



Proven Distributed Memory Parallelization of Particle Methods

JOHANNES PAHLKE, Faculty of Computer Science, Dresden University of Technology, Dresden, Germany and Center for Systems Biology Dresden, Max Planck Institute of Molecular Cell Biology and Genetics, Dresden, Germany

IVO F. SBALZARINI, Faculty of Computer Science, Dresden University of Technology, Dresden, Germany and Center for Systems Biology Dresden, Max Planck Institute of Molecular Cell Biology and Genetics, Dresden, Germany

We provide a mathematically proven parallelization scheme for particle methods on distributed-memory computer systems. Particle methods are a versatile and widely used class of algorithms for computer simulations and numerical predictions in various applications, ranging from continuum fluid dynamics to discrete granular flows and molecular dynamics simulations. Particle methods naturally lend themselves to implementation on parallel computing systems. So far, however, a mathematical proof of correctness and equivalence to sequential implementations has only been available for shared-memory parallelism. Here, we leverage a formal definition of the algorithmic class of particle methods to provide a proven parallelization scheme for distributed-memory computers. We prove that thus parallelized particle methods on distributed-memory computers are formally equivalent to their sequential counterpart for a well-defined class of particle methods, and we provide analytical expressions for the speedup and scalability bounds of this class of algorithms in function of their parameters. The parallelization scheme analyzed here is the basis of many real-world software designs for parallel particle methods. The present analysis is, therefore, of direct relevance to existing and new parallel implementations of particle methods and places them on solid theoretical grounds, rationalizing best practices and providing useful scalability and speedup bounds for benchmarking.

CCS Concepts: • **Theory of computation** → **Parallel algorithms; Distributed algorithms**; • **Computing methodologies** → **Distributed algorithms; Parallel algorithms**;

Additional Key Words and Phrases: Particle methods, simulation algorithms, formal definition, algorithmics, software engineering, parallelization, distributed memory, meshfree methods, scalability bounds

ACM Reference Format:

Johannes Pahlke and Ivo F. Sbalzarini. 2024. Proven Distributed Memory Parallelization of Particle Methods. *ACM Trans. Parallel Comput.* 11, 4, Article 17 (November 2024), 45 pages. <https://doi.org/10.1145/3696189>

This work was funded in parts by the German Research Foundation (Deutsche Forschungsgemeinschaft, DFG) within the Research Training Group “Role-based Software Infrastructures for continuous-context-sensitive Systems” (GRK 1907).

Authors’ Contact Information: Johannes Pahlke, Faculty of Computer Science, Dresden University of Technology, Dresden, Germany and Center for Systems Biology Dresden, Max Planck Institute of Molecular Cell Biology and Genetics, Dresden, Germany; e-mail: pahlke@mpi-cbg.de; Ivo F. Sbalzarini, Faculty of Computer Science, Dresden University of Technology, Dresden, Germany and Center for Systems Biology Dresden, Max Planck Institute of Molecular Cell Biology and Genetics, Dresden, Germany; e-mail: ivos@mpi-cbg.de.



This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2024 Copyright held by the owner/author(s).

ACM 2329-4949/2024/11-ART17

<https://doi.org/10.1145/3696189>

1 Introduction

High-performance computing (HPC) is becoming increasingly important in research, especially in the life sciences [19, 32], where many physical experiments are not feasible due to ethical reasons or technical limitations in control and observation. At the same time, the power of computer hardware is increasing, mainly due to parallelism. This computing power enables the simulation of increasingly complex models but demands elaborate code, which typically incurs long development times. Therefore, generic software frameworks are actively developed in order to bridge the gap between accessible programming and heterogeneous hardware parallelization [27]. At the foundation of these frameworks are generic and parallelizable numerical algorithms.

Prominent examples of parallelizable numerical algorithms belong to the algorithmic class of particle methods, which has widespread use in scientific computing. Applications range from computational fluid dynamics [7] over computational chemistry [23] and plasma physics [15] to particle-based image processing [1, 6], encompassing many well-known numerical methods, such as **Discrete-Element Methods (DEMs)** [33], **Molecular Dynamics (MD)** [2], **Particle-In-Cell (PIC)** methods [15], **Reproducing Kernel Particle Methods (RKPMs)** [20], **Particle Strength Exchange (PSE)** [8, 9], **Discretization-Corrected PSE (DC-PSE)** [5, 30], and **Smoothed Particle Hydrodynamics (SPH)** [10, 22]. Particle methods have been efficiently parallelized on both shared and distributed-memory systems [16–18, 26, 29].

Recently, particle methods were mathematically defined [24], enabling their formal study across applications. Leveraging this formal definition, particle methods have been proven to be parallelizable on shared-memory systems under certain conditions [24]. However, no such results have been available for the distributed-memory parallelism prevalent in HPC. It was, therefore, not known under which conditions a distributed-memory implementation of a particle method is formally equivalent to its sequential counterpart, i.e., is correct in the sense that it computes the same results for any possible input.

Here, we provide a proof of equivalence for a basic distributed-memory parallelization scheme for particle methods and analyze the conditions under which it is valid. Our analysis is independent of a specific application or a specific numerical method. The proof covers a broad and well-defined class of particle-based algorithms. Moreover, the parallelization scheme we analyze is the basis of a wide range of practically used distributed-memory implementations. Hence, it provides the theoretical foundation for more specialized parallelization schemes [21] that can be reduced to or composited from the base case considered here.

The present base case builds on the classic cell-list algorithm [15] and a checkerboard-like domain decomposition. We formalize this scheme in mathematical equations as well as a Nassi–Shneiderman diagram. We then provide an exhaustive list of conditions a particle method must fulfill for the scheme to be correct. Under these conditions, we prove the equivalence of the presented parallelization scheme to the underlying sequential particle method. Furthermore, we use the presented formal analysis to infer the scheme's time complexity and parallel scalability bounds.

We hope this work provides a starting point for proving complexity bounds and correctness of distributed parallel codes in scientific HPC. In the long run, such efforts could lead to the development of provably correct software frameworks for scientific computing and enable formal analysis of distributed numerical simulation codes.

2 Prerequisites

Before analyzing shared-memory parallelization of particle methods, we introduce the notation used in this article and recall the classic cell-list data structure, the concept of domain

decomposition, and the formal definition of the algorithmic class of particle methods for which the statements are going to be derived.

2.1 Terminology and Notation

We start by introducing the notation and terminology used and defining the underlying mathematical concepts.

Definition 1. The Kleene star A^* is the set of all finite tuples of elements of a set A , including the empty tuple $()$. It is defined using the Cartesian product as

$$A^0 := \{()\}, A^1 := A, A^{n+1} := A^n \times A \text{ for } n \in \mathbb{N}_{>0} \quad (1)$$

$$A^* := \bigcup_{j \in \mathbb{N}_0} A^j. \quad (2)$$

In this article, we use the Kleene star to describe the set of all particle tuples. A particle tuple is the mathematical representation of an array of particles in the computer memory.

Notation 1. We use *bold symbols* for tuples of arbitrary length, e.g.,

$$\mathbf{p} \in P^*. \quad (3)$$

This is used to denote a specific particle tuple out of the set of all possible ones.

Notation 2. We use *regular symbols with subscript indices* to denote the elements of a tuple, e.g.,

$$\mathbf{p} = (p_1, \dots, p_n). \quad (4)$$

The elements of a particle tuple are the individual particles of a particle method.

Notation 3. We use *regular symbols* for tuples of determined length with specific elements, e.g.,

$$p = (a, b, c) \in A \times B \times C. \quad (5)$$

This is used, for example, to describe the properties of a particle p .

Notation 4. We use the same *indices* for tuples of determined length and their named elements, e.g.,

$$p_j = (a_j, b_j, c_j). \quad (6)$$

Notation 5. We use *underlined symbols* for vectors, e.g.,

$$\underline{v} \in A^n. \quad (7)$$

Vectors are a third kind of tuple whose length is fixed and all elements are from the same space. We use this to describe vectorial particle properties and cell lists.

Definition 2. Let $a \in \mathbb{R}$, $a = z + r$ with $z \in \mathbb{Z}$, $r \in \mathbb{R}$, and $0 \leq r < 1$. Then *rounding down* of a real number $a \in \mathbb{R}$ is defined as

$$\lfloor a \rfloor := z. \quad (8)$$

Definition 3. *Rounding down* of a real-valued vector $\underline{v} \in \mathbb{R}^d$ is defined element-wise:

$$\lfloor \underline{v} \rfloor := \begin{pmatrix} \lfloor v_1 \rfloor \\ \vdots \\ \lfloor v_d \rfloor \end{pmatrix}. \quad (9)$$

Definition 4. The number of elements in a tuple $\mathbf{p} = (p_1, \dots, p_n) \in P^*$ is defined as

$$|\mathbf{p}| := n. \quad (10)$$

Definition 5. The Euclidean length of a vector $\underline{v} \in \mathbb{R}^d$ is defined by the ℓ^2 -norm

$$|\underline{v}| = \left| \begin{array}{c} v_1 \\ \vdots \\ v_d \end{array} \right| := \sqrt{v_1^2 + \dots + v_d^2}. \quad (11)$$

Definition 6. The composition operator [24] $*_h$ of a binary function $h : A \times B \rightarrow A$ is recursively defined as

$$*_h : A \times B^* \rightarrow A \quad (12)$$

$$a *_h () := a \quad (13)$$

$$a *_h (b_1, b_2, \dots, b_n) := h(a, b_1) *_h (b_2, \dots, b_n). \quad (14)$$

We use the composition operator to shorten and simplify equations, e.g., for the interaction of a particle with all of its neighbors.

Definition 7. The concatenation $\circ : A^* \times A^* \rightarrow A^*$ of tuples $(a_1, \dots, a_n), (b_1, \dots, b_m) \in A^*$ is defined as

$$(a_1, \dots, a_n) \circ (b_1, \dots, b_m) := (a_1, \dots, a_n, b_1, \dots, b_m). \quad (15)$$

The big concatenation \bigcirc of tuples $\mathbf{a}_1, \dots, \mathbf{a}_n \in A^*$ is defined as

$$\bigcirc_{j=1}^n \mathbf{a}_j := \mathbf{a}_1 \circ \dots \circ \mathbf{a}_n. \quad (16)$$

We use the concatenation operator, for instance, to add particles from one tuple to another tuple.

Definition 8. We define the construction of a subtuple $\mathbf{b} \in A^*$ of $\mathbf{a} \in A^*$. Let $f : A^* \times \mathbb{N} \rightarrow \{\top, \perp\}$ ($\top = \text{true}$, $\perp = \text{false}$) be the condition for an element a_j of the tuple \mathbf{a} to be in \mathbf{b} . Then, $\mathbf{b} = (a_j \in \mathbf{a} : f(\mathbf{a}, j))$ defines a subtuple of \mathbf{a} as

$$\begin{aligned} \mathbf{b} = (a_j \in \mathbf{a} : f(\mathbf{a}, j)) &:= (a_{j_1}, \dots, a_{j_n}) \\ \longleftrightarrow \quad \mathbf{a} &= (a_1, \dots, a_{j_1}, \dots, a_{j_2}, \dots, a_{j_n}, \dots, a_m) \\ \wedge \quad \forall k \in \{1, \dots, n\} &: f(\mathbf{a}, j_k) = \top \\ \wedge \quad \forall l \in \{1, \dots, m\} \setminus \{j_1, \dots, j_n\} &: f(\mathbf{a}, l) = \perp. \end{aligned} \quad (17)$$

Subtuples are used in the following to describe the neighborhood (i.e., the interaction partners) of a particle.

Definition 9. A permutation σ is a bijective function mapping the finite set A ($|A| < \infty$) to itself:

$$\sigma : A \rightarrow A. \quad (18)$$

Definition 10. Let $\mathbf{a} = (a_1, \dots, a_n) \in A^n$ a tuple and $\sigma : \{a_1, \dots, a_n\} \rightarrow \{a_1, \dots, a_n\}$ a permutation. Then, the permutation of a tuple is defined as

$$\sigma(\mathbf{a}) := (\sigma(a_1), \dots, \sigma(a_n)). \quad (19)$$

The permutation of a tuple is useful to keep track of potential reordering of particles throughout the algorithm.

