ matrix and purple for small square matrices $C[t, s]$). Nonadmissible cell pairs are represented using light gray.

convergence. The Ewald summation technique¹⁴ transforms eq 1 into $\mathcal{E} = \mathcal{E}_{\text{real}} + \mathcal{E}_{\text{rec}} + \mathcal{E}_{\text{self}}$ where

$$\begin{aligned} \mathcal{E}_{\text{real}} &:= \sum_{\mathbf{t}} \sum_{\mathbf{x}} \sum_{\mathbf{y}} \mathcal{D}_{\mathbf{x}} \mathcal{D}_{\mathbf{y}} \left(\frac{\text{erfc}(\xi \|\mathbf{x} - \mathbf{y} + \mathbf{t}\|)}{\|\mathbf{x} - \mathbf{y} + \mathbf{t}\|} \right), \\ \mathcal{E}_{\text{rec}} &:= \frac{1}{2\pi(2r_B)^3} \sum_{\mathbf{m} \neq 0} \frac{e^{-2(\pi r_B / \xi)^2 \mathbf{m} \cdot \mathbf{m}}}{\mathbf{m} \cdot \mathbf{m}} S(\mathbf{m}) S(-\mathbf{m}), \\ S(\mathbf{m}) &:= \sum_{\mathbf{x}} e^{i[\pi/r_B] \mathbf{m} \cdot \mathbf{x}} (q_{\mathbf{x}} + 2i\pi r_B \mu_{\mathbf{x}} \cdot \mathbf{m} - (2\pi r_B)^2 \Theta_{\mathbf{x}} : (\mathbf{m} \mathbf{m}^T)), \\ \mathcal{E}_{\text{self}} &:= -\frac{\xi}{\sqrt{\pi}} \sum_{\mathbf{x}} \left(q_{\mathbf{x}}^2 + \frac{2\xi^2}{3} \mu_{\mathbf{x}} \cdot \mu_{\mathbf{x}} + \frac{8\xi^4}{5} \Theta_{\mathbf{x}} : \Theta_{\mathbf{x}} \right) \end{aligned} \quad (10)$$

in which all the series are absolutely convergent, $\xi \in \mathbb{R}^{+*}$ and $\text{erfc}(x) = 1 - \frac{2}{\pi} \int_0^x e^{-t^2} dt$ denotes the complementary error function. The main point is that erfc quickly decays for sufficiently large ξ , implying that only a small amount of images (i.e., of different \mathbf{t} 's) in $\mathcal{E}_{\text{real}}$ have to be considered to reach a given accuracy. Applying a cutoff with radius δ on the sum $\sum_{\mathbf{x}} \sum_{\mathbf{y}}$ transforms it into $\sum_{\mathbf{x}} \sum_{\mathbf{y} \in \mathcal{B}(\mathbf{x}, \delta)}$ accelerating its evaluation, using the notation $\mathcal{B}(\mathbf{x}, \delta) := \{\mathbf{z} \in \mathbb{R}^3 \mid \|\mathbf{z} - \mathbf{x}\| \leq \delta\}$. \mathcal{E}_{rec} also is a quickly converging sum and only a limited number of \mathbf{m} 's have to be considered to reach a given accuracy, provided that ξ is sufficiently large. Usually, ξ is chosen to minimize the evaluation cost according to the possible cutoffs, balancing the computation costs of $\mathcal{E}_{\text{real}}$ and \mathcal{E}_{rec} . The Particle Mesh Ewald¹⁴ (PME) algorithm uses Fast Fourier Transforms

(FFT) in the computation of $S(\mathbf{m})$'s, further reducing the complexity.

III.A. A New Alternative to PME. On one side, the parallelization of the PME algorithm in distributed memory can be a practical bottleneck due to the scalability of the FFT.⁴ On the other side, there exists fast scalable methods in distributed memory for the evaluation of N -body problems involving asymptotically smooth kernels, such as (IB)FMM (see Sect. II.B). We thus propose a new alternative to PME, based on a very simple idea. Indeed, the kernel

$$H(\mathbf{x}, \mathbf{y}) := \frac{\operatorname{erfc}(\xi|\mathbf{x} - \mathbf{y}|)}{|\mathbf{x} - \mathbf{y}|} \quad (11)$$

is asymptotically smooth,⁵ as well as is all its derivatives. For sufficiently large p , thanks to the absolute convergence of the terms in eq 10, we have

$$\begin{aligned} \mathcal{E}_{\text{real}} &\approx \sum_{\mathbf{t} \in 2r_B Z_p} \sum_{\mathbf{x}} \sum_{\mathbf{y}} \mathcal{D}_{\mathbf{x}} \mathcal{D}_{\mathbf{y}} \left(\frac{\operatorname{erfc}(\xi|\mathbf{x} - \mathbf{y} + \mathbf{t}|)}{|\mathbf{x} - \mathbf{y} + \mathbf{t}|} \right) \\ &= \sum_{\mathbf{t} \in 2r_B Z_p} \sum_{\mathbf{x}} \sum_{\mathbf{y}} \mathcal{D}_{\mathbf{x}} \mathcal{D}_{\mathbf{y}} H(\mathbf{x}, \mathbf{y} + \mathbf{t}) =: \tilde{\mathcal{E}}_{\text{real}} \end{aligned} \quad (12)$$

with $Z_p := \{(z_0, z_1, z_2) \in \mathbb{Z}^3 | |z_k| \leq p, k \in [0, 2]\}$ and p to be fixed by user.

Our approach is thus very simple: we consider small ξ (typically $\xi = 0.01$), so that $\mathcal{E}_{\text{rec}} \approx 0$ and can be neglected for our targeted accuracies. In this case, the counterpart relies in the computation of $\tilde{\mathcal{E}}_{\text{real}}$ that requires a relatively large amount of periodic images (i.e., a large p) to reach the numerical convergence. Since $\mathcal{E}_{\text{self}}$ involves $O(N)$ flops to be computed (and is embarrassingly parallel), only this $\tilde{\mathcal{E}}_{\text{real}}$ has a corresponding intensive computational cost. Hopefully, eq 12 allows the computation of $\tilde{\mathcal{E}}_{\text{real}}$ to be handled through fast methods for N -body problems with kernel $\mathcal{D}_{\mathbf{x}} \mathcal{D}_{\mathbf{y}} H(\mathbf{x}, \mathbf{y})$.

III.B. Reformulation. In this section, we consider that $p = 0$. Extension to other p 's will be done in Sect. III.C. Let

$$\tilde{\mathcal{E}}_{\text{real}}(t, s) := \sum_{\mathbf{x} \in t} \sum_{\mathbf{y} \in s} \mathcal{D}_{\mathbf{x}} \mathcal{D}_{\mathbf{y}} H(\mathbf{x}, \mathbf{y})$$

where $t, s \in \mathcal{L}(\mathcal{T})$. We can now divide $\tilde{\mathcal{E}}_{\text{real}}$ into two parts:

$$\underbrace{\sum_{\substack{t, s \in \mathcal{L}(\mathcal{T}) \\ \text{dist}(t, s) = 0}} \tilde{\mathcal{E}}_{\text{real}}(t, s)}_{=: \mathcal{N}_{\text{real}}} + \underbrace{\sum_{\substack{t, s \in \mathcal{L}(\mathcal{T}) \\ \text{dist}(t, s) \neq 0}} \tilde{\mathcal{E}}_{\text{real}}(t, s)}_{=: \mathcal{F}_{\text{real}}}$$

where $\text{dist}(t, s)$ denotes the distance between t and s . Hence, by increasing the octree depth, one decreases the number of interactions computed inside $\mathcal{N}_{\text{real}}$ (near field) while increasing the number of interactions in $\mathcal{F}_{\text{real}}$ (far field).

Each $(t, s) \in \mathcal{L}(\mathcal{T})$ such that $\text{dist}(t, s) \neq 0$ actually is a well-separated pair of cells. Hence, by means of multivariate polynomial interpolation as in eq 6 on t and s , one may accurately approximate $\tilde{\mathcal{E}}_{\text{real}}(t, s)$ as

$$\begin{aligned} &\sum_{\mathbf{x} \in t} \sum_{\mathbf{y} \in s} \mathcal{D}_{\mathbf{x}} \mathcal{D}_{\mathbf{y}} \left(\sum_k S_k[t](\mathbf{x}) \sum_l H(\mathbf{x}_k^t, \mathbf{y}_l^s) S_l[s](\mathbf{y}) \right) \\ &= \sum_k \left(\sum_{\mathbf{x} \in t} \mathcal{D}_{\mathbf{x}} S_k[t](\mathbf{x}) \right) \sum_l H(\mathbf{x}_k^t, \mathbf{y}_l^s) \left(\sum_{\mathbf{y} \in s} \mathcal{D}_{\mathbf{y}} S_l[s](\mathbf{y}) \right) \\ &= \sum_k Q_t(\mathbf{x}_k^t) \sum_l H(\mathbf{x}_k^t, \mathbf{y}_l^s) Q_s(\mathbf{y}_l^s) \end{aligned} \quad (13)$$

These $Q_c(\mathbf{u})$, with \mathbf{u} any interpolation node in c , can be interpreted as modified charges associated with this interpolation node, which can itself be seen as a particle. Hence, one can perform an IBFMM in which the multipole/local expansions are computed using the modified charges while the near field (i.e., $\mathcal{N}_{\text{real}}$) is evaluated using the explicit formula for the complete interaction between particles with respect to the kernel H (i.e., involving charges, dipoles, and quadrupoles here). In other words, this method (that we named ANKH) consists in the following steps:

1. Compute the tree decomposition \mathcal{T} of the simulation box B ,
2. For each leaf $c \in \mathcal{L}(\mathcal{T})$, compute the modified charges in c (using the same interpolation grid up to a translation to the center of c),
3. Compute $\mathcal{N}_{\text{real}}$ using direct computation,
4. Compute $\mathcal{F}_{\text{real}}$ solving the N -body problem defined on interpolation nodes and modified charges,
5. Compute $\mathcal{E}_{\text{self}}$ using eq 10,
6. Return $\mathcal{N}_{\text{real}} + \mathcal{F}_{\text{real}} + \mathcal{E}_{\text{self}}$.

The N -body problem mentioned in step 4 can be written this way:

$$\sum_{\substack{t \in \mathcal{L}(\mathcal{T}), \\ \mathbf{u} \text{ interpolation} \\ \text{node on } t}} Q_t(\mathbf{u}) \left(\sum_{\substack{s \in \mathcal{L}(\mathcal{T}), \\ \mathbf{v} \text{ interpolation} \\ \text{node on } s}} H_{t,s}(\mathbf{u}, \mathbf{v}) Q_s(\mathbf{v}) \right) \quad (14)$$

N -body problem as in eq (3)

where

$$H_{t,s} = \begin{cases} H & \text{if } \text{dist}(t, s) > 0, \\ 0 & \text{otherwise} \end{cases}$$

According to eq 14, the quantity in parentheses can be efficiently computed in linear time by means of IBFMM applied on kernel H and in which we simply avoid the P2P operators. We name this approach ANKH-FMM.

In the case of quasi-uniform particle distribution (actually the kind of particle distributions we are targeting), a perfect octree with depth $\log(N)$ generates leaves with $O(1)$ particles. Hence, in this situation, step 3 of the ANKH method costs $O(N)$ flops. For such perfect trees, the tree construction (step 1) also is $O(N)$ and computation of $\mathcal{E}_{\text{self}}$ is, of course, also linear in complexity (step 5). Since it is obvious that step 6 is $O(1)$, the overall complexity is $O(N)$.

Two last questions remain: what about the periodicity and the computation of modified charges? The first of these questions is the purpose of Sect. III.C and the second one is the purpose of Sect. III.D.

III.C. Efficient Handling of Periodicity. Various formulations of the FMM provide efficient ways of dealing with periodicity.^{11,32,35} We consider a slightly different approach in our method, well-suited for IBFMM. One idea is to perform a classical IBFMM on the computational box as well as on the direct $3^3 - 1$ adjacent images of it. Then, all other distant images are well-separated from the computational box and their interaction can be approximated using only the multipole expansion \mathcal{M}_B of this box. More precisely, denoting $C[B, I]$ as the M2L matrix between the computational box B and its image I , one wants to rapidly compute

$$\sum_{\substack{I \in Z_p \\ p \geq |I| > 1}} C[B, I] \mathcal{M}_B = \left(\sum_{\substack{I \in Z_p \\ p \geq |I| > 1}} C[B, I] \right) \mathcal{M}_B \quad (15)$$

The remaining task is thus to sum the M2L matrices between images of B and B itself efficiently. This can be done by means of interpolation over the equispaced grid.^{6,13}

Indeed, such a tool allows us to express $C[B, I]$ as a product $\chi^* \mathbb{F}^* D[B, I] \mathbb{F} \chi$, where χ is a zero-padding matrix that increases the size of the inner matrix by slightly less than 8, \mathbb{F} denotes the discrete Fourier matrix (that can be applied through FFT), and $D[B, I]$ is a diagonal matrix that can be computed in linearithmic time with respect to the size of $C[B, I]$.

