

AN OPTIMAL-TRANSPORT FINITE-PARTICLE METHOD FOR MASS DIFFUSION

A. PANDOLFI¹, L. STAINIER² AND M. ORTIZ³

ABSTRACT. We formulate a class of velocity-free finite-particle methods for mass transport problems based on a time-discrete incremental variational principle that combines entropy and the cost of particle transport, as measured by the Wasserstein metric. The incremental functional is further spatially discretized into finite particles, i. e., particles characterized by a fixed spatial profile of finite width, each carrying a fixed amount of mass. The motion of the particles is then governed by a competition between the cost of transport, that aims to keep the particles fixed, and entropy maximization, that aims to spread the particles so as to increase the entropy of the system. We show how the optimal width of the particles can be determined variationally by minimization of the governing incremental functional. Using this variational principle, we derive optimal scaling relations between the width of the particles, their number and the size of the domain. We also address matters of implementation including the acceleration of the computation of diffusive forces by exploiting the Gaussian decay of the particle profiles and by instituting fast nearest-neighbor searches. We demonstrate the robustness and versatility of the finite-particle method by means of a test problem concerned with the injection of mass into a sphere. These test results demonstrate the meshless character of the method in any spatial dimension, its ability to redistribute mass particles and follow their evolution in time, its ability to satisfy flux boundary conditions for general domains based solely on a distance function, and its robust convergence characteristics.

1. INTRODUCTION

Particle methods have attained considerable acceptance in solid and fluid mechanics, e. g. the Smoothed Particle Hydrodynamics method (SPH) [1, 2], the Material Point Method [3], the Reproducing Kernel Particle method [4], the Corrective Smoothed Particle Method [5], the Modified Smoothed Particle Hydrodynamics method [6], the Optimal-Transportation Meshfree method (OTM) [7, 8], the dislocation monopole method [9, 10], and others. Whereas the initial development of particle methods was often *ad hoc*, recently there has been increasing awareness of the synergistic connection between particle methods and optimal transport (cf., e. g., [11, 12, 13] for background), in as much as the theory of optimal transport provides a suitable mathematical foundation for formulating particle methods and,

conversely, particle methods supply a natural discretization of problems of optimal transport.

The essential structure of *flow* problems is exemplified by Benamou and Brennier’s reformulation of the compressible Euler flow problem [14, 7] in terms of two fields, the mass density and the velocity field, which are required to jointly minimize an action integral subject to a continuity equation constraint. In particular, the discretization of the problem requires the formulation of discretization rules for the mass density and for the velocity field. In the OTM method [7, 8], the mass density is discretized into point masses, or *particles*, and the velocity field is discretized by means of conforming interpolation from nodal values. The convergence properties of OTM have been established in [15, 16, 17] and applications to solid and fluid flows, including plasticity and fragmentation, have been presented in [7, 18, 19, 20, 21], among others.

Fedeli *et al.* [22] applied the OTM method to advection-diffusion problems. The connection between diffusion problems and optimal transport was elucidated in the seminal work of Jordan, Kinderlehrer and Otto [23, 24, 25]. In the approach of Fedeli *et al.* [22] the OTM two-field structure is retained, with the mass density discretized by means of point masses and the velocity field, which for diffusion problems is entropic in nature, by means of conforming interpolation. A particularly powerful property of optimal transport methods as they bear on advection-diffusion problems is that they are *geometrically exact*, in the sense that all advection and volume updates are computed exactly from the incremental transport mapping. This property is in contrast with upwinding schemes and Eulerian treatments of the volume constraint, which are inexact, costly and prone to instabilities. The geometrical structure of optimal transport (cf., e. g., [26, 13] for relevant background) is, in fact, a main source of its computational power.

In contrast to the OTM two-field representation of diffusion, so-called *blob methods* work solely with the mass density field, which is discretized into smoothed mass particles, or *blobs*, of a certain finite width [27, 28]. For diffusion problems, these finite-particle methods have the appeal of setting forth single-field numerical schemes, at great gain in simplification. However, they also raise a number of numerical analysis questions, such as the consistent variational foundation of the schemes, the enforcement of boundary conditions, the optimal selection of particle thickness and the overall accuracy and convergence of the schemes.

In the present work, we resort to a time-discretized variational formulation of mass diffusion, as in [22], but discretize the mass density by means of finite particles, or blobs, as in [27, 28], in order to formulate a class of velocity-free finite-particle methods. The immediate benefit of using finite-width particles, as opposed to Dirac deltas, is that their entropy can be computed directly from the relative entropy, or Kullback-Leibler, functional (cf., e. g., [29]). This is in contrast to mass discretization based on point masses, such as used in [22], for which the entropy functional needs to be carefully

rephrased in a weak form. The motion of the finite particles then follows as the result of a competition between entropy, as given by the Kullback-Leibler functional, and mobility, as represented by the Wasserstein distance (cf., e. g., [12]) between consecutive configurations. Thus, entropy works to disperse the particles uniformly over the domain, whereas mobility works to hinder their motion.

In Section 3, we show how this competition results in discrete Euler-Lagrange equations that, in combination with suitable quadrature rules, are explicit in the positions of the particles. A numerical issue that is central to finite-particle schemes such as considered here concerns the efficient choice of particle width. In the present work, we optimize the particle widths *variationally*, i. e., we determine them by minimizing the same incremental functional that determines the motion of the particles. We show that this optimality criterion results in particle widths on an intermediate scale between the mean-free-distance between particles and the size of the domain, as expected. Another critical issue concerns the formulation of boundary conditions. We show that flux boundary conditions can be conveniently formulated in terms of injection/ejection of particles into/from the domain. We also show that zero-flux conditions can be formulated by means of a barrier potential based solely on a distance function to the domain, which greatly facilitates consideration of general domains defined, e. g., by means of CAD tools.

In Section 4, a number of implementational issues are discussed, including the efficient calculation of particle widths and the use of fast searches to accelerate the calculation of pairwise interaction forces. These issues become critical when the number of particles becomes large and need to be addressed carefully in order for the calculations to be feasible. In Section 5, we illustrate the scope and properties of the method by means of a test problem concerned with the injection of mass into a sphere. The ability to inject finite particles into the domain and follow the diffusion-driven trajectories in time, including convergence to a uniform steady-state density after the injection ends, is quite remarkable. The weak convergence of the scheme is monitored through the moment of inertia of the particles at steady state. The test is indicative of robust convergence with respect to the number of particles.

2. PROBLEM DEFINITION

We consider the advection-diffusion initial-boundary-value problem

$$\begin{aligned}
 (1a) \quad & \partial_t \rho + \nabla \cdot (\rho u) = \kappa \Delta \rho + \rho s, & \text{in } \Omega \times [0, T], \\
 (1b) \quad & \rho = g, & \text{on } \Gamma_D \times [0, T], \\
 (1c) \quad & (\kappa \nabla \rho - \rho u) \cdot n = f, & \text{on } \Gamma_N \times [0, T], \\
 (1d) \quad & \rho(x, 0) = \rho_0(x), & \text{in } \Omega,
 \end{aligned}$$

where ρ is the unknown density or concentration, ρ_0 is its initial value, u is a given advection velocity field, s is a source density, g is a prescribed density, f is a prescribed outward mass flux, κ is a diffusion coefficient, Ω is a bounded domain in \mathbb{R}^d , $\Gamma = \partial\Omega$ is the boundary of Ω , n is the outward unit normal at the boundary, Γ_D is the Dirichlet boundary and Γ_N the Neumann boundary, with $\Gamma_D \cap \Gamma_N = \emptyset$ and $\Gamma_D \cup \Gamma_N = \Gamma$.

Problem (1a) can be equivalently reformulated as the transport problem

$$(2a) \quad \partial_t \rho + \nabla \cdot (\rho v) = \rho s,$$

$$(2b) \quad \rho v = \rho u - \kappa \nabla \rho,$$

where v is a velocity field that results from the combined effect of advection and diffusion, eq. (2a) is the continuity equation of Eulerian continuum mechanics and (2b) is a mobility law combining the effects of advection and diffusion. The corresponding Lagrangian formulation is

$$(3a) \quad \rho(\varphi(x, t), t) = \frac{\rho_0(x)}{\det \nabla \varphi(x, t)} \exp \left(\int_0^t s(\varphi(x, \tau), \tau) d\tau \right)$$

$$(3b) \quad \partial_t \varphi(x, t) = v(\varphi(x, t), t),$$

where $\varphi : \Omega \times [0, T] \rightarrow \Omega$ is the transport map. The mass density

$$(4) \quad \rho_0(x) \exp \left(\int_0^t s(\varphi(x, \tau), \tau) d\tau \right)$$

in (3a) gives the initial mass density at material point x adjusted for the mass inserted or removed by the sources. Eq. (3a) states that $\rho(x, t)$ is the push-forward of (4) by the transport map $\varphi(\cdot, t)$ whereas eq. (3b) relates the transport map to the particle velocity.