Definition 11. Let $\alpha = (a_1, \dots, a_n) \in A_1 \times \dots \times A_n$ a tuple. Then, an *element* $\langle \alpha \rangle_j$ of a tuple is defined as

$$\langle \alpha \rangle_j := a_j, \quad (20)$$

and a *collection tuple* $\langle \alpha \rangle_{(j_1, \dots, j_m)}$ of a tuple with $j_1, \dots, j_m \in \{1, \dots, n\}$ is defined as

$$\langle \alpha \rangle_{(j_1, \dots, j_m)} := (a_{j_1}, \dots, a_{j_m}). \quad (21)$$

The collection tuple is used to select specific elements of a tuple, e.g., the neighbor particles in a particle tuple or a specific cell from a cell list.

Definition 12. A *subresult* ${}_k f$ of a function

$$f : A_1 \times \dots \times A_n \rightarrow B_1 \times \dots \times B_m \quad (22)$$

is defined as

$${}_k f(a_1, \dots, a_n) := \langle f(a_1, \dots, a_n) \rangle_k \quad \text{with } k \in \{1, \dots, m\}. \quad (23)$$

If a function has a tuple as a result, the subresult is used to obtain a specific element of the result tuple.

Notation 6. We follow the notational convention of [24], according to which functions have uppercase superscripts to specify the level of abstraction, e.g., $copy^{ALL}$, and subscripts to specify fixed parameters, e.g., $copy_g^{ALL}$.

Notation 7. We use a *bold number with over- and underbars* to denote a vector containing the same number for all entries, e.g.,

$$\bar{\mathbf{1}} := (1, \dots, 1)^T \in \mathbb{N}^d. \quad (24)$$

This notation is used to shift vectorial indices, especially in cases where the dimension is a variable.

Definition 13. Let d be the dimension of a discrete space with *vectorial index space*

$$\mathbb{N}^d \cap \left[\bar{\mathbf{1}}, \bar{\mathbf{1}} \right] \quad \text{with } \bar{\mathbf{I}} = (I_1, \dots, I_d)^T \in \mathbb{N}_{>0}^d \quad (25)$$

and corresponding *scalar index space*

$$\left\{ 1, \dots, \prod_{t=1}^d I_t \right\}. \quad (26)$$

Then, the *translation of a scalar index to a vectorial index* is

$$\bar{I}_t : \left\{ 1, \dots, \prod_{t=1}^d I_t \right\} \rightarrow \mathbb{N}^d \cap \left[\bar{\mathbf{1}}, \bar{\mathbf{1}} \right] \quad (27)$$

and defined as

$$\bar{I}_l(j) := \begin{pmatrix} (j-1) - \left\lfloor \frac{j-1}{I_1} \right\rfloor I_1 + 1 \\ \left\lfloor \frac{j-1}{I_1} \right\rfloor - \left\lfloor \frac{j-1}{I_1 I_2} \right\rfloor I_2 + 1 \\ \vdots \\ \left\lfloor \frac{j-1}{\prod_{t=1}^{l-1} I_t} \right\rfloor - \left\lfloor \frac{j-1}{\prod_{t=1}^l I_t} \right\rfloor I_l + 1 \\ \vdots \\ \left\lfloor \frac{j-1}{\prod_{t=1}^{d-2} I_t} \right\rfloor - \left\lfloor \frac{j-1}{\prod_{t=1}^{d-1} I_t} \right\rfloor I_{d-1} + 1 \\ \left\lfloor \frac{j-1}{\prod_{t=1}^{d-1} I_t} \right\rfloor + 1 \end{pmatrix}. \quad (28)$$

The *backward translation of a vectorial index to a scalar index* is

$$\bar{I}_l^{-1} : \mathbb{N}^d \cap \left[\bar{\mathbf{1}}, \bar{\mathbf{l}} \right] \rightarrow \left\{ 1, \dots, \prod_{t=1}^d I_t \right\} \quad (29)$$

and defined as

$$\bar{I}_l^{-1}(j) := 1 + (j_1 - 1) + (j_2 - 1)I_1 + (j_3 - 1)I_1 I_2 + \dots + (j_l - 1) \prod_{t=1}^{l-1} I_t + \dots + (j_d - 1) \prod_{t=1}^{d-1} I_t. \quad (30)$$

$[\bar{\mathbf{1}}, \bar{\mathbf{l}}]$ is the d -dimensional interval from $\bar{\mathbf{1}}$ to $\bar{\mathbf{l}}$. The vectorial index space is used, for instance, to index the cells of a cell list along each spatial dimension. This is possible since they form a regular Cartesian grid. The scalar index space is the linearized index, which can be used to store the cells from the d -dimensional grid in memory or iterate over them.

Notation 8. We use *square brackets* $[\cdot]$ in two situations: first, to describe closed intervals with defined borders that are included in the interval, and second, to combine multiple variables into a *single* variable, e.g., $[g, \mathbf{p}]$. The meaning is always clear from the data type of the arguments.

2.2 Cell Lists

In particle simulations, cell lists [15] are a classic data structure for neighbor access. Since particles can change their location during a simulation, cell lists are used to efficiently find the neighboring particles around any given particle in order to, e.g., compute particle interactions. The size of the neighborhood is defined by a parameter r_c , called the *cutoff radius*. Cell lists are then constructed by partitioning the simulation domain into uniform and equilateral Cartesian cells, where each edge length of each cell is equal to r_c (see Figure 1). Each particle is then, based on its position \underline{x}_p , assigned to the one cell with vectorial index $\underline{j}_{\text{cell}}$:

$$\underline{j}_{\text{cell}} = \left\lfloor \frac{\underline{x}_p}{r_c} \right\rfloor + \bar{\mathbf{1}}. \quad (31)$$

This is done by adding the index p to the end of a linked list of particle indices that is stored in each cell. This construction ensures that particles within the same cell and its neighboring cells are the only possible interaction partners. Using the cell list as a particle-index lookup table when evaluating all pairwise particle interactions within an interaction radius of r_c reduces the time complexity of the interaction evaluation from $O(N^2)$ to $\Theta(N)$ on average for uniformly distributed particles, where N is the total number of particles. The worst-case complexity remains $O(N^2)$ when




Fig. 1. Illustration of a cell-list grid with some particles. If the red particle needs to interact with all other particles within the red circle, searching over the blue cells is sufficient to find all interaction partners.

all particles cluster in a single cell. That case, however, is not of practical relevance. Moreover, we use cell lists here as they also simplify the data distribution and inter-processor communication in a distributed-memory parallel particle method.

2.3 Domain Decomposition

In order to parallelize a particle method onto a distributed-memory parallel computer, the simulation domain needs to be partitioned among the processes. In that way, each processor, or more precisely its local memory address space, only stores a part of the data of the overall particle method. This ensures that the algorithms also scale in the weak sense when using more processors. It fundamentally assumes that the simulation domain, i.e., the space in which the particles are distributed, possesses a bounding box. This assumption is not limiting, since a bounding box can always be chosen as large as needed to leave room for the free movement of the particles within the finite time of a simulation. Moreover, on a digital computer, positions cannot go to infinity, so the domain is trivially bounded by the limits of the internal number representation. Assuming the existence of a bounding box for domain decomposition, however, does not restrict how the particles can be distributed within the domain. Particle distributions are routinely irregular in order to, e.g., represent complex or moving and deforming geometries [4, 28].

The most common way of distributing data across memory address spaces is to compute a domain decomposition, assigning distinct segments to individual processes. A *process* is an autonomous and self-contained unit of computation with its own memory address space [14]. If a process needs to access data that is stored in the memory of another process, that data needs to be explicitly communicated, usually over a network interconnect. The collective amount of all communication performed by a distributed-memory parallel application, i.e., its *communication overhead*, limits the performance scalability of the application.

In particle methods, particles interact with their neighbors within a given cutoff radius. Therefore, physical-space domain decompositions (over particle locations) are used, as they lead to lower communication overhead than index-space decompositions (over particle indices). A physical-space domain decomposition partitions the simulation domain into subdomains as illustrated in Figure 2 on the left. While arbitrary partitions of unity are possible, software implementations are often restricted to aligned cuboidal subdomains, which is also what we consider here.

For all particle interactions on a given process to be computed independent of and in parallel with those on other processes, each process must possess information about particles from neighboring (in the domain decomposition) processes within the cutoff radius from its domain boundary. Most practical implementations of parallel particle methods exchange this data between processes




Fig. 2. Illustration of domain decomposition for distributed-memory particle methods. The simulation domain on the left is partitioned into cuboidal subdomains. Each subdomain (color) is assigned to one process. Each individual process (right inset for the example of the blue subdomain) extends its subdomain by a ghost layer of one cell-list cell thickness all around. This layer is filled with copies of the particle data from the respective neighboring processes (indicated by color). This defines the communication overhead of the parallel implementation and allows subsequent computation of pairwise particle interactions on each process without further communication.

before any interactions are computed. The boundary particles from neighboring processes are then stored in a *ghost layer*, also known as a padding or halo layer, which is an additional layer of cell-list cells outside the subdomain of the given process (Figure 2, right). Therefore, when using cell lists, domain decomposition typically operates on units of complete cells to minimize the communication overhead.

2.4 Mathematical Definition of Particle Methods

We briefly recall the formal definition of the algorithmic class of particle methods as presented in more detail elsewhere [24]. This mainly serves to remind of the core concepts for the derivations that follow. In any particle method, particles are collections of properties, such as position, velocity, mass, color, and so forth. Each particle interacts pairwise with other particles in order to change its properties and independently evolves. For convenience, and in order to better reflect practical implementations of particle methods, the definition also includes a global variable, which usually is a collection of aggregate properties, such as simulation time, total energy, total number of particles, and so forth.

In scientific simulation applications, particles can be used either to represent discrete model entities or to mathematically discretize continuous fields. A prominent example of the discrete case is molecular dynamics simulations, where particles represent atoms or molecules that interact with each other according to physical force potentials and evolve according to Newton’s second law. Other discrete examples include traffic simulations with particles representing traffic participants (cars, pedestrians, bicycles, etc.) and granular flow simulations where particles represent individual grains (e.g., of sand, salt, or any other granular material). When numerically simulating models involving continuous mathematical fields, such as the electric fields in Maxwell’s equations or the fluid pressure and velocity fields in the Navier–Stokes equations, particles represent mathematical discretization points of some numerical scheme. Since these discretization points are not required to lie on a grid or mesh, and they can move during the simulation (e.g., with the flow velocity in Lagrangian fluid simulations or to locally adapt the numerical resolution in adaptive solvers),

particle methods possess certain beneficial numerical properties, such as improved linear stability. Examples of widely used numerical particle schemes for continuous models include SPH for fluid mechanics and PIC for plasma physics. While these examples serve to illustrate what particle methods are, we note that equivalent algorithms are also used in fields other than scientific computing. This includes point-based graphics in computer graphics, particle-based computer vision, particle filters in statistics, and so forth, which are all subsumed in our formal definition of particle methods [24].

Formally, the definition of any particle method consists of three parts: the algorithm, the instance, and the state transition.

Particle Method Algorithm. The definition of a particle method algorithm encapsulates the structural elements of its implementation in a small set of data structures and functions that need to be specified at the onset.

In detail, the components are defined as follows:

Definition 14. A *particle method algorithm* is a 7-tuple $(P, G, u, f, i, e, \dot{e})$, consisting of the two data structures

$$P := A_1 \times A_2 \times \cdots \times A_n \quad \text{the particle space,} \quad (32)$$

$$G := B_1 \times B_2 \times \cdots \times B_m \quad \text{the global variable space,} \quad (33)$$

such that $[G \times P^*]$ is the *state space* of the particle method, and five functions:

$$u : [G \times P^*] \times \mathbb{N} \rightarrow \mathbb{N}^{*0} \quad \text{the neighborhood function,} \quad (34)$$

$$f : G \rightarrow \{\top, \perp\} \quad \text{the stopping condition,} \quad (35)$$

$$i : G \times P \times P \rightarrow P \times P \quad \text{the interact function,} \quad (36)$$

$$e : G \times P \rightarrow G \times P^{*0} \quad \text{the evolve function,} \quad (37)$$

$$\dot{e} : G \rightarrow G \quad \text{the evolve function of the global variable.} \quad (38)$$

Particle Method Instance. A particle method instance describes the initial state of the particle method, instantiating the data structures of a particle method algorithm.

