

Accelerating the calculation of dipolar interactions in particle based simulations with open boundary conditions by means of the P²NFFT method

Rudolf Weeber^{1,*}, Franziska Nestler², Florian Weik¹, Michael Pippig², Daniel Potts², and Christian Holm¹

¹ *Institute for Computational Physics, University of Stuttgart, Allmandring 3, Germany and*

² *Chemnitz University of Technology, Faculty of Mathematics, 09107 Chemnitz, Germany*

(Dated: July 25, 2018)

Magnetic gels are soft elastic materials consisting of magnetic particles embedded in a polymer network. Their shape and elasticity can be controlled by an external magnetic field, which gives rise to both, engineering and biomedical applications. Computer simulations are a commonly used tool to study these materials. A well-known bottleneck of these simulations is the demanding calculation of dipolar interactions. Under periodic boundary conditions established algorithms are available for doing this, however, at the expense of restricting the way in which the gels can deform in an external magnetic field. Moreover, the magnetic properties depend on the sample shape, ruling out periodic boundary conditions entirely for some research questions. In this article we will employ the recently developed dipolar variant of the P²NFFT method that is able to calculate dipolar interactions under open boundary conditions with an $N \log N$ scaling in the number of particles, rather than the expensive N^2 scaling of a direct summation of pair forces. The dipolar P²NFFT method has been implemented within the ScaFaCoS library. The molecular dynamics software ESPResSo has been extended to make use of the library.

After a short summary of the method, we will discuss its value for studying magnetic soft matter systems. A particular focus is put on developing a tuning strategy to reach the best performance of the method at a predefined accuracy, and lastly applying the method to a magnetic gel model. Here, adapting to the gel's change in shape during the course of a simulation is of particular interest.

Keywords: simulations, magnetic gels, dipolar interactions, open boundary conditions, nonequispaced fast Fourier transform, NFFT, P²NFFT

I. INTRODUCTION

In this article we report on the application of a novel scheme called particle-particle NFFT method (P²NFFT) to calculate the dipolar interactions in soft magnetic materials for the case of open boundary conditions. This is a relevant use case, in particular, for magnetic gels. These materials consist of magnetic nanoparticles embedded into a gel matrix [1, 2]. This allows to control their shape and elasticity via external magnetic fields [3–6] giving rise to applications such as actuation [7–9] or biomedical applications, since the magnetic field controlling their behaviour does not significantly interact with living matter. A review on magnetic gels can be found in Ref. [10], a discussion of their study using particle-based simulations in Ref. [11].

Simulations of magnetic gels have been performed mostly using molecular dynamics [12–17] as well as finite-elements and finite-volume methods [18–22].

Two interesting aspects in the simulation of a magnetic gel are the calculation of the elasticity, – and sometimes the structure – of the hydrogel matrix, and the calculation of the interactions between the embedded magnetic particles. The excluded volume interaction of the magnetic

particles is often approximated by hard or soft sphere potentials, while the magnetic interaction is simplified to a point dipole approximation, with the dipoles sitting at the particles center. Hence, the task is to calculate the dipolar interaction energy

$$-\frac{\mu_0}{4\pi} \left[3 \frac{\langle \vec{\mu}_i, \vec{r}_{ij} \rangle \langle \vec{\mu}_j, \vec{r}_{ij} \rangle}{r_{ij}^5} - \frac{\langle \vec{\mu}_i, \vec{\mu}_j \rangle}{r_{ij}^3} \right], \quad (1)$$

where μ_0 is the magnetic vacuum permittivity, and $\vec{\mu}_i, \vec{\mu}_j \in \mathbb{R}^3$ denote the particles' dipole moments. Furthermore, we denote by $\vec{r}_i \in \mathbb{R}^3$ and $\vec{r}_j \in \mathbb{R}^3$ the positions of the dipoles as well as by

$$r_{ij} := \|\vec{r}_{ij}\| := \|\vec{r}_i - \vec{r}_j\|$$

their distance.

Which of the two parts, elasticity or dipolar interactions, is most computing intensive, depends on the modeling approach. In those cases, where the hydrogel or its elasticity is simulated in great detail, either via the inclusion of explicit polymer chains or via a continuum model, these calculations dominate the computing time. The calculation of the dipolar interactions, on the other hand, is most relevant for models which model the hydrogel's elasticity in a very coarse-grained fashion, e.g., by entropic springs connecting the magnetic particles. These

* weeber@icp.uni-stuttgart.de

models can then contain a large number of magnetic particles. In case of Ref. [16], e.g., 10 000 dipolar particles are simulated.

It is important to note that dipolar interactions are long-ranged, as the interaction scales like r^{-3} , where r denotes the distance between two dipoles. A simple cut-off scheme for reducing the number of dipole pairs to be considered results in significant errors. Under periodic boundary conditions the computational complexity can be reduced to $O(N \log N)$, where N denotes the number of dipoles in the system, using particle mesh Ewald methods like the dipolar P³M method [45]. However, the behaviour of a magnetic gel is also determined by its shape [16, 24, 25], due to the shape-dependent demagnetization energy. Hence, when this effect is important and needs to be taken into account, open boundary conditions have to be applied.

While it is possible to calculate the dipolar interactions for an open system by considering all pairs of particles, the calculations scale like $O(N^2)$, making the simulation for large N computationally expensive. In this contribution, the dipolar P²NFFT method [26, 27] is applied to solve this problem. By using a Fourier-space based calculation for the long range part and applying a regularization at the boundaries to obey the open boundaries, the method achieves an $O(N \log N)$ scaling with respect to the number of dipoles in the system.

Support for the dipolar P²NFFT method with open, periodic, and mixed boundary conditions has been added to the ScaFaCoS library [28, 29]. The implementation is parallelized by means of MPI. Support for using the ScaFaCoS library for long range interactions has been added to the ESPResSo molecular dynamics software [30, 31]. Furthermore, a tuning scheme has been developed that allows to find those method parameters resulting in a given accuracy. The paper is structured as follows. In Sec. II, a brief overview of the method is presented, in Sec. IV the tuning procedure is described. The model system, on which the method and tuning are evaluated, is described in Sec. III, tuning results as well as timings and the method's scaling behaviour are then presented in Sec. V. Based on this, a heuristic scheme for selecting method parameters is discussed in Sec. VI. Lastly, we apply the method to a magnetic gel simulation (Sec. VII) and conclude with a summary.

II. DIPOLAR P²NFFT FOR OPEN BOUNDARY CONDITIONS

Given a system consisting of N dipoles at positions $\vec{r}_j \in \mathbb{R}^3$, each possessing a dipole moment $\vec{\mu}_j \in \mathbb{R}^3$, the

energy of this system is given by

$$U_{\text{dd}} = -\frac{\mu_0}{4\pi} \sum_{j=1}^N \sum_{i=1, i \neq j}^N \frac{\langle \vec{\mu}_i, \nabla_i \rangle \langle \vec{\mu}_j, \nabla_j \rangle}{r_{ij}}. \quad (2)$$

Here, ∇ denotes the gradient operator and $\langle \cdot, \cdot \rangle$ represents the Euclidean inner product in \mathbb{R}^3 . Further, the single summands are a short form notation of the expressions

$$\left\langle \vec{\mu}_i, \nabla_{\vec{x}} \left\langle \vec{\mu}_j, \nabla_{\vec{y}} \frac{1}{\|\vec{x} - \vec{y}\|} \right\rangle \right\rangle$$

evaluated at $\vec{x} = \vec{r}_i$ and $\vec{y} = \vec{r}_j$. Calculating the partial derivatives results in a representation of the single summands as introduced in Eqn. 1.