Eqs. (3a) and (3b) may be thought as jointly defining an evolution for both the measure ρdx and the transport map φ . In particular, we note that this extension to measures requires consideration of the transport map φ as an additional unknown of the problem.

3. APPROXIMATION

We aim to approximate the solutions of (1) variationally through a combination of time discretization followed by spatial discretization of the mass density into particles. To this end, we resort to an incremental variational principle that sets forth an optimality criterion for the approximations.

3.1. Fractional steps. We begin by discretizing time into a sequence of discrete times $t_0, t_1, \dots, t_k, t_{k+1} \dots$ and approximating the corresponding mass densities $\rho_0, \rho_1, \dots, \rho_k, \rho_{k+1} \dots$. To this end, we exploit the additive structure of (1a) to decompose incremental updates by the method of fractional steps [22]. The advective fractional step is governed by the pure advection equation

$$(5) \quad \partial_t \rho + \nabla \cdot (\rho u) = 0,$$

which follows formally from (2a) by setting $\kappa = 0$ and $s = 0$ and it be conveniently solved exactly following the approach described in [22]. This exact geometrical treatment of advection sets particle methods apart from upwinding methods and confer them great advantage and strength. The source fractional step

$$(6) \quad \partial_t \rho = \rho s,$$

is likewise amenable to an exact treatment, cf. (4). Therefore, in the remainder of the section we focus on the pure diffusional fractional step, obtained by formally dropping the advection and source terms in (2a). Further detail on the convection step and examples of application may be found in [22].

3.2. Variational formulation of the diffusional fractional step. Following [23, 24, 25, 22], we characterize the incremental evolution of the density by means of the Jordan-Kinderlehrer-Otto (JKO) functional

$$(7) \quad F(\rho_{k+1}) = \frac{1}{2} \frac{d_W^2(\rho_k, \rho_{k+1})}{t_{k+1} - t_k} + \int_{\Omega} \kappa \rho_{k+1} \log \left(\frac{\rho_{k+1}}{\rho_{\infty}} \right) dx \rightarrow \inf!,$$

where d_W is the Wasserstein distance (cf., e. g., Appendix A of [22]) and ρ_{∞} is a reference density. For definiteness, we specifically choose $\rho_{\infty} = M/|\Omega|$, where M is the total mass of the system at steady state and $|\Omega|$ is the volume of the domain.

It is verified that this formulation indeed supplies a time discretization of the mobility law (2b) in the limit of zero advection. Thus, a straightforward calculation (cf., e. g., Appendix A of [22]) gives the Euler-Lagrange equation

$$(8) \quad \rho_{k+1}(x) \frac{x - \varphi_{k+1 \rightarrow k}(x)}{t_{k+1} - t_k} = -\kappa \nabla \rho_{k+1}(x),$$

where $\varphi_{k+1 \rightarrow k}$ is the incremental transport map, $\varphi_{k \rightarrow k+1} = \varphi_{k+1 \rightarrow k}^{-1}$ and ρ_{k+1} follows from $\varphi_{k+1 \rightarrow k}$ through the geometrically exact pushforward operation

$$(9) \quad \rho_{k+1} \circ \varphi_{k \rightarrow k+1} = \rho_k / \det(\nabla \varphi_{k \rightarrow k+1}).$$

We verify that (8) is indeed a time discretization of (2b). A rigorous analysis of convergence with respect to the time step may be found in [23, 24].

3.3. Spatial discretization. From the standpoint of weak convergence of measures, it would be natural to approximate $\rho(x, t)$ by concentrating mass on a finite collection of points $\{x_p(t)\}_{p=1}^n$, i. e., by setting

$$(10) \quad \rho(x, t) \sim \sum_{p=1}^n m_p \delta(x - x_p(t)),$$

where δ is the Dirac delta and m_p is the mass carried by particle p . However, the functional (7) is not defined for trial densities of this type. In [22], this difficulty is circumvented by replacing the incremental functional (7) with an alternative weak statement, modeled after (2), in which the mass density appears linearly and can therefore be approximated by the *ansatz* (10).

In the present work, following Carrillo *et al.* [27, 28] we instead spread, or 'fatten', the particles over a finite width, to be determined. The corresponding approximation is

$$(11) \quad \rho(x, t) \sim \sum_{p=1}^n m_p \varphi_p(x - x_p(t), t),$$

where the function $\varphi_p(x - x_p(t); t)$ represents the time-dependent profile of particle x_p . The form of the profile $\varphi_p(x - x_p(t), t)$ may be parametrized within a convenient class of functions and the explicit dependence on t then represents the time-dependence of the parameters. The discrete system is then obtained by inserting the *ansatz* (11) into the functional (7) and rendering the resulting discrete functional stationary with respect to all degrees of freedom.

3.4. Gaussian particle profile. We specifically choose a Gaussian particle profile of the form

$$(12) \quad \varphi_p(x - x_p(t), t) = \left(\frac{\beta_p(t)}{\pi} \right)^{d/2} e^{-\beta_p(t) \|x - x_p(t)\|^2} := N(x - x_p(t), \beta_p(t)),$$

normalized so that

$$(13) \quad \int_{\mathbb{R}^d} N(x - x_p(t), \beta_p(t)) dx = 1.$$

We note that $N(x - x_p(t), \beta_p(t))$ has dimensions of an inverse volume, carries unit mass and has a width set by the parameter $\beta_p(t)$.

3.5. Wasserstein distance. The Gaussian profile of the particles and the absence of mass redistribution renders the calculation of the Wasserstein distance $d_W^2(\rho_k, \rho_{k+1})$ in (7) straightforward. Thus, for particles of constant mass $d_W^2(\rho_k, \rho_{k+1})$ simply records the total cost of transportation of all the particles from their initial to their final positions, namely,

$$(14) \quad d_W^2(\rho_k, \rho_{k+1}) \sim \sum_{p=1}^n d_W^2(m_p N_{p,k}(x), m_p N_{p,k+1}),$$

where we write $N_{p,k}(x) := N(x - x_{p,k}, \beta_{p,k})$, $N_{p,k+1}(x) := N(x - x_{p,k+1}, \beta_{p,k+1})$, and neglect the effect of the clipping of the tails of the Gaussians by a finite domain. Conveniently, the Wasserstein distance between Gaussian measures is known explicitly [30, 31, 32, 33], whence we obtain

$$(15) \quad d_W^2(m_p N_{p,k}(x), m_p N_{p,k+1}) \sim \left(\|x_{p,k+1} - x_{p,k}\|^2 + |\beta_{p,k+1}^{-1/2} - \beta_{p,k}^{-1/2}|^2 \right) m_p.$$

We note that $d_W^2(\rho_k, \rho_{k+1})$ quantifies not only the extent of particle displacement during a time step but also the change in the width of the particles, if any.

3.6. Hermitian quadrature. The choice of Gaussian profiles also facilitates the computation of integrals by numerical quadrature. Thus, for densities of the form (11) and (12) we have

$$(16) \quad \int_{\Omega} \kappa \rho_{k+1} \log \rho_{k+1} dx = \sum_{p=1}^n \int_{\Omega} \kappa \log \left(\frac{\rho_{k+1}(x)}{\rho_{\infty}} \right) \times \\ \left(\frac{\beta_{p,k+1}}{\pi} \right)^{d/2} \exp \left(-\beta_{p,k+1} \|x - x_{p,k+1}\|^2 \right) m_p dx.$$

The Gaussian weight in the integral suggests approximation by Hermitian quadrature [34]. The one-point quadrature rule is particularly simple and takes the form

$$(17) \quad \int_{\Omega} \kappa \rho_{k+1}(x) \log \left(\frac{\rho_{k+1}(x)}{\rho_{\infty}} \right) dx \sim \sum_{p=1}^n \kappa \log \left(\frac{\rho_{k+1}(x_{p,k+1})}{\rho_{\infty}} \right) m_p.$$

Higher accuracy can be achieved by recourse to quadrature rules of higher order, albeit at increased computational cost and complexity.