Definition 15. An initial state defines a *particle method instance* for a given particle method algorithm $(P, G, u, f, i, e, \dot{e})$:

$$[g^1, \mathbf{p}^1] \in [G \times P^*]. \quad (39)$$

The instance consists of an initial value for the global variable $g^1 \in G$ and an initial tuple of particles $\mathbf{p}^1 \in P^*$.

Particle State Transition Function. The particle method state transition function describes how a particle method proceeds from the instance to the final state by using the particle method algorithm. The state transition function consists of a series of state transition steps. This series ends when the stopping condition f returns *true* (\top). The state transition steps use the functions u, i, e, \dot{e} in a way that is the same for all particle methods, regardless of their specific implementation.

Definition 16. The *state transition function* $S : [G \times P^*] \rightarrow [G \times P^*]$ is defined with the following Nassi–Shneiderman diagram:




Fig. 3. Nassi–Shneiderman diagram of the state transition function S .

The Nassi–Shneiderman diagram corresponds to equations that mathematically define the state transition function S . It is divided into sub-functions here for better readability. We explicitly list these equations hereafter, indicating for each one of them which line (1–12) in the above the Nassi–Shneiderman diagram it corresponds to.

Notation 9. In the following, we use superscripts to indicate over how many particles or pairs of particles the superscribed function iterates. These are only names/labels and not variables. Specifically, I indicates that only one particle or pair is involved, whereas N indicates iterations over all particles. $I \times U$ indicates iterations over all pairs of one particle with each of its neighbors, and $N \times U$ indicates iterations over all particles and for each particle with all of its respective neighbors.

Notation 10. An *overbar* indicates a variable that stores the change to the value of the overbared variable. It therefore denotes a different variable and is not a unary operator.

All interaction sub-functions return an altered particle tuple \mathbf{p} . The first interaction sub-function i^I calculates one single interaction and formalizes line 6:

$$i^I([g, \mathbf{p}], j, k) := (p_1, \dots, p_{j-1}, \bar{p}_j, p_{j+1}, \dots, p_{k-1}, \bar{p}_k, p_{k+1}, \dots, p_{|\mathbf{p}|}),$$

$$\text{with } (p_1, \dots, p_{|\mathbf{p}|}) = \mathbf{p}, \quad (\bar{p}_j, \bar{p}_k) := i(g, p_j, p_k). \quad (40)$$

The second interaction sub-function $i^{I \times U}$ calculates the interaction of one particle with *all* its interaction partners and formalizes lines 4 to 6:

$$i^{I \times U}([g, \mathbf{p}], j) := \mathbf{p} *_{i^I} u([g, \mathbf{p}], j) \quad \text{with } i^I_{(g, j)}(\mathbf{p}, k) := i^I([g, \mathbf{p}], j, k). \quad (41)$$

The third interaction sub-function $i^{N \times U}$ calculates the interactions of *all* particles with *all* of their respective interaction partners and formalizes lines 3 to 6:

$$i^{N \times U}([g, \mathbf{p}]) := \mathbf{p} *_{i^{I \times U}} (1, \dots, |\mathbf{p}|) \quad \text{with } i^{I \times U}_g(\mathbf{p}, j) := i^{I \times U}([g, \mathbf{p}], j). \quad (42)$$

¹ \mathbf{q} is an intermediate result.

² $\bar{\mathbf{q}}$ is an intermediate result.

The first evolution sub-function ϵ^1 calculates the evolution of one particle and stores the result in an intermediate particle tuple \mathbf{q} and in the global variable. It formalizes line 9 and 10:

$$\epsilon^1([g, \mathbf{p}], \mathbf{q}, j) := [\bar{g}, \mathbf{q} \circ \bar{\mathbf{q}}] \quad \text{with} \quad (\bar{g}, \bar{\mathbf{q}}) := e(g, p_j). \quad (43)$$

The second evolution sub-function ϵ^N calculates the evolution of *all* particles and returns the result in a new particle tuple \mathbf{q} and the global variable. It formalizes lines 7 to 10:

$$\epsilon^N([g, \mathbf{p}]) := [g, ()] *_{\epsilon_p^1} (1, \dots, |\mathbf{p}|) \quad \text{with} \quad \epsilon_p^1([g, \mathbf{q}], j) := \epsilon^1([g, \mathbf{p}], \mathbf{q}, j). \quad (44)$$

The state transition step s combines all sub-functions, formalizing lines 3 to 12:

$$s([g, \mathbf{p}]) := [\hat{e}(\bar{g}), \bar{\mathbf{p}}] \quad \text{with} \quad [\bar{g}, \bar{\mathbf{p}}] := \epsilon^N([g, i^{\mathbb{N} \times \mathbb{U}}([g, \mathbf{p}])]). \quad (45)$$

Finally, the state transition function S advances the instance to the final state, formalizing the complete diagram:

$$S([g^1, \mathbf{p}^1]) = [g^T, \mathbf{p}^T] \quad \longleftrightarrow \\ f(g^T) = \top \quad \wedge \quad \forall t \in \{2, \dots, T\} : [g^t, \mathbf{p}^t] = s([g^{t-1}, \mathbf{p}^{t-1}]) \wedge f(g^{t-1}) = \perp. \quad (46)$$

3 Distributed Pull Particle Methods without Global Operations

While most parallelization strategies tend to be algorithm specific, the formal definition of particle methods [24] enables us to formulate, prove, and analyze a parallelization scheme for an entire class of algorithms. We consider the parallelization of particle methods on distributed-memory computers. Shared-memory parallelism of particle methods has previously been considered [24].

The present parallelization scheme is valid for particle methods with pull interactions and without global operations. Pull interactions denote the case where only the interacting particle is changed and there is no change to any of its neighbor particles and interaction partners. This is also sometimes referred to as *asymmetric particle interactions*. The parallelization scheme analyzed here is based on cell lists and domain decomposition, as described in the previous sections and often used in practical code implementations.

Contrary to practical implementations, however, we assume a checkerboard inter-process communication schedule in order to avoid conflicting or colliding communications. This is not usually implemented in practice but rather left to the scheduler of the network communication library used. Nevertheless, we assume it here for simplicity of the mathematical notation. In addition, we assume that the particle method contains no distributed operations on the global variable, as that would require additional inter-process communication (reduction followed by broadcast).

In the resulting parallelization scheme, each process concurrently executes the steps illustrated in Figure 4 from top to bottom for each iteration (time step) of the particle method. This is the minimal parallelization scheme for particle methods, which we study here in order to prove its correctness and to derive bounds on its complexity and scalability. Any more sophisticated or elaborate scheme would have to be within the resulting bounds and be reducible to this minimal scheme for correctness.

Formally, cell lists require that each particle p has a position \underline{x} :

$$p = (\dots, \underline{x}, \dots) \in P, \quad (47)$$

and that the neighborhood function u is based on a cutoff radius r_c and not directly on particle indices. This makes it, in some sense, order independent. Hence, it is restricted to the form

$$u([g, \mathbf{p}], j) := \left(k \in (1, \dots, |\mathbf{p}|) : p_k, p_j \in \mathbf{p} \wedge |\underline{x}_k - \underline{x}_j| \leq r_c \wedge \Omega(g, j, k, p_j, p_k) \right) \\ \text{with } \Omega : G \times \mathbb{N} \times \mathbb{N} \times P \times P \rightarrow \{\top, \perp\}, \quad (48)$$

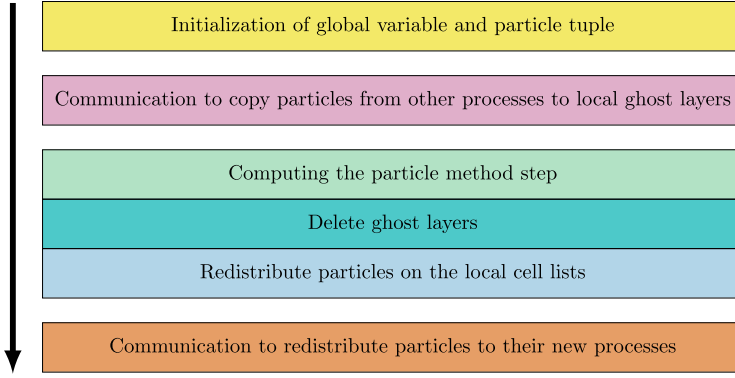


Fig. 4. Overview of the minimal parallelization scheme for distributed-memory particle methods with pull interactions and without global operations. In each iteration of the particle method, each process concurrently executes the steps from top to bottom.

where the function Ω is an additional constraint that can be used to encode more complex interaction neighborhoods, e.g., to exclude self-interactions ($\Omega(g, j, k, p_j, p_k) := (j \neq k)$) or to express anisotropic or adaptive interaction neighborhoods.

The position \underline{x} of a particle is limited to a cuboidal spatial domain of dimension d for all states/times t :

$$\forall t \in \{1, \dots, T\} \forall p_j^t \in \mathbf{p}^t : \underline{x}_j^t \in [D_{\min}, D_{\max}] \subset \mathbb{R}^d. \quad (49)$$

This is usually achieved with some boundary conditions, e.g., periodic boundary conditions, at the edges of the domain. In addition, the position is not allowed to change by more than a cutoff radius r_c in a single state transition step:

$$\forall t \in \{1, \dots, T-1\} \forall p_j^t \in \mathbf{p}^t : |\underline{x}_j^t - \underline{x}_j^{t+1}| \leq r_c. \quad (50)$$

This latter condition is, in most cases, not practically limiting, since the particle displacements are usually even more strongly restricted by the consistency and stability requirements of the numerical method, e.g., MD, PSE, DEM. Here, we require this restriction because we are going to consider the limit of the finest possible domain decomposition, where each subdomain is of smallest possible edge length r_c as well as the ghost layer. In practice, unless the numerical scheme requires otherwise, the condition can be relaxed but must be accounted for.

Further, the interaction is limited to a pull interaction:

$$\forall p_j, p_k \in P, g \in G : i(g, p_j, p_k) = (\bar{p}_j, p_k), \quad (51)$$

independent of previous interactions of the interaction partner to avoid chain dependencies of interaction results:

$$\forall p_j, p_k, p_{k'} \in P, g \in G : {}_1i_g(p_j, {}_1i_g(p_k, p_{k'})) = {}_1i_g(p_j, p_k) \quad \text{with} \quad {}_1i_g(p_j, p_k) := \langle i(g, p_j, p_k) \rangle_1, \quad (52)$$

and be order independent:

$$\forall p_j, p_k, p_{k'} \in P, g \in G : {}_1i_g({}_1i_g(p_j, p_k), p_{k'}) = {}_1i_g({}_1i_g(p_j, p_{k'}), p_k). \quad (53)$$

Order independence is not strictly required, but it is not limiting in practice and avoids re-sorting the particles. The neighborhood function u needs to be independent of previous interactions so the cell list can be prepared before the interactions happen. Here, $*_{\mu}$ is the composition operator

defined in Definition 6 and used on l^I , whereas l^I is the first transition sub-function defined in Equation (40):

$$\forall j, k', k'' \in \{1, \dots, |\mathbf{p}|\}, [g, \mathbf{p}] \in [G \times P^*] : u([g, \mathbf{p}], j) = u\left(\left[g, \mathbf{p} *_{l^I_{(g, k')}}(k'')\right], j\right). \quad (54)$$

Finally, we assume that the evolve function e does not change the global variable g :

$$\forall g \in G, p \in P : e(g, p) = (g, \mathbf{q}), \quad (55)$$

which avoids global operations and their $O(\log(\text{number of processes}))$ communication complexity [11].

Under these constraints on a particle method, we formulate the parallelization of any particle method onto multiple distributed-memory processes. The parallelization scheme is based on the cell-list algorithm [15]. Cell lists require the definition of a cutoff radius r_c to divide the domain into equal-sized Cartesian cells. The number of cells along each dimension of the computational domain is represented by the vector

$$\bar{\mathbf{I}} = \begin{pmatrix} I_1 \\ \vdots \\ I_d \end{pmatrix} := \left\lfloor \frac{1}{r_c} (D_{\max} - D_{\min}) + \bar{\mathbf{1}} \right\rfloor. \quad (56)$$

The total number of cells is

$$N_{\text{cell}} := \prod_{l=1}^d I_l. \quad (57)$$

To prove correctness of the parallel scheme, we consider the case with the highest degree of parallelism for this scheme. The highest degree of parallelism is reached when assigning each cell-list cell to a separate process. This domain decomposition of finest granularity also avoids load balancing. Therefore, it results in a simpler overall scheme and, hence, proof. Each process then has to exchange information with its direct (face-, edge-, and corner-connected) neighbors in the cell-list grid. We assume that each process can only communicate with one other process at a time, and that the network behaves as a fully connected graph. In order to guarantee that communications never overlap, we assume that any two communicating processes have at least two inactive processes between them. This results in a checkerboard-like inter-process communication pattern, as illustrated in Figure 5, which is the Cartesian special case of the graph-coloring communication scheduling often implemented in practice [29].