Instead of computing only the overall energy U_{dd} , one is often interested in computing the single potentials

$$\phi_j := -\frac{\mu_0}{4\pi} \sum_{i=1, i \neq j}^N \frac{\langle \vec{\mu}_i, \nabla_i \rangle}{r_{ij}}, \quad (3)$$

the fields

$$\vec{E}_j := -\nabla_j \phi_j, \quad (4)$$

the torques

$$\vec{T}_j := \vec{\mu}_j \times \vec{E}_j,$$

as well as the acting forces

$$\vec{F}_j := -[\nabla_j \nabla_j^T \phi_j] \vec{\mu}_j. \quad (5)$$

A direct evaluation of the underlying sums of length N , cf. Eqn. 2 or 3 for instance, results in an arithmetic complexity of $O(N^2)$, which is impractical for large particle systems. The dipolar P²NFFT method, as introduced in the following, enables an efficient approximation of the introduced quantities with only $O(N \log N)$ arithmetic operations.

In order to derive fast approximation methods in the field of molecular dynamics simulations, a common approach is to apply the so-called Ewald splitting. This is especially done in the case that periodic boundary conditions are desired. Although we are interested in open boundary conditions, we still apply the Ewald splitting in order to separate the interactions into short range and long range parts, see Sec. II A.

The so-called long range part is approximated by a trigonometric polynomial and evaluated based on the well known fast Fourier transform (FFT). Since the particles have arbitrary positions $\vec{r}_j \in \mathbb{R}^3$, we make use of the fast Fourier transform for nonequispaced data (NFFT), which is introduced in Sec. II B. Finally, the NFFT based fast summation approach, which combines different NFFT modules, is applied in order to approximate the long range potentials and forces. We give a short introduction to the fast summation approach in Sec. II C.

A. Ewald splitting

The function $f(r) = r^{-1}$, where r denotes a distance between two particles, takes a central role in the field of dipolar interactions, see Eqns. 2 and 3. We split the function into two parts via

$$\frac{1}{r} = \frac{\text{erf}(\alpha r)}{r} + \frac{\text{erfc}(\alpha r)}{r}, \quad (6)$$

where $\alpha > 0$ is the Ewald or splitting parameter, $\text{erf}(\cdot)$ denotes the error function given by

$$\text{erf}(x) := \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt$$

and $\text{erfc}(\cdot) := 1 - \text{erf}(\cdot)$ is the complementary error function, see [32]. Note that the second summand in Eqn. 6 is still singular at $r = 0$, whereas the first part has a finite limit for $r \rightarrow 0$, which is given by

$$\lim_{r \rightarrow 0} \frac{\text{erf}(\alpha r)}{r} = \frac{2\alpha}{\sqrt{\pi}}.$$

Furthermore, the erf-term is approaching the value zero very slowly for $r \rightarrow \infty$, whereas the second summand tends to zero exponentially fast.

The potentials can now be computed based on Eqn. 3 by applying the introduced Ewald splitting (Eqn. 6) to the single summands with $r := r_{ij}$. The included partial derivatives are computed later on. In case of the potentials we obtain

$$\phi_j = \phi_j^{\text{short}} + \phi_j^{\text{long}},$$

where we define the short range part

$$\phi_j^{\text{short}} := -\frac{\mu_0}{4\pi} \sum_{i=1, i \neq j}^N \langle \vec{\mu}_i, \nabla_i \rangle \frac{\text{erfc}(\alpha r_{ij})}{r_{ij}} \quad (7)$$

and the long range part

$$\phi_j^{\text{long}} := -\frac{\mu_0}{4\pi} \sum_{i=1, i \neq j}^N \langle \vec{\mu}_i, \nabla_i \rangle \frac{\text{erf}(\alpha r_{ij})}{r_{ij}} \quad (8)$$

$$= -\frac{\mu_0}{4\pi} \sum_{i=1}^N \langle \vec{\mu}_i, \nabla_i \rangle \frac{\text{erf}(\alpha r_{ij})}{r_{ij}}. \quad (9)$$

Note that the function $r^{-1}\text{erf}(\alpha r)$ is continuous at $r = 0$ and, thus, we allow $r_{ij} = 0$. In addition, the gradient of the radial function vanishes at the origin. Consequently, the contribution for $i = j$ is equal to zero, i.e., no self-potential is computed. In contrast, in order to compute

the fields, as defined in Eqn. 4, we apply a further differential operator to the expressions given in Eqns. 7 and 8. Thus, we compute a self-field given by

$$\vec{E}_j^{\text{self}} = \frac{\mu_0}{4\pi} \cdot \frac{4\alpha^3}{3\sqrt{\pi}} \vec{\mu}_j.$$

Analogously to the self-potential, the self-force is equal to zero.

B. FFT for nonequispaced data

We give a brief introduction to fast Fourier transforms for nonequispaced data (NFFT) in three dimensions, cf. [33–36]. The NFFT is going to be applied in order to approximate the long range interactions in Eqn. 9. Therefor, the interaction kernel $r^{-1}\text{erf}(\alpha r)$ will be approximated by a trigonometric polynomial. Finally, a combination of different NFFT modules [37] will be used in order to evaluate the long range interactions, see Secs. II C and II D.

In the following we define for some $\vec{m} = (m_1, m_2, m_3) \in 2\mathbb{N}^3$ the index set

$$\mathcal{I}_{\vec{m}} := \left(\left[-\frac{m_1}{2}, \frac{m_1}{2} \right) \times \dots \times \left[-\frac{m_3}{2}, \frac{m_3}{2} \right) \right) \cap \mathbb{Z}^3.$$

The NFFT realizes an efficient computation of the sums

$$f(\vec{x}_j) = \sum_{\vec{k} \in \mathcal{I}_{\vec{m}}} \hat{f}_{\vec{k}} e^{2\pi i \langle \vec{k}, \vec{x}_j \rangle}, \quad j = 1, \dots, N, \quad (10)$$

i.e., the efficient evaluation of a trigonometric polynomial f at given nonequispaced nodes $\vec{x}_j \in [-1/2, 1/2]^3$. Note that the components of the vector \vec{m} represent the number of Fourier coefficients present in each single dimension. In the context of the dipolar P²NFFT method, $|\vec{m}| = m_1 m_2 m_3$ is the total number of Fourier coefficients used for the approximation of the function $r^{-1}\text{erf}(\alpha r)$. Later in this article, see Section V, we will restrict our considerations to meshes of size $\vec{m} = (m, m, m)$, i.e., we use the same number of mesh points $m \in 2\mathbb{N}$ in each single dimension.

The basic idea of the NFFT is explained as follows. We simply apply the ordinary inverse FFT on the given equispaced mesh in Fourier space. Afterward, the values of f in the nonequispaced nodes \vec{x}_j are recovered based on the computed equispaced data via a so called window function.

In other words, the function values $f(\vec{x}_j)$ are finally approximated via

$$f(\vec{x}_j) \approx \sum_{\vec{l} \in \mathcal{I}_{\vec{m}}} g_{\vec{l}} \tilde{\varphi} \left(\vec{x}_j - \frac{\vec{l}}{m} \right), \quad (11)$$

where we set $\vec{l} := \left(\frac{l_1}{m_1}, \frac{l_2}{m_2}, \frac{l_3}{m_3}\right) \in \mathbb{R}^3$. The function $\tilde{\varphi}$ is in general the 1-periodization of a compactly supported window function φ with small support, for example a B-spline. Thus, the sums in Eqn. 11 consist only of a small number of nonzero summands.