3.7. The discrete problem. Combining the preceding approximations, the reduced JKO functional takes the form

$$(18) \quad F(\rho_{k+1}) \sim \sum_{p=1}^n \frac{1}{2} \frac{\|x_{p,k+1} - x_{p,k}\|^2 + (\beta_{p,k+1}^{-1/2} - \beta_{p,k}^{-1/2})^2}{t_{k+1} - t_k} m_p \\ + \sum_{p=1}^n \kappa \log \left(\frac{\rho_{k+1}(x_{p,k+1})}{\rho_{\infty}} \right) m_p \rightarrow \min!$$

with

$$(19) \quad \rho_{k+1}(x) := \sum_{p=1}^n \left(\frac{\beta_{p,k+1}}{\pi} \right)^{d/2} e^{-\beta_{p,k+1} \|x - x_{p,k+1}\|^2} m_p.$$

We note that the incremental variational principle (18) determines both the updated position of the particles and their optimal widths. The corresponding Euler-Lagrange equations are

$$(20) \quad \frac{\partial F}{\partial x_{r,k+1}} = 0, \quad \frac{\partial F}{\partial \beta_{r,k+1}} = 0,$$

which evaluate to

$$(21a) \quad m_r \frac{x_{r,k+1} - x_{r,k}}{t_{k+1} - t_k} + \frac{\partial}{\partial x_{r,k+1}} \sum_{p=1}^n \kappa \log \left(\frac{\rho_{k+1}(x_{p,k+1})}{\rho_{\infty}} \right) m_p = 0,$$

$$(21b) \quad -\frac{m_r}{2} \frac{\beta_{r,k+1}^{-1/2} - \beta_{r,k}^{-1/2}}{\beta_{r,k+1}^{3/2} (t_{k+1} - t_k)} + \frac{\partial}{\partial \beta_{r,k+1}} \sum_{p=1}^n \kappa \log \left(\frac{\rho_{k+1}(x_{p,k+1})}{\rho_{\infty}} \right) m_p = 0.$$

We note that these equations represent an implicit forward-Euler update for the positions and widths of the particles.

3.8. Boundary conditions. Neumann, or flux, boundary conditions can be readily implemented within the particle representation just outlined. Thus, a prescribed inward flux can be built into the calculations by injecting particles through the boundary of the domain so as to achieve the desired mass input rate per unit area. A prescribed outward flux can be implemented either by removing particles crossing the boundary, or by inserting particles with negative mass [35], or a combination of both, again so as to achieve the desired mass input rate per unit area.

A zero flux condition can be implemented by means of a barrier potential, i. e., a continuous differentiable function $V(x)$ that is zero on the domain, including its boundary, and grows steeply away from it. For instance, in calculations we choose

$$(22) \quad V(x) = \frac{C}{2} \text{dist}^2(x, \Omega),$$

where

$$(23) \quad \text{dist}(x, \Omega) = \min\{\|x - y\| : y \in \Omega\}$$

is the distance function to Ω and the constant C is to be calibrated. The incremental problem (18) is then augmented to

$$(24) \quad F(\rho_{k+1}) + \sum_{p=1}^n V(x_p) m_p \rightarrow \min!$$

Evidently, the effect of the barrier potential is to apply a restoring diffusional force $DV(x_p) m_p$ to particles x_p that wander out of the domain, causing them to return to it.

Conveniently, the distance function $d(x, \Omega)$ is a functionality that is provided by most Computer-Aided Design (CAD) codes, which facilitates the implementation of the finite-particle method based on CAD models of the domain of analysis.

A number of extensions of optimal transport and the JKO functional have been proposed for problems with Dirichlet boundary conditions that can be applied to particle methods [36, 37, 38, 35]. A possible simple implementation of Dirichlet boundary conditions consists of fixing particles on the Dirichlet boundary at the desired areal mass density. However, such extensions are beyond the scope of this paper and will not be addressed further.

4. NUMERICAL IMPLEMENTATION

In this section, we discuss ways in which the discrete problem (21) can be simplified, and the solution accelerated, without significantly compromising accuracy or convergence.

4.1. Computation of diffusive forces. It is possible to derive a particular expression for the diffusive forces in (21a) within the accuracy of the Hermitian quadrature rule. To this end, we write

$$(25) \quad \begin{aligned} & \frac{\partial}{\partial x_{r,k+1}} \sum_{p=1}^n \kappa \log \left(\frac{\rho_{k+1}(x_{p,k+1})}{\rho_\infty} \right) m_p = \\ & \sum_{p=1}^n \kappa \left[\frac{\partial}{\partial x_{r,k+1}} \log \left(\frac{\rho_{k+1}(x)}{\rho_\infty} \right) \right]_{x=x_{p,k+1}} m_p + \kappa \nabla \rho_{k+1}(x_{r,k+1}) m_r, \end{aligned}$$

with $\rho_{k+1}(x)$ as in (19). Within the accuracy of one-point Hermitian interpolation, we have

$$(26) \quad \begin{aligned} & \sum_{p=1}^n \kappa \left[\frac{\partial}{\partial x_{r,k+1}} \log \left(\frac{\rho_{k+1}(x)}{\rho_\infty} \right) \right]_{x=x_{p,k+1}} m_p = \\ & \sum_{p=1}^n \kappa \left[\frac{1}{\rho_{k+1}(x)} \frac{\partial}{\partial x_{r,k+1}} \rho_{k+1}(x) \right]_{x=x_{p,k+1}} m_p \sim \\ & \int_{\Omega} \kappa \frac{\partial}{\partial x_{r,k+1}} \rho_{k+1}(x) dx = \frac{\partial}{\partial x_{r,k+1}} \int_{\Omega} \kappa \rho_{k+1}(x) dx \sim 0. \end{aligned}$$

To within this approximation, (25) therefore simplifies to the particularly simple expression

$$(27) \quad \frac{\partial}{\partial x_{r,k+1}} \sum_{p=1}^n \kappa \log \left(\frac{\rho_{k+1}(x_{p,k+1})}{\rho_\infty} \right) m_p \sim \kappa \nabla \rho_{k+1}(x_{r,k+1}) m_r,$$

which we use in all calculations.

4.2. Uniform particle width. Considerable simplification is achieved if the width of the particles is assumed to be uniform, i. e., $\beta_{r,k+1} = \beta_{k+1}$. Inserting this assumption into (18) and rendering the functional stationary with respect to β_{k+1} , (21b) simplifies to the single equation

$$(28) \quad -\frac{M}{2} \frac{\beta_{k+1}^{-1/2} - \beta_k^{-1/2}}{\beta_{k+1}^{3/2} (t_{k+1} - t_k)} + \frac{\partial}{\partial \beta_{k+1}} \sum_{p=1}^n \kappa \log \left(\frac{\rho_{k+1}(x_{p,k+1})}{\rho_\infty} \right) m_p = 0,$$

where $M = \sum_{r=1}^n m_r$ is the total mass of the system and

$$(29) \quad \rho_{k+1}(x) := \sum_{p=1}^n \left(\frac{\beta_{k+1}}{\pi} \right)^{d/2} e^{-\beta_{k+1} \|x - x_{p,k+1}\|^2} m_p.$$

in place of (19). Evaluating the derivative in (28) using (29) yields

$$(30) \quad \frac{\partial}{\partial \beta_{k+1}} \sum_{p=1}^n \kappa \log \left(\frac{\rho_{k+1}(x_{p,k+1})}{\rho_\infty} \right) m_p = M \kappa \frac{d}{2} \frac{1}{\beta_{k+1}} - \sum_{p=1}^n \kappa \left(\frac{\sum_{q=1}^n \|x_{p,k+1} - x_{q,k+1}\|^2 e^{-\beta_{k+1} \|x_{p,k+1} - x_{q,k+1}\|^2} m_q}{\sum_{q=1}^n e^{-\beta_{k+1} \|x_{p,k+1} - x_{q,k+1}\|^2} m_q} \right) m_p.$$

A further simplification is obtained if we assume that the particle widths are at steady state at all times. Under this assumption, (30) further reduces to the algebraic equation

$$(31) \quad M \frac{d}{2} \frac{1}{\beta_{k+1}} = \sum_{p=1}^n \left(\frac{\sum_{q=1}^n \|x_{p,k+1} - x_{q,k+1}\|^2 e^{-\beta_{k+1} \|x_{p,k+1} - x_{q,k+1}\|^2} m_q}{\sum_{q=1}^n e^{-\beta_{k+1} \|x_{p,k+1} - x_{q,k+1}\|^2} m_q} \right) m_p,$$

which gives the instantaneous width of the particles as a function of their positions.