The number of active cells in each dimension is given by the vector

$${}^k \bar{\mathbf{I}}^* = \begin{pmatrix} {}^k I_1^* \\ \vdots \\ {}^k I_d^* \end{pmatrix} := \left\lfloor \frac{1}{3} \left(\bar{\mathbf{I}} - \bar{\mathfrak{3}}_I(k) + \bar{\mathfrak{3}} \right) \right\rfloor, \quad (58)$$

where k represents one of the 3^d different communicating process configurations $k \in \{1, \dots, 3^d\}$. In total, the number of communicating processes for each k is

$${}^k N_{\text{cell}}^* := \prod_{l=1}^d {}^k I_l^*. \quad (59)$$

We address the communicating processes for each k by

$$\gamma(k, j) := \bar{\mathbf{I}}^{-1} \left({}^k \bar{\mathbf{I}}^*(j) \cdot 3 + \bar{\mathfrak{3}}_I(k) - \bar{\mathfrak{3}} \right). \quad (60)$$




Fig. 5. Illustration of the checkerboard-like inter-process communication strategy. Each colored square represents a different process. Processes of the same color communicate concurrently, while the others wait. Once all pairwise communication has happened, each process (right inset for an example) has its ghost layers populated with data from all neighboring processes.

The l th neighbor cell of the t th process is

$$\beta(t, l) := \bar{\Gamma}_t^{-1} \left(\bar{\Gamma}_t(t) + \bar{\Xi}_t(l) - \bar{\mathbf{2}} \right). \quad (61)$$

The result of β is undefined unless the result of $\bar{\Gamma}_t^{-1}$, Definition 13 and Equation (56), is defined. We note that this checkerboard-like pattern is not how codes are usually implemented in practice. There, communications would all be launched at once, leaving their scheduling to the networking sub-system. For the sake of theoretical analysis, however, the two are interchangeable, and it separates the consideration of the algorithm from the properties of the interconnect.

We formulate the standard procedure of distributing particles into a cell list according to Figure 1 as a condition for the initial particle distribution. Hence, each cell k initially contains a tuple of particles \mathbf{p}_k^1 . To ensure no particle is lost, we require that the concatenation of these particle tuples is a permutation π of the initial particle tuple \mathbf{p}^1 :

$$\mathbf{p}_1^1 \circ \dots \circ \mathbf{p}_{N_{\text{cell}}}^1 = \pi(\mathbf{p}^1), \quad (62)$$

and that particles are distributed according to their position:

$$\forall p_j^1 \in \mathbf{p}^1 : p_j^1 \in \mathbf{p}_w^1 \quad \text{with} \quad w = \bar{\Gamma}_t^{-1} \left(\left\lfloor \frac{1}{r_c} (\mathbf{x}_j^1 - \underline{D}_{\min}) \right\rfloor + \bar{\mathbf{1}} \right). \quad (63)$$

Each process has its own memory address space, containing its global variable storage g and its particle storage $\underline{\mathbf{p}}$. The latter stores the particles within the cell-list cell assigned to that process and copies of the particles from neighboring cells, as illustrated in Figure 5. Therefore, the particle storage of each process is compartmentalized cell-wise:

$$\underline{\mathbf{p}} = \begin{pmatrix} \langle \underline{\mathbf{p}} \rangle_1 \\ \vdots \\ \langle \underline{\mathbf{p}} \rangle_{3^d} \end{pmatrix}^T \in (P^*)^{3^d}, \quad (64)$$

where the center entry (i.e., location $\frac{3^d+1}{2}$) contains the “real” particles of the cell assigned to that process. The other entries contain the copies of the cells from the neighboring processes in the order corresponding to their position in the cell list. Two processes are neighbors if the corresponding cells are direct (face-, edge-, and corner-connected) neighbors in the cell-list grid.

The **particle storages of all processes (PROC)** together is

$$\mathcal{P} = \begin{pmatrix} \text{PROC}_1 \underline{\mathbf{p}} \\ \vdots \\ \text{PROC}_{N_{\text{cell}}} \underline{\mathbf{p}} \end{pmatrix}^\top \in \left((P^*)^{3^d} \right)^{N_{\text{cell}}}. \quad (65)$$

We use “PROC” with an index to emphasize that the variable lives on a specific process. The initial particle tuples in each process’s particle storage are

$$\text{PROC}_n \underline{\mathbf{p}}^1 = \begin{pmatrix} 0, \dots, 0, \underbrace{\mathbf{p}_n^1}_{\frac{3^d+1}{2}\text{-th entry}}, 0, \dots, 0 \end{pmatrix}. \quad (66)$$

Then, the initial particle storage of all processes is

$$\mathcal{P}^1 := \begin{pmatrix} \text{PROC}_1 \underline{\mathbf{p}}^1 \\ \vdots \\ \text{PROC}_{N_{\text{cell}}} \underline{\mathbf{p}}^1 \end{pmatrix}^\top. \quad (67)$$

The global variable storage of all processes is

$$\mathcal{G}^1 := (g^1, \dots, g^1) \in G^{N_{\text{cell}}}. \quad (68)$$

We introduce a second global variable \tilde{g} to store the cell-list-specific parameters from Equations (49) and (56). This is different from the global variable g of the particle method. We therefore use a tilde to distinguish this second global variable from the standard one:

$$\tilde{g} = (\underline{D}_{\min}, \underline{D}_{\max}, d, \bar{\mathbf{I}}). \quad (69)$$

The function $\text{copy}_{(\tilde{g}, \mathcal{P})}$ copies the center particle storage compartment of a process $\text{PROC}_{\beta(w,l)}$ to the specified compartment l in the storage $\underline{\mathbf{p}}$ of another process:

$$\text{copy}_{(\tilde{g}, \mathcal{P}, w)}(\underline{\mathbf{p}}, l) := \left(\langle \underline{\mathbf{p}} \rangle_1, \dots, \langle \underline{\mathbf{p}} \rangle_{l-1}, \langle \text{PROC}_{\beta(w,l)} \underline{\mathbf{p}} \rangle_{\frac{3^d+1}{2}}, \langle \underline{\mathbf{p}} \rangle_{l+1}, \dots, \langle \underline{\mathbf{p}} \rangle_{3^d} \right). \quad (70)$$

This coping is done for every k th cell in the checkerboard-like configuration by

$$copy_{\underline{g}}^{ACTIVE}(\mathcal{P}, k) := \left(\begin{array}{c} PROC_1 \underline{\mathbf{p}} \\ \vdots \\ PROC_{\gamma(k,1)-1} \underline{\mathbf{p}} \\ PROC_{\gamma(k,1)} \underline{\mathbf{p}} * copy_{(\underline{g}, \mathcal{P}, \gamma(k,1))} (1, \dots, 3^d) \\ PROC_{\gamma(k,1)+1} \underline{\mathbf{p}} \\ \vdots \\ PROC_{\gamma(k,l)-1} \underline{\mathbf{p}} \\ PROC_{\gamma(k,l)} \underline{\mathbf{p}} * copy_{(\underline{g}, \mathcal{P}, \gamma(k,l))} (1, \dots, 3^d) \\ PROC_{\gamma(k,l)+1} \underline{\mathbf{p}} \\ \vdots \\ PROC_{\gamma(k, k_{N^*_{cell}})-1} \underline{\mathbf{p}} \\ PROC_{\gamma(k, k_{N^*_{cell}})} \underline{\mathbf{p}} * copy_{(\underline{g}, \mathcal{P}, \gamma(k, k_{N^*_{cell}}))} (1, \dots, 3^d) \\ PROC_{\gamma(k, k_{N^*_{cell}})+1} \underline{\mathbf{p}} \\ \vdots \\ PROC_{N_{cell}} \underline{\mathbf{p}} \end{array} \right)^T. \quad (71)$$

Doing so for all distinct checkerboard-like configurations results in

$$copy_{\underline{g}}^{ALL}(\mathcal{P}) := \mathcal{P} * copy_{\underline{g}}^{ACTIVE} (1, \dots, 3^d). \quad (72)$$

After the function $copy_{\underline{g}}^{ALL}$, each process has (copies of) all particles required to calculate the interactions and evolutions of the particles in its cell without any further inter-process communication. We define for this scheme the interaction of all particles with their respective interaction partners by the function

$$interaction_{\underline{g}}(\underline{\mathbf{p}}) := \left(\begin{array}{c} \langle \mathbf{q} \rangle_1 *_{1 i_g} \left\langle \bigcirc_{w=1}^{3^d} \langle \underline{\mathbf{p}} \rangle_w \right\rangle_u \left(\left[g, \bigcirc_{w=1}^{3^d} \langle \underline{\mathbf{p}} \rangle_w \right], z+1 \right) \\ \vdots \\ \langle \mathbf{q} \rangle_{|q|} *_{1 i_g} \left\langle \bigcirc_{w=1}^{3^d} \langle \underline{\mathbf{p}} \rangle_w \right\rangle_u \left(\left[g, \bigcirc_{w=1}^{3^d} \langle \underline{\mathbf{p}} \rangle_w \right], z+|q| \right) \end{array} \right)^T, \quad (73)$$

where $z = \left\lfloor \frac{3^d+1}{2} - 1 \right\rfloor$ and $\mathbf{q} = \langle \underline{\mathbf{p}} \rangle_w$ and $\mathbf{q} = \langle \underline{\mathbf{p}} \rangle_w$. The step function $step_{(\underline{g}, g)}$ uses the function $interaction_{\underline{g}}$ to compute the state transition step (mostly simulation time step) including the

evolutions of the particle properties and positions:

$$step_{(\tilde{g},g)}(\underline{\mathbf{p}}) := \begin{pmatrix} \langle \underline{\mathbf{p}} \rangle_1 \\ \vdots \\ \langle \underline{\mathbf{p}} \rangle_{\frac{3^d+1}{2}-1} \\ 2\epsilon^N \left(\left[g, interaction_g(\underline{\mathbf{p}}) \right] \right) \\ \langle \underline{\mathbf{p}} \rangle_{\frac{3^d+1}{2}+1} \\ \vdots \\ \langle \underline{\mathbf{p}} \rangle_{3^d} \end{pmatrix}^\top. \quad (74)$$

Calculating the state transition step on all processes results in

$$step_{(\tilde{g},\mathcal{G})}^{ALL}(\mathcal{P}) := \left(step_{(\tilde{g}, [PROC1]g)} \left([PROC1] \underline{\mathbf{p}} \right), \dots, step_{(\tilde{g}, [PROC|\mathcal{P}]g)} \left([PROC|\mathcal{P}] \underline{\mathbf{p}} \right) \right), \quad (75)$$

where $[PROCK]g := \langle \mathcal{G} \rangle_k$ and $[PROCK]\underline{\mathbf{p}} := \langle \mathcal{P} \rangle_k$.