We remark that the NFFT method allows the application of other window functions, such as Bessel functions or Gaussians, for instance. Furthermore, a so called oversampling factor may be applied in Eqn. 11. For the sake of simplicity we restrict our considerations in the present paper to the above mentioned B-spline window function without oversampling and refer to Ref. [38], for instance, for further discussions. We denote by $a \in 2\mathbb{N}$ the order of the B-spline and refer to a as the assignment order in the following.

The coefficients $g_{\vec{l}}$ in Eqn. 11 are obtained by applying a classical inverse FFT on the given regular grid in Fourier space. In order to correct for the convolution with the window function afterward, the given Fourier coefficients $\hat{f}_{\vec{k}}$ are deconvolved with the Fourier coefficients of the window function before. Finally, the NFFT basically consists of the following three steps:

1. Deconvolution in Fourier space:

$$\hat{g}_{\vec{k}} := \frac{\hat{f}_{\vec{k}}}{c_{\vec{k}}(\tilde{\varphi})} \quad (12)$$

2. Inverse FFT: $(\hat{g}_{\vec{k}})_{\vec{k} \in \mathcal{I}_{\vec{m}}} \mapsto (g_{\vec{l}})_{\vec{l} \in \mathcal{I}_{\vec{m}}}$.
3. Approximate the function values by evaluating the short sums 11.

The efficient computation of the sums

$$h(\vec{k}) = \sum_{j=1}^N f_j e^{-2\pi i \langle \vec{k}, \vec{x}_j \rangle}, \quad \vec{k} \in \mathcal{I}_{\vec{m}},$$

is realized similarly based on the FFT. The efficient algorithm is widely known as adjoint NFFT.

In addition, the NFFT has been generalized in order to evaluate the gradients $\nabla f(\vec{x}_j) \in \mathbb{C}^3$, see [39], as well as the Hessians $\nabla \nabla^\top f(\vec{x}_j) \in \mathbb{C}^{3 \times 3}$, as presented in [27]. We refer to these variants as the gradient NFFT and the Hessian NFFT, respectively.

A further variant is called adjoint gradient NFFT enabling the efficient computation of

$$h(\vec{k}) := \sum_{j=1}^N \langle \vec{f}_j, \nabla_j \rangle e^{2\pi i \langle \vec{k}, \vec{x}_j \rangle} \in \mathbb{C}, \quad (13)$$

where $\vec{k} \in \mathcal{I}_{\vec{m}}$. For more details we refer to [27]. The introduced variants of the NFFT are applied in order to compute the considered dipolar interactions efficiently, see Sec. IID.

C. Fast summation for radial kernels

We aim to approximate sums of the form

$$f_j := \sum_{i=1}^N \langle \vec{c}_i, \nabla_i \rangle K(\|\vec{x}_i - \vec{x}_j\|), \quad j = 1, \dots, N, \quad (14)$$

where $K : [0, L] \rightarrow \mathbb{R}$ is a continuously differentiable function, $\vec{c}_i \in \mathbb{C}^3$ are some given vectors and the nodes \vec{x}_j , $j = 1, \dots, N$, are supposed to be arbitrarily distributed in a ball of radius $L/2$, i.e., we have $\|\vec{x}_i - \vec{x}_j\| \leq L$ for all i, j .

Note that the long range parts of the potentials (Eqn. 9) are exactly of the form given in Eqn. 14, where we have $\vec{c}_i = \vec{\mu}_i$ and $K(r) = -\mu_0(4\pi r)^{-1} \text{erf}(\alpha r)$.

In Sec. IIB we introduced the NFFT algorithm, enabling an efficient evaluation of trigonometric polynomials (Eqn. 10) at arbitrary nodes $\vec{x}_j \in \mathbb{T}^3$. Now we consider expressions as given in Eqn. 14, which are in general non-periodic. A second difficulty results from the fact that a separation of the source nodes \vec{x}_i and the target nodes \vec{x}_j is not readily possible, and, hence, a naive evaluation would require $\mathcal{O}(N^2)$ arithmetic operations. To overcome these two problems, we extend the non-periodic kernel function K to a smooth periodic function, which we realize based on the so called regularization technique. Next, we approximate the resulting periodic function by a trigonometric polynomial, which finally enables an efficient evaluation of Eqn. 14 via separating the source and target nodes. This ansatz is closely related to the classical NFFT based fast summation approach [40].

If we are able to calculate the derivatives of the function K , we may construct a regularized function $K_{\text{R}} : [-h/2, h/2]^3 \rightarrow \mathbb{R}$ defined via

$$K_{\text{R}}(\vec{x}) := \begin{cases} K(\|\vec{x}\|) & : \|\vec{x}\| \leq L, \\ K_{\text{B}}(\|\vec{x}\|) & : L < \|\vec{x}\| \leq h/2, \\ K_{\text{B}}(h/2) & : \text{else,} \end{cases} \quad (15)$$

with some $h \geq 2L$. The function $K_{\text{B}} : [L, h/2] \rightarrow \mathbb{R}$ is constructed such that the values of the first $p-1$ derivatives coincide with those of the kernel function K at $x = L$ and that the first $p-1$ derivatives vanish at $x = h/2$. The constant continuation with value $K_{\text{B}}(h/2)$ makes the function K_{R} smooth on the cube $[-h/2, h/2]^3$. The polynomial K_{B} , which fulfills the given interpolation conditions, is computed via the so called two point Taylor interpolation, see [41, 42].

Note that there are no further restrictions to the period $h > 0$, except for $h \geq 2L$. In order to describe the size of the regularization domain relative to the size of the whole interval, we introduce the variable $\epsilon \in [0, 1/2)$ and write

$$L = h(1/2 - \epsilon) \iff \epsilon = \frac{h - 2L}{2h}.$$

By construction, the h -periodic continuation (regarding all three dimensions) of the regularization K_R is smooth. Thus, it is well approximated by a three-dimensional trigonometric polynomial

$$K_R(\vec{x}) \approx \sum_{\vec{k} \in \mathcal{I}_{\vec{m}}} \hat{b}_{\vec{k}} e^{2\pi i \langle \vec{k}, h^{-1} \vec{x} \rangle}, \quad (16)$$

where $h^{-1} \vec{x} \in [-1/2, 1/2]^3$. The Fourier coefficients $\hat{b}_{\vec{k}}$ can be computed via sampling the regularized function K_R at equispaced nodes and applying the FFT.

Since $\|\vec{x}_i - \vec{x}_j\| \leq L$ we obtain

$$f_j \approx \sum_{\vec{k} \in \mathcal{I}_{\vec{m}}} \hat{b}_{\vec{k}} S_{\vec{k}} e^{-2\pi i \langle \vec{k}, h^{-1} \vec{x}_j \rangle}, \quad (17)$$

where we define

$$S_{\vec{k}} := \sum_{i=1}^N \langle \vec{c}_i, \nabla_i \rangle e^{2\pi i \langle \vec{k}, h^{-1} \vec{x}_i \rangle}. \quad (18)$$

An efficient approximation of the values f_j , as defined in Eqn. 14, is now possible as follows. The sums $S_{\vec{k}}$, $\vec{k} \in \mathcal{I}_{\vec{m}}$, are evaluated via the adjoint gradient NFFT, cf. Eqn. 13, followed by a simple multiplication with the coefficients $\hat{b}_{\vec{k}}$, and, finally, the sums in Eqn. 17 are computed via the NFFT, cf. Eqn. 10.

D. The P²NFFT method

In the following we assume that the particles are located in a ball of radius $L/2$ with $L > 0$. Since the complementary error function tends to zero exponentially fast, we may reduce the computational effort needed in order to compute the short range potentials (Eqn. 7) by only considering distances $r_{ij} \leq r_c$ with some appropriate near field cutoff radius $r_c < L$. Of course, such a near field radius only exists in the case that the splitting parameter α is also reasonably chosen. The same applies to the short range parts of the fields and the forces, which are defined analogously containing further partial derivatives.