More precisely, let $\Xi_B := \{\Xi_B(i, j, k) := (-r_B + 2r_{Bi}/(L-1), -r_B + 2r_{Bj}/(L-1), -r_B + 2r_{Bk}/(L-1)) | i, j, k \in [[0, L-1]]\}$ be an equispaced grid over B with L nodes in each direction (i.e., with L^3 nodes in total) and let Ξ_I be its translation from the origin to the center of the image I of B . The matrix $C[B, I]$, in this case, has a block-Toeplitz structure.^{6,13} Any such block-Toeplitz matrix can be embedded into a circulant one of size $(2L-1)^3$ whose first row is composed of the $(2L-1)^3$ different entries of $C[B, I]$. $\chi \in \{0, 1\}^{(2L-1)^3 \times L^3}$ and $D[B, I] \in \mathbb{C}^{(2L-1)^3 \times (2L-1)^3}$ are then defined as follows:

$$\chi_{q,p} := \begin{cases} 1 & \text{if } i_0 = i_1, j_0 = j_1; \\ & k_0 = k_1; i_1, j_1, k_1 < L \\ 0 & \text{otherwise} \end{cases}$$

$$D[B, I] := \text{diag}(\mathbb{F} R[B, I]) \quad (16)$$

where $R[B, I]$ is the first row of the circulant matrix formed by embedding of the Toeplitz matrix $C[B, I]$, i.e.,

$$R[B, I]_{i(2L-1)^2 + j(2L-1) + k} := H(\mathbf{z}_{i,j,k}^{(I)}, \mathbf{0}),$$

$$\mathbf{z}_{i,j,k}^{(I)} := \Xi_B(\tilde{i}, \tilde{j}, \tilde{k}) - \Xi_I(\hat{i}, \hat{j}, \hat{k}),$$

$$\tilde{i} := \begin{cases} 0 & \text{if } i < L \\ 2L - 1 - i & \text{otherwise} \end{cases}$$

$$\hat{i} := \begin{cases} i & \text{if } i < L \\ 0 & \text{otherwise} \end{cases} \quad (17)$$

Since \mathbb{F} is practically performed through FFT, computation of $D[B, I]$ results in $O(L^3 \log(L))$ flops.

We can thus reformulate eq 15 as

$$\chi^* \mathbb{F}^* \underbrace{\left(\sum_{\substack{I \in Z_p \\ p \geq |I| > 1}} D[B, I] \right)}_{=: D_B} \mathbb{F} \chi \mathcal{M}_B \quad (18)$$

As a sum of diagonal matrix, D_B clearly is diagonal. We can switch the sum in eq 18 to the vector $R[B, I]$ thanks to eq 16, resulting into

$$D_B = \text{diag} \left(\mathbb{F} \sum_{\substack{I \in Z_p \\ p \geq |I| > 1}} R[B, I] \right) \quad (19)$$

Since the image I of B is a translation of B by $\mathbf{t} = 2r_B I$, we have that $\Xi_I(\hat{i}, \hat{j}, \hat{k}) = \Xi_B(\hat{i}, \hat{j}, \hat{k}) + \mathbf{t}$, so that, thanks to eq 17 and the translational invariance of H ,

$$R[B, I]_{i(2L-1)^2 + j(2L-1) + k} = H(\mathbf{z}_{i,j,k}^{(0)} + \mathbf{t}, \mathbf{0})$$

$$= H(\mathbf{z}_{i,j,k}^{(0)}, -\mathbf{t}) \quad (20)$$

Notice that $-I$ also lies into the set of images of B so that one can eliminate the minus sign in eq 20. Also, $\{\mathbf{z}_{i,j,k}^{(0)} | i, j, k \in [[0, (2L-1)]]\}$ can be interpreted as a particle distribution over B and its direct neighbors. Moreover, applying the sum over I of eq 18 to eq 20 gives

$$\sum_{\substack{I \in Z_p \\ p \geq |I| > 1}} H(\mathbf{z}_{i,j,k}^{(0)}, \mathbf{t}) \quad (21)$$

which can be interpreted as a N -body problem with all charges set to 1 over two particle distributions: the set P_{target} of $(2L-1)^3$ different $\mathbf{z}_{ijk}^{(0)}$'s and the set P_{source} of \mathbf{t} 's. For large p , using a hierarchical method for such problem, one may compute these quantities very efficiently since the first of these particle distribution is already well-separated from each \mathbf{t} . In addition, both H , P_{target} and P_{source} are invariant under the action of the rotation group that preserves the cube. This can be exploited to speed-up the computation.¹³ See Figure 3 for a schematic representation.

A critical point here (that actually motivates this entire subsection) is that D_B does not depend on the particle distribution nor its modified charges, but only relies on the size of the simulation box B . Hence, the computation of D_B can be done only once inside a precomputation step, not at each IBFMM call. This has a particular interest since in daylife molecular dynamics applications, such IBFMM calls on the same B would be numerous, but needing only the single precomputation of D_B with our method.

In terms of complexity, at runtime, $D[B, I]$ is already computed, so the cost of handling the periodicity this way is the one of two FFTs on a small grid of size $(2L-1)^3$. Since L is a (small) constant, these FFTs cost $O(L^3 \log(L)) = O(1)$ flops.

III.D. Numerical Differentiation. The idea of switching derivatives to the interpolation polynomials has already found applications in the literature.³⁰ In our case, things are slightly more difficult because we consider differential operators instead of simple derivatives. To do so, we consider a Lagrange interpolation over a product of 1D Chebyshev grids in each leaf cell as well as all along the tree structure with a fixed order L in




Figure 3. 1D schematic representation of the handling of far periodic images through hierarchical methods. Octree \mathcal{T} in the simulation box B is represented in gray. Its two direct (non-well-separated) neighbor images (in red) directly interact with \mathcal{T} . Red and blue nodes in these trees respectively correspond to non-well-separated and well-separated nodes of the images regarding nodes of \mathcal{T} at the same tree level. On the root of \mathcal{T} (i.e., on B), interpolation nodes over an equispaced grid of size S are represented (in black). P_{target} is drawn in green. Blue nodes at the same tree level than B are elements of P_{source} . Upper nodes can be interpreted as aggregation of some of them during a hierarchical procedure to solve the N -body problem of eq 21.

each direction. Notice that this interpolation is different from the one of Sect. III.C that performs over a equispaced grid and only for the root. However, this does not impact the method at all since the root is only involved in interactions between nonadjacent images of B (see Figure 3) and because the definition of the M2M/L2L operators is valid no matter what the interpolation grid type is.

The choice of multivariate Lagrange interpolation over products of 1D Chebyshev nodes is motivated by the numerical instability that arises on large interpolation order and at fixed arithmetic precision when dealing with equispaced grids. Moreover, this choice of interpolation is known to be particularly good in the context of multilevel summation schemes.

As presented in Sect. III.B, in the ANKH approach, the differential operators apply on interpolation polynomials directly. Any Lagrange polynomial over a product grid is the product of the Lagrange polynomials in each one-dimensional grid. Assuming for the sake of clarity (one only has to introduce a scaling function in order to recover the general case^{6,30}) that the one-dimensional grids are defined on $[-1, 1]$, we thus have the following decomposition of such a polynomial S :

$$S_{iL^2+jL+k}(\mathbf{x}) = S_i(x_0)S_j(x_1)S_k(x_2) \quad (22)$$

where S_{iL^2+jL+k} can be either $S_k[t]$ or $S_i[s]$ following the notations of eq 6, S_i 's are one-dimensional Lagrange polynomials over Chebyshev (i.e., Lagrange-Chebyshev polynomials) nodes and $\mathbf{x} = (x_0, x_1, x_2) \in \mathbb{R}^3$. Explicit forms¹⁸ are known for these polynomials S_i :

$$S_k(x) = \frac{1}{L} + \frac{2}{L} \sum_{m=1}^{L-1} T_m(r_k) T_m(x) \quad (23)$$

where T_m is the m^{th} Chebyshev polynomial and $r_k := \cos\left(\frac{2k-1}{2L}\pi\right)$ is the k^{th} Chebyshev node. $T_m(x)$ can be computed using the recurrence relation:

$$\begin{aligned} T_{m+2}(x) &= 2xT_{m+1}(x) - T_m(x), \\ T_0(x) &= 1, T_1(x) = x \end{aligned} \quad (24)$$

The problem is that the explicit derivation of eq 23 leads to

$$S'_k(x) = \frac{2}{L\sqrt{1-x^2}} \sum_{m=1}^{L-1} mT_m(x_k) \sin(m \arccos(x)) \quad (25)$$

which is numerically unstable near the bounds of $[-1, 1]$. This can here be solved by using Chebyshev polynomials of the second kind in the derivation, but the same kind of problem appears when trying to derive S_k once again (which is needed to take into account quadrupoles). We thus opted for another way of deriving these polynomials: there exists³ an upper triangular matrix \mathcal{D} such that

$$\mathbf{T}'(x) = \mathcal{D}\mathbf{T}(x), \quad \mathbf{T}(x) := \begin{bmatrix} T_0(x) \\ \vdots \\ T_{L-1}(x) \end{bmatrix} \quad (26)$$

This can trivially be generalized to any derivation order, giving:

$$\mathbf{T}^{(N)}(x) = \mathcal{D}^N \mathbf{T}(x) \quad (27)$$

This last form does not suffer from numerical instability near the bounds of $[-1, 1]$. To exploit it in ANKH, we start by reformulating the set of Lagrange-Chebyshev polynomials of eq 23 in matrix form

$$\begin{aligned} \mathbf{S}(x) &= \mathbb{H}\mathbf{T}(x), \\ \mathbb{H} &:= \frac{1}{L} \begin{bmatrix} 1 & 2T_1(r_0) & \dots & 2T_{L-1}(r_0) \\ \vdots & \vdots & \ddots & \vdots \\ 1 & 2T_1(r_{L-1}) & \dots & 2T_{L-1}(r_{L-1}) \end{bmatrix} \end{aligned} \quad (28)$$

where \mathbb{H} does not depend on x and can be precomputed only once. We can now combine eq 27 and eq 28 in order to obtain our general formula:

$$\mathbf{S}^{(N)}(x) = \mathbb{H}\mathcal{D}^N \mathbf{T}(x) \quad (29)$$

To extend this technique to the three-dimensional case, a direct application of eq 22 provides expressions of the multivariate polynomials.

We summarize using C++-like pseudocode in Alg. 2 the steps needed to compute the modified charges for multipoles up to quadrupoles (trivial generalizations can be extrapolated from this pseudocode) using numerical differentiation through derivation of Chebyshev polynomials. Notice that the different \mathbf{t} 's computed in Alg. 2 could be stacked in practice inside a matrix with 9 columns before applying the product by \mathbb{H} , which allows us to benefit from matrix–matrix products instead of matrix–vector ones (i.e., BLAS3 instead of BLAS2), that should better perform.