After the function $step_{(\tilde{g},\mathcal{G})}^{ALL}$, all particles in the center storage compartments of all processes have the correct positions and properties. However, they may be in the wrong cell/storage compartment after moving. Therefore, the particles in the center storage compartments must be reassigned into the compartments and communicated to the new process if they have moved to another cell/compartment. Cell/compartment assignment of each particle q is done by

$$dist_{(\tilde{g},g,j)}(\underline{\mathbf{p}},q) := \begin{pmatrix} \langle \underline{\mathbf{p}} \rangle_1 \\ \vdots \\ \langle \underline{\mathbf{p}} \rangle_{\alpha-1} \\ \langle \underline{\mathbf{p}} \rangle_{\alpha} \circ (q) \\ \langle \underline{\mathbf{p}} \rangle_{\alpha+1} \\ \vdots \\ \langle \underline{\mathbf{p}} \rangle_{3^d} \end{pmatrix}^\top, \quad (76)$$

where

$$q = (\dots, \underline{\mathbf{x}}, \dots) \in P, \quad \alpha := \bar{3}_l^{-1} \left(\left\lfloor \frac{1}{r_c} (\underline{\mathbf{x}} - D_{\min}) \right\rfloor - \bar{1}_l(j) + \bar{3} \right). \quad (77)$$

For all particles in a center storage compartment \mathbf{q} , redistribution is done by the function

$$dist_{(\tilde{g},g,j)}^N(\mathbf{q}) := ((), \dots, ()) * dist_{(\tilde{g},g,j)} \mathbf{q}, \quad (78)$$

where $((), \dots, ())$ is the tuple of 3^d empty tuples, representing the empty particle storage of a process. For all processes, the redistribution procedure is

$$dist_{(\tilde{g},\mathcal{G})}^{ALL}(\mathcal{P}) := \begin{pmatrix} dist_{(\tilde{g}, [PROC1]g,1)}^N \left(\left\langle [PROC1] \underline{\mathbf{p}} \right\rangle_{\frac{3^d+1}{2}} \right) \\ \vdots \\ dist_{(\tilde{g}, [PROC|\mathcal{P}]g,|\mathcal{P}|)}^N \left(\left\langle [PROC|\mathcal{P}] \underline{\mathbf{p}} \right\rangle_{\frac{3^d+1}{2}} \right) \end{pmatrix}^\top, \quad (79)$$

where $\text{PROC}_k g := \langle \mathcal{G} \rangle_k$ and $\text{PROC}_k \underline{\mathbf{p}} := \langle \mathcal{P} \rangle_k$. After the function $\text{dist}_{(\underline{g}, \underline{\mathcal{G}})}^{ALL}$ the particles on all processes are correctly assigned to their respective storage compartments/cell-list cells, but those particles may belong to another process now, except for the particles in the center storage compartment. Therefore, communication between processes is required for the particles that have moved to a different storage compartment/cell. Collecting the particles that moved from the $\beta(w, l)$ -th process to the w th process is done by the function

$$\text{collect}_{(\underline{g}, \mathcal{P}, w)}(\underline{\mathbf{p}}, l) := \begin{pmatrix} \langle \underline{\mathbf{p}} \rangle_1 \\ \vdots \\ \langle \underline{\mathbf{p}} \rangle_{\frac{3^{d+1}}{2}-1} \\ \langle \underline{\mathbf{p}} \rangle_{\frac{3^{d+1}}{2}} \circ \langle \text{PROC}_{\beta(w, l)} \underline{\mathbf{p}} \rangle_{\frac{3}{2}l-1} \left(\frac{3}{4} - \frac{3}{2}l(l) \right) \\ \langle \underline{\mathbf{p}} \rangle_{\frac{3^{d+1}}{2}+1} \\ \vdots \\ \langle \underline{\mathbf{p}} \rangle_{3^d} \end{pmatrix}. \quad (80)$$

All processes collecting particles from the other processes simultaneously could lead to overlapping communications. As discussed above, we aim to prevent this. Therefore, the collection procedure again follows the checkerboard-like pattern. For the k th checkerboard-like pattern, the collection function is

$$\text{collect}_{\underline{g}}^N(\mathcal{P}, k) := \begin{pmatrix} \text{PROC}_1 \underline{\mathbf{p}} \\ \vdots \\ \text{PROC}_{\gamma(k, 1)-1} \underline{\mathbf{p}} \\ \text{PROC}_{\gamma(k, 1)} \underline{\mathbf{p}} * \text{collect}_{(\underline{g}, \mathcal{P}, \gamma(k, 1))} (1, \dots, 3^d) \\ \text{PROC}_{\gamma(k, 1)+1} \underline{\mathbf{p}} \\ \vdots \\ \text{PROC}_{\gamma(k, l)-1} \underline{\mathbf{p}} \\ \text{PROC}_{\gamma(k, l)} \underline{\mathbf{p}} * \text{collect}_{(\underline{g}, \mathcal{P}, \gamma(k, l))} (1, \dots, 3^d) \\ \text{PROC}_{\gamma(k, l)+1} \underline{\mathbf{p}} \\ \vdots \\ \text{PROC}_{\gamma(k, k N_{\text{cell}}^*)-1} \underline{\mathbf{p}} \\ \text{PROC}_{\gamma(k, k N_{\text{cell}}^*)} \underline{\mathbf{p}} * \text{collect}_{(\underline{g}, \mathcal{P}, \gamma(k, k N_{\text{cell}}^*))} (1, \dots, 3^d) \\ \text{PROC}_{\gamma(k, k N_{\text{cell}}^*)+1} \underline{\mathbf{p}} \\ \vdots \\ \text{PROC}_{|\mathcal{P}|} \underline{\mathbf{p}} \end{pmatrix}, \quad (81)$$

where $\text{PROC}_k \underline{\mathbf{p}} := \langle \mathcal{P} \rangle_k$. The complete collection procedure is sequential over the 3^d checkerboard-like patterns. Hence, it takes 3^d steps to finish. The complete collection is

$$\text{collect}_{\tilde{g}}^{ALL}(\mathcal{P}) := \mathcal{P} *_{\text{collect}_{\tilde{g}}^N} (1, \dots, 3^d), \quad (82)$$

which results in each central particle storage compartment of all processes containing the correct particles. The copies from the neighboring processes will then be updated in the subsequent iteration, again by the function $\text{copy}_{\tilde{g}}^{ALL}$.

Taking all functions together, the parallelized state transition step for a distributed-memory particle method can be formally expressed as

$$\tilde{s}_{\tilde{g}}([\mathcal{G}, \mathcal{P}]) = \left[\begin{array}{c} \circ \text{e}(\langle \mathcal{G} \rangle_1) \\ \vdots \\ \circ \text{e}(\langle \mathcal{G} \rangle_{N_{\text{cell}}}) \end{array} \right]^T, \text{collect}_{\tilde{g}}^{ALL} \left(\text{dist}_{(\tilde{g}, \mathcal{G})}^{ALL} \left(\text{step}_{\mathcal{G}}^{ALL} \left(\text{copy}_{\tilde{g}}^{ALL}(\mathcal{P}) \right) \right) \right). \quad (83)$$

We use this parallel state transition step $\tilde{s}_{\tilde{g}}$ to define the parallel state transition function similar to Equation (46):

$$\begin{aligned} \tilde{S}([\mathcal{G}^1, \mathcal{P}^1]) &:= [\mathcal{G}^T, \mathcal{P}^T] \quad \longleftrightarrow \\ f(\langle \mathcal{G}^T \rangle_1) &= \top \quad \wedge \quad \forall t \in \{2, \dots, T\} : [\mathcal{G}^t, \mathcal{P}^t] = \tilde{s}_{\tilde{g}}([\mathcal{G}^{t-1}, \mathcal{P}^{t-1}]) \quad \wedge \quad f(\langle \mathcal{G}^{t-1} \rangle_1) = \perp. \end{aligned} \quad (84)$$

The complete parallelization scheme is summarized by the Nassi–Shneiderman diagram in Figure 6, where the blocks are colored according to the steps from Figure 4.

The algorithm starts by initializing the global variable storage (line 1) and the particle storage (line 3). The global variable storage is filled with the initial global variable, and the center particle storage compartment is filled with the particles from the corresponding cell-list cell. The remaining particle storage compartments are filled with copies of the particles from directly adjacent neighboring cells by copying them from the respective process (lines 5–8), which is repeated for each state (line 4). Then, each process evaluates the state transition step of the particle method (lines 10 and 11), which is only guaranteed to return the correct result for the particles of the center particle storage compartment. The results of the remaining particles may be corrupted by missing interactions with particles outside the process storage. Therefore, the algorithm proceeds to store the particles of the center storage compartment in \mathbf{q} (line 12) and deletes all particles in the particle storage \mathbf{p} (line 13). The particles of the center storage compartment, now in \mathbf{q} , are redistributed to the process's particle storage compartments according to their new positions (lines 14–16). Finally, the algorithm collects for each process all particles that newly belong to it and which are in the corresponding particle storage compartments on other processes (lines 17–21).

3.1 Lemmata

We aim to formally prove that the above distributed-memory parallelization of a particle method is equivalent to the original sequential algorithm. To prepare the proof, we first derive a couple of useful lemmata. The proofs for all lemmata can be found in Appendix A.

LEMMA 1. \bar{I}_l and \bar{I}_l^{-1} are bijections and mutual functional inverses, i.e., $(\bar{I}_l^{-1})^{-1} = \bar{I}_l$.

LEMMA 2. Order independence of a series of interactions follows from the order independence of the interact function:

$${}_1 i_g({}_1 i_g(p_j, p_k), p_{k'}) = {}_1 i_g({}_1 i_g(p_j, p_{k'}), p_{k''}) \quad (85)$$

$$\rightarrow \tilde{p} *_{i_g} \sigma(p_1, \dots, p_n) = \tilde{p} *_{i_g} (p_1, \dots, p_n). \quad (86)$$

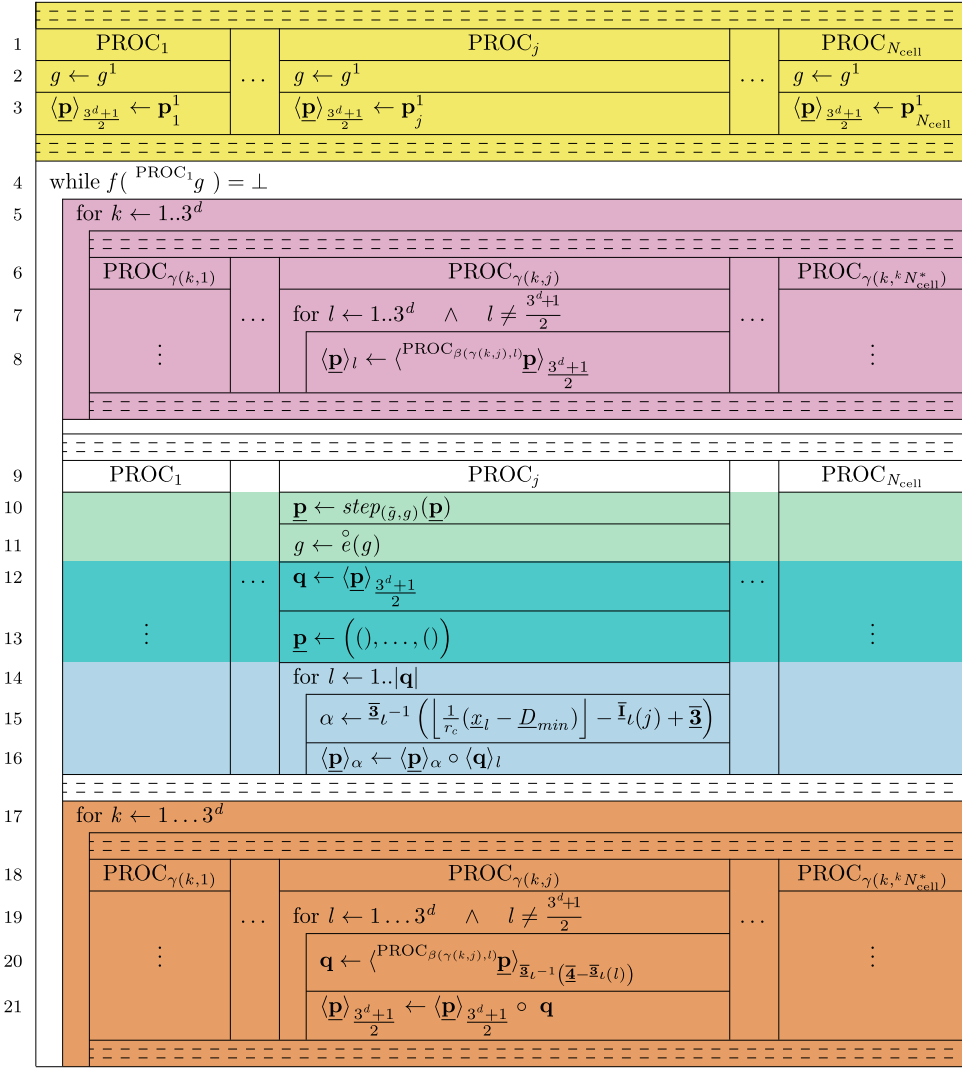


Fig. 6. Nassi-Shneiderman diagram of the distributed-memory parallelization of the outer loop of a particle method with pull interactions. The dashed double lines mark parallel sections of the algorithm. The diagram is colored according to the six basic steps of the algorithm as defined in Figure 4.