Provided that $\|\vec{r}_j\| \leq L/2$ for all $j = 1, \dots, N$, we can simply apply the presented NFFT based fast summation approach in order to compute the long range parts of the potentials (Eqn. 9), i.e., we set $\vec{c}_i = \vec{\mu}_i$ and $K(r) = -\mu_0(4\pi r)^{-1} \text{erf}(\alpha r)$ in Eqn. 14. We start with approximating the sums in Eqn. 18 based on the adjoint gradient NFFT, cf. Eqn. 13. Afterward, the sums in Eqn. 17 are approximated via the NFFT, cf. Eqn. 10.

In an analog manner, the long range parts of the fields and the forces may be approximated. Whereas the long range potentials are obtained by applying the NFFT to

sums in Eqn. 17, the long range parts of the fields and the forces are obtained via

$$\vec{E}_j^{\text{long}} = -\nabla_j f_j \quad \text{and} \quad \vec{F}_j^{\text{long}} = -[\nabla_j \nabla_j^T f_j] \vec{\mu}_j,$$

respectively. Here, f_j are approximately given via Eqn. 17, with $\vec{c}_i = \vec{\mu}_i$ and $K(r) = -\mu_0(4\pi r)^{-1} \text{erf}(\alpha r)$. Instead of the NFFT we hereby apply the gradient NFFT and the Hessian NFFT, respectively.

Error control

The error behaviour of dipolar particle mesh methods is already well studied regarding fully periodic boundary conditions, cf. [43–45]. This allows for a fast tuning of the method parameters such that a given accuracy can be reached. In contrast, there are no error estimation formulae for the presented method under mixed periodic and open boundary conditions. For open boundary conditions, however, the exact results are easily calculable by means of direct summation. The method parameters of P²NFFT can then be tuned by numerically minimizing the error of the P²NFFT calculations, depending on the method parameters. This is explained in Sec. IV.

Recent results indicate that for charge-charge systems, the parameters can be tuned heuristically based on the error estimates known for 3d-periodic constraints [42]. This approach may be transferrable to the dipole-dipole case in the future, particularly for 1d and 2d periodic geometries, where the exact result cannot be easily calculated.

III. MODEL SYSTEM

As this paper is focussed on the calculation of dipolar interactions, we use a dipolar hard sphere system for illustrating the tuning procedure and timing of the method. I.e., soft repulsion between the coatings of magnetic particles as well as the elasticity of the hydrogel matrix in a magnetic gel is not included, see Sec. VII for the simulation of an actual magnetic gel system.

A fluid of dipolar hard spheres is characterized by two dimensionless quantities. First, the volume fraction

$$\phi = \frac{1}{6} \frac{N\pi\sigma^3}{V}, \quad (19)$$

measuring the ratio of the volume covered by the hard spheres of diameter σ , and the total simulation volume V . Second,

$$\lambda = \frac{\mu_0 \mu^2}{4\pi \sigma^3 kT} \quad (20)$$

is the ratio of the maximum dipolar energy per particle of two touching particles (Eqn. 1), divided by the thermal energy kT . The number of particles N and their diameter σ can be chosen arbitrarily in a system with periodic boundaries. In the non-periodic case discussed here, however, the number of particles N is an additional parameter, since it controls the relative influence of bulk and surface properties. We use reduced units and set the particle diameter σ and the thermal energy kT equal to 1. For the magnetic fluid, experimentally achievable parameters of $\phi = 0.05$ and $\lambda = 4$ are assumed. Furthermore, model systems with $\lambda \in \{2, 4, 6, 8\}$ at $\phi = 0.05$ and for $\phi \in \{0.01, 0.05, 0.09, 0.13, 0.17\}$ at $\lambda = 4$ are considered for some system sizes.

The test system is generated by randomly placing N particles into the simulation box. Then, overlap between the particles is removed as follows. A purely repulsive Lennard-Jones potential

$$U(r) = \begin{cases} 4 \left[\left(\frac{\sigma}{r}\right)^{12} - \left(\frac{\sigma}{r}\right)^6 \right] & \text{for } r < \sigma, \\ 0 & \text{otherwise,} \end{cases} \quad (21)$$

is applied to the particles. Then, the interaction energy is minimized by the steepest descent method. At the end of this procedure, the minimum distance between any two particles in the system is larger or equal to σ . Note that this is different from the Weeks-Chandler-Andersen potential, which is cut at $2^{\frac{1}{6}}\sigma$. During the tuning and timing of the P²NFFT method, the particles do not move, so the Lennard-Jones interaction is not necessary here. We consider system sizes of 5 000, 10 000, 20 000, 40 000, 80 000, 160 000, and 320 000 particles. The orientations of all magnetic dipole moments are chosen randomly. By selecting a suitable unit for the electric current, the vacuum permittivity in Eqn. 1 can be set to $\frac{\mu_0}{4\pi} = 1$. Inserting the definition of λ from Eqn. 20 into the dipole-dipole interaction (Eqn. 1), setting the thermal energy $kT = 1$ and the particle distance to σ yields a value of

$$\mu = \sqrt{\lambda\sigma^3} \quad (22)$$

for the magnetic moment.

IV. TUNING PROCEDURE

At this point, there are no error estimation formulae for the method. However, in contrast to periodic ones, for open boundary conditions the exact result can be obtained from direct summation. The error for a given set of parameters is then obtained by comparing the P²NFFT result to the exact one. As error measure we use the average between the root mean square force error and root

mean square torque error

$$\Delta = \frac{1}{2} \left(\sqrt{\sum_{i=1}^N |\vec{F}_i - \vec{F}_i^{\text{exact}}|^2} + \sqrt{\sum_{i=1}^N |\vec{T}_i - \vec{T}_i^{\text{exact}}|^2} \right), \quad (23)$$

where N denotes the number of particles, and \vec{F}_i and \vec{T}_i denote the force and torque on particle i . Exact results are obtained from direct summation.

When applying the method in simulations, the set of parameters has to be found which gives the fastest computation time at the desired level of accuracy. The parameter space is high-dimensional, as it consists of number the of mesh points m , the Ewald splitting parameter α , the real-space cutoff for calculating short range interactions r_c , the assignment order a and the parameter ϵ controlling the regularization at the boundaries. Therefore not all possible combinations of parameters can be tested. We developed a tuning procedure, which consists of two phases. First, the set of parameters giving the fastest calculation which still fulfill the accuracy requirement are generated for all combinations of mesh sizes m and assignment orders a to be considered. In the second phase, the calculations for these parameter sets are timed on the desired number of processor cores, and the parameter set providing the fastest calculation is selected.

Let us now examine the first phase in more detail. It is important to note that while all parameters can affect the accuracy of the result, the Ewald splitting parameter α and the parameter ϵ controlling the regularization do not influence the run-time of the algorithm. Hence in the first step, for a given number of mesh points, assignment order, and real space cutoff, we select the Ewald splitting parameter α and the regularization parameter ϵ which minimize the error (Eqn. 23) by means of a numerical minimization using the LBFGS.B method from the SciPy package (www.scipy.org). The optimization is constrained: for the Ewald splitting parameter α , values between $0.4/r_c$ and $5/r_c$ are considered, where r_c is the real space cutoff. The regularization parameter ϵ can adopt values between 0.001 and 0.15. The optimization is then repeated for different real space cutoffs. For a fixed number of mesh points and assignment order, the smallest real space cutoff which allows the accuracy target to be met results in the fastest simulations. It is determined by bisecting the interval between the minimum and maximum real space cutoff considered (five and 20 particle diameters (σ), respectively). Finally, this procedure is performed for all combinations of the number of mesh points and assignment order, which are to be considered. As the results do not depend on each other, they can be executed in parallel.