Regarding the total complexity, the computation of Q inside a leaf cell s has a cost of $O(L^3N_s)$, where N_s denotes the number of particles in s . Since L is a constant, one may drop it from the big O and the total complexity becomes

$$\sum_{s \in \mathcal{L}(\mathcal{T})} O(N_s) = O\left(\sum_{s \in \mathcal{L}(\mathcal{T})} N_s\right) = O(N) \quad (30)$$

Algorithm 2 Compute \mathcal{Q}_s in leaf cell s

```

// Initialize modified charges with zero
 $\mathcal{Q}_s = 0$ ;
// Loop over particles in  $s$  (order does not matter)
for  $\mathbf{y} \in Y_s$  do
  // Notation:  $\mathbf{y} = (y_0, y_1, y_2)$ 
  // Allocate arrays storing Chebyshev polynomials
   $S_0^{(0)}, S_0^{(1)}, S_0^{(2)} = \text{new } \mathbb{R}[L]$ ;
   $S_1^{(0)}, S_1^{(1)}, S_1^{(2)} = \text{new } \mathbb{R}[L]$ ;
   $S_2^{(0)}, S_2^{(1)}, S_2^{(2)} = \text{new } \mathbb{R}[L]$ ;
   $\mathbf{T} = \text{new } \mathbb{R}[L]$ ;
  // Decompose the computation per direction
  for  $k \in \llbracket 0, 2 \rrbracket$  do
    // Initialize the first two polynomials, see Eq. (24)
     $\mathbf{T}[0] = 1$ ;
     $\mathbf{T}[1] = y_k$ ;
    // Compute recurrence following Eq. (24)
    for int  $i = 2; i < L; i++$  do
       $\mathbf{T}[i] = 2y_k\mathbf{T}[i-1] - \mathbf{T}[i-2]$ ;
    end for
    // Apply differentiation matrix following Eq. (26)
    for int  $j = 0; j < 3; j++$  do
       $\mathbf{T} = \mathbf{D} \cdot \mathbf{T}$ ;
       $S_k^{(p)} = \mathbb{H} \cdot \mathbf{T}$ ;
    end for
  end for
  // Add all possible combinations scaled with charge...
   $\mathcal{Q}_s = \mathcal{Q}_s + S_0^{(0)} S_1^{(0)} S_2^{(0)} q_{\mathbf{y}}$ ;
  // ... dipoles...
   $\mathcal{Q}_s = \mathcal{Q}_s + S_0^{(1)} S_1^{(0)} S_2^{(0)} (\mu_{\mathbf{y}})_0$ ;
   $\mathcal{Q}_s = \mathcal{Q}_s + S_0^{(0)} S_1^{(1)} S_2^{(0)} (\mu_{\mathbf{y}})_1$ ;
   $\mathcal{Q}_s = \mathcal{Q}_s + S_0^{(0)} S_1^{(0)} S_2^{(1)} (\mu_{\mathbf{y}})_2$ ;
  // ... and quadrupoles, exploiting symmetry
   $\mathcal{Q}_s = \mathcal{Q}_s + S_0^{(2)} S_1^{(0)} S_2^{(0)} (\Theta_{\mathbf{y}})_{0,0}$ ;
   $\mathcal{Q}_s = \mathcal{Q}_s + S_0^{(1)} S_1^{(1)} S_2^{(0)} (\Theta_{\mathbf{y}})_{0,1} \times 2$ ;
   $\mathcal{Q}_s = \mathcal{Q}_s + S_0^{(0)} S_1^{(0)} S_2^{(1)} (\Theta_{\mathbf{y}})_{0,2} \times 2$ ;
   $\mathcal{Q}_s = \mathcal{Q}_s + S_0^{(0)} S_1^{(2)} S_2^{(0)} (\Theta_{\mathbf{y}})_{1,1}$ ;
   $\mathcal{Q}_s = \mathcal{Q}_s + S_0^{(0)} S_1^{(1)} S_2^{(1)} (\Theta_{\mathbf{y}})_{1,2} \times 2$ ;
   $\mathcal{Q}_s = \mathcal{Q}_s + S_0^{(0)} S_1^{(0)} S_2^{(2)} (\Theta_{\mathbf{y}})_{2,2}$ ;
end for

```

III.E. Overall Complexity and Advantages. At the end of the day, ANKH-FMM defines a linear method since the overall complexity, thanks to all previous sections, can be counted this way:

$$\underbrace{O(N)}_{\text{step 1}} + \underbrace{O(N)}_{\text{step 2}} + \underbrace{O(N)}_{\text{step 3}} + \underbrace{O(N)}_{\text{step 4}} + \underbrace{O(N)}_{\text{step 5}} + \underbrace{O(1)}_{\text{step 6}} \quad (31)$$

where the steps are those presented in Sect. III.B. This has to be compared with the linearithmic complexity of PME. Hence, our method is asymptotically less costly than PME. However, the FMM (so as IBFMM) is known to suffer from an important prefactor, while the FFT has a (very) small one and may almost reach CPU peak performance. This means that the theoretical complexities are not sufficient to compare the two approaches, and numerical tests (that will be provided in the following) have to verify the efficiency of ANKH.

However, as we mentioned in Sect. III.A, there is another strong advantage of using ANKH-FMM: its distributed parallel potential. Indeed, computation of modified charges, as local operations, are easy to parallelize in distributed memory context. Moreover, IBFMM has shown impressive parallel performance in this context, and general parallel strategies aiming at reaching exascale are already developed for FMMs. Parallel scaling of FMMs can be put in comparison with ones of FFTs on which PME relies. In addition, efficient implementations of IBFMM on

GPU⁴² were proposed in the literature for perfect trees (i.e., exactly our application case). Hence, the ANKH strategy appears as a serious theoretical candidate for exascale hybrid computations.

We have already mentioned that GPU implementations of IBFMM are technical. This is why we wanted to devise a simple and portable scheme to run on GPU while keeping the parallel structure of the FMM. This is the purpose of Sect. IV.

III.F. Mutual Interactions. In the particular context of energy computation, the global algorithmic discussed in Sect. III.E still is suboptimal. Indeed, coming back to eq 5, we locally want to compute $q_{t'}^T \phi_{t',s'}$ instead of $\phi_{t'}$ only (that would rather be suited for forces computations). Hence, we would have

$$\begin{aligned}
 q_{t'}^T \phi_{t',s'} &\approx (q_{t'}^T U^*[t']) U^*[t', t] C[t, s] V[s, s'] (V[s'] q_{s'}) \\
 &= (q_{t'}^T V[t']^*) V[t, t']^* C[t, s] V[s, s'] \mathcal{M}_{s'} \\
 &= (V[t'] q_{t'})^T V[t, t']^* C[t, s] V[s, s'] \mathcal{M}_s \\
 &= \mathcal{M}_t^T V[t', t]^* C[t, s] V[s, s'] \mathcal{M}_{s'} \\
 &= \underbrace{(\mathcal{M}_t^T V[t, t']^T)}_{=(V[t', t'] \mathcal{M}_t)^T = \mathcal{M}_t^T} C[t, s] \underbrace{(V[s, s'] \mathcal{M}_{s'})}_{=\mathcal{M}_s} \\
 &= \mathcal{M}_t^T C[t, s] \mathcal{M}_s
 \end{aligned} \quad (32)$$

since $U^*[t'] = V[t']^* = V[t']^T$ because IBFMM L2P and P2M operators are adjoint matrices with real entries (so as IBFMM L2L and M2M operators). This provides a new formula for $\mathcal{F}_{\text{real}}$:

$$\begin{aligned}
 \mathcal{F}_{\text{real}} &\approx \sum_{\text{target cell } t} \mathcal{M}_t^T \sum_{s \in \Lambda(s)} C[t, s] \mathcal{M}_s \\
 &= \sum_{\text{target cell } t} \sum_{s \in \Lambda(s)} (\mathcal{M}_t^T C[t, s] \mathcal{M}_s)
 \end{aligned} \quad (33)$$

This does not reduce the theoretical complexity but strongly simplifies the method. In addition, many optimizations may be derived from this formula.

First, when performing M2L, we can directly compute the quantity into parentheses in eq 33, i.e., adding the left product by \mathcal{M}_t^T . Second, since if two cells $t, s \subseteq B$ are such that $s \in \Lambda(t)$, then $t \in \Lambda(s)$, meaning that the two interactions will be computed. This can be simplified because $C[t, s] = C[s, t]^T$:

$$\begin{aligned}
 &\mathcal{M}_t^T C[t, s] \mathcal{M}_s + \mathcal{M}_s^T C[s, t] \mathcal{M}_t \\
 &= \mathcal{M}_t^T C[s, t]^T \mathcal{M}_s + \mathcal{M}_s^T C[s, t] \mathcal{M}_t \\
 &= (\mathcal{M}_s^T C[s, t] \mathcal{M}_t)^T + \mathcal{M}_s^T C[s, t] \mathcal{M}_t \\
 &= 2\mathcal{M}_s^T C[s, t] \mathcal{M}_t
 \end{aligned} \quad (34)$$

This implies that only half of the M2L interactions need to be computed. In addition, since the M2L matrices are practically approximated by low-rank ones¹⁸ of the form $C[t, s] \approx \mathbf{L}[t, s] \mathbf{R}[t, s]$ with $\mathbf{L}[t, s], \mathbf{R}[t, s] \in \mathbb{R}^{k \times L^3}$, $k \ll L^3$, we can exploit this factorization to further accelerate the M2L evaluation in eq 34. This takes the following form:

$$\begin{aligned}
 \mathcal{M}_s^T C[s, t] \mathcal{M}_t &\approx \mathcal{M}_s^T \mathbf{L}[t, s]^* \mathbf{R}[t, s] \mathcal{M}_t \\
 &= \underbrace{(\mathbf{L}[t, s] \mathcal{M}_s)}_{O(kL^3)} \underbrace{(\mathbf{R}[t, s] \mathcal{M}_t)}_{O(kL^3)} \\
 &\quad \underbrace{\quad}_{O(kL^3)} \quad (35)
 \end{aligned}$$

Different techniques can be used to obtain such low-rank approximations. We compared two of them: partial Adaptive Cross Approximation² (pACA) and Singular Value Decomposition (SVD). pACA benefits from a linear algorithmic complexity while SVD is cubic cost. However, the sizes of the M2L matrices are relatively small and thanks to the use of symmetry,¹³ it is possible to only compute a few dozens of such approximations per level (and only once during the precomputation step, not at runtime). Hence, because numerical ranks obtained using pACA are higher than ranks obtained through SVD (that are optimal), we measured a slightly higher cost when evaluating the FMM using pACA approximations. For these reasons, we decided to rely on SVDs for performance tests.

Similarly, because our kernels K are translationally invariant, we have

$$\mathcal{D}_x \mathcal{D}_y K(\mathbf{x}, \mathbf{y}) + \mathcal{D}_y \mathcal{D}_x K(\mathbf{y}, \mathbf{x}) = 2 \mathcal{D}_x \mathcal{D}_y K(\mathbf{x}, \mathbf{y})$$

meaning that only half of the direct interactions have to be computed.

III.G. Global ANKH-FMM algorithm. For the sake of clarity, we summarize all the ANKH-FMM steps in Alg. 3. More precisely, ANKH-FMM first step consists in computing the modified charges in each leaf cell, once the (perfect) octree space decomposition is known. These modified charges are interpreted as terms of a IBFMM multipole expansion associated with interpolation nodes in these cells. This thus leads to a computation of other multipole expansions in the non-leaf-tree levels. Notice that this step seems quite easy to parallelize in a shared-memory context due to the independence of the modified charges of two different leaves. The second ANKH-FMM step corresponds to the M2L evaluation involved in IBFMM, except that in our case mutual interactions give a simpler expression of it (following Sect. III.F). Since we are focused on energy computation and thanks to the adjoint form of the L2L/M2M and L2P/P2M operators, results of each mutual M2L can be added to a scalar variable. Thus, no L2L/L2P application is needed in practice.

An important point to mention at this presentation stage is that, even if the IBFMM part of ANKH-FMM also involves cells of the direct neighbor images of the simulation box B , there is no need of computing its multipole expansions. Indeed, there are equal to the multipole expansion in the original box up to a translation, which only appears in the M2L formula since all other formula are in practice taken respectively to cell centers.¹²

Then, the third step of ANKH-FMM is just a direct computation of the near field $\mathcal{N}_{\text{real}}$. To achieve it in practice, we used explicit derivations⁴⁰ of the kernel K . Fourth step corresponds to the handling of periodicity (i.e., here the influence on B of nondirect neighbor images of B) using the technique presented in Sect. III.C. Finally, the two last steps are just direct computation of the self-energy in eq 10 and simple sum of all the computed quantities.