LEMMA 3. *The function $\text{copy}_{\bar{g}}^{ALL}$ does not cause overlapping communications. This implies that two processes can only communicate to the same third process or to each other.*

LEMMA 4. *Each process “owns” one distinct cell-list cell. The $\text{copy}_{\bar{g}}^{ALL}$ function copies the particles from all neighbor cells/processes. After the copy procedure, the particle storage compartments of all processes contain the particles of the corresponding cell and copies of the particles from all neighboring cells. Therefore, after each process has executed the copy function, all interaction partners of all particles in the center cell are in process-local memory.*

LEMMA 5. *The function $\text{dist}_{(\bar{g}, \mathcal{G}^1)}^{ALL}$ redistributes particles from the center particle storage compartment to only the other storage compartments on the same process. Therefore, it does not try to place*

them into non-existing storage compartments. Hence,

$$\begin{aligned} \forall w \in \{1, \dots, N_{\text{cell}}\} \forall p_j^1 \in \left\langle \langle \mathcal{P}^1 \rangle_w \right\rangle_{\frac{3^{d+1}}{2}} : p_j^2 := \left\langle \left\langle \text{step}_{\mathcal{G}^1}^{\text{ALL}} \left(\text{copy}_{\bar{g}}^{\text{ALL}}(\mathcal{P}^1) \right) \right\rangle_w \right\rangle_{\frac{3^{d+1}}{2}} \\ \longrightarrow \alpha = \bar{3}_l^{-1} \left(\left\lfloor \frac{1}{r_c} (x_j^2 - D_{\min}) \right\rfloor - \bar{1}_l(w) + \bar{3} \right) \in \{1, \dots, 3^d\}. \end{aligned} \quad (87)$$

LEMMA 6. The function $\text{dist}_{(\bar{g}, \mathcal{G}^1)}^{\text{ALL}}$ places particles only into storage compartments that represent cells inside the computational domain. Hence, processes at the domain border do not have particles in their “outer” storage compartments. Let

$$\underline{w} = (w_1, \dots, w_d)^T := \bar{1}_l(w) \quad (88)$$

$$\underline{\alpha} = (\alpha_1, \dots, \alpha_d)^T := \left\lfloor \frac{1}{r_c} (x_j^2 - D_{\min}) \right\rfloor - \underline{w} + \bar{3}, \quad (89)$$

then

$$\begin{aligned} \forall w \in \{1, \dots, N_{\text{cell}}\} \quad \forall p_j^2 := \left\langle \left\langle \text{step}_{\mathcal{G}^1}^{\text{ALL}} \left(\text{copy}_{\bar{g}}^{\text{ALL}}(\mathcal{P}^1) \right) \right\rangle_w \right\rangle_{\frac{3^{d+1}}{2}} : \\ (k \in \{1, \dots, d\} \wedge w_k = 1) \rightarrow \alpha_k \in \{2, 3\} \wedge (k \in \{1, \dots, d\} \wedge w_k = I_k) \rightarrow \alpha_k \in \{1, 2\}. \end{aligned} \quad (90)$$

LEMMA 7. $\bar{\mathcal{P}}^1 := \text{step}_{\mathcal{G}^1}^{\text{ALL}}(\text{copy}_{\bar{g}}^{\text{ALL}}(\mathcal{P}^1))$ does not necessarily fulfill the condition, Equation (63), to iterate $\text{step}_{\mathcal{G}^2}^{\text{ALL}}(\text{copy}_{\bar{g}}^{\text{ALL}}(\bar{\mathcal{P}}^1))$. To avoid interact_g missing interactions due to incomplete neighborhoods, the particles in $\bar{\mathcal{P}}^1$ need to be redistributed such that Equation (63) is fulfilled. The decomposition of the initial permuted particle tuple is stored in the center storage compartment of the processes. Hence, Equation (63) can be rewritten for all state transition steps as

$$\forall p_j^t \in \bigcirc_{v=1}^{N_{\text{cell}}} \left\langle \langle \mathcal{P}^t \rangle_v \right\rangle_{\frac{3^{d+1}}{2}} : p_j^t \in \left\langle \langle \mathcal{P}^t \rangle_w \right\rangle_{\frac{3^{d+1}}{2}}, \quad (91)$$

where

$$w = \bar{1}_l^{-1} \left(\left\lfloor \frac{1}{r_c} (x_j^t - D_{\min}) \right\rfloor + \bar{1} \right). \quad (92)$$

This is achieved by the two functions $\text{dist}_{(\bar{g}, \mathcal{G}^1)}^{\text{ALL}}$ and $\text{collect}_{\bar{g}}^{\text{ALL}}$.

3.2 Proof of Correctness and Equivalence with Sequential Algorithm

THEOREM 1. On the particles stored in the center storage compartments of the processes, and under the assumptions stated at the beginning of this section, the presented parallelization scheme for particle methods on distributed-memory computers in Figure 6 computes the same results, except for particle ordering, as the underlying sequential particle method in Figure 3. Hence,

$$S([\mathcal{G}^1, \mathcal{P}^1]) = \left[\langle \mathcal{G}^T \rangle_1, \sigma^* \left(\bigcirc_{w=1}^{N_{\text{cell}}} \left\langle \langle \mathcal{P}^T \rangle_w \right\rangle_{\frac{3^{d+1}}{2}} \right) \right], \quad (93)$$

where

$$[\mathcal{G}^T, \mathcal{P}^T] = \tilde{s}([\mathcal{G}^1, \mathcal{P}^1]). \quad (94)$$

We prove that each sequential state transition step s , Equation (45), is equivalent to each parallel state transition step \tilde{s} , Equation (83), and that both algorithms terminate after the same number of state transitions. The sequential state transition step was defined as follows:

$$s([g, \mathbf{p}]) := \left[e^\circ(\bar{g}), \bar{\mathbf{p}} \right] \quad \text{with } [\bar{g}, \bar{\mathbf{p}}] = \epsilon^N \left(g, \iota^{N \times U}([g, \mathbf{p}]) \right). \quad (\text{Equation (45)})$$

Under the condition that the interact function, Equation (52), and the neighborhood function, Equation (54), are independent of previous interactions, we can rewrite the third interact subfunction $t^{N \times U}$ for pull interaction particle methods, Equation (51), to [24]

$$t^{N \times U}([g, \mathbf{p}]) = \begin{pmatrix} p_1 *_{i_g} \langle \mathbf{p} \rangle_{u(g, \mathbf{p}, 1)} \\ \vdots \\ p_{|\mathbf{p}|} *_{i_g} \langle \mathbf{p} \rangle_{u(g, \mathbf{p}, |\mathbf{p}|)} \end{pmatrix}^\top. \quad (95)$$

We also know [24] that if the evolve function does not change the global variable, Equation (55), we can write the state transition step as

$$s([g, \mathbf{p}]) := \left[\overset{\circ}{e}(g), {}_2\epsilon^N \left(\begin{pmatrix} p_1 *_{i_g} \langle \mathbf{p} \rangle_{u(g, \mathbf{p}, 1)} \\ \vdots \\ p_{|\mathbf{p}|} *_{i_g} \langle \mathbf{p} \rangle_{u(g, \mathbf{p}, |\mathbf{p}|)} \end{pmatrix}^\top \right) \right]. \quad (96)$$

We use proof by induction to prove equivalence for the global variable and the particles separately. We start with the global variable. Following the construction of \mathcal{G}^1 , we directly get

$$\begin{pmatrix} g^2 \\ \vdots \\ g^2 \end{pmatrix}^\top = \begin{pmatrix} \overset{\circ}{e}(g^1) \\ \vdots \\ \overset{\circ}{e}(g^1) \end{pmatrix}^\top = \begin{pmatrix} \overset{\circ}{e}(\langle \mathcal{G}^1 \rangle_1) \\ \vdots \\ \overset{\circ}{e}(\langle \mathcal{G}^1 \rangle_{N_{\text{cell}}}) \end{pmatrix}^\top = \begin{pmatrix} \langle \mathcal{G}^2 \rangle_1 \\ \vdots \\ \langle \mathcal{G}^2 \rangle_{N_{\text{cell}}} \end{pmatrix}^\top = \mathcal{G}^2. \quad (97)$$

This is the base case for the proof by induction. For the induction step, we start from

$$\begin{pmatrix} g^t \\ \vdots \\ g^t \end{pmatrix}^\top = \mathcal{G}^t \quad (98)$$

and get

$$\begin{pmatrix} g^{t+1} \\ \vdots \\ g^{t+1} \end{pmatrix}^\top = \begin{pmatrix} \overset{\circ}{e}(g^t) \\ \vdots \\ \overset{\circ}{e}(g^t) \end{pmatrix}^\top = \begin{pmatrix} \overset{\circ}{e}(\langle \mathcal{G}^t \rangle_1) \\ \vdots \\ \overset{\circ}{e}(\langle \mathcal{G}^t \rangle_{N_{\text{cell}}}) \end{pmatrix}^\top = \begin{pmatrix} \langle \mathcal{G}^{t+1} \rangle_1 \\ \vdots \\ \langle \mathcal{G}^{t+1} \rangle_{N_{\text{cell}}} \end{pmatrix}^\top = \mathcal{G}^{t+1}. \quad (99)$$

This completes the part for the global variable.

We now prove that the sequential state transition function S and the distributed-memory state transition function \tilde{S} stop after the same number of states. For the state transition S , the stop function f only depends on g^t . For \tilde{S} , f depends only on $\langle \mathcal{G}^t \rangle_1$. From this and $g^t = \langle \mathcal{G}^t \rangle_1$ follows

$$\underline{{}_1S([g^1, \mathbf{p}^1])} = g^T = \langle \mathcal{G}^T \rangle_1 = \langle \tilde{S}([g^1, \mathbf{p}^1]) \rangle_1 \quad (100)$$

$${}_2S([g^1, \mathbf{p}^1]) = {}_2\epsilon^N \left(\begin{pmatrix} p_1^1 *_{i_{g^1}} \langle \mathbf{p}^1 \rangle_{u(g^1, \mathbf{p}^1, 1)} \\ \vdots \\ p_{|\mathbf{p}^1|}^1 *_{i_{g^1}} \langle \mathbf{p}^1 \rangle_{u(g^1, \mathbf{p}^1, |\mathbf{p}^1|)} \end{pmatrix}^\top \right). \quad (101)$$

The interact function is order independent, Equation (53), and the permutation $\pi(\mathbf{p}^1)$, Equation (62), of the initial particle tuple sorts the particles into cells. The neighborhood function is restricted

to not using indices, Equation (48). Hence, for a permuted particle tuple, the neighborhood function returns the same result as for the unpermuted particle tuple, with the same permutation also applied to the result. This leads to

$${}_2s([g^1, \mathbf{p}^1]) = {}_2\epsilon^N \left(\pi^{-1} \left(\left(\begin{array}{c} p_{\pi(1)}^1 *_{i_{g^1}} \langle \pi(\mathbf{p}^1) \rangle_{u(g^1, \pi(\mathbf{p}^1), \pi(1))} \\ \vdots \\ p_{\pi(|\mathbf{p}^1|)}^1 *_{i_{g^1}} \langle \pi(\mathbf{p}^1) \rangle_{u(g^1, \pi(\mathbf{p}^1), \pi(|\mathbf{p}^1|))} \end{array} \right)^\top \right) \right). \quad (102)$$