Once viable parameter sets are collected, the calculation is timed on the target architecture with the desired

number of processor cores using these parameter sets obtained in the previous step, and the fastest one is selected for production simulations. In practice, both, the root-mean-square force and torque error, and the timing results slightly depend on the configuration of particles in the system. When tuning for an accuracy of 10^{-4} , we therefore averaged over four configurations. Timing is conducted for eight configurations.

V. SCALING BEHAVIOUR AND METHOD PARAMETERS

For the purpose of obtaining the scaling behaviour of the method and studying the optimal values of its parameters, tuning was carried out on systems of 5 000 to 320 000 particles. Assignment orders a of two to five were considered. For the number of mesh points, candidates were evaluated which are multiples of small prime factors, as this results in faster Fourier transforms. The full list is

$$m = \{64, 80, 96, 104, 112, 120, 128, 136, 144, 152, 160, 176, 192, 224, 256, 288, 320, 352, 384, 416, 448, 480, 512\}. \quad (24)$$

Mesh sizes below $60(N/10000)^{\frac{1}{3}}$ were not considered. Furthermore, a time limit of 16 hours on four processor cores was applied to the tuning of any individual combination of number of mesh points and assignment order. Parameter sets which could not be tuned in that time will not provide good simulation performance, in the first place. Applying this limitation, even for the largest system containing 320 000 particles, 39 viable parameter sets were obtained, with the number of mesh points ranging from 192 to 512. As mentioned in Sec. IV, the sets of parameters that satisfy the accuracy requirements have to be timed to pick the fastest one on a given architecture. Timings were run on 1, 2, 4, 8, and 16 cores on a node containing two Intel Xenon E2630 processors with 8 cores, each. In Fig. 1, results for the strong scaling are shown. I.e., we plot the time for a calculation versus the number of particles N . Data is shown for different numbers of processor cores. It can be seen that the computing time scales almost linearly in the number of particles. The logarithmic component is not significant when increasing the number of particles by a factor of 64 (5 000 to 320 000 particles).

Results for weak scaling are shown in Fig. 2. Here, the computing time for a fixed number of particles per processor core is shown versus the number of processor cores. It can be seen that up to four cores, the loss in efficiency is small. For eight and 16 cores, however, the




FIG. 1. Strong scaling for calculations on 5 000 to 320 000 particles performed on one to 16 cores. It can be seen that the scaling is close to linear over the complete range of system sizes

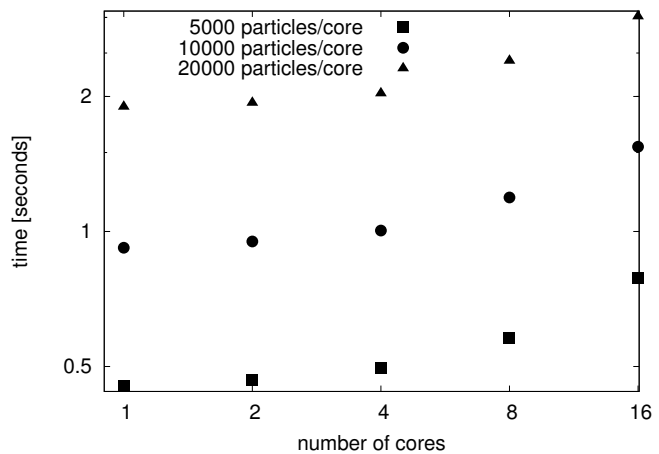


FIG. 2. Weak scaling for calculations on 5 000, 10 000 and 20 000 particles per core for one to 16 cores. The computation time increases with the number of cores, in particular, for 8 and 16 cores. This might be caused by the large amount of data being communicated between the cores.

calculations take significantly more time. This might be related to the rather high communication volume caused by a large real space cutoff and a large FFT grid.

Let us now turn to the parameter sets obtained by the tuning procedure. The timings and associated method parameters for one and 16 cores and the full range of system sizes can be found in Tbl. I. It is worth mentioning that in many cases, several sets of method parameters lead to very similar computation times. Often, the five most favorable sets of parameters result in computation time differences of only a few percent. This is due

N	1 core						16 cores					
	t [s]	m	a	r_c	α	ϵ	t [s]	m	a	r_c	α	ϵ
5000	0.453	80	8	11.445	0.262	0.060	0.070	96	8	9.922	0.308	0.051
10000	0.920	112	8	10.391	0.293	0.041	0.141	112	8	10.391	0.293	0.041
20000	1.897	144	8	10.039	0.304	0.030	0.291	144	10	9.570	0.322	0.029
40000	3.860	176	8	10.156	0.302	0.021	0.579	192	8	9.453	0.330	0.018
80000	7.805	224	8	9.922	0.310	0.012	1.191	224	8	9.922	0.310	0.012
160000	15.674	288	8	9.688	0.320	0.001	2.402	256	8	10.625	0.287	0.006
320000	31.726	352	10	9.336	0.331	0.001	4.741	352	10	9.336	0.331	0.001

TABLE I. Timings and method parameters for the fastest calculations meeting the accuracy target of 10^{-4} on one and 16 cores.




FIG. 3. Number of mesh sites m in each Cartesian direction versus the number of particles N . Results are shown for calculations on one and 16 cores, respectively. It can be seen that m scales like $N^{1/3}$, resulting in an approximately constant ratio of mesh sites per particle m^3/N in the simulation box. The fluctuations are, in part, caused because for some parameters, an assignment order of 10 rather than 8 yielded the fastest calculations.

to trade-offs, e.g. between a higher assignment order or number of mesh points on the one hand and a lower real-space cutoff on the other hand. In Fig. 3, the number of mesh sites per direction m is plotted against the number of particles N . Results are shown for calculations on one and 16 cores, respectively. They can be fitted with $m \sim N^{1/3}$. This implies that the ratio of mesh sites per particle m^3/N stays roughly constant for all system sizes. Because for some cases an assignment order of 10 rather than 8 yielded the fastest calculations, resulting in a slightly smaller real space cutoff, we observe some fluctuations of the timing data around the fit in Fig. 3. For the dipolar interaction parameter $\lambda = 4$ and volume fraction $\phi = 0.05$, we also find that the real space cutoff has to be tuned to $\approx 10\sigma \pm 1.5\sigma$ for all system sizes and number of processor cores. Furthermore, the Ewald splitting parameter scales roughly as $\alpha \sim 1/r_c$, as long as




FIG. 4. Real space cutoff, r_c , versus the mesh size per Cartesian direction, m , for the parameter sets resulting in the fastest calculation while meeting the accuracy requirement for 10 000 and 320 000 particles. The results can be fitted via $r_c \sim m^{-0.75}$.

only the system size and the number of processor cores is changed.

When choosing the method parameters, a trade-off is made between reducing the effort of the real and Fourier space parts, respectively. The computation time of the real space part scales like $O(r_c^3)$, where r_c is the real-space cutoff, whereas the Fourier space part scales like $O(m^3 \log m)$. In Fig. 4, we show the cutoff versus the number of mesh points per Cartesian direction, for the fastest calculations meeting the accuracy target. It can be seen that the cutoff scales approximately as $r_c \sim m^{-0.75}$.