III.H. A First Conclusion. We presented in this section a new method, based on both the Ewald summation and the interpolation-based FMM, aiming at efficiently treating N -body problems arising in the energy computation context in

Algorithm 3 ANKH-FMM

```

// Step 1: get modified charges, multipole expansions
for each leaf cell  $s \in \mathcal{L}(T)$  do
  Compute  $\mathcal{Q}_s$  using Alg. 2
   $\mathcal{M}_s = \mathcal{Q}_s$ 
end for
for each non leaf level  $e$  starting from the deepest one do
  for each cell  $s$  at level  $e$  do
     $\mathcal{M}_s = \mathbf{0}$ 
    for each daughter  $s'$  of  $s$  do
       $\mathcal{M}_s = \mathcal{M}_s + V[s, s'] \mathcal{M}_{s'}$ 
    end for
  end for
end for
// Step 2: get  $\mathcal{F}_{\text{real}}$  using mutual M2L
 $\mathcal{F}_{\text{real}} = \mathbf{0}$ 
for each target cell  $t$  do
  for each source cell  $s \in \Lambda(t)$  do
    if interaction of  $s$  and  $t$  has not been yet computed then
       $\mathcal{F}_{\text{real}} = \mathcal{F}_{\text{real}} + 2 \mathcal{M}_t^T C[t, s] \mathcal{M}_s$ 
    end if
  end for
end for
// Step 3: get  $\mathcal{N}_{\text{real}}$  through mutual P2P
 $\mathcal{N}_{\text{real}} = \mathbf{0}$ 
for each leaf cell  $t \in \mathcal{L}(\mathcal{T})$  do
  for each source cell  $s$  adjacent to  $t$  do
    if interaction of  $s$  and  $t$  has not been yet computed then
      for  $\mathbf{x} \in X|_t$  do
        for  $\mathbf{y} \in Y|_s$  do
           $\mathcal{N}_{\text{real}} = \mathcal{N}_{\text{real}} + 2 \mathcal{D}_x \mathcal{D}_y G(\mathbf{x}, \mathbf{y})$ 
        end for
      end for
    end if
  end for
end for
// Step 4: get periodic influence
Let  $\mathcal{M}_B$  be the root multipole expansion
 $\mathcal{F}_{\text{per}} = \mathcal{M}_B^T \chi^* \mathbb{F}^* D_B \mathbb{F} \chi \mathcal{M}_B$ 
// Step 5: compute self energy
 $\mathcal{E}_{\text{self}} = \mathbf{0}$ 
for each  $\mathbf{x} \in X$  do
   $\mathcal{E}_{\text{self}} = \mathcal{E}_{\text{self}} + q_{\mathbf{x}}^2 \frac{2\epsilon^2}{3} \mu_{\mathbf{x}} \cdot \mu_{\mathbf{x}} + \frac{8\epsilon^4}{5} \Theta_{\mathbf{x}} : \Theta_{\mathbf{x}}$ 
end for
// Step 6: return energy  $\mathcal{E}$ 
 $\mathcal{E} = -\frac{\epsilon}{\sqrt{\pi}} \mathcal{E}_{\text{self}} + \mathcal{F}_{\text{real}} + \mathcal{N}_{\text{real}} + \mathcal{F}_{\text{per}}$ 

```

molecular dynamics and in a way that should efficiently perform in parallel (mainly distributed memory) context. Hence, since FMMs benefit from highly scalable strategies,^{22,28,45} this new ANKH-FMM approach may pave the way for a linear complexity and scalable family of methods suited for molecular large scale dynamics computations. We proved, under molecular dynamics hypotheses, that our method is actually $O(N)$, even when exploiting periodicity with large amount of images, thanks to a new way of taking into account the influence of the far images. We introduced the mutual interactions in the context of M2L evaluations for this application case and we presented a numerical differentiation through polynomial interpolation in order to deal with differential operators arising when considering dipoles and quadrupoles (our methodology can be extended to any multipole order with low effort). Finally, we provided explicit algorithms detailing our method.

Among the direct perspective, the ideas behind ANKH-FMM can be used to compute forces, even if some optimizations have to be changed to this purpose (especially in the mutual M2L interactions). The parallel implementation of ANKH-FMM as well as real tests on supercomputers also counts among the important and relatively short-term perspectives.

III.I. Important Remarks. Before presenting our second approach, we dwell on a few important details that could have aroused the curiosity of an informed reader.

1. For the sake of clarity, we considered only cubical box B . However, the method described in Sect. III as well as the one presented in the following Sect. IV directly extend to cuboids.
2. The computation of self-energy $\mathcal{E}_{\text{self}}$ does not involve particular treatment in our approach and its magnitude is small compared to the real energy one. However, because the decay of the term $\left| \frac{\xi}{\sqrt{\pi}} \right|$ in $\mathcal{E}_{\text{self}}$ with respect to decreasing ξ is far slower than the $e^{-2(\pi r_B/\xi)^2 \mathbf{m} \cdot \mathbf{m}}$ one (involved in \mathcal{E}_{rec}), we decided to keep $\mathcal{E}_{\text{self}}$ anyway and to neglect only \mathcal{E}_{rec} for our choices of very small ξ . Numerical tests suggest that for targeted precisions, this choice is justified since, for instance, $\mathcal{E}_{\text{self}}/\mathcal{E}_{\text{real}} \approx 3.8 \times 10^{-5}$ in the case of 96000 atoms (see Sect. V.D for detailed test cases) and using $\xi = 0.01$. We are aware that a careful analysis on this point should be proposed to provide precise mathematical error bounds.
3. One may wonder why we did not consider even smaller ξ than in the second remark. The reason is that a small ξ implies a large amount of far images of B to be considered before reaching convergence. Hence, even if this cost can be mitigated thanks to precomputations in Sect. III.C, we decided to keep ξ in a value range that is numerically both sufficient to neglect \mathcal{E}_{rec} and ensure a convergence using (very) large but still fixed number of far images of B even for high required precisions. Once again, we believe that there is a point here to be covered by a mathematical theoretical study in order to provide explicit bounds.

IV. ANKH-FFT

Despite the sequential and distributed efficiency of the interpolation-based FMM, limitations arise in the context of GPU computing. Indeed, if the near-field part is well suited to GPU, the M2L evaluation suffers from data loading in shared memory that mitigates the performance on this type of architecture. However, the perfectness of the 2^d -trees in molecular dynamics allows us to derive efficient schemes, as those proposed in the literature.⁴² In this section, we introduce a new method, also based on interpolation, and designed to bypass the limitations arising in GPU context.

Starting from the interpolation of the particle multipoles at leaves on equispaced grids (instead of Chebyshev ones), one may construct a single level FMM as the evaluation of the matrix concatenating all the M2L matrices between nonadjacent or equal cells. This matrix can be proven to be a block-Toeplitz matrix (we will focus on the theory in a forthcoming mathematical article). Such a matrix can be embedded in a larger circulant matrix. Because any circulant matrix can be explicitly diagonalized in the Fourier domain, and because the product by a discrete Fourier basis can be efficiently processed through FFTs, one may compute this diagonalization in a linearithmic time. Doing so, we obtain a fast linearithmic scheme to deal with the far-field of N-body problems on quasi-uniform particle distributions (this assumption allowing to obtain perfect trees). Of course, the near-field is treated as in the FMM, i.e. by performing local direct computation between nonseparated cells.

However, in our case, since we want to perform interpolation on Chebyshev polynomials in order to mitigate the precision loss of the multipole operator evaluation without much numerical stability issue, we start by performing such a Chebyshev interpolation at leaves, directly followed by a reinterpolation of the data now defined at Chebyshev nodes to a local equispaced grid, using the same grid (up to translation) for every leaf (since they all have the same radius).

Once in the Fourier domain, the energy computation problem is nothing more than a scalar product between two vectors: the diagonal of the Fourier diagonalization of the circulant matrix and the modified charges vector Fourier transform squared modulus after zero-padding (to have the same size than the circulant matrix). Since the circulant matrix has $O(N)$ rows and columns ($O(N)$ leaves with $O(1)$ interpolation nodes in each), this scalar product costs $O(N)$ floating point operations (flops), while the FFT costs $O(N \log N)$ flops. Since the interpolation phase is analogue to the application of all P2M operators in ANKH-FMM, this step can also be performed in $O(N)$ flops. Hence, we end up with a linearithmic method. This new method is named ANKH-FFT.

ANKH-FFT is theoretically more costly than ANKH-FMM, but ANKH-FFT benefits from highly optimized FFT implementations, also available on GPU, and the interpolation phase is a highly parallel one. Hence, ANKH-FFT seems to be a better candidate than ANKH-FMM in a GPU context. We should still compare the performance of these two methods, which is done in the following. These performances are comparable, and ANKH-FFT behaves as a linear method in the tested range (i.e., from 1000 to 100 million atoms). This leads us to the conclusion about the use of this new approach: ANKH-FFT appears as an entirely new and efficient approach to deal with local N-body problems with an algorithmic structure well-suited for GPU programming. This should thus be an excellent alternative to local interpolation-based FMM used in one core of the distributed memory framework of the parallel interpolation-based FMM.

We may here insist on the fact that ANKH-FFT requires a strong mathematical background to be entirely presented, hence we postpone this presentation to a forthcoming mathematical paper, including proofs. ANKH-FFT is based on a matrix factorization presented in Sect. IV.A. We then show how this factorization leads to a diagonalization in Sect. IV.B and we present the fast periodic condition handling in Sect. IV.C. The important role of mutual interactions in ANKH-FFT is presented in Sect. IV.D. Finally, we summarize the main algorithm in Sect. IV.E, followed in Sect. IV.F by a complexity analysis. The numerical comparison between ANKH-FFT and ANKH-FMM is provided in Sect. IV.G.

IV.A. Matrix Factorization. We start by considering all the leaf cells $\mathcal{L}(\mathcal{T})$ of \mathcal{T} (that are all at the same tree level since \mathcal{T} is perfect). Let us consider the *extended interaction list* Λ^+ of any leaf cell t defined as the set of leaf cells with a strictly positive distance with t :

$$\Lambda^+(t) := \{s \in \mathcal{L}(\mathcal{T}) | \text{dist}(t, s) > 0\}$$

We then decompose $\tilde{\mathcal{E}}_{\text{real}}$ of eq 12 as

$$\tilde{\mathcal{E}}_{\text{real}} = \underbrace{\sum_{t \in \mathcal{I}(\mathcal{T})} \sum_{\substack{\mathbf{x} \in t, \mathbf{y} \in s \\ s \in \Lambda^+(t)}} \mathcal{D}_{\mathbf{x}} \mathcal{D}_{\mathbf{y}} H(\mathbf{x}, \mathbf{y})}_{=: \mathcal{F}} + \underbrace{\sum_{t \in \mathcal{I}(\mathcal{T})} \sum_{\substack{\mathbf{x} \in t, \mathbf{y} \in s \\ s \notin \Lambda^+(t)}} \mathcal{D}_{\mathbf{x}} \mathcal{D}_{\mathbf{y}} H(\mathbf{x}, \mathbf{y})}_{=: \mathcal{N}} \quad (36)$$

As in our ANKH-FMM algorithm, we propose to compute \mathcal{N} directly and to approximate \mathcal{F} by means of interpolation. Indeed, following Sect. III.F, we can write

$$\mathcal{F} \approx \sum_{t \in \mathcal{I}(\mathcal{T})} \sum_{s \in \Lambda^+(t)} \mathcal{M}_t^T C[t, s] \mathcal{M}_s \quad (37)$$

Now, let \mathbf{E}_c for any $c \in \mathcal{L}(\mathcal{T})$ be a reinterpolation matrix (see the M2M definition in eq 8) from the (Chebyshev) interpolation nodes of c into a equispaced grid \mathcal{J}_L defined by

$$\mathcal{J}_L := \left\{ r_B 2^{-(1+E)} \left(\begin{bmatrix} -1 \\ -1 \\ -1 \end{bmatrix} + \frac{2\mathbf{k}}{L-1} \right) \mid \mathbf{k} \in [[0, L-1]]^3 \right\} \quad (38)$$

i.e., \mathbf{E}_c interpolates the modified charges associated with Chebyshev nodes in a cell c to new modified charges associated with the nodes of a equispaced grid in c . This interpolation grid \mathcal{J}_L has to be centered in each leaf cell center. Said differently, the same equispaced interpolation grid \mathcal{J}_L is used in each leaf cell, up to a translation (see Figure 4).

Each M2L matrix $C[t, s]$ can be written as

$$C[t, s] \approx \mathbf{E}_t^T G[t, s] \mathbf{E}_s \quad (39)$$




Figure 4. Schematic representation of 2D interpolation on tensorized equispaced grids (as in eq 38) with $L = 5$ on four different leaf cells (in red, orange, green and gray). Interpolation nodes belonging on a unique cell are represented in black on the top figure, nodes belonging in two cells are represented in purple, and the node belonging in the four cells is represented in cyan. In the bottom, we represented the actual grid used in ANKH-FFT, with a slight node translation in order to exhibit duplications: purple nodes are duplicated once and the cyan one is duplicated four times. Assuming that we only have these four leaf cells (i.e., a tree depth of only one), this cyan node is indexed (following the multi-indexing of Sect. IV.B) as $(0, 4, 0, 0)$ in the red cell, as $(1, 0, 0, 0)$ in the gray cell, as $(0, 4, 0, 4)$ in the orange cell, and as $(1, 0, 1, 4)$ in the green cell.

where $(G[t, s])_{k,l} := H(ctr(t) + \tilde{\mathbf{x}}_{k,ctr}(s) + \tilde{\mathbf{x}}_l)$, $ctr(c)$ denoting the center of the cell c and $\tilde{\mathbf{x}}_k$ referring to the k^{th} elements of \mathcal{J}_L . In other terms, eq 39 simply switches between Chebyshev nodes in a cell and equispaced nodes.