4.3. Optimal particle width at steady state. Suppose that, at steady state, the mass density attains the constant value $\rho_\infty = M/|\Omega|$, where M is the total mass and $|\Omega|$ the volume of the domain, assumed to be finite. Suppose further that this steady state is approximated by a collection of n particles of uniform width,

$$(32) \quad \rho(x) := \sum_{p=1}^n \left(\frac{\beta}{\pi} \right)^{d/2} e^{-\beta \|x - x_p\|^2} m_p,$$

as in (19). The corresponding relative entropy then follows as

$$(33) \quad F_\infty(\beta) = \int_\Omega \kappa \rho(x) \log \left(\frac{\rho(x)}{\rho_\infty} \right) dx,$$

to be minimized with respect to β . To this end, we rewrite (33) as

$$(34) \quad F_\infty(\beta) = \int_{\mathbb{R}^d} \kappa \rho(x) \log \left(\frac{\rho(x)}{\rho_\infty} \right) dx - \int_{\mathbb{R}^d \setminus \Omega} \kappa \rho(x) \log \left(\frac{\rho(x)}{\rho_\infty} \right) dx,$$

and proceed to estimate each term in turn. We note that in the first term in (34) the tails of the finite particles are allowed to extend outside the domain, as tacitly assumed in the quadrature rule (17). The second term in (34) then corrects for the incurred mass leakoff.

For large β the leakoff mass is confined to a narrow layer of thickness $1/\sqrt{\beta}$ in the vicinity of the boundary $\partial\Omega$ of the domain. For sufficiently regular domains, therefore,

$$(35) \quad \int_{\mathbb{R}^d \setminus \Omega} \kappa \rho(x) \log \left(\frac{\rho(x)}{\rho_\infty} \right) dx \sim -\kappa \rho_\infty |\partial\Omega| \beta^{-1/2},$$

where $|\partial\Omega|$ is the area of the domain boundary. We note that the effect of mass leakoff vanishes in the limit of $\beta \rightarrow +\infty$, as expected. To estimate the first term of (34), we partition the domain Ω into n subdomains of volume $|\Omega|/n$, each containing one particle. Again, for sufficiently large β , we have $\log(\rho(x)) \sim -\beta\|x - x_p\|^2$, in the area containing particle x_p , and we compute

$$(36) \quad \int_{\mathbb{R}^d} \kappa \rho(x) \log\left(\frac{\rho(x)}{\rho_\infty}\right) dx \sim -\kappa \rho_\infty |\Omega| \beta \left(\frac{|\Omega|}{n}\right)^{2/d}.$$

Inserting these estimates into (34) and minimizing, we obtain

$$(37) \quad \beta \sim \left(\frac{|\partial\Omega|}{|\Omega|^{(d+2)/d}}\right)^{2/3} N^{4/3d},$$

modulo a multiplicative constant, to be calibrated. In (37), the first factor is of the form R^{-2} , where R is an effective radius of the domain. The second factor defines a precise scaling $\beta^{-1/2} \sim N^{-2/3d}$ for the particle width with respect to the number of particles. Remarkably, the optimal particle size takes values on a scale intermediate between the size of the particle neighborhoods, $(|\Omega|/n)^{1/d}$ and the size of the domain $|\Omega|^{1/d}$.

4.4. Local-neighborhood searches. A naïve evaluation of the diffusive forces (27) is an $O(n^2)$ operation which, for large n , becomes the main computational bottleneck. The complexity of the operation can be reduced—and the calculations accelerated—by noting that, due to the Gaussian decay of the particle profiles, points $x_{p,k+1}$ in the sum (19) for the calculation of $\nabla \rho_{k+1}(x_{r,k+1})$ at distances much greater than $1/\sqrt{\beta_{r,k+1}}$ from the $x_{r,k+1}$ contribute negligibly to the sum. Therefore, the sum can be restricted to a local neighborhood of $x_{r,k+1}$ on the scale of $1/\sqrt{\beta_{r,k+1}}$ without appreciable loss of accuracy.

The rate-limiting step in this approach is the operation of searching for the local neighborhoods as they evolve during in calculations as a result of the relative shuffling of the particles. Numerous techniques have been developed for solving the local-neighborhood problem. One approach makes use of space partitioning algorithms where the branch and bound method [39] is applied. Early examples in this class are the k -d tree [40, 41] and the k -means [42] algorithms, with many variants thereof. A deficiency of exact methods is that they cannot guarantee logarithmic search time [43, 44]. Conveniently, for many applications approximate nearest-neighbor (ANN) algorithms [45, 46], which trade off accuracy for efficiency, suffice. More recent work has focused on proximity graph-based methods [47, 48], which are well-suited to high-dimensionality problems arising in image recognition, computational linguistics, product recommendation, and others.

The best choice of search algorithm depends on data-related properties such as dimension, data-set size, correlations, density distribution, and others, and performance metrics such as query time, accuracy, query workload, building time and memory usage. In this work, we employ a simple *ad hoc* algorithm based on a regular one-level subdivision of the computational

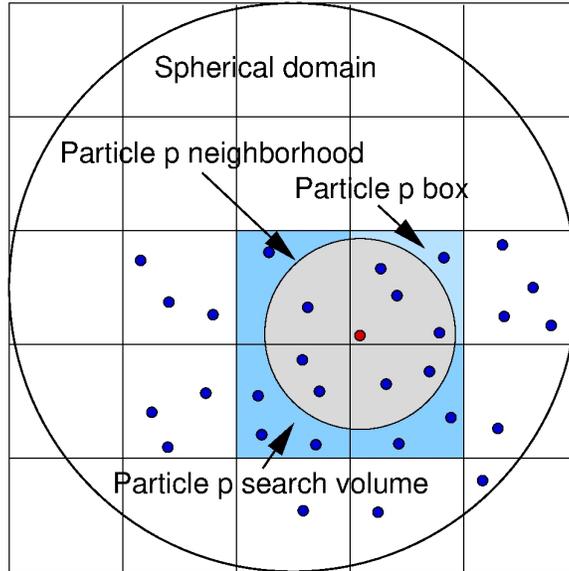


FIGURE 1. Schematic of the search algorithm with reference to the problem of an injected sphere. The spherical domain is subdivided in same-size cubic cells, or boxes, and each particle is assigned to a cell according to its current position. A cell-connectivity array lists the particles contained in each cell. For each particle, a second nearest-neighbor array records the particles contained within a spherical neighborhood of radius $\sim 1/\sqrt{\beta}$. At every time step, each particle interacts only with the particles in its neighborhood. Periodically, the nearest-neighbor array is updated by means of searches restricted to the cells that intersect the neighborhood of each particle.

domain into cubic cells, or boxes, Fig. 1. For simplicity, we consider the case of uniform and constant β pegged, e. g., to the estimate (37). In this case, the size of cells is taken as a multiple of $1/\sqrt{\beta}$, with factors to be calibrated numerically. Thus, the algorithm requires two connectivity arrays. At initialization, as well as for any new particle entering the domain as a result of a prescribed flux, the particles are assigned to a specific cell according to their coordinates. Each particle is outfitted with the connectivity list of interacting particles, and each cell is outfitted with the list of the particles contained in the cell. In shared memory platforms, the arrays may be conveniently implemented as C++ STL containers. The algorithm keeps the cell array up-to-date by transferring the particles between cells at every time step as needed. By contrast, the particle connectivity lists are updated only after a prescribed number of time steps. The update is carried out by searching only the particles contained in the cells intersected by the spherical neighborhood of the query particle.

5. EXAMPLE: MASS INJECTION INTO A SPHERE

We demonstrate the properties of the finite-particle scheme defined in the foregoing by means of a test case concerned with the injection of mass into a spherical domain through a small orifice in its surface.

5.1. Problem definition. We consider a spherical domain of radius R , initially empty, which is filled with mass injected through a small orifice of area A at the North pole at constant mass flux f during an injection time interval $[0, T_0]$. Subsequently, the particles diffuse freely, eventually filling the domain and attaining a uniform steady-state density $\rho_\infty = M/|\Omega|$. The total injected mass $M = fAT_0$ is divided into n particles of equal mass $m = M/n$, each inserted at time intervals of T_0/n to a position below the North pole offset by a distance d from the boundary. The velocity of injection of the particles is, therefore, $v = d/(T_0/n)$. The problem data (dimensionless units) are collected in table 1. To analyze the performance and the convergence properties of the method, we consider a total diffusion time $T > T_0$ and numbers of particles ranging from a minimum of 125 to a maximum of 128000. The parameters used in calculations are listed in Table 1.