Let us finally examine the influence of the strength of the dipolar interactions λ (Eqn. 20) and the volume fraction ϕ (Eqn. 19) on the optimal calculation time and corresponding method parameters. Results for a system containing 10 000 particles with a dipolar interaction parameter of $\lambda = 4$ and for different densities can be found in Tbl. II. It can be seen that the calculation time increases with increasing density. While the calculations

ρ	t [s]	m	a	r_c	α	ϵ
0.01	0.649	96	8	15.430	0.175	0.041
0.05	0.920	112	8	10.391	0.293	0.041
0.09	1.045	112	10	8.867	0.356	0.048
0.13	1.104	112	10	8.164	0.394	0.049
0.17	1.171	120	10	7.344	0.446	0.049

TABLE II. Optimal calculation times and associated method parameters for various volume fractions ϕ (Eqn. 19) for a system of 10000 particles with a dipolar interaction parameter $\lambda = 4$ simulated on a single cpu core.

λ	t [s]	m	a	r_c	α	ϵ
2.0	0.813	112	8	9.453	0.312	0.036
4.0	0.920	112	8	10.391	0.293	0.041
6.0	1.002	112	8	10.977	0.283	0.043
8.0	1.042	112	10	10.742	0.296	0.048

TABLE III. Optimal calculation times and associated method parameters for various dipolar interaction strengths λ (Eqn. 20) for a system of 10000 particles with a volume fraction $\phi = 0.05$ simulated on a single cpu core.

for low densities are carried out largely in real space, the effort is shifted to the long range calculations in Fourier space for higher densities. It is worth mentioning that the effort for the short range calculations also increases with density. This effort is proportional to the number of particles n_c in a sphere of radius r_c , where r_c is the short range cutoff, we therefore have

$$n_c \sim \phi r_c^3. \quad (25)$$

The ratio between n_c at $\phi = 0.17$ and $\phi = 0.01$ is approximately 1.83.

Results for varying dipolar interaction parameters λ are presented in Tbl. III. The volume fraction is kept constant at $\phi = 0.05$. While varying λ changes the strength of the dipolar interaction, we keep the target accuracy constant at 10^{-4} . This is justified, because the relative importance of interactions in a soft matter system is measured by their relative strength compared to the thermal energy kT . Hence, the required accuracy is determined by this energy scale rather than by the strength of a particular interaction. From the table, it can be seen that the computation time increases with increasing interaction strength, both, due to an increase of the short range cutoff and the number of mesh points. So, in contrast to the case of a varying volume fraction, the relative importance of short and long range calculations does not change strongly. The increase in computation time is likely mostly due to the error being proportional to λ .

VI. TUNING HEURISTICS

The full tuning procedure described in Sec. IV requires a very high computational effort. To make the use of the P²NFFT method practical in simulations, a faster approach is needed. If tabulated tuning results are available for a similar system, this can be done by extrapolating. This is possible based on the $m \sim N^{\frac{1}{3}}$ scaling of the number of mesh points m with the number of particles N (Fig. 3). Moreover, while in some cases an assignment order of $a = 10$ produced slightly faster calculations, the difference to $a = 8$ is not large. Based on these observations, an extensive tuning only has to be carried out for a single number of particles. From those tuning results, the number of mesh points for other system sizes can be extrapolated. We tested this scheme as follows. The number of mesh points ($m_0 = 112$) for $N_0 = 10000$ particles, a dipolar interaction parameter of $\lambda = 4$, and a volume fraction of $\phi = 0.05$ was taken as basis. For all system sizes N , considered (5000 to 320000 particles), the number of mesh points was obtained as

$$m_N = m_0 \left(\frac{N}{N_0} \right)^{\frac{1}{3}}. \quad (26)$$

Based on this estimate, the closest m_N was chosen from the list of mesh sizes considered (Eqn. 24). The timing for this m and an assignment order of $a = 8$ was then compared against the fastest one obtained for the same system size in the full tuning. For $\lambda = 4$ and $\phi = 0.05$, and for all system sizes above 10000 particles, the computation time using the estimated m was found to be within five percent of the optimal one obtained in the full tuning for the respective system size. The scheme can even be used for varying volume fractions ϕ and dipolar interaction parameters λ . Then, however, not only the mesh constant from Eqn. 24 closest to the extrapolated one, but the three closest ones should be considered. In that case, the calculation time was within 15 percent of the optimal one even for the systems with different volume fractions and interaction parameters (Tbl. II and Tbl. III) and for the systems with less than 10000 particles discussed in the previous section.

In practice, after the number of mesh points has been estimated, only a tuning of the Ewald splitting parameter α and the regularization parameter ϵ has to be run, which typically takes less than an hour. This is a low effort compared to a typical soft matter simulation, which runs for millions of time steps.

If no tabulated tuning results are available for extrapolation, it is also possible to employ a simplified version of the full tuning which is based on the assumption that the real space cutoff decreases monotonously with increasing mesh size. This is applied in the next section.




FIG. 5. Initial configuration of a simple model ferrogel. It consists of a number of flexible chains of magnetic particles. In each node of the network, the ends of six chains are attached.

VII. APPLICATION TO MAGNETIC GELS

Let us now consider the simulation of a magnetic gel. Since these materials can deform quite strongly in an external magnetic field, and since the simulation is run under open boundary conditions, special care has to be taken. Since the P²NFFT method uses a mesh covering the simulation box, the box geometry has to be adapted, when the gel changes its shape. A too large simulation box, where a lot of the mesh does not contain particles, slows down the calculation. On the other hand, the sample should also not leave the simulation box. Moreover, the P²NFFT method parameters need to be adapted, when the simulation box is changed. Lastly, the local structure of the material changes considerably in an external magnetic field. This can result in the need for a complete re-tuning of the method after an initial reshaping of the gel.

To illustrate how to cope with these requirements, in this section, we describe a sample simulation protocol. The scripts to run this protocol are provided with this article and at http://github.com/RudolfWeeber/scafacos_espressomd_dipoles. They can be used as a starting point for new projects. The ‘dipoles’ branch of ScaFaCoS [28] with the commit hash ‘066f753f0572c7397508231cb4fc9432d5aeaf04’ is used. The commit hash for ESPResSo is ‘7a0ec981f24de839689f61781581ce6194ecf51d’.

The three-dimensional gel model used is loosely based on “Model I” in Refs.[13, 14]. Please note that the model used here is intended as an example application for P²NFFT, not as a model suitable for study of ferrogels.

The model gel consists of a network of flexible chains of 14 particles, each. Additional particles are placed in the nodes of the network, so that the ends of six chains are connected to each node. As depicted in Fig. 5, the initial arrangement is part of a simple cubic lattice. The gel is cut in a spherical shape and any dangling chains at the surface are removed. In total, there are 17957 particles in the system. Both, node particles and the particles making up the chains, carry a magnetic moment giving a dipolar interaction parameter of $\lambda = 2$ (Eqn. 20). The particles also interact via a purely repulsive Lennard-Jones potential (WCA-potential [46])

$$U(r) = \left\{ \begin{array}{ll} 4\epsilon \left[\left(\frac{\sigma}{r}\right)^{12} - \left(\frac{\sigma}{r}\right)^6 \right] + \frac{1}{4} & r < 2^{\frac{1}{6}}\sigma \\ 0 & \text{otherwise,} \end{array} \right\}, \quad (27)$$

with $\epsilon = \sigma = 1$. Neighboring particles in the chains are connected by a harmonic bond

$$U(r) = \frac{1}{2}kr^2, \quad (28)$$

with $k = 200$. The particles are thermalized by a Langevin thermostat, which applies random forces and a velocity dependent friction to the particles [47]. Thermal energy and friction are set to $kT = \gamma = 1$. During the simulation, both, the particles forming the chains and the particles in the network’s nodes can move. The bonds are maintained throughout the simulation.

a. Setup of the simulation model The initial state shown in Fig. 5 is easily constructed but is not yet a realistic representation of a gel in thermal equilibrium. Due to the entropy introduced at thermal energies $kT > 0$, the polymer chains coil and the network as a whole shrinks. This initial relaxation ($1.5 \cdot 10^6$ time steps of $dt = 0.015$) is done without considering the dipolar interactions between the particles, to save computation time.

b. Obtaining suitable P²NFFT parameters First it is necessary to calculate the exact dipolar forces and torques for the thermalized system via direct summation. This is then used as reference to calculate the error of the P²NFFT calculation for a given set of method parameters. Then the tuning is run. For the model gel, we use a simplified version of the tuning procedure described in Sec. IV, which avoids the complex workflow and makes use of a single script. Rather than tuning for different mesh sizes in parallel, this is done serially. Assuming a monotone relationship between mesh size and real space cutoff (Fig. 4), the highest and lowest candidate mesh sizes m are tuned first to determine the highest and lowest relevant real space cutoffs r_c . Further mesh sizes are attempted in ascending order, such that the maximum real space cutoff to consider can be decreased further. Thus, tuning runs for later in the process will take less time.