Remark 1 In practice, the order of the Chebyshev rules are chosen so that the derivation of the associated polynomials results in the targeted error. Nevertheless, after modified charges computation, the equispaced interpolation does not have to consider the same interpolation order. This last order can be chosen without consideration about the derivatives, which practically implies lower orders. Hence, another compression level may be induced by this reinterpolation, provided that the interpolation order on equispaced grids is chosen to be smaller than the one on Chebyshev grids.

Now, we aim at efficiently computing the interactions between a cell t and all the cells in its extended interaction list. Let us extend the definition of $G[t, s]$ in order to cover the blanks induced by pairs of leaf cells (t, s) such that $s \notin \Lambda^+(t)$, introducing

$$\tilde{G}[t, s] = \begin{cases} G[t, s] & \text{if } s \in \Lambda^+(t), \\ 0 & \text{otherwise} \end{cases} \quad (40)$$

We thus obtain

$$\mathcal{F} \approx \sum_{t \in \mathcal{L}(\mathcal{T})} \sum_{s \in \mathcal{L}(\mathcal{T})} \mathcal{M}_t^T \mathbf{E}_t^T \tilde{G}[t, s] \mathbf{E}_s \mathcal{M}_s \quad (41)$$

that corresponds to the product of a matrix \mathbb{A} formed by the concatenation of the blocks $\tilde{G}[t, s]$ with a left and a right vector formed by the concatenation of $\mathbf{E}_c \mathcal{M}_c$'s, $c \in \mathcal{L}(\mathcal{T})$. The blocks are ordered in 3D (in both rows and column) following a lexicographic order on their center difference:

$$c < c' \Leftrightarrow [r_0 < 0 \vee (r_0 = 0 \wedge r_1 < 0) \vee (r_0 = 0 \wedge r_1 = 0 \wedge r_2 < 0)]$$

for any $c, c' \in \mathcal{L}(\mathcal{T})$ and with $(r_0, r_1, r_2) := ctr(c) - ctr(c')$. The result of this concatenation is a vector $\mathbf{v} \in \mathbb{R}^{8^E L^3}$, where 8^E corresponds to the number of leaf cells in a perfect octree with depth E . Using these notations, we obtain the following approximation:

$$\mathcal{F} \approx \mathbf{v}^T \mathbb{A} \mathbf{v} \quad (42)$$

We are thus now interested by an efficient computation of the product by \mathbb{A} . This is the purpose of Sect. IV.B.

Remark 2 In \mathbb{A} , a given target cell t may interact with various source cells s that spatially share interpolation nodes. However, each equispaced interpolation grid and each multipole expansion \mathcal{M}_s associated with it has to be considered independently because, depending on t , \mathcal{M}_s may be masked or used, thanks to possible products by zeros in eq 40. This means that we do not see the concatenation of equispaced grids as a global equispaced grid, but rather as a collection of such small grids with nodes duplication on the cell's edges. This consideration is illustrated in Figure 4

IV.B. Diagonalization. In this section, we provide an explicit diagonalization of an embedding of \mathbb{A} in a Fourier basis. Let \mathbf{k}, \mathbf{l} be multi-indices in

$$\begin{pmatrix} \left[\frac{[[0, n(E)]]}{\text{cell on } x\text{-axis}} \times \frac{[[0, L-1]]}{\text{interpolation node on } x\text{-axis}} \right] \\ \times \left(\frac{[[0, n(E)]]}{\text{cell on } y\text{-axis}} \times \frac{[[0, L-1]]}{\text{interpolation node on } y\text{-axis}} \right) \\ \times \left(\frac{[[0, n(E)]]}{\text{cell on } z\text{-axis}} \times \frac{[[0, L-1]]}{\text{interpolation node on } z\text{-axis}} \right) \end{pmatrix}$$

where $n(E) := 2^E - 1$ (see Figure 4 for graphical 2D representation). Let $\mathbf{u} = (\tilde{u}_0, \tilde{u}_1, \tilde{u}_2) \in \{0, 1\}^6$ be a boolean vector and let $\tilde{\mathbf{k}} = (\tilde{k}_0, \tilde{k}_1, \tilde{k}_2)$, $\tilde{\mathbf{l}} = (\tilde{l}_0, \tilde{l}_1, \tilde{l}_2)$ be defined as

$$\tilde{k}_p = \begin{cases} k_p + u_p & \text{if } p \text{ is even and } k_p \neq 2^E - 1, \\ k_p + u_p & \text{if } p \text{ is odd and } k_p \neq L - 1, \\ k_p & \text{otherwise} \end{cases} \quad (43)$$

In our study case, \mathbb{A} verifies a strong property which is given in Theorem 1.

Theorem 1 If H is a radial kernel, then $\mathbb{A}_{\mathbf{k}, \mathbf{l}} = \mathbb{A}_{\tilde{\mathbf{k}}, \tilde{\mathbf{l}}}$, where $\mathbb{A}_{\mathbf{k}, \mathbf{l}}$ refers to the element of \mathbb{A} at row $k_0(2^E L^3)^2 + k_1 2^E L^3 + k_2$ and column $l_0(2^E L^3)^2 + l_1 2^E L^3 + l_2$.

Actually, Theorem 1 exhibits a Toeplitz structure for \mathbb{A} in six dimensions (that are the dimensions of the indexing at the beginning of this section). The direct consequence of Theorem 1 is summarized in Corollary 1.

Corollary 1 Let $\chi \in \{0, 1\}^{((2^{E+1}-1) \times (2L-1))^3 \times ((2^E \times L)^3)}$ be defined as

$$\chi_{\mathbf{k}, \mathbf{l}} = \begin{cases} 1 & \text{if } \mathbf{k} = \mathbf{l}, \\ 0 & \text{otherwise} \end{cases} \quad (44)$$

for any $\mathbf{k} \in ([0, 2^E - 1] \times [0, L - 1])^3$, $\mathbf{l} \in ([0, 2^{E+1} - 2] \times [0, 2L - 2])^3$ and let \mathbb{F}_6 be the discrete Fourier transform matrix of dimensions $(2^E, L, 2^E, L, 2^E, L)$. We have

$$\mathbb{A} = \chi^* \mathbb{F}_6^* \mathbf{D}(\mathbb{A}) \mathbb{F}_6 \chi \quad (45)$$

where $\mathbf{D}(\mathbb{A})$ is a diagonal matrix.

The main idea is that \mathbb{A} is a block-Toeplitz matrix according to Theorem 1. Such a matrix can be embedded into a block-circulant matrix one. Any block-circulant matrix can be diagonalized in a Fourier basis (here \mathbb{F}_6), giving the factorization of Corollary 1. Here, this block-circulant matrix is never built because the diagonal of $\mathbf{D}(\mathbb{A})$ can be computed through a single application of \mathbb{F}_6 to a vector of correctly ordered elements of the first row and column of \mathbb{A} . Similar tools were used in Sect. III.C in a much simpler context.

The expression 45 allows to derive a fast scheme for the product by \mathcal{F} in eq 42 since thanks to Theorem 1,

$$\mathcal{F} \approx \mathbf{v}^T \chi^* \mathbb{F}_6^* \mathbf{D}(\mathbb{A}) \mathbb{F}_6 \chi \mathbf{v} \quad (46)$$

in which products by χ^* , χ can be evaluated with linear complexity (because of their sparse structure) and products by \mathbb{F}_6^* , \mathbb{F}_6 can be processed with linearithmic complexity through FFTs (regarding the size of \mathbf{v}).

We refer to Alg. 5 in Appendix A for the explicit computation of the diagonal of $\mathbf{D}(\mathbb{A})$.

IV.C. Efficient Handling of the Periodicity. When large numbers of periodic images are considered, the computation of $\sum_{\mathbf{t} \in 2r_B Z_p} H(\mathbf{x}, \mathbf{y} + \mathbf{t})$ involved in Alg. 5 becomes too costly to be used. Fortunately, this step can be accelerated through another level of interpolation. Exploiting translational invariance of H , one gets

$$\sum_{\mathbf{t} \in 2r_B Z_p} H(\mathbf{x}, \mathbf{y} + \mathbf{t}) = \sum_{\mathbf{t} \in 2r_B Z_p} \underbrace{H(\mathbf{x} - \mathbf{y}, \mathbf{t})}_{=: \mathbf{z}} \quad (47)$$

that can be separated into two sums

$$\left(\sum_{\mathbf{t} \in 2r_B Z_2} H(\mathbf{z}, \mathbf{t}) \right)_{=: \mathcal{N}_{\text{per}}(\mathbf{z})} + \left(\sum_{\mathbf{t} \in 2r_B Z_p \setminus Z_2} H(\mathbf{z}, \mathbf{t}) \right)_{\mathcal{F}_{\text{per}}(\mathbf{z})} \quad (48)$$

Since both $\mathbf{x}, \mathbf{y} \in B$ lie inside the (centered at zero) simulation box of radius r_B , \mathbf{z} belongs inside a box \tilde{B} of radius $2r_B$. Each $\mathbf{t} \in 2r_B Z_p \setminus Z_2$ has distance at least $2r_B$ from \tilde{B} , hence, each such \mathbf{t} is well-separated from \tilde{B} , using a treecode-like criterion.¹⁵ Hence, performing an interpolation over a equispaced grid on \tilde{B} , we obtain

$$\mathcal{F}_{\text{per}}(\mathbf{z}) \approx \sum_k U_k[\tilde{B}](\mathbf{z}) \mathcal{F}_{\text{per}}(\mathbf{z}_k) \quad (49)$$

U_k being the Lagrange polynomial associated with the k^{th} node of the equispaced interpolation grid over \tilde{B} .

Combining eq 48 and eq 49, we get

$$\mathcal{F}_{\text{per}}(\mathbf{z}) \approx \sum_k U_k[\tilde{B}](\mathbf{z}) \left(\sum_{\mathbf{t} \in 2r_B Z_p \setminus Z_2} H(\mathbf{z}_k, \mathbf{t}) \right)_{=: h_k} \quad (50)$$

in which the sum into parentheses corresponds to an N -body problem on two point clouds: the Cartesian grid over \tilde{B} (with L^3 nodes) and the set $2r_B Z_p \setminus Z_2$ (with all charges equal to 1). This N -body problem can be solved efficiently using a treecode or a IBFMM. Then, once all h_k 's are known, any $\mathcal{F}_{\text{per}}(\mathbf{z})$ can be retrieved in $O(L^3) = O(1)$. In addition, these h_k 's can be computed once during the precomputation step since they do not depend on the modified charges but only on the kernel H .

The term $\mathcal{N}_{\text{per}}(\mathbf{z})$ involves only a sum over $S^3 = O(1)$ elements, meaning that using this interpolation, one may deduce the result of eq 48 in $O(1)$ flops.

IV.D. Mutual Interactions. The evaluation of eq 46 involves two FFTs if computed naively. Since they are the most consuming part of the evaluation time, we may wonder how to mitigate it. This can be done by exploiting once again, as in Sect. III.F, the mutual interactions. Our idea is based on the following result:

Corollary 2 If H has real values, then $\mathbf{D}(\mathbb{A}) \in \mathbb{R}^{[(2^{E+1}-1)(2L-1)]^3 \times [(2^{E+1}-1)(2L-1)]^3}$.

This can be proven from Theorem 1 noticing that radial property of H implies symmetry of \mathbb{A} . Hence, the spectrum of \mathbb{A} is real.

As a direct consequence, following eq 46, we have

$$\mathbf{v}^T \chi^* \mathbb{F}_6^* \mathbf{D}(\mathbb{A}) \mathbb{F}_6 \chi \mathbf{v} = (\mathbb{F}_6 \chi \mathbf{v})^* \mathbf{D}(\mathbb{A}) \underbrace{\mathbb{F}_6 \chi \mathbf{v}}_{=: \mathbf{w}} \quad (51)$$

with $\mathbf{w} = (w_0, \dots, w_{[(2^{E+1}-1)(2L-1)]^3-1})^T$, and which reformulates as

$$\begin{aligned} \mathbf{w}^* \mathbf{D}(\mathbb{A}) \mathbf{w} &= \sum_{k=0}^{[(2^{E+1}-1)(2L-1)]^3-1} \overline{w_k} \mathbf{D}(\mathbb{A})_{k,k} w_k \\ &= \sum_{k=0} \mathbf{D}(\mathbb{A})_{k,k} |w_k|^2 \end{aligned} \quad (52)$$

Hence, only \mathbf{w} has to be computed, meaning that only a single FFT needs to be performed during evaluation of eq 46. As another important consequence of Corollary 2, once this FFT has been performed, one can only keep the real part of the output (the imaginary one being equal to 0). In our implementation, we even rely on real-to-complex FFTs from FFTW3¹⁹ since the modified charges (as well as the entries of \mathbf{w}) are real provided that the charges, dipoles, and quadrupoles also are.