R	M	κ	T	T_0	v
1	200	1	120	25	1

TABLE 1. Parameters used in the sphere-injection test calculations.

The discrete incremental problem (21) is solved by a sequential Runge-Kutta four-point scheme (RK4) for small systems, $n \leq 8,000$, and a parallel implementation of forward-Euler on 1 to 32 cores for larger systems. The diffusional forces are computed as in (27). For simplicity, the particle width $1/\sqrt{\beta}$ is assumed to be equal for all particles and constant in time, with value given by the scaling relation (37), calibrated as $\beta = 10n^{4/9}/R^2$, as function of the number of particles. The subdivision scheme described in Section 4.4 is then used with tolerance 0.01 to accelerate the calculations.

Owing to the explicit nature of the calculations, the stable time step is constrained by the diffusional CFL condition $\Delta t \sim \Delta x^2/\kappa$, where we take $\Delta x \sim (|\Omega|/n)^{1/d}$ as an estimate of the distance between particles. Conveniently, the stable time step is always smaller than the injection time step. We additionally enforce the zero-flux boundary condition on the surface of the sphere by means of the barrier potential (22), with the highest constant C that does not decrease the stable time step.

5.2. Particle flow visualizations. The evolution of the particle distribution during and after injection is visualized by a sequence of successive snapshots. Fig. 2 shows the three-dimensional distributions of the centers of the particles at the end of the injection time T_0 for four different values

of n . Every particle is colored according to the number of particles contained in its $1/\sqrt{\beta}$ -neighborhood and its size is related to the actual mass m . The distribution of particles is characterized by short-range disorder, or randomness, and long-range order, or smoothness, and a general trend towards convergence is evident from the figures. As dictated by optimal scaling, the particle width increases sublinearly with the number of particles and the size of the particle neighborhoods also increases concomitantly, up to a value in the order of 50 particles for $n = 32,000$. The value of the variational principle in elucidating these intricate tradeoffs is remarkable.

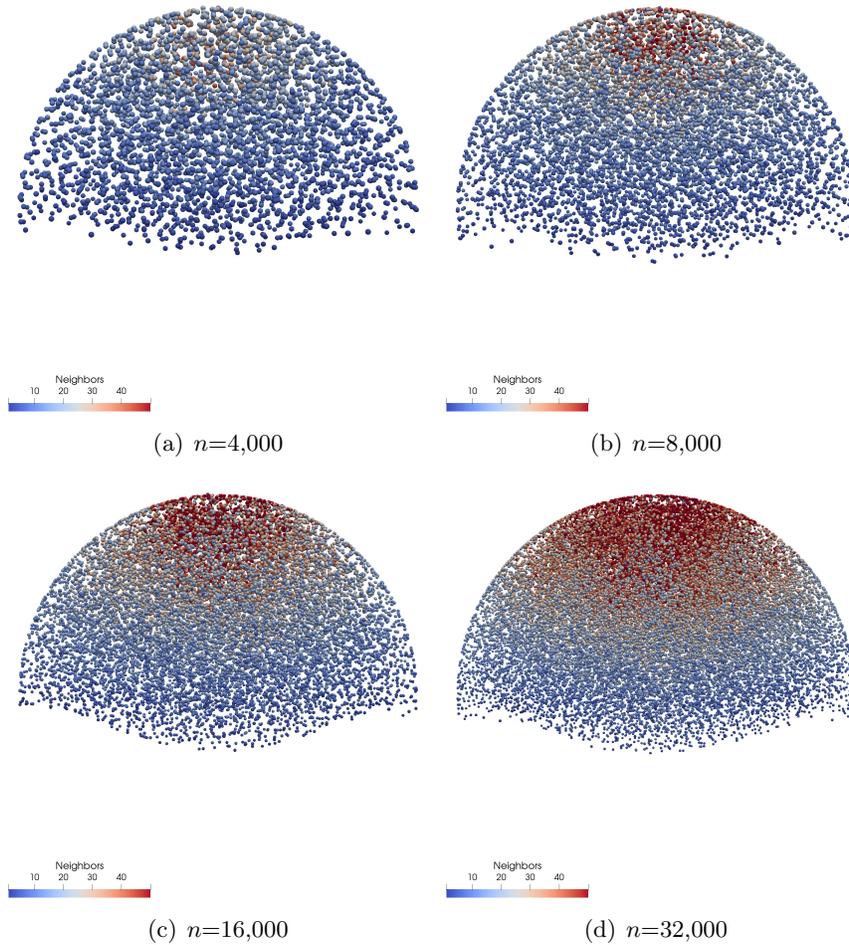


FIGURE 2. Visualization of the 3D distribution of the centers of the particles at the end of the injection phase T_0 for different number of particles. The analysis adopts a constant value of β pegged through the optimal scaling relation to the number of particles n . The color of each particle relates to the number of particles in its neighborhood. The size of the particles is related to their mass $m = M/n$.

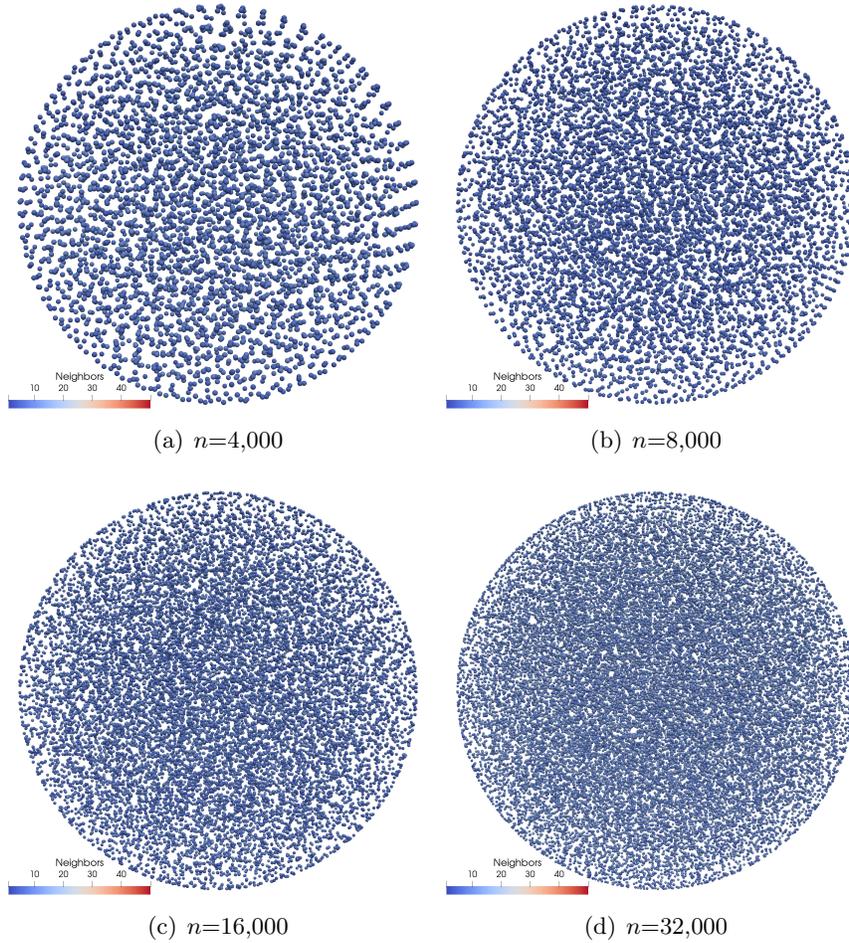


FIGURE 3. Visualization of the 3D distribution of the centers of the particles at the end of the simulation time T for different number of particles. The analysis adopts a constant value of β pegged through the optimal scaling relation to the number of particles n . The color of each particle relates to the number of particles in its neighborhood. The size of the particles is related to their mass $m = M/n$.

Fig. 3 shows the same distributions of the centers of the particles at the end of the simulation time T for different values of n . We recall that the slowest relaxation time of the diffusional problem scales as $R^2/\kappa = 1$. At $T = 120$, we may therefore expect the system to be at steady state. Indeed, an ostensibly uniform steady-state density is clear from the figures, as expected. The uniformity of the steady-state density is borne out by a different view of the system, Fig. 4, which depicts a meridional section of the sphere in the form of a disk of width 0.2. The long-term stability of the finite-particle scheme is again remarkable.