FIG. 6. Snapshot of the model ferrogel at the end of the simulation. **Left:** no external field is applied, the gel has a roughly spherical shape. **Right:** A field of $H = 10$ is applied along the x -axis. The gel elongates parallel to the field direction and shrinks in the perpendicular one.

c. Running and re-tuning the model Based on the P²NFFT parameters obtained by the tuning, the actual simulation – including dipolar interactions –, is run. This is done once for the case without an external magnetic field applied, and once for a field of $H = 10$. Every ten time steps ($dt = 0.015$), it is checked whether the gel is too close to the boundary of the simulation box. If there is a layer of less than two or more than six particle diameters (σ in Eqn. 27) around the gel, the box is reshaped to a boundary of four diameters. To maintain a good accuracy of the P²NFFT calculation, the mesh size m in each Cartesian coordinate is adjusted to recover the mesh density obtained in the tuning process. Without this adjustment, the accuracy of the P²NFFT calculations deteriorate significantly.

Due to the presence of the external field and the dipole-dipole interactions, the gel changes its shape and inner structure. So, even with the adjustment of the mesh size as the box is re-shaped, there is no guarantee that the initial choice of P²NFFT parameters are still a good choice. Hence, the forces and torques obtained by P²NFFT have to be compared to an exact calculation using direct summation during a long simulation run.

When no external magnetic field is applied, and the gel stays roughly spherical (left part of Fig. 6), we found that the accuracy after 100 000 time steps is still comparable to the accuracy target of the P²NFFT tuning. When, on the other hand, the gel is deformed by an external field, the deviations turn out to be significant, making a retuning of the method after 50 000 time steps and a simulation of the remaining 50 000 time steps with the new parameters necessary. Images of the model gel both,

with and without an external field applied, are shown in Fig. 6.

d. Scripts provided with this article The following scripts are provided with this article and can serve as a starting point for new models

- **run-test.sh**
Run a short test version of the simulation protocol described above
- **run-full.sh**
Run the full simulation protocol described above. This is used to obtain the samples shown in Fig. 6.
- **model.py**
Model class
- **gen-system.py**
Sets up a model gel and thermalizes it
- **add-reference-forces-torques.py**
Adds reference forces and torques to a stored particle configuration by means of direct summation of the dipolar interactions
- **tune.py**
Tunes the P²NFFT method for a given particle configuration including reference forces and torques
- **run.py**
Runs the model, calculating dipolar interactions by means of P²NFFT

- **get-accuracy.py**
Compares the P²NFFT results for dipolar interaction to an exact calculation by means of direct summation
- **p2nfft_common.py**
Support methods for P²NFFT, e.g., for the re-tuning and mesh size adapting

VIII. CONCLUSION AND OUTLOOK

In this article, we provided a brief overview of the P²NFFT method for the calculation of dipolar interactions under open boundary conditions. The method is based on the fast Fourier transform for nonequispaced data (NFFT), and yields a scaling of $N \log N$ in the number of particles, as compared to the N^2 scaling of summing up the dipolar interactions directly. In order to make the NFFT algorithms applicable, the involved long ranged non-periodic functions are periodized based on a polynomial interpolation and finally approximated by trigonometric polynomials.

Furthermore, we have demonstrated the usefulness of the algorithm to magnetic soft matter by showing a sample case, namely the application to a ferrogel simulation. We have developed a tuning procedure based on a comparison to exact results obtained by direct summation, provide a simulation protocol for a magnetic gel model,

and explain how to cope with a strong change of sample shape and structure during the simulation.

In summary, we have demonstrated that the P²NFFT method is well suited for the simulation of large magnetic soft matter systems with open boundaries. However, due to the effort needed for interpolating onto a regular grid, the method will outperform direct summation for system sizes on the order of 10 000 particles and above. For the application to magnetic gels, large simulations are important, because the material properties in the bulk and surface regions may be different. Hence, the surface to bulk ratio of simulations should be comparable to that in experiments. In the future, it may be of interest to off-load part of the calculations, such as the Fourier transforms, onto an accelerator such as a graphics card. Since Fourier transforms perform very well on graphics cards, a considerable speed-up can be expected. Lastly, the P²NFFT method is also applicable to systems with mixed periodicity. Then, a different tuning scheme is required, as the exact result can no longer be calculated directly.

ACKNOWLEDGEMENTS

RW, FW, and CH are grateful for financial support from the DFG through the SPP 1681, SFB 716 and cluster of excellence Simtech EXC 310, and to the BW-Unicluster for computing resources. FN, DP and CH gratefully acknowledge support by the German Research Foundation (DFG), project PO 711/12-1 and HO 1108/25-1.

-
- [1] L. Barsi, A. Büki, D. Szabo, and M. Zrinyi. Gels with magnetic properties. *Progress in Colloid and Polymer Science*, 102:57, 1996.
 - [2] Z. Varga, J. Feher, G. Filipcsei, and M. Zrinyi. Smart nanocomposite polymer gels. *Macromolecular Symposia*, 200:93, 2003.
 - [3] D. Szabó, G. Szeghy, and M. Zrínyi. Shape transition of magnetic field sensitive polymer gels. *Macromolecules*, 31:6541–6548, September 1998.
 - [4] C. Gollwitzer, A. Turanov, M. Krekhova, G. Lattermann, I. Rehberg, and R. Richter. Measuring the deformation of a ferrogel in a homogeneous magnetic field. *Journal of Chemical Physics*, 128:164709, 2008.
 - [5] T. Mitsumata and S. Otori. Magnetic polyurethane elastomers with wide range modulation of elasticity. *Polymer Chemistry*, 2(5):1063–1067, 2011.
 - [6] S. Odenbach. Microstructure and rheology of magnetic hybrid materials. *Archive of Applied Mechanics*, 86:1–11, 2016.
 - [7] R. Ramanujan and L. Lao. The mechanical behavior of smart magnet–hydrogel composites. *Smart materials and structures*, 15(4):952, 2006.
 - [8] S. Monz, A. Tschöpe, and R. Birringer. Magnetic properties of isotropic and anisotropic CoFe₂O₄-based ferrogels and their application as torsional and rotational actuators. *Physical Review E*, 78(2):021404, 2008.
 - [9] K. Zimmermann, V. Bohm, and I. Zeidis. Vibration-driven mobile robots based on magneto-sensitive elastomers. In *Advanced Intelligent Mechatronics (AIM), 2011 IEEE/ASME International Conference on*, pages 730–735. IEEE, 2011.
 - [10] R. Weeber, M. Hermes, A. M. Schmidt, and C. Holm. Polymer architecture of magnetic gels: a review. *Journal of Physics: Condensed Matter*, 30(6):063002, 2018.
 - [11] R. Weeber, P. Kreissl, and C. Holm. Studying the field-controlled change of shape and elasticity of magnetic gels using particle-based simulations. *Archive of Applied Mechanics*, pages 1–14, 2018.
 - [12] M. Dudek, B. Grabiec, and K. Wojciechowski. Molecular dynamics simulations of auxetic ferrogel. *Reviews on Advanced Materials Science*, 14:167–173, 2007.
 - [13] R. Weeber, S. Kantorovich, and C. Holm. Deformation mechanisms in 2d magnetic gels studied by computer simulations. *Soft Matter*, 8:9923–9932, 2012.