From the construction of the initial particle storages of all processes \mathcal{P}^1 , Equations (63) to (67), as well as Lemmata 3 and 4, we derive

$${}_2s([g^1, \mathbf{p}^1]) = {}_2\epsilon^N \left(\pi^{-1} \left(\bigcirc_{w=1}^{N_{\text{cell}}} \left(\begin{array}{c} {}^w p_1^1 *_{i_{g^1}} {}^w \mathbf{q}^1_{u(g^1, {}^w \mathbf{q}^1, {}^w z+1)} \\ \vdots \\ {}^w p_n^1 *_{i_{g^1}} {}^w \mathbf{q}^1_{u(g^1, {}^w \mathbf{q}^1, {}^w z+{}^w n)} \end{array} \right)^\top \right) \right), \quad (103)$$

where the tuple of all particles in the storage of the w th process is

$${}^w \mathbf{q}^1 := \bigcirc_{l=1}^{3^d} \left\langle \left\langle \text{copy}_{\tilde{g}}^{ALL}(\mathcal{P}^1) \right\rangle_w \right\rangle_l, \quad (104)$$

the number of particles in ${}^w \mathbf{q}^1$ in front of the particles of the central storage compartment is

$${}^w z := \sum_{l=1}^{\frac{3^d+1}{2}-1} \left| \left\langle \left\langle \text{copy}_{\tilde{g}}^{ALL}(\mathcal{P}^1) \right\rangle_w \right\rangle_l \right|, \quad (105)$$

the number of particles in the central storage compartment is

$${}^w n = \left| \left\langle \left\langle \text{copy}_{\tilde{g}}^{ALL}(\mathcal{P}^1) \right\rangle_w \right\rangle_{\frac{3^d+1}{2}} \right|, \quad (106)$$

and the particles in the central storage compartment are

$$({}^w p_1^1, \dots, {}^w p_n^1) = \left\langle \left\langle \text{copy}_{\tilde{g}}^{ALL}(\mathcal{P}^1) \right\rangle_w \right\rangle_{\frac{3^d+1}{2}}. \quad (107)$$

The interactions of the particles with their neighbors in Equation (103) resemble the *interaction_g* function (Equation (3)):

$${}_2s([g^1, \mathbf{p}^1]) = {}_2\epsilon^N \left(\pi^{-1} \left(\bigcirc_{w=1}^{N_{\text{cell}}} \text{interaction}_{g^1} \left(\left\langle \left\langle \text{copy}_{\tilde{g}}^{ALL}(\mathcal{P}^1) \right\rangle_w \right\rangle \right) \right) \right). \quad (108)$$

The evolution function cannot change the global variable (condition in Equation (55)). Hence, the second evolution subfunction ϵ^N cannot change the global variable either and is, therefore, independent of the ordering of the particles. Since the evolve function e can create or destroy particles, the permutation π^{-1} cannot rearrange the result of ϵ^N to the sequential state transition result. But the calculation on the particles is identical. Hence, there exists a permutation $\tilde{\pi}^{-1}$ that rearranges the result such that

$${}_2s([g^1, \mathbf{p}^1]) = \tilde{\pi}^{-1} \left({}_2\epsilon^N \left(\bigcirc_{w=1}^{N_{\text{cell}}} \text{interaction}_{g^1} \left(\left\langle \left\langle \text{copy}_{\tilde{g}}^{ALL}(\mathcal{P}^1) \right\rangle_w \right\rangle \right) \right) \right). \quad (109)$$

All processes calculate the $interaction_g$ function independently. Hence, the ϵ^N function can also be executed on each process independently. This means that

$${}_2s([g^1, \mathbf{p}^1]) = \tilde{\pi}^{-1} \left(\bigcirc_{w=1}^{N_{\text{cell}}} {}_2\epsilon^N \left(interaction_{g^1} \left(\left\langle copy_{\tilde{g}}^{ALL}(\mathcal{P}^1) \right\rangle_w \right) \right) \right). \quad (110)$$

The combination of ${}_2\epsilon^N$ and $interaction_g$ is the same as the $step_{(\tilde{g}, g^1)}$ function, Equation (74), at the center storage compartment:

$${}_2s([g^1, \mathbf{p}^1]) = \tilde{\pi}^{-1} \left(\bigcirc_{w=1}^{N_{\text{cell}}} \left\langle step_{(\tilde{g}, g^1)} \left(\left\langle copy_{\tilde{g}}^{ALL}(\mathcal{P}^1) \right\rangle_w \right) \right\rangle_{\frac{3^d+1}{2}} \right). \quad (111)$$

The function $step_{\mathcal{G}^1}^{ALL}$, Equation (75), calculates the step for all processes, leading to

$${}_2s([g^1, \mathbf{p}^1]) = \tilde{\pi}^{-1} \left(\bigcirc_{w=1}^{N_{\text{cell}}} \left\langle \left\langle step_{\mathcal{G}^1}^{ALL} \left(copy_{\tilde{g}}^{ALL}(\mathcal{P}^1) \right) \right\rangle_w \right\rangle_{\frac{3^d+1}{2}} \right). \quad (112)$$

$\overline{\mathcal{P}}^1 := step_{\mathcal{G}^1}^{ALL}(copy_{\tilde{g}}^{ALL}(\mathcal{P}^1))$ does not necessarily fulfill the condition, Equation (63), to iterate $step_{\mathcal{G}^2}^{ALL}(copy_{\tilde{g}}^{ALL}(\overline{\mathcal{P}}^1))$. For $interact_g$ to not miss interactions due to incomplete neighborhoods, the particles in $\overline{\mathcal{P}}^1$ need to be redistributed such that the condition in Equation (63) is fulfilled. This is achieved by the functions $dist_{(\tilde{g}, \mathcal{G}^1)}^{ALL}$, Equation (79), and $collect_{\tilde{g}}^{ALL}$, Equation (82), as proven in the Lemmata 5, 6, and 7. Hence,

$${}_2s([g^1, \mathbf{p}^1]) = \tilde{\pi}'^{-1} \left(\bigcirc_{w=1}^{N_{\text{cell}}} \left\langle \left\langle collect_{\tilde{g}}^{ALL} \left(dist_{(\tilde{g}, \mathcal{G}^1)}^{ALL} \left(step_{\mathcal{G}^1}^{ALL} \left(copy_{\tilde{g}}^{ALL}(\mathcal{P}^1) \right) \right) \right) \right\rangle_w \right\rangle_{\frac{3^d+1}{2}} \right), \quad (113)$$

where $\tilde{\pi}'^{-1}$ is a new permutation. We insert the definition of ${}_2\tilde{s}_{\tilde{g}}$ to get

$${}_2s([g^1, \mathbf{p}^1]) = \tilde{\pi}'^{-1} \left(\bigcirc_{w=1}^{N_{\text{cell}}} \left\langle \left\langle \underbrace{{}_2\tilde{s}_{\tilde{g}}([G^1, \mathcal{P}^1])}_{=: \mathcal{P}^2} \right\rangle_w \right\rangle_{\frac{3^d+1}{2}} \right), \quad (114)$$

$$[g^2, \mathbf{p}^2] = \tilde{\pi}'^{-1} \left(\bigcirc_{w=1}^{N_{\text{cell}}} \left\langle \langle \mathcal{P}^2 \rangle_w \right\rangle_{\frac{3^d+1}{2}} \right). \quad (115)$$

This is the base case for the proof by induction. For the induction step, we start from

$$[g^t, \mathbf{p}^t] = \tilde{\pi}''^{-1} \left(\bigcirc_{w=1}^{N_{\text{cell}}} \left\langle \langle \mathcal{P}^t \rangle_w \right\rangle_{\frac{3^d+1}{2}} \right). \quad (116)$$

We can assume that \mathcal{P}^t fulfills the same conditions as \mathcal{P}^2 and, hence, also as \mathcal{P}^1 , especially the condition in Equation (63) (or Equations (103) and (104)). Then, we can define a new particle method where $[g^t, \mathbf{p}^t]$ is the instance and \mathcal{P}^t its corresponding cell-list-based distribution onto processes. Using Lemmata 5, 6, and 7, we derive that

$${}_2s([g^t, \mathbf{p}^t]) = \tilde{\pi}'''^{-1} \left(\bigcirc_{w=1}^{N_{\text{cell}}} \left\langle \left\langle \underbrace{{}_2\tilde{s}_{\tilde{g}}([G^t, \mathcal{P}^t])}_{=: \mathcal{P}^{t+1}} \right\rangle_w \right\rangle_{\frac{3^d+1}{2}} \right), \quad (117)$$

and hence,

$$[g^{t+1}, \mathbf{p}^{t+1}] = \tilde{\pi}'''^{-1} \left(\bigcirc_{w=1}^{N_{\text{cell}}} \left\langle \langle \mathcal{P}^{t+1} \rangle_w \right\rangle_{\frac{3^d+1}{2}} \right). \quad (118)$$

We do this now for all t until $f(g^t) = \top$. Together with the derivation of the proof for the global variable, this leads to

$$S([g^1, \mathbf{p}^1]) = \left[\langle \mathcal{G}^T \rangle_1, \sigma^* \left(\bigcirc_{w=1}^{N_{\text{cell}}} \langle \langle \mathcal{P}^T \rangle_w \rangle_{\frac{3^{d+1}}{2}} \right) \right], \quad (119)$$

where $[\mathcal{G}^T, \mathcal{P}^T] = \tilde{S}([\mathcal{G}^1, \mathcal{P}^1])$.

Hence, the present parallelization scheme produces the same result up to a different ordering of the particles. The particles are permuted by an unknown permutation σ^* . This proves that the distributed-memory parallelization scheme is correct for order-independent particle methods.

4 Bounds on Time Complexity and Parallel Scalability

For the parallelization scheme considered in the previous section, we derive bounds on the time complexity and the parallel scalability in both the weak and strong scaling cases. Since these bounds are for a minimal parallelization scheme, they should provide useful limits for more elaborate parallelization schemes, and they indicate the relevance of the considered parallelization scheme. The detailed derivations and closed-form formulas can be found in Appendix B.

An algorithm's time complexity describes the runtime required by a machine to execute that algorithm. It depends on the input size of the algorithm. For a particle method, the input size is the length of the initial tuple \mathbf{p}^1 . We assume that constants bound the sizes of the global variable and of each particle. We further assume that the algorithm terminates in finite time. Hence, upper bounds exist for all functions' time complexities.

Since particles can move, the neighborhood search function checks each pair of particles to see if they are neighbors. This has a complexity in $O(n^2)$, where n is the number of particles. Fast neighbor list algorithms like cell lists [15] reduce this to $O(n)$ under the assumption of even and constant particle density. We can confirm this by comparing the simplified time complexity of the cell-list-based scheme, Equation (267), with the time complexity of the sequential state transition, Equation (260), where we eliminated the dependency of the time complexity of the neighborhood function on the particle number. Our cell-list-based scheme scales linearly with the number of particles ($N_{\text{cell}} n_{\text{max}} \approx N_{\text{p}}^{\text{max}}$) and not quadratically as it would without the assumption of even and constant particle density. Hence, the speedup is $O(N_{\text{p}}^{\text{max}})$ as derived in Equation (277) and visualized in Figure 7(a).

Second, Amdahl's law [3] provides an upper bound on the speedup of the cell-list scheme on multiple processors when the problem size is fixed for increasing processors n_{CPU} (strong scaling). In this case, we can increase n_{CPU} until we reach the number of cells N_{cell} . After that, there will be no further speedup. But also, until then, we find a step-like behavior with increasing n_{CPU} because cells cannot be split across CPUs, as visualized in Figure 7(b).

Third, Gustafson's law [12] provides an upper bound on the speedup of the cell-list scheme on multiple processors when the ratio of problem size to process number is constant while increasing the number of processors (weak scaling); $\frac{N_{\text{cell}}}{n_{\text{CPU}}} = \text{const}$. For a perfectly fitting processor interconnect network topology, we predict a linear speedup on average with steps as visualized in Figure 7(c).

Overall, the scheme behaves as expected for cell-list algorithms.

5 Conclusions and Discussion

Particle methods encompass a wide range of computer simulation algorithms, such as DEMs [33], MD [2], RKPMs [20], PSE [8, 9], DC-PSE [5, 30], and SPH [10, 22]. Although particle methods are widely used, mathematically defined, and practically parallelized in software frameworks such as the PPM Library [29], OpenFPM [16], POOMA [26], or FDPS [17], little work has been

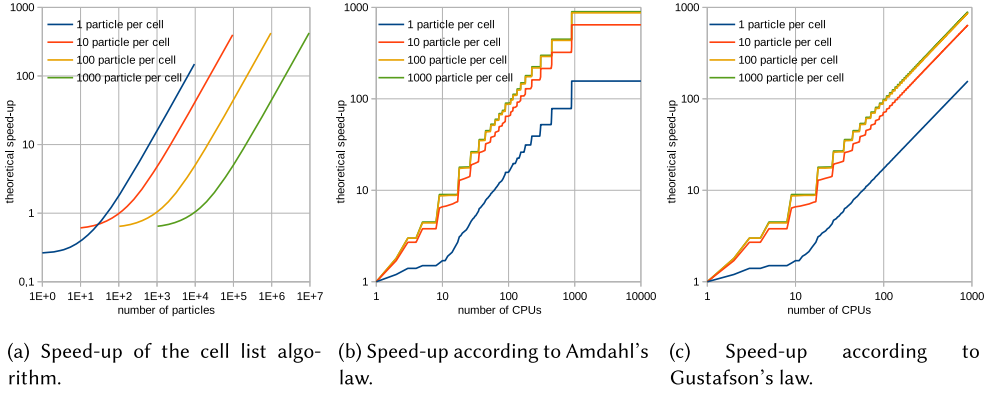


Fig. 7. Theoretical speedup bounds. The constants are chosen to be $d = 2$, $C_u = C_\alpha = C_\beta = C_\gamma = 1$, $\tau_i = \tau_e = 3$, and $\tau_f = \tau_\epsilon = 1$. For Amdahl's and Gustafson's laws, $N_{\text{cell}} = 900$.

done to formally investigate generic parallelization schemes for particle methods and prove their correctness, complexity, and scalability.