IV.E. Global Algorithmic. It is now possible to provide the overall algorithm for the method derived from our ANKH methodology and exploiting the diagonalization of the interpolated Ewald summation over local equispaced grids. Namely, this corresponds to ANKH-FFT, and Alg. 4 summarizes all its steps.

Algorithm 4 ANKH-FFT

```
// Step 1: get modified cartesian multipole expansions
for each leaf cell  $s \in \mathcal{L}(T)$  do
  Compute  $Q_s$  using Alg. 2
   $\mathcal{M}_s = \mathbf{E}_s Q_s$ 
end for
// Step 2: get  $\mathcal{F}_{real}$  using FFT
Let  $\mathbf{v}$  be defined as in Eq. (46)
 $\mathbf{w} = \mathbb{F}_6 \chi \mathbf{v}$ 
 $\mathcal{F}_{real} = \mathbf{w}^T \mathbf{D}(\mathbb{A}) \mathbf{w}$  using Eq. (52)
// Step 3: get  $\mathcal{N}_{real}$  through mutual P2P
 $\mathcal{N}_{real} = 0$ 
for each leaf cell  $t \in \mathcal{L}(\mathcal{T})$  do
  for each source cell  $s$  adjacent to  $t$  do
    if interaction of  $s$  and  $t$  has not been yet computed then
      for  $\mathbf{x} \in X_t$  do
        for  $\mathbf{y} \in Y_s$  do
           $\mathcal{N}_{real} = \mathcal{N}_{real} + 2\mathfrak{D}_{\mathbf{x}} \mathfrak{D}_{\mathbf{y}} G(\mathbf{x}, \mathbf{y})$ 
        end for
      end for
    end if
  end for
end for
// Step 4: compute self energy
 $\mathcal{E}_{self} = 0$ 
for each  $\mathbf{x} \in X$  do
   $\mathcal{E}_{self} = \mathcal{E}_{self} + q_{\mathbf{x}}^2 \frac{2\epsilon^2}{3} \mu_{\mathbf{x}} \cdot \mu_{\mathbf{x}} + \frac{8\epsilon^4}{5} \Theta_{\mathbf{x}} : \Theta_{\mathbf{x}}$ 
end for
// Step 5: return energy  $\mathcal{E}$ 
 $\mathcal{E} = -\frac{\epsilon}{\sqrt{\pi}} \mathcal{E}_{self} + \mathcal{F}_{real} + \mathcal{N}_{real}$ 
```

ANKH-FMM and ANKH-FFT share strong similarities: both use the same computation of modified charges, both compute the near field part (\mathcal{N}_{real}) using direct interactions, and both compute the self-energy in the same way. However, the computation of the far field part (\mathcal{F}_{real}) is strongly simplified in ANKH-FFT.

Among the differences, we decided to take into account the influence of the periodicity in the diagonalization process of ANKH-FFT while we opted for an external handling in ANKH-FMM, meaning that the evaluation of ANKH-FFT does not involve any periodic influence (already contained in $\mathbf{D}(\mathbb{A})$). The reason behind this choice is that in ANKH-FFT, the far field computation is done in the Fourier domain directly, in which the particle cannot be interpreted as easily as in the Cartesian domain, while in ANKH-FMM, the multilevel structure

provides a way of measuring the influence of the far images efficiently *on-the-fly*. This points out the main difference between the two approaches: ANKH-FFT is not multilevel but is built on a global operation (FFT).

IV.F. Complexity. There are two complexities to count in ANKH-FFT: the precomputation complexity and the evaluation one. The precomputation step involves the solving of the N -body problem of eq 50, which may be done in a linear time with respect to the number of far images, i.e., in $O(p^3)$. In practice, $p = O(1)$ is small. Then, the diagonal of \mathbb{A} is computed using Alg. 5. It is easy to see from this algorithm that, thanks to Sect. IV.C, its complexity is linearithmic with respect to the size of \mathbf{v} .

When considering quasi-uniform particle distributions, the total number of leaf cells can be considered as $O(N)$. In this case, the size of \mathbf{v} is $O(NL^3)$ with a constant L corresponding to the number of interpolation nodes in one direction of the equispaced grid, which is a small constant. Hence, the size of \mathbf{v} is $O(N)$ and the FFTs perform in $O(N \log N)$.

The precomputation cost of the modified charges is linear with respect to the number of particles in the simulation box, therefore, the overall precomputation complexity is

$$O(N) + O(N) + O(N \log N) = O(N \log N) \quad (53)$$

Since the product by a diagonal matrix is also $O(N)$, the overall complexity of the evaluation of \mathcal{F} complexity though eq 46 is $O(N \log N)$ (because of the FFT).

IV.G. Comparison between ANKH-FMM and ANKH-FFT and Partial Conclusions. This section is dedicated to the comparison between ANKH-FMM and ANKH-FFT. As we already discussed, the theoretical complexities (respectively linear and linearithmic) are not sufficient to conclude in the practical efficiency. One of the main reasons behind that is the arithmetic intensity of the *state-of-the-art* FFT implementations¹⁹ that should be higher than our IBFMM one. Hence, we compared running times for the evaluation of both methods (i.e., excluding precomputations) on the same particle distributions, using charges only (since the near field part is the same between the two methods) and no periodic condition (since its handling has a very small impact on the timings and we want to compare the main core of these methods). Timings are provided in Figure 5. To do these tests, we run our (sequential) codes on a single Intel(R) Xeon(R) Gold 6152 CPU.

Performance of the two methods are comparable in the (large) tested range of systems, i.e., from 1000 atoms to 10 M atoms: ANKH-FFT performs slightly better on the case with




Figure 5. Timing comparison of the two ANKH variants with respect to the number of particles (atoms) in the simulation box.

10000 atoms, but not in a significant way. Since ANKH-FMM follows the theoretical estimates of linear complexity, ANKH-FFT also behaves as a linear method in practice.

We may still insist on the fact that the precomputations involved in ANKH-FFT are more costly than the ones of ANKH-FMM (but in the same order of magnitude than the evaluation step). These precomputation timings are however quite relative to the application: in the MD context, the same simulation box is used for a massive amount of evaluations, hence these precomputations are performed only once for many evaluations, which mitigates their costs.

To summarize, we list the main difference, advantages, and drawbacks of the two ANKH approaches:

- ANKH-FMM is a truly linear method while ANKH-FFT only is a linearithmic one. Both achieve fast computations with comparable running times (on CPU) and can take into account periodic boundary conditions as well as distributed point multipoles in a convergent way.
- ANKH-FFT is easier to implement than ANKH-FMM, especially in projection for the GPU context while ANKH-FMM, as a IBFMM-based approach, directly benefits from the literature on HPC for this kind of method.
- Since ANKH-FFT relies mainly on FFTs, it may be easier to provide portable codes for it (most modern architectures provide efficient FFT implementations).
- The drawback, however, is the precomputation timings of ANKH-FFT are more costly than the ANKH-FMM ones.

One of the main motivation behind providing this ANKH approach is to design methods able to efficiently scale on modern HPC architectures and (very) large systems of particles. Hence, it is important to realize that *state-of-the-art* techniques (mainly relying on Local Essential Trees^{1,16,28,34,45}) for parallel computations using FMMs can be used in MD exploiting the ANKH framework. More specifically, Local Essential Trees may also be combined with the ANKH-FFT approach, since local subproblems can be treated with the latter (only the periodic boundary condition handling has to be slightly adapted). We believe that this last idea, exploiting GPU implementation of ANKH-FFT, can lead to highly scalable MD simulations on modern and hybrid HPC architectures.

V. NUMERICAL RESULTS

In this section, we provide numerical experiments aiming at validating our ANKH approach. Following Sect. IV.G, since the two ANKH variants (ANKH-FMM and ANKH-FFT) provide almost the same timing results, the numerical timing results in this section are based on ANKH-FFT only. Our code is sequential and still run on a single Intel(R) Xeon(R) Silver 4116 CPU. We used the intel oneAPI Math Kernel Library for BLAS calls and FFTW3 for FFT computations. All the code uses double precision arithmetic. The Ewald Summation parameter ξ (see eq 10) is always fixed at 0.01.

V.A. Interpolation Errors. First, we want to verify that our way of treating the point multipole expansions results in admissible errors. To do so, for each partial derivative of H (as in eq 11) in situations corresponding to well-separated pair of cells (t, s), we numerically measured the relative error on the interpolation:

$$\frac{\max_{\substack{\text{on 100 randomly generated} \\ (\mathbf{x}, \mathbf{y}) \in t \times s}} \{ \partial_{\mathbf{x}}^{\alpha} \partial_{\mathbf{y}}^{\beta} (H(\mathbf{x}, \mathbf{y}) - I(\mathbf{x}, \mathbf{y})) \}}{\max_{\substack{\text{on 100 randomly generated} \\ (\mathbf{x}, \mathbf{y}) \in t \times s}} \{ |H(\mathbf{x}, \mathbf{y})| \}} \quad (54)$$

where we considered the maximum error, the same set of 100 test points on both numerator and denominator. The important point is that these errors are given relative to the maximum magnitude of the (nondifferentiated) kernel H . This is motivated by the form of our differential operators $\mathcal{D}_{\mathbf{x}}$, $\mathcal{D}_{\mathbf{y}}$ that are linear combinations of partial derivatives, hence influenced by the maximal magnitude of these partial derivatives, which is obtained on the kernel evaluation (without derivatives) directly. In addition, because we do not consider scaling by charges, dipoles and quadrupoles moment, and because in our applications the magnitudes of these moments are far smaller than the charged one, the results of Figure 6 should overestimate




Figure 6. Relative error (according to eq 54) of kernel interpolation and numerical differentiation: q, charges; d, dipoles; T, quadrupoles. Error types $i-j$ correspond to derivatives used for i and j applied on kernel H .

the numerical error induced by numerical differentiation on practical cases. Notice that similar results are obtained using more test points (but prohibit the tests for high orders). We tested all possible partial derivatives: since if $|\alpha_0| = |\alpha_1|$ and $|\beta_0| = |\beta_1|$ the errors using $\partial_{\mathbf{x}}^{\alpha_0} \partial_{\mathbf{y}}^{\beta_0}$ and $\partial_{\mathbf{x}}^{\alpha_1} \partial_{\mathbf{y}}^{\beta_1}$ have the same order, we only report one of them in the figure.

According to Figure 6, the numerical differentiation described in Sect. III.D results in errors on derivatives always smaller than the error on the kernel interpolation itself (relative to this kernel direct evaluation). The numerical error of this interpolation process on the (nondifferentiated) kernel follows a $O(\nu^L)$ shape, with 1D interpolation order L and $\nu \in (0, 1)$, so do all the tested partial derivatives (with the error defined in eq 54).

As a sum of these ($1 + 3 + 6 = 10$, using quadrupole matrix symmetry) partial derivatives, the numerical differentiation error using operators $\mathcal{D}_{\mathbf{x}}$ and $\mathcal{D}_{\mathbf{y}}$ (as in Sect. I) should also behave as $O(\nu^L)$. This last point is investigated in Sect. V.B.

V.B. Overall Error in Nonperiodic Configuration. In this section, we measure the relative error r_{ANKH} of ANKH-FFT results $\mathcal{E}_{\text{ANKH}}$ (relative to exact solution) on a small water test case without periodic boundary conditions and with respect to

1D (Chebyshev) interpolation order. This relative error is given by

$$\mathcal{E}_{\text{exa}} = \frac{1}{2} \sum_{\mathbf{x} \in X} \sum_{\mathbf{y} \in Y} \mathcal{D}_{\mathbf{x}} \mathcal{D}_{\mathbf{y}} \left(\frac{1}{|\mathbf{x} - \mathbf{y} + \mathbf{t}|} \right),$$

$$r_{\text{ANKH}} = \frac{|\mathcal{E}_{\text{exa}} - \mathcal{E}_{\text{ANKH}}|}{|\mathcal{E}_{\text{exa}}|} \quad (55)$$

According to Remark 1, the 1D interpolation order L_u on equispaced grids does not have to fit with the Chebyshev interpolation one L_c , hence we considered different L_u for the tested L_c in order to check the compression error that may be induced by choosing $L_u < L_c$. Results are provided in Figure 7.




Figure 7. Relative error on water box with 648 atoms (without periodic boundary condition) with respect to the 1D Chebyshev interpolation order L_c and according to various 1D equispaced interpolation order L_u .