- [14] R. Weeber, S. Kantorovich, and C. Holm. Ferrogels cross-linked by magnetic nanoparticles – deformation mechanisms in two and three dimensions studied by means of computer simulations. *Journal of Magnetism and Magnetic Materials*, 383:262–266, 2015.
- [15] R. Weeber, S. Kantorovich, and C. Holm. Ferrogels cross-linked by magnetic particles: Field-driven deformation and elasticity studied using computer simulations. *Journal of Chemical Physics*, 143, 2015.
- [16] R. Weeber and C. Holm. Interplay between particle microstructure, network topology and sample shape in magnetic gels – a molecular dynamics simulation study. *Arxiv preprint 1704.06578*, 2017.
- [17] E. Minina, P. Sánchez, C. Likos, and S. Kantorovich. The influence of the magnetic filler concentration on the properties of a microgel particle: zero-field case. *Journal of Magnetism and Magnetic Materials*, 2017.
- [18] O. V. Stolbov, Y. L. Raikher, and M. Balasoiu. Modelling of magnetodipolar striction in soft magnetic elastomers. *Soft Matter*, 7:8484–8487, 2011.
- [19] P. Cremer, H. Löwen, and A. M. Menzel. Tailoring superelasticity of soft magnetic materials. *Applied Physics Letters*, 107(17):171903, 2015.
- [20] A. Attaran, J. Brummund, and T. Wallmersperger. Development of a continuum model for ferrogels. *Journal of Intelligent Material Systems and Structures*, page 1045389X16672564, 2016.
- [21] P. Metsch, K. A. Kalina, C. Spieler, and M. Kästner. A numerical study on magnetostrictive phenomena in magnetorheological elastomers. *Computational Materials Science*, 124:364–374, 2016.
- [22] K. A. Kalina, P. Metsch, and M. Kästner. Microscale modeling and simulation of magnetorheological elastomers at finite strains: A study on the influence of mechanical preloads. *International Journal of Solids and Structures*, 102:286–296, 2016.
- [23] J. J. Cerdà, V. Ballenegger, and C. Holm. Particle-particle particle-mesh method for dipolar interactions: On error estimates and efficiency of schemes with analytical differentiation and mesh interlacing. *Journal of Chemical Physics*, 135:184110, 2011.
- [24] Y. L. Raikher and O. V. Stolbov. Magnetodeformational effect in ferrogel samples. *JMMM*, 258/259:477, 2003.
- [25] A. Y. Zubarev. On the theory of the magnetic deformation of ferrogels. *Soft Matter*, 8(11):3174–3179, 2012.
- [26] F. Nestler. An NFFT based approach to the efficient computation of dipole-dipole interactions under various periodic boundary conditions. *Appl. Numer. Math.*, 105:25–46, 2016.
- [27] M. Hofmann, F. Nestler, and M. Pippig. NFFT based Ewald summation for electrostatic systems with charges and dipoles. *Appl. Numer. Math.*, 122:39–65, 2017.
- [28] M. Bolten, F. Fahrenberger, R. Halver, F. Heber, M. Hofmann, I. Kabadshow, O. Lenz, M. Pippig, and G. Sutmann. ScaFaCoS, C subroutine library. <http://scafacos.github.com/>.
- [29] A. Arnold, F. Fahrenberger, C. Holm, O. Lenz, M. Bolten, H. Dachsels, R. Halver, I. Kabadshow, F. Gähler, F. Heber, J. Iseringhausen, M. Hofmann, M. Pippig, D. Potts, and G. Sutmann. Comparison of scalable fast methods for long-range interactions. *Physical Review E*, 88:063308, December 2013.
- [30] H. J. Limbach, A. Arnold, B. A. Mann, and C. Holm. ESPResSo – an extensible simulation package for research on soft matter systems. *Computer Physics Communications*, 174(9):704–727, May 2006.
- [31] A. Arnold, O. Lenz, S. Kesselheim, R. Weeber, F. Fahrenberger, D. Röhm, P. Košován, and C. Holm. ESPResSo 3.1 — molecular dynamics software for coarse-grained models. In M. Griebel and M. A. Schweitzer, editors, *Meshfree Methods for Partial Differential Equations VI*, volume 89 of *Lecture Notes in Computational Science and Engineering*, pages 1–23. Springer Berlin Heidelberg, 2013.
- [32] P. P. Ewald. Die Berechnung optischer und elektrostatischer Gitterpotentiale. *Ann. Phys.*, 369:253–287, 1921.
- [33] A. Dutt and V. Rokhlin. Fast Fourier transforms for nonequispaced data. *SIAM J. Sci. Stat. Comput.*, 14:1368–1393, 1993.
- [34] G. Beylkin. On the fast Fourier transform of functions with singularities. *Appl. Comput. Harmon. Anal.*, 2:363–381, 1995.
- [35] G. Steidl. A note on fast Fourier transforms for nonequispaced grids. *Adv. Comput. Math.*, 9:337–353, 1998.
- [36] A. J. W. Duijndam and M. A. Schonewille. Nonuniform fast Fourier transform. *Geophysics*, 64:539–551, 1999.
- [37] M. Pippig. PNFFT - Parallel Nonequispaced FFT software library. <https://github.com/mpip/pnfft>.
- [38] F. Nestler. Automated parameter tuning based on RMS errors for nonequispaced FFTs. *Adv. Comput. Math.*, 42(4):889–919, 2016.
- [39] M. Pippig. *Massively Parallel, Fast Fourier Transforms and Particle-Mesh Methods*. Dissertation. Universitätsverlag Chemnitz, 2015.
- [40] D. Potts and G. Steidl. Fast summation at nonequispaced knots by NFFT. *SIAM J. Sci. Comput.*, 24:2013–2037, 2003.
- [41] M. Fenn and G. Steidl. Fast NFFT based summation of radial functions. *Sampl. Theory Signal Image Process.*, 3:1–28, 2004.
- [42] F. Nestler, M. Pippig, and D. Potts. Fast Ewald summation based on NFFT with mixed periodicity. *J. Comput. Phys.*, 285:280–315, 2015.
- [43] Z. W. Wang and C. Holm. Estimate of the cutoff errors in the Ewald summation for dipolar systems. *J. Chem. Phys.*, 115:6277–6798, 2001.
- [44] J. J. Cerdà, V. Ballenegger, O. Lenz, and C. Holm. P3M algorithm for dipolar interactions. *J. Chem. Phys.*, 129:234104, 2008.
- [45] J. J. Cerdà, V. Ballenegger, and C. Holm. Particle-particle particle-mesh method for dipolar interactions: On error estimates and efficiency of schemes with analytical differentiation and mesh interlacing. *J. Chem. Phys.*, 135:184110, 2011.
- [46] J. D. Weeks, D. Chandler, and H. C. Andersen. Role of repulsive forces in determining the equilibrium structure of simple liquids. *The Journal of Chemical Physics*, 54:5237, 1971.

- [47] R. Kubo. The fluctuation-dissipation theorem. *Reports on Progress in Physics*, 29(1):255–284, 1966.