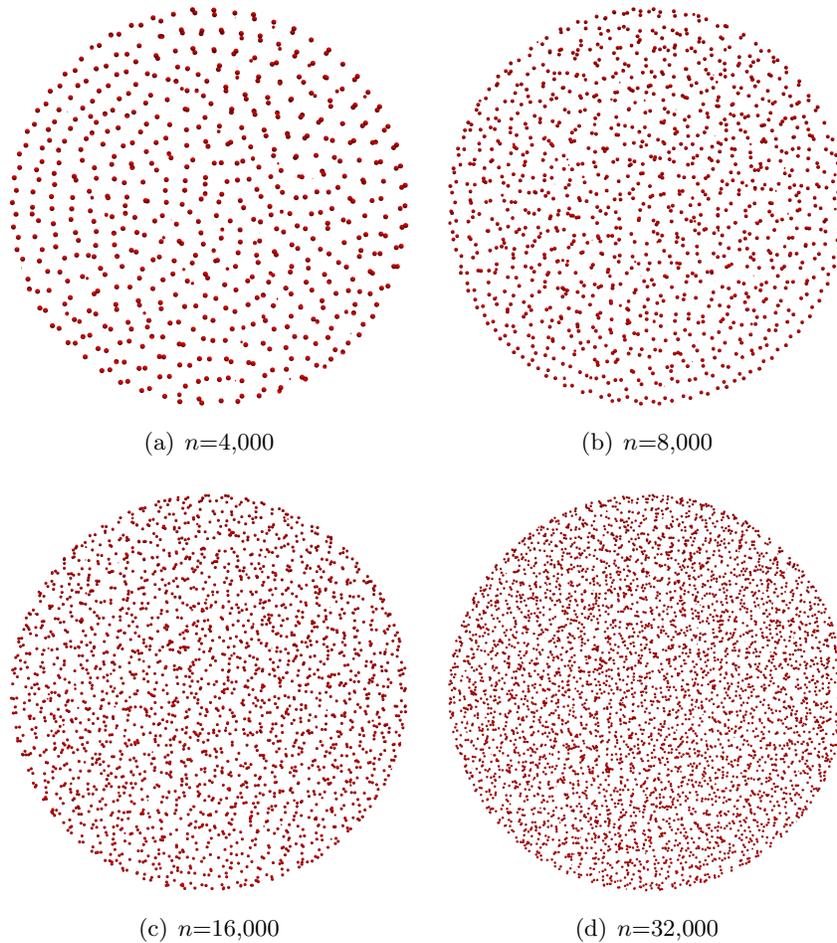


FIGURE 4. Visualization of the distribution of the particles at the end of the analysis for different number of particles, $C_3 = 10$. The image shows all the particles contained within a meridian disk-like slice of the sphere, 0.2 thick, taken at the maximum diameter. The particles have the same size, but the associated mass is different in the four cases.

A different type of visualization, which supplies a sense of the motion of the particles, consists of simulating a long-exposure picture of the flow showing the particle streamers, or paths. In Fig. 5, these streamers are shown for $n = 1,000$ and $4,000$ only, since larger numbers of particles result in excessive clutter. It is noteworthy from the figures how the particles spread out from the injection point and gradually fill the sphere. In particular, the outermost particles trace trajectories tangential to the spherical surface of the domain, as enforced by the barrier potential, thus abiding by the zero-flux boundary condition.

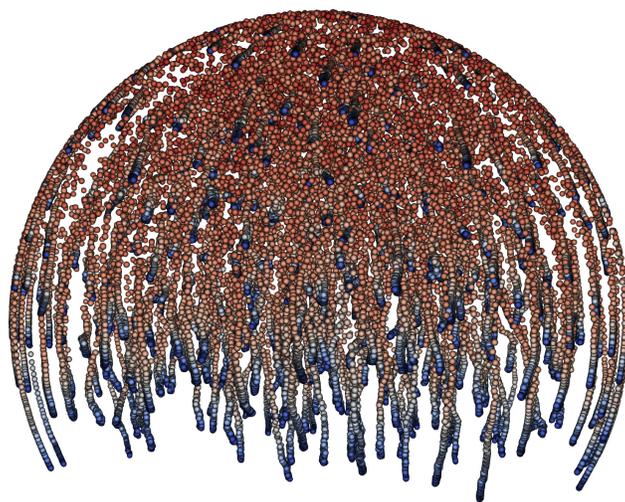
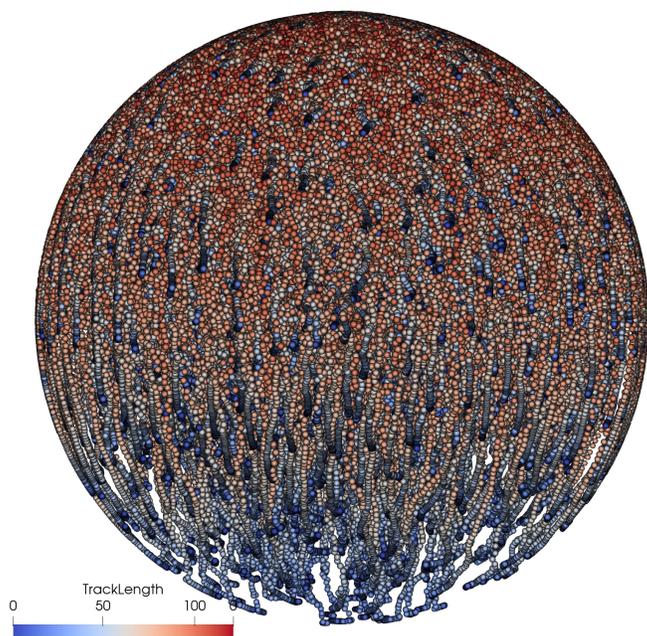
(a) $n=1000$ (b) $n=4000$

FIGURE 5. Long-exposure picture of the flow showing the particle streamers, or paths. The particles fill the domain following injection and trace trajectories tangential to the spherical surface of the domain, in accordance with the zero-flux boundary condition.

5.3. Analysis of convergence. We assess the (weak) convergence of the method by computing the polar moment of inertia of the distribution at steady-state,

$$(38) \quad J_G = \frac{3}{5}MR^2, \quad M = \frac{4}{3}\pi R^3\rho_\infty.$$

The approximate value from the particle discretization is

$$(39) \quad J_G = \sum_{p=1}^n J_p,$$

where

$$(40) \quad J_p = m_p \int_{\Omega} r^2 N(x - x_p; \beta) dx = m_p \left(r_p^2 + \frac{2}{3\beta} \right)$$

and we write $r_p = \|x_p\|$.

The time evolution of the numerical polar moment of inertia normalized with respect to the steady-state value (38) is shown in Fig. 6 as a function of the number of particles. The initial injection phase and the subsequent diffusion phase are clearly evident in the figure. The polar moment of inertia overshoots the steady state value at the end of the injection phase and then converges to the steady state from above in at large times, as expected. A general trend towards convergence of the time evolution—and towards the steady-state a large times—with respect to the number of particles is also clearly visible in the figure.

The relative error of the numerical polar moment of inertia at steady-state is plotted in Fig. 7 as a function of the number of particles. This error effectively quantifies the spatial accuracy of the finite-particle discretization, including the clipping of the particle tail profiles at the boundary, for a random Poisson distribution of particle centers. As expected (cf., e. g., [49, 17]), the plots show a clear trend towards convergence at a rate of convergence of 1/4.