First, we observe that the error geometrically decreases with respect to the 1D Chebyshev interpolation order, provided that the equispaced one equals this order. Hence, the overall ANKH-FFT approach numerically converges in accordance with the theory. According to the geometric convergence, when fixing the equispaced order to 4, the relative error does not decrease under $\approx 10^{-5}$, while it stagnates at $\approx 10^{-6}$ and $\approx 10^{-7}$ for equispaced orders fixed to 6 and 8, respectively. Moreover, when fixing to a constant the equispaced 1D interpolation order, the convergence is still observed up to the precision allowed by the latter (and then stagnates). This thus validates the trick we presented in Remark 1: we can further compress the Chebyshev interpolation process using possibly smaller equispaced interpolation orders.

V.C. Overall Error in Periodic Boundary Configuration.

The tested relative error is the same than in eq 55 except that periodic boundary conditions are used to \mathcal{E}_{exa} , fitting with the target quantity defined in eq 1:

$$\mathcal{E}_{\text{exa}} = \frac{1}{2} \sum_{\mathbf{t} \in 2r_b \mathbb{Z}^3} \sum_{\mathbf{x} \in X} \sum_{\mathbf{y} \in Y} \mathcal{D}_{\mathbf{x}} \mathcal{D}_{\mathbf{y}} \left(\frac{1}{|\mathbf{x} - \mathbf{y} + \mathbf{t}|} \right),$$

$$r_{\text{ANKH}} = \frac{|\mathcal{E}_{\text{exa}} - \mathcal{E}_{\text{ANKH}}|}{|\mathcal{E}_{\text{exa}}|} \quad (56)$$

The corresponding results are provided in Figure 8, in which we considered $31^3 - 1$ images of the simulation box, an 1D equispaced interpolation order equal to the Chebyshev one L_c and interpolation order for the far images handling (see Sect. IV.C) equal to $L_c + 1$ (which was empirically fixed and may not




Figure 8. Relative error of ANKH-FFT de with respect to the 1D Chebyshev interpolation order.

be optimal). The test case is the same as in Sect. V.B. This choice of image number is also empirical: we observed in our tests that such a parameter allowed to fully converge in most cases.

According to these errors, our ANKH approach converges in the periodic boundary condition configuration, still geometrically (hence as the nonperiodic case of Sect. V.B). The relative error value for $L_c = 8$ is surprisingly low, but we believe that this is just a particular case on which more digits are luckily correct than expected in this particular test case and using this particular parameter choice. This validates our ANKH approach when using periodic boundary conditions.

Unfortunately, due to the cost of direct computation of exact reference solution using large number of far images, we cannot directly check our code with exact solution when considering important numbers of atoms. However, to emphasize the conclusion in this section, as discussed (and done) in Sect. V.D, we can verify that our ANKH approach is able to reach the same accuracy as the provably convergent code on larger cases. Notice that we also retrieve the geometrical convergence under periodic boundary conditions when using few numbers of images (to perform the direct computation) on test cases with ≈ 100000 particles (hence, since the conclusions are the same, we do not report them in this paper).

V.D. Performance Test. In this section, we compare the performance of both ANKH-FFT and Smooth Particle Mesh Ewald (SPME) using Tinker-HP on various systems. Their sizes range from a small water box of 216 water molecules (648 atoms) up to the Satellite Tobacco Mosaic Virus in water (1066624 atoms), covering the typical scale of systems of interest in biochemistry even if both methods could naturally handle larger cases. In more detail these are

- water box of 648 atoms, size 18.643^3 \AA ,
- dhfr protein in water, 23558 atoms, size 62.23^3 \AA ,
- water box of 96000 atoms, size 98.65^3 \AA ,
- water box of 288000 atoms, size 142.17^3 \AA ,
- water box of 864000 atoms, size 205.19^3 \AA ,
- Satellite Tobacco Mosaic Virus in water, 1066624 atoms, 223.0^3 \AA .

Most of these test cases can be found in the Tinker-HP git repository: <https://github.com/TinkerTools/tinker-hp/tree/master/v1.2/example>

Let us recall that we compare sequential execution of both applications without including precomputations (see Sect. IV.G). We use Tinker-HP results with default SPME parameters as a reference, which we observed to be of order 10^{-4} compared

to pure Ewald summation. Thus, we tuned the ANKH-FFT parameters to reach similar relative errors. We considered, for these tests, a 1D equispaced interpolation order equal to 2.

On small test cases (between few hundreds and few thousands), ANKH-FFT and SPME show similar timings (Figure 9). Medium test cases appear to be the most favorable




Figure 9. Comparison of application timings (in seconds) for ANKH-FFT (red) and SPME (purple) with regard to the number of atoms in *B*.

to SPME here, where it performs 20% faster than ANKH-FFT (for 96000 atoms). Notice that this number of particles is quite small for hierarchical methods and that performance still is very similar (same magnitude order each time), which allows us to claim that our method almost performs the same way as *state-of-the-art* approaches on this type of particle distributions. On larger test cases, ANKH-FFT outperforms SPME, providing timings up to 40% smaller than the last.

Hence, ANKH-FFT shows important performance on a wide range of particle distribution sizes and for various types of typical systems arising in biochemistry. This direct code performance comparison however does not allow us to conclude on the interest of the optimizations we provided. We thus focus on ANKH-FFT in Sect. V.E.

V.E. Split Timings and Details. In this section we provide additional details on the tests of Sect. V.D. First, we present in Figure 10 split timings in order to compare the running times of the various evaluation steps. There mainly are three of them:

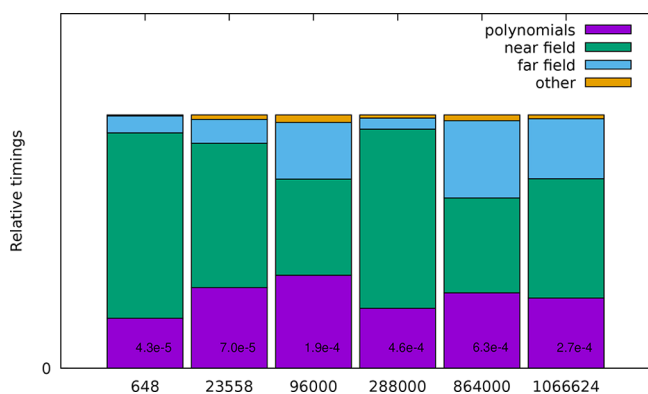


Figure 10. Relative timings (with respect to the entire running time of each test case) of the different steps of ANKH-FFT application. Relative errors (compared to the results of Tinker-HP) are indicated in black, and the number of atom in each corresponding system corresponds to the abscissa.

- the computation of interpolation polynomials at leaves and transformation to modified charges,
- the far field computation (using FFT),
- the near field computation (using direct evaluation).

Clearly, the most time-consuming part is dedicated to near field computations. This step, as well as the polynomial handling, has the strong advantage of being a local one. In addition, tree structure localizes the particles in an efficient way, avoiding the computation of neighbor lists. We thus expect these two steps to fully benefit from parallelization.

Since the far-field computation time does not exceed $\approx 6\%$ of the overall application time, these results tend to validate our fast handling of far interactions. Notice that this far field still involves most of the interactions (at least for medium and large test cases). However, these relatively small timings are mainly due to the important recompression, using 1D equispaced interpolation order equal to 2. This was sufficient to reach the Tinker-HP SPME implementation precision (see relative errors on Figure 10). For higher required precision, the balance between near and far fields is better preserved in overall timings. Nevertheless, this recompression allows us to strongly mitigate the hierarchical methods prefactor.

As stated in Sect. IV.F, ANKH-FFT precomputation complexity is linearithmic. These costs are negligible in applications according to Sect. IV.G. We still present in Figure 11 the precomputation timings of ANKH-FFT on the different test cases of Sect. V.D.




Figure 11. Precomputation timings with respect to number of atoms for fixed 10^{-4} output precision and using $31^3 - 1$ images of the simulation box *B* (in red), compared to $O(N \log N)$ curve (in black).

According to these timings, the precomputation step behaves as our theoretical linearithmic estimates. This is a consequence of the interpolation used for the handling of far field images since we also compared to a noninterpolated version that was already not able to run in medium test cases. In terms of magnitude order, ANKH-FFT precomputation timings are up to few hundred seconds on larger test cases. Regarding the entire cost of MD simulation, since these precomputations are done only once; the conclusion proposed in Sect. IV.G is numerically verified.

VI. CONCLUSION

In this paper, we introduced ANKH, a new polynomial interpolation-based accelerated Ewald Summation for energy computations in molecular dynamics. Our method mainly exploits the (absolutely convergent) real part of the Ewald Summation in order to express the energy computation problem

as a generalized N -body problem on which a fast kernel-independent hierarchical method can perform. We then proposed two fast approaches, the first one based on IBFMM (accordingly named ANKH-FMM) will be well-suited for distributed memory parallelization and the second, built on a fast matrix diagonalization using FFTs (named ANKH-FFT), will be a good candidate for GPU computations. ANKH-FMM has a linear complexity while ANKH-FFT has a linearithmic one but practically behaves as a linear complexity approach. We demonstrated the numerical convergence of our interpolation-based method, and we compared our implementation to a Smooth Particle Mesh Ewald production code (Tinker-HP), exhibiting comparable performance on small and medium test systems and superior performances on large test cases. A public code is provided at <https://github.com/IChollet/ankh-fft> whereas a version of it will also be available within a future Tinker-HP code release (<https://github.com/TinkerTools/tinker-hp>).

We believe that ANKH is an interesting approach for MD on modern HPC architectures, especially because it theoretically solved the scalability issues of Particle-Mesh methods in distributed memory context. However, our current implementation does not use shared or distributed memory parallelism and our intuition still has to be verified in practice. Hence, our priority is now to develop a scalable ANKH-based code (MPI + GPU) and to test it on early exascale architectures.

Among other short-term direct perspectives of the work we presented here, we are considering the (almost) straightforward extension to forces computations. Indeed, the main difference between the algorithm we presented and a general force algorithm belongs in the mutual interactions that cannot be exploited the same way. We would also like to look at application of ANKH-FFT to other scientific fields on which N -body problems on quasi-uniform particle distributions have to be solved.

Considering long-term perspectives, the adaptation of the algorithms we presented to noncubic simulation boxes may also find applications in material science.

COMPUTATION OF $\mathbf{D}(\mathbb{A})$

In this algorithm, we used the notations $\mathbf{I} = (I_0, \dots, I_5)^T$, $\mathbf{J} = (J_0, \dots, J_5)^T$, $\mathbf{R} = (R_0, \dots, R_5)^T$. Mainly, this algorithm generates a vector composed of the (correctly ordered for the circulant embedding) elements of the first row and column of \mathbb{A} , followed by the discrete Fourier transform of this vector. This results in the diagonal elements of the diagonalization of \mathbb{A} in the Fourier domain, i.e., $\mathbf{D}(\mathbb{A})$.

The last if test ($t > n$) is needed to avoid near field interactions. This algorithm is implemented in the C++ library we developed and is available at <https://github.com/IChollet/ankh-fft>. Especially, it can be found in the following file <https://github.com/IChollet/ankh-fft/blob/main/include/circulant.hpp>.

Algorithm 5 Get ANKH-FFT diagonal matrix

```

 $u = \frac{(2^E - 1)r_B}{2^E}$ 
 $v = \frac{r_B}{2^E}$ 
 $P_{cell} = 2^{E+1} - 1$ 
 $P_{ntrp} = 2L - 1$ 
 $\mathbf{Z} := \begin{bmatrix} \frac{2u}{2^E-1} & \frac{2v}{L-1} & 0 & 0 & 0 & 0 \\ 0 & 0 & \frac{2u}{2^E-1} & \frac{2v}{L-1} & 0 & 0 \\ 0 & 0 & 0 & 0 & \frac{2u}{2^E-1} & \frac{2v}{L-1} \end{bmatrix}$ 
for  $\mathbf{R} \in ([0, P_{cell} - 1] \times [0, P_{ntrp} - 1])^3$  do
  for  $k = 0, 1, 2$  do
    if  $R_{2k} < 2^E$  then
       $I_{2k} = R_{2k}$  and  $J_{2k} = 0$ 
    else
       $I_{2k} = 0$  and  $J_{2k} = P_{cell} - R_{2k}$ 
    end if
    if  $R_{2k+1} < L$  then
       $I_{2k+1} = R_{2k+1}$  and  $J_{2k+1} = 0$ 
    else
       $I_{2k+1} = 0$  and  $J_{2k+1} = P_{cell} - R_{2k+1}$ 
    end if
     $\mathbf{x} = \mathbf{Z}\mathbf{I}$ 
     $\mathbf{y} = \mathbf{Z}\mathbf{J}$ 
  end for
   $t = 0$ 
   $n = d$ 
  for  $k = 0, 1, 2$  do
     $t = t + |I_{2k} - J_{2k}|$ 
    if  $(I_{2k} - J_{2k}) \neq 0$  then
       $n = n - 1$ 
    end if
  end for
   $i = \sum_{k=0}^2 (R_{2k}P_{ntrp} + R_{2k+1})P_{cell}^{2-k}P_{ntrp}^{2-k}$ 
  if  $t > n$  then
     $\mathbf{v}_i = \sum_{\mathbf{t} \in 2r_B Z_p} H(\mathbf{x}, \mathbf{y} + \mathbf{t})$ 
  else
     $\mathbf{v}_i = 0$ 
  end if
end for
 $\text{diag}(\mathbf{D}(\mathbb{A})) = \mathbb{F}_6 \mathbf{v}$ 

```

AUTHOR INFORMATION

Corresponding Authors

Igor Chollet – LAGA, Université Sorbonne Paris Nord, UMR 7539, Villetaneuse, France and LCT, Sorbonne Université, UMR 7616, Paris 75006, France; Email: chollet@math.univ-paris13.fr

Louis Lagardère – LCT, Sorbonne Université, UMR 7616, Paris 75005, France; Email: louis.lagardere@sorbonne-universite.fr

Jean-Philip Piquemal – LCT, Sorbonne Université, UMR 7616, Paris 75005, France; orcid.org/0000-0001-6615-9426; Email: jean-philip.piquemal@sorbonne-universite.fr

Complete contact information is available at: <https://pubs.acs.org/10.1021/acs.jctc.3c00015>

Notes

The authors declare no competing financial interest.