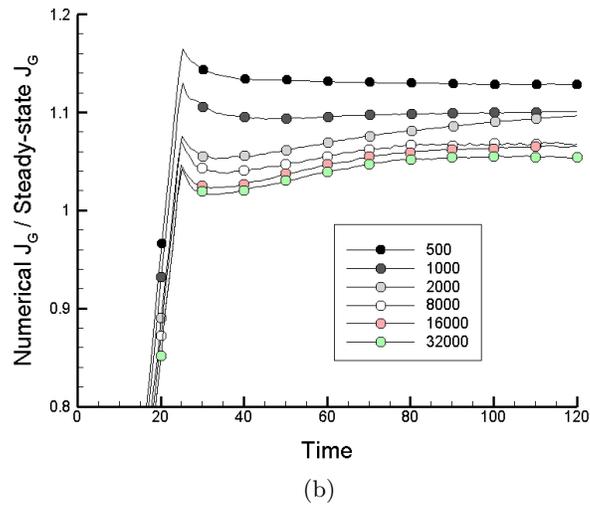
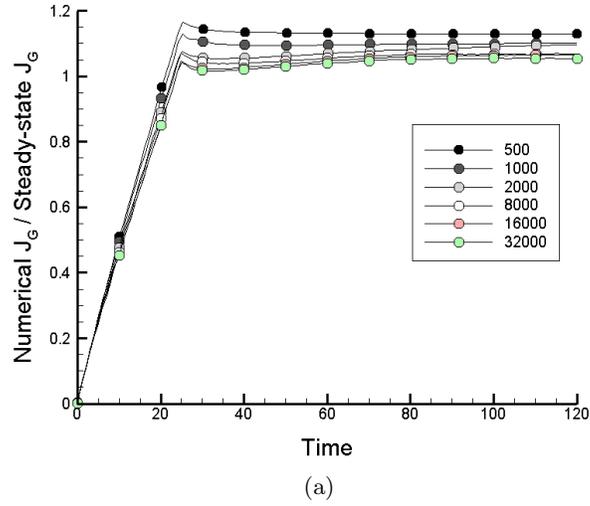
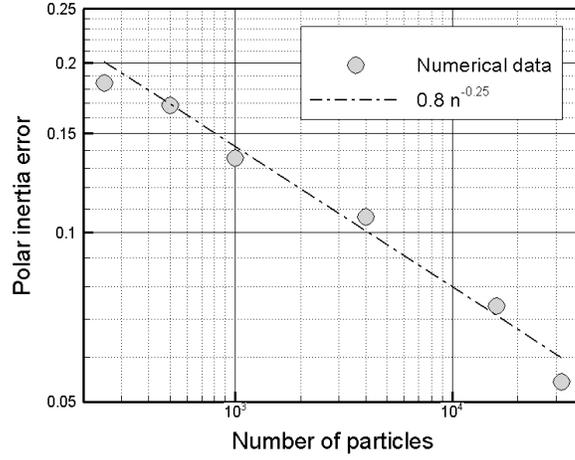
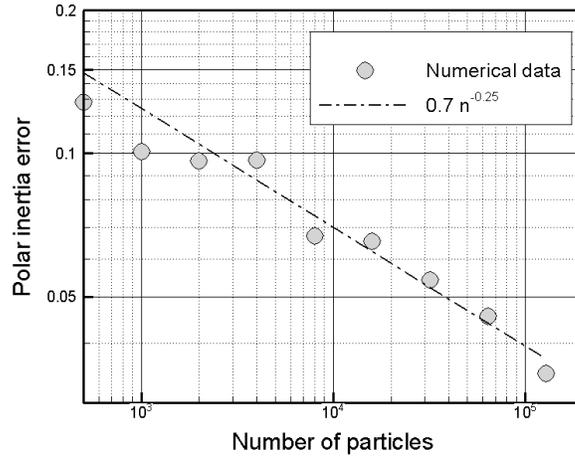


FIGURE 6. Time evolution of the numerical value of the polar moment of inertia, normalized with respect to the steady-state value, for various numbers of particles. (a) Full history. (b) Detail of the attainment of the steady state condition.



(a) Runge-Kutta



(b) Forward-Euler

FIGURE 7. Relative error of the polar inertia as a function of the number of particles and linear regression. a) Results using Runge-Kutta integration, sequential. b) Results using the Forward Euler integration, parallel, with increasing number of cores. The inferred rate of convergence is $1/4$ for both cases.

6. SUMMARY AND CONCLUDING REMARKS

We have formulated a class of finite-particle methods for mass transport problems based on a time-discrete incremental variational principle that combines entropy and the cost of particle transport, as measured by the Wasserstein metric. The incremental functional is further spatially discretized into finite particles, i. e., particles characterized by a fixed spatial profile of finite width, each carrying a fixed amount of mass. The motion of the particles is then governed by a competition between the cost of transport, that aims to keep the particles fixed, and entropy maximization, that aims to spread the particles so as to increase the entropy of the system. Unlike other particle schemes based on optimal transport and intended for deformable bodies [7, 22, 20, 19], the formulation is velocity-free, i. e., it relies solely on the particle representation and eschews the need to define a velocity field in addition, which confers the approach great advantage.

A critical issue in formulating finite-particle methods concerns the optimal choice of particle width. We have shown how the optimal width of the particles can be determined variationally by minimization of the governing incremental functional. In this manner, the incremental functional determines not only the motion of the particles but also the evolution of their profiles. Using this extended variational principle, we have in particular derived optimal scaling relations between the width of the particles, their number and the size of the domain. It is found that the optimal particle width lies in an intermediate scale between the particle distance and the size of the domain, with exponents that are precisely determined by the theory.

We have also addressed sundry matters of implementation including the acceleration of the computation of diffusive forces by exploiting the Gaussian decay of the particle profiles and by instituting fast nearest-neighbor searches. We have demonstrated how a simple domain subdivision scheme, with particle-neighborhood size and cell sizes subordinate to the optimal particle profile width, greatly accelerates the computation of diffusional forces, which is the main computational bottleneck, and allows consideration of large systems of particles.

Finally, we have demonstrated the robustness and versatility of the finite-particle method by means of a test problem concerned with the injection of mass into a sphere. There test results demonstrate the meshless character of the method in any spatial dimension, its ability to redistribute mass particles and follow their evolution in time, its ability to satisfy flux boundary conditions for general domains based solely on a distance function, and its robust convergence characteristics.

We close by remarking that the finite-particle method presented in this work, and others like it, is not restricted to mass transport and can in principle be applied to other scalar transport problems. However, whereas Dirichlet boundary conditions are somewhat inimical to mass transport, they play a central role in other transport problems such as heat conduction.

In the present work, we have eschewed consideration of Dirichlet boundary conditions in the interest of conciseness, but it remains of interest to elucidate efficient implementations thereof. In addition, while the simple one-level subdivision scheme employed in this work has sufficed to explore the convergence properties of the method, it remains of interest to explore other fast-search algorithms, including hierarchical k -means, approximate nearest-neighbor (ANN) algorithms and graph-based methods. These and other extensions and enhancements of the approach suggest themselves as worthwhile directions for further research.

ACKNOWLEDGEMENTS

AP is grateful for support of the Italian National Group of Physics-Mathematics (GNFM) of the Italian National Institution of High Mathematics “Francesco Severi” (INDAM). MO gratefully acknowledges support from the Deutsche Forschungsgemeinschaft (DFG) through the Schwerpunktprogramme (SPP) 2311.

REFERENCES

- [1] L. B. Lucy. Numerical approach to testing the fission hypothesis. *Astronomical Journal*, 82:1013–1024, 1977.
- [2] J. J. Monaghan. Why particle methods work. *SIAM Journal on Scientific and Statistical Computing*, 3(4):422–433, 1982.
- [3] D. Sulsky, Z. Chen, and H. L. Schreyer. A particle method for history-dependent materials. *Computer Methods in Applied Mechanics and Engineering*, 118:179–186, 1994.
- [4] W. K. Liu, S. Jun, and Y. F. Zhang. Reproducing kernel particle methods. *International Journal for Numerical Methods in Fluids*, 20(8–9):1081–1106, 1995.
- [5] J. K. Chen, J. E. Beraun, and C. J. Jih. An improvement for tensile instability in smoothed particle hydrodynamics. *Computational Mechanics*, 23(4):279–287, 1999.
- [6] G. M. Zhang and R. C. Batra. Modified smoothed particle hydrodynamics method and its application to transient problems. *Computational mechanics*, 34(2):137–146, 2004.
- [7] B. Li, F. Habbal, and M. Ortiz. Optimal transportation meshfree approximations schemes for fluid and plastic flows. *International Journal for Numerical Methods in Engineering*, 83(12):1541–1579, 2010.
- [8] C. Weißenfels and P. Wriggers. Stabilization algorithm for the optimal transportation meshfree approximation scheme. *Computer Methods in Applied Mechanics and Engineering*, 329:421–443, 2018.
- [9] A. Deffo, M. P. Ariza, and M. Ortiz. A line-free method of monopoles for 3d dislocation dynamics. *Journal of the Mechanics and Physics of Solids*, 122:566–589, 2019.
- [10] M. P. Ariza and M. Ortiz. A semi-discrete line-free method of monopoles for dislocation dynamics. *Extreme Mechanics Letters*, 45:101267, 2021.
- [11] L. C. Evans. *Partial differential equations and Monge-Kantorovich mass transfer*. Current Developments in Mathematics. International Press: Cambridge, MA, 1997.
- [12] C. Villani. *Topics in Optimal Transportation Theory*, volume 58 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, Rhode Island, USA, 2003.