ACKNOWLEDGMENTS

This work has also received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No 810367), project EMC2 (JPP).

REFERENCES

- (1) Abduljabbar, M.; Farhan, M. A.; Al-Harthi, N.; Chen, R.; Yokota, R.; Bagci, H.; Keyes, D. Extreme scale fmm-accelerated boundary

integral equation solver for wave scattering. *SIAM Journal on Scientific Computing* **2019**, *41* (3), C245–C268.

(2) Aimi, A.; Desiderio, L.; Di Credico, G. Partially pivoted aca based acceleration of the energetic bem for time-domain acoustic and elastic waves exterior problems. *Computers & Mathematics with Applications* **2022**, *119*, 351–370.

(3) Amiraslani, A.; Corless, R. M.; Gunasingam, M. Differentiation matrices for univariate polynomials. *Numerical Algorithms* **2020**, *83*, 1–31.

(4) Ayala, A.; Tomov, S.; Stoyanov, M.; Dongarra, J. In *Parallel Computing Technologies*; Malyshev, V., Ed.; Springer International Publishing, 2021; pp 279–287.

(5) Badreddine, S.; Chollet, I.; Grigori, L. Factorized structure of the long-range two-electron integrals tensor and its application in quantum chemistry. *SSRN Journal* **2022**, 4270096.

(6) Blanchard, P. *Fast hierarchical algorithms for the low-rank approximation of matrices, with applications to materials physics, geostatistics and data analysis*. Ph.D. Thesis, Université de Bordeaux, Feb. 2017.

(7) Board, J. A.; Humphres, C. W.; Lambert, C. G.; Rankin, W. T.; Toukmaji, A. Y. Ewald and multipole methods for periodic n-body problems. In *Computational Molecular Dynamics: Challenges, Methods, Ideas*; Deuffhard, P., Hermans, J., Leimkuhler, B., Mark, A. E., Reich, S., Skeel, R. D., Eds.; Springer Berlin Heidelberg, 1999; pp 459–471.

(8) Börm, S. *Efficient Numerical Methods for Non-local Operators: H2-matrix Compression, Algorithms and Analysis*; EMS tracts in mathematics; European Mathematical Society, 2010.

(9) Brooks, B. R.; Brooks, C. L., III; Mackerell, A. D., Jr; Nilsson, L.; Petrella, R. J.; Roux, B.; Won, Y.; Archontis, G.; Bartels, C.; Boresch, S.; et al. Charmm: the biomolecular simulation program. *Journal of computational chemistry* **2009**, *30* (10), 1545–1614.

(10) Case, D. A.; Cheatham, T. E., III; Darden, T.; Gohlke, H.; Luo, R.; Merz, K. M., Jr; Onufriev, A.; Simmerling, C.; Wang, B.; Woods, R. J. The amber biomolecular simulation programs. *Journal of computational chemistry* **2005**, *26* (16), 1668–1688.

(11) Challacombe, M.; White, C.; Head-Gordon, M. Periodic boundary conditions and the fast multipole method. *J. Chem. Phys.* **1997**, *107*, 10131.

(12) Chollet, I. *Symmetries and Fast Multipole Methods for Oscillatory Kernels*. Ph.D. Thesis, Sorbonne Université, Mar. 2021.

(13) Chollet, I.; Claeys, X.; Fortin, P.; Grigori, L. A Directional Equispaced interpolation-based Fast Multipole Method for oscillatory kernels. *SIAM J. Sci. Comput.* **2023**, *45*, C20.

(14) Darden, T.; York, D.; Pedersen, L. Particle mesh ewald: An $n \log(n)$ method for ewald sums in large systems. *J. Chem. Phys.* **1993**, *98* (12), 10089–10092.

(15) Dehnen, W. A hierarchical (n) force calculation algorithm. *J. Comput. Phys.* **2002**, *179* (1), 27–42.

(16) Dubinski, J. A parallel tree code. *New Astronomy* **1996**, *1* (2), 133–147.

(17) Essmann, U.; Perera, L.; Berkowitz, M.; Darden, T.; Lee, H.; Pedersen, L. A smooth particle mesh ewald method. *J. Chem. Phys.* **1995**, *103* (11), 8577.

(18) Fong, W.; Darve, E. The black-box fast multipole method. *J. Comput. Phys.* **2009**, *228* (23), 8712–8725.

(19) Frigo, M.; Johnson, S. G. The design and implementation of FFTW3. *Proceedings of the IEEE* **2005**, *93* (2), 216–231. Special issue on “Program Generation, Optimization, and Platform Adaptation”.

(20) Greengard, L.; Rokhlin, V. A fast algorithm for particle simulations. *J. Comput. Phys.* **1987**, *73* (2), 325–348.

(21) Greengard, L.; Rokhlin, V. The rapid evaluation of potential fields in three dimensions In *Vortex Methods*; Anderson, C., Greengard, C., Eds.; Springer Berlin Heidelberg, pp 121–141.

(22) Hamada, T.; Narumi, T.; Yokota, R.; Yasuoka, K.; Nitadori, K.; Taiji, M. 42 tflops hierarchical n-body simulations on gpus with applications in both astrophysics and turbulence. *Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis* **2013**, 1654123.

(23) Hardy, D. J.; Wu, Z.; Phillips, J. C.; Stone, J. E.; Skeel, R. D.; Schulten, K. Multilevel summation method for electrostatic force evaluation. *J. Chem. Theory Comput.* **2015**, *11* (2), 766–779. PMID: 25691833.

(24) Jing, Z.; Liu, C.; Cheng, S. Y.; Qi, R.; Walker, B. D.; Piquemal, J.-P.; Ren, P. Polarizable force fields for biomolecular simulations: Recent advances and applications. *Annual Review of Biophysics* **2019**, *48* (1), 371–394. PMID: 30916997.

(25) Kohnke, B.; Kutzner, C.; Grubmüller, H. A gpu-accelerated fast multipole method for gromacs: Performance and accuracy. *J. Chem. Theory Comput.* **2020**, *16* (11), 6938–6949. PMID: 33084336.

(26) Lagardère, L.; Jolly, L.-H.; Lipparini, F.; Aviat, F.; Stamm, B.; Jing, Z. F.; Hager, M.; Torabifard, H.; Cisneros, G. A.; Schnieders, M. J.; et al. Tinker-hp: a massively parallel molecular dynamics package for multiscale simulations of large complex systems with advanced point dipole polarizable force fields. *Chemical science* **2018**, *9* (4), 956–972.

(27) Lambert, C. G.; Darden, T. A.; Board, J. A., Jr. A multipole-based algorithm for efficient calculation of forces and potentials in macroscopic periodic assemblies of particles. *J. Comput. Phys.* **1996**, *126* (2), 274–285.

(28) Lashuk, I.; Chandramowlishwaran, A.; Langston, H.; Nguyen, T.-A.; Sampath, R.; Shringarpure, A.; Vuduc, R.; Ying, L.; Zorin, D.; Biros, G. A massively parallel adaptive fast multipole method on heterogeneous architectures. *Commun. ACM* **2012**, *55* (5), 101–109.

(29) Melcr, J.; Piquemal, J.-P. Accurate biomolecular simulations account for electronic polarization. *Frontiers in Molecular Biosciences* **2019**, *6*, No. 143, DOI: 10.3389/fmolb.2019.00143.

(30) Messner, M. *Fast Boundary Element Methods in Acoustics*. Dissertation, Graz University of Technology, 2011.

(31) Ohno, Y.; Yokota, R.; Koyama, H.; Morimoto, G.; Hasegawa, A.; Masumoto, G.; Okimoto, N.; Hirano, Y.; Ibeid, H.; Narumi, T.; Taiji, M. Petascale molecular dynamics simulation using the fast multipole method on k computer. *Comput. Phys. Commun.* **2014**, *185* (10), 2575–2585.

(32) Rokhlin, V.; Wandzura, S. The fast multipole method for periodic structures. *Proceedings of IEEE Antennas and Propagation Society International Symposium and URSI National Radio Science Meeting* **1994**, *1*, 424–426.

(33) Sagui, C.; Pedersen, L. G.; Darden, T. A. Towards an accurate representation of electrostatics in classical force fields: Efficient implementation of multipolar interactions in biomolecular simulations. *J. Chem. Phys.* **2004**, *120* (1), 73–87.

(34) Salmon, J. K. *Parallel hierarchical N-body methods*. Ph.D. Thesis, California Institute of Technology, 1990.

(35) Schmidt, K. E.; Lee, M. A. Implementing the fast multipole method in three dimensions. *J. Stat. Phys.* **1991**, *63*, 1223–1235.

(36) Shi, Y.; Ren, P.; Schnieders, M.; Piquemal, J.-P. *Polarizable Force Fields for Biomolecular Modeling*. John Wiley & Sons, Ltd, 2015; Chapter 2, pp 51–86, DOI: 10.1002/9781118889886.ch2.

(37) Shimada, J.; Kaneko, H.; Takada, T. Performance of fast multipole methods for calculating electrostatic interactions in biomacromolecular simulations. *J. Comput. Chem.* **1994**, *15* (1), 28–43.

(38) Simmonett, A.; Brooks, B.; Darden, T. Efficient and scalable electrostatics via spherical grids and treecode summation. Working paper or preprint, 2022.

(39) Simmonett, A. C.; Brooks, B. R. A compression strategy for particle mesh ewald theory. *J. Chem. Phys.* **2021**, *154* (5), 054112.

(40) Smith, W. Point Multipoles in the Ewald Summation (Revisited). *CCPS Newsletter* **2007**, *46*, 18.

(41) Stamm, B.; Lagardère, L.; Polack, É.; Maday, Y.; Piquemal, J.-P. A coherent derivation of the ewald summation for arbitrary orders of multipoles: The self-terms. *J. Chem. Phys.* **2018**, *149* (12), 124103.


(42) Takahashi, T.; Cecka, C.; Fong, W.; Darve, E. Optimizing the multipole-to-local operator in the fast multipole method for graphical processing units. *International Journal for Numerical Methods in Engineering* **2012**, *89* (1), 105–133.

(43) Toukmaji, A.; Sagui, C.; Board, J.; Darden, T. Efficient particle-mesh ewald based approach to fixed and induced dipolar interactions. *J. Chem. Phys.* **2000**, *113* (24), 10913–10927.

(44) Van Der Spoel, D.; Lindahl, E.; Hess, B.; Groenhof, G.; Mark, A. E.; Berendsen, H. J. Gromacs: fast, flexible, and free. *Journal of computational chemistry* **2005**, 26 (16), 1701–1718.

(45) Warren, M. S.; Salmon, J. K. Astrophysical n-body simulations using hierarchical tree data structures. *Proceedings of the 1992 ACM/IEEE Conference on Supercomputing* **1992**, 570–576, DOI: 10.1109/SUPERC.1992.236647.

(46) Ying, L.; Biros, G.; Zorin, D. A kernel-independent adaptive fast multipole algorithm in two and three dimensions. *J. Comput. Phys.* **2004**, 196 (2), 591–626.



CAS BIOFINDER DISCOVERY PLATFORM™

BRIDGE BIOLOGY AND CHEMISTRY FOR FASTER ANSWERS

Analyze target relationships,
compound effects, and disease
pathways

Explore the platform