- [13] S. Daneri and G. Savaré. Lecture notes on gradient flows and optimal transport. In *Optimal Transportation: Theory and Applications*, pages 100–144. Cambridge University Press, 2014.
- [14] J.-D. Benamou and Y. Brenier. A computational fluid mechanics solution to the monge-kantorovich mass transfer problem. *Numerische Mathematik*, 84(3):375–393, 2000.
- [15] A. Bompadre, B. Schmidt, and M. Ortiz. Convergence analysis of meshfree approximation schemes. *SIAM Journal of Numerical Analysis*, 50(3):1344–1366, 2012.
- [16] A. Bompadre, L. E. Perotti, C. J. Cyron, and M. Ortiz. Convergent meshfree approximation schemes of arbitrary order and smoothness. *Computer Methods in Applied Mechanics and Engineering*, 221–222:83–103, 2012.
- [17] B. Schmidt. On the infinite particle limit in lagrangian dynamics and convergence of optimal transportation meshfree methods. *Multiscale Modeling & Simulation*, 12:265–289, 2014.
- [18] B. Li, A. Pandolfi, and M. Ortiz. Material-point erosion simulation of dynamic fragmentation of metals. *Mechanics of Materials*, 80:288–297, 2015.
- [19] H. Wessels, c. Weißenfels, and P. Wriggers. Metal particle fusion analysis for additive manufacturing using the stabilized optimal transportation meshfree method. *Computer Methods in Applied Mechanics and Engineering*, 339:91–114, 2018.
- [20] P. Navas, S. López-Querol, R. C. Yu, and M. Pastor. Optimal transportation meshfree method in geotechnical engineering problems under large deformation regime. *International Journal for Numerical Methods in Engineering*, 115(10):1217–1240, 2018.
- [21] H. Wang, H. Liao, Z. Fan, J. Fan, L. Stainier, X. Li, and B. Li. The hot optimal transportation meshfree (hotm) method for materials under extreme dynamic thermomechanical conditions. *Computer Methods in Applied Mechanics and Engineering*, 364:112958, 2020.
- [22] L. Fedeli, A. Pandolfi, and M. Ortiz. Geometrically exact time-integration mesh-free schemes for advection-diffusion problems derived from optimal transportation theory and their connection with particle methods. *International Journal for Numerical Methods in Engineering*, 112(9):1175–1193, 2017.
- [23] R. Jordan, D. Kinderlehrer, and F. Otto. Free energy and the Fokker-Planck equation. *Physica D*, 107(2–4):265–271, 1997.
- [24] R. Jordan, D. Kinderlehrer, and F. Otto. The variational formulation of the Fokker-Planck equation. *SIAM Journal on Mathematical Analysis*, 29(1):1–17, 1998.
- [25] R. Jordan, D. Kinderlehrer, and F. Otto. Dynamics of the Fokker-Planck equation. *Phase Transitions - B*, 69(3):271–288, 1999.
- [26] W. Gangbo and McCann R. J. The geometry of optimal transportation. *Acta Mathematica*, 177(2):113–161, 1996.
- [27] J. A. Carrillo, Y. Huang, F. S. Patacchini, and G. Wolansky. Numerical study of a particle method for gradient flows. *Kinetic and Related Models*, 10(3):613–641, 2017.
- [28] J. A. Carrillo, K. Craig, and F. S. Patacchini. A blob method for diffusion. *Calculus of Variations and Partial Differential Equations*, 58(2), 2019.
- [29] F. J. Pinski, G. Simpson, A. M. Stuart, and H. Weber. Kullback–Leibler approximation for probability measures on infinite dimensional spaces. *SIAM Journal on Mathematical Analysis*, 47(6):4091–4122, 2015.
- [30] D. C. Dowson and B. V. Landau. The Fréchet distance between multivariate normal distributions. *Journal of Multivariate Analysis*, 12(3):450–455, 1982.
- [31] I. Olkin and F. Pukelsheim. The distance between two random vectors with given dispersion matrices. *Linear Algebra and its Applications*, 48:257–263, 1982.
- [32] C. R. Givens and R. M. Shortt. A class of Wasserstein metrics for probability distributions. *Michigan Mathematical Journal*, 31(2):231–240, 1984.
- [33] M. Knott and C. S. Smith. On the optimal mapping of distributions. *Journal of Optimization Theory and Applications*, 43(1):39–49, 1984.

- [34] F. B. Hildebrand. *Introduction to Numerical Analysis: Second Edition*. Dover Books on Mathematics. Dover Publications, 2013.
- [35] A. Profeta and K.-T. Sturm. Heat flow with dirichlet boundary conditions via optimal transport and gluing of metric measure spaces. *Calculus of Variations and Partial Differential Equations*, 59:117, 2020.
- [36] A. Figalli and N. Gigli. A new transportation distance between non-negative measures, with applications to gradients flows with dirichlet boundary conditions. *Journal de Mathématiques Pures et Appliquées*, 94(2):107–130, 2010.
- [37] S. Dweik and F. Santambrogio. Summability estimates on transport densities with dirichlet regions on the boundary via symmetrization techniques. *ESAIM: Control, Optimisation and Calculus of Variations*, 24(3):1167–1180, 2018.
- [38] S. Dweik. Optimal transportation with boundary costs and summability estimates on the transport density. *Journal of Convex Analysis*, 25(1):135–160, 2018.
- [39] A. H. Land and A. G. Doig. An automatic method for solving discrete programming problems. In *50 Years of Integer Programming 1958-2008*, pages 105–132. Springer, 2010.
- [40] J. L. Bentley. Multidimensional binary search trees used for associative searching. *Communications of the ACM*, 18(9):509–517, 1975.
- [41] J. H. Friedman, J. L. Bentley, and R. A. Finkel. An algorithm for finding best matches in logarithmic expected time. *ACM Transactions on Mathematical Software (TOMS)*, 3(3):209–226, 1977.
- [42] K. Fukunaga and P. M. Narendra. A branch and bound algorithm for computing k-nearest neighbors. *IEEE transactions on computers*, 100(7):750–753, 1975.
- [43] J. K. Uhlmann. Satisfying general proximity/similarity queries with metric trees. *Information processing letters*, 40(4):175–179, 1991.
- [44] R. F. Sproull. Refinements to nearest-neighbor searching ink-dimensional trees. *Algorithmica*, 6(1-6):579–589, 1991.
- [45] M. Muja and D. Lowe. Flann-fast library for approximate nearest neighbors user manual. *Computer Science Department, University of British Columbia, Vancouver, BC, Canada*, 5, 2009.
- [46] M. Muja and D. G. Lowe. Scalable nearest neighbor algorithms for high dimensional data. *IEEE transactions on pattern analysis and machine intelligence*, 36(11):2227–2240, 2014.
- [47] W. Li, Y. Zhang, Y. Sun, W. Wang, M. Li, W. Zhang, and X. Lin. Approximate nearest neighbor search on high dimensional data-experiments, analyses, and improvement. *IEEE Transactions on Knowledge and Data Engineering*, 32(8):1475–1488, 2019.
- [48] F. Groh, L. Ruppert, P. Wieschollek, and H. P. A. Lensch. GGNN: Graph-based GPU Nearest Neighbor Search. *IEEE Transactions on Big Data*, 9(1):267–279, 2022.
- [49] M. Huesmann and K.-T. Sturm. Optimal transport from Lebesgue to Poisson. *The Annals of Probability*, 41(4):2426–247, 12 2013.

¹POLITECNICO DI MILANO, CIVIL AND ENVIRONMENTAL ENGINEERING DEPARTMENT, 20133 MILANO, ITALY, ²INSTITUT DE RECHERCHE EN GÉNIE CIVIL ET MÉCANIQUE (GEM), ECOLE CENTRALE DE NANTES, 4321 NANTES CEDEX 3, FRANCE, ³CALIFORNIA INSTITUTE OF TECHNOLOGY, ENGINEERING AND APPLIED SCIENCE DIVISION, PASADENA CA, 91125, USA

Email address: anna.pandolfi@polimi.it, laurent.stainier@ec-nantes.fr, ortiz@caltech.edu