Here, we took a first step in this direction by providing a distributed-memory parallelization scheme for particle methods independent of specific applications. We proved the correctness of the presented cell-list parallelization scheme by showing equivalence to the sequential particle methods definition under certain assumptions, which we also defined here. Finally, we derived upper bounds on the time complexity of the proposed scheme executed on both sequential and parallel computers, and we discussed the parallel scalability limits in the weak and strong scaling case.

The presented parallelization scheme is not novel, and it is, in fact, a consensus of what is commonly implemented in software. Therefore, our analysis is of immediate practical relevance for general-purpose particle methods frameworks, even for critical applications, as the proof guarantees algorithmic correctness, leaving room only for errors induced by finite-precision arithmetics.

The present analysis was deliberately done for a minimal, simple parallelization scheme. Practically implemented schemes tend to be more elaborate and, e.g., deviate from our assumption of a uniform Cartesian domain decomposition, isotropic and uniform particle interaction radii, or the collision-free checkerboard communication pattern. They are, however, similar in structure, such that the bounds derived here can still provide useful limits. Also, many more elaborate schemes can be reduced to the present minimal scheme to show their correctness. This would not yield conceptually different or novel results, albeit the mathematical notation of the proofs would become a lot more involved.

We focused here on the distributed-memory parallelization of particle methods. The shared-memory case has already been proven earlier [24]. We also did not consider hardware-specific parallelization strategies, like streaming multi-threading on GPUs [13, 16], nor instruction-level parallelism in SIMD processor architectures [25, 31]. Another limitation of the present parallelization scheme is that it considers only pull interactions between the particles. This neglects the potential runtime benefits of symmetric interaction evaluations, where both interaction partners are changed simultaneously. However, pull interactions are suitable for more computer architectures, and they are more general, encompassing also particle methods that do not possess symmetric interaction functions (e.g., DC-PSE [30]). Pull interactions also have lower communication overhead, since no changes to ghost particles need to be sent back to the source processes [29]. In this sense, the scalability bounds derived here are valid also for push-pull schemes, which are expected to scale worse but potentially have lower overall runtime. We also restricted the neighborhood

function such that the cell-list strategy became applicable. This limits the expressiveness of the particle method, for example, neglecting spatially adaptive or anisotropic interaction neighborhoods. However, more complex neighborhoods could be modeled using the function Ω in the neighborhood step, maintaining the validity of the present correctness proof also in this case.

Further, the constraint that particles are not allowed to leave the domain keeps domain decomposition simple. Otherwise, subdomains would dynamically need to be added or removed as necessary, resulting in a more complex and dynamic mapping of subdomains to processes. In practice, however, this assumption is not limiting, since a bounding box can always be chosen as large as needed to leave room for the free movement of the particles within the finite time of a simulation. Moreover, it does not preclude irregular particle distributions within the domain to, e.g., represent complex or moving and deforming geometries [4, 28]. We also restricted particles to not move further than the cutoff radius in a single-state iteration or time step of the algorithm. Since the cutoff radius determines the smallest possible cell size and individual cells cannot be split across multiple processes, this guarantees that processes only need to communicate with their immediately adjacent neighbors. In practice, this constraint can be relaxed to require particles to never move further than one subdomain per state iteration. Our choice, however, enabled us to prove correctness in the limit case of the highest degree of parallelism, where each cell-list cell is assigned to a different process. Additionally, we restricted the global variable to only be changed by the evolve function of the global variable and not by the evolve function of any particle. Therefore, no global operations were allowed that would require additional communication. Still, global variable changes can be independently computed locally, keeping them synchronized without communication. Finally, the time complexity of the checkerboard-like communication scheme does not have optimal pre-factor, leaving many processes inactive. Nevertheless, it scales linearly with the number of processes and abstracts the internal scheduling of the network subsystem, providing a correct asymptotic bound on the scalability while permitting correctness proofs.

Notwithstanding these limitations, the present parallelization scheme, with its proof of equivalence to the sequential particle methods definition, stands in contrast with algorithm-specific or empirically tested parallelization schemes. We therefore hope that our analysis paves the way for future research on the theory of parallel scientific simulation algorithms and the engineering of provably correct parallel software implementations.

Future theoretical work could optimize the presented scheme for computer architectures with parallel or asynchronously clocked communication. Global operations could then also be allowed, and the checkerboard-like communication pattern could be relaxed, leading to an improved scalability pre-factor. Also, the network topology of the machine's interconnect could be explicitly incorporated into the parallelization scheme. Furthermore, proofs for push and symmetric interaction schemes could be beneficial for specific use cases, as well as combining parallelization schemes for shared and distributed memory to better match the heterogeneous architecture of modern supercomputers. On the software engineering side, future work could leverage the presented parallelization scheme and proofs to design a new generation of theoretically founded software frameworks. They would potentially be more predictable, suitable for security- and safety-critical applications, and more maintainable and understandable as they are based on a common formal framework [27].

Overall, the presented proven parallelization scheme provides a way to parallelize a class of particle methods on distributed-memory systems with full knowledge of its validity, performance, and assumptions. We proved that it computes the same result as the underlying sequential particle method. Therefore, using it in a general framework for particle methods is suitable, even for critical computations, since the proof guarantees that the parallelization does not change the results. We therefore hope that the present work will generate downstream investigation and studies in the theory of algorithms for scientific computing.

Appendices

A Proofs of Lemmata

We prove the helping lemmata for the proof that the presented parallelization scheme is equivalent to the original particle method.

PROOF OF LEMMA 1. We need to prove that \bar{I}_l and \bar{I}_l^{-1} are bijective, and we need to prove that \bar{I}_l and \bar{I}_l^{-1} are the inverse functions of one another.

$$\forall j \in \left\{ 1, \dots, \prod_{\delta=1}^d I_\delta \right\} : \bar{I}_l^{-1} \left(\bar{I}_l(j) \right) = j \quad (120)$$

$$\wedge \forall \underline{j} \in \mathbb{N}^d \cap \left[\bar{\mathbf{1}}, \bar{\mathbf{1}} \right] : \bar{I}_l \left(\bar{I}_l^{-1}(\underline{j}) \right) = \underline{j}. \quad (121)$$

First, we prove Equation (120):

$$\bar{I}_l^{-1} \left(\bar{I}_l(j) \right) = 1 + \left((j-1) - \left\lfloor \frac{j-1}{I_1} \right\rfloor I_1 + 1 \right) + \left(\left(\left\lfloor \frac{j-1}{I_1} \right\rfloor - \left\lfloor \frac{j-1}{I_1 I_2} \right\rfloor I_2 + 1 \right) - 1 \right) I_1 \quad (122)$$

$$+ \left(\left(\left\lfloor \frac{j-1}{I_1 I_2} \right\rfloor - \left\lfloor \frac{j-1}{I_1 I_2 I_3} \right\rfloor I_3 + 1 \right) - 1 \right) I_1 I_2 + \dots + \left(\left(\left\lfloor \frac{j-1}{\prod_{t=1}^{l-1} I_t} \right\rfloor - \left\lfloor \frac{j-1}{\prod_{t=1}^l I_t} \right\rfloor I_l + 1 \right) - 1 \right) \prod_{t=1}^{l-1} I_t \quad (123)$$

$$+ \left(\left(\left\lfloor \frac{j-1}{\prod_{t=1}^l I_t} \right\rfloor - \left\lfloor \frac{j-1}{\prod_{t=1}^{l+1} I_t} \right\rfloor I_{l+1} + 1 \right) - 1 \right) \prod_{t=1}^l I_t \quad (124)$$

$$+ \dots + \left(\left(\left\lfloor \frac{j-1}{\prod_{t=1}^{d-2} I_t} \right\rfloor - \left\lfloor \frac{j-1}{\prod_{t=1}^{d-1} I_t} \right\rfloor I_{d-1} + 1 \right) - 1 \right) \prod_{t=1}^{d-2} I_t + \left(\left(\left\lfloor \frac{j-1}{\prod_{t=1}^{d-1} I_t} \right\rfloor + 1 \right) - 1 \right) \prod_{t=0}^{d-1} I_t \quad (125)$$

$$= 1 + (j-1) - \underbrace{\left\lfloor \frac{j-1}{I_1} \right\rfloor I_1 + \left\lfloor \frac{j-1}{I_1} \right\rfloor I_1}_{=0} - \underbrace{\left\lfloor \frac{j-1}{I_1 I_2} \right\rfloor I_2 I_1 + \left\lfloor \frac{j-1}{I_1 I_2} \right\rfloor I_1 I_2}_{=0} \quad (126)$$

$$- \left\lfloor \frac{j-1}{I_1 I_2 I_3} \right\rfloor I_3 I_1 I_2 + \dots + \left\lfloor \frac{j-1}{\prod_{t=1}^{l-1} I_t} \right\rfloor \prod_{t=1}^{l-1} I_t - \underbrace{\left\lfloor \frac{j-1}{\prod_{t=1}^l I_t} \right\rfloor \prod_{t=1}^l I_t + \left\lfloor \frac{j-1}{\prod_{t=1}^l I_t} \right\rfloor \prod_{t=1}^l I_t}_{=0} \quad (127)$$

$$- \left\lfloor \frac{j-1}{\prod_{t=1}^{l+1} I_t} \right\rfloor \prod_{t=1}^{l+1} I_t + \dots + \left\lfloor \frac{j-1}{\prod_{t=1}^{d-2} I_t} \right\rfloor \prod_{t=1}^{d-2} I_t - \underbrace{\left\lfloor \frac{j-1}{\prod_{t=1}^{d-1} I_t} \right\rfloor I_{d-1} \prod_{t=1}^{d-2} I_t + \left\lfloor \frac{j-1}{\prod_{t=1}^{d-1} I_t} \right\rfloor \prod_{t=1}^{d-1} I_t}_{=0} \quad (128)$$

$$= j. \quad (129)$$

Second, we prove Equation (121) by subdividing the resulting vector in its entries, starting with the first entry:

$$\left\langle \bar{I}_l \left(\bar{I}_l^{-1}(\underline{j}) \right) \right\rangle_1 = \left(\left(1 + (j_1 - 1) + (j_2 - 1) I_1 + \dots + (j_d - 1) \prod_{t=1}^{d-1} I_t \right) - 1 \right) \quad (130)$$

$$- \left\lfloor \frac{\left(1 + (j_1 - 1) + \dots + (j_d - 1) \prod_{t=1}^{d-1} I_t \right) - 1}{I_1} \right\rfloor I_1 + 1 \quad (131)$$

$$\begin{aligned}
& \left[\begin{aligned}
& = \left[\frac{\left(1+(j_1-1)+\dots+(j_d-1)\prod_{t=1}^{d-1} I_t\right)-1}{I_1} \right] I_1 \\
& = \left[\frac{j_1-1}{I_1} + \frac{(j_2-1)I_1+\dots+(j_d-1)\prod_{t=1}^{d-1} I_t}{I_1} \right] I_1 \\
& = \left[\underbrace{\frac{j_1-1}{I_1}}_{1 \leq j_1 \leq I_1} + \underbrace{(j_2-1) + \dots + (j_d-1) \prod_{t=2}^{d-1} I_t}_{\in \mathbb{N}_0} \right] I_1 \\
& = \left((j_2-1) + \dots + (j_d-1) \prod_{t=2}^{d-1} I_t \right) I_1 \\
& = (j_2-1)I_1 + \dots + (j_d-1) \prod_{t=1}^{d-1} I_t
\end{aligned} \right] \tag{132} \\
& = (j_1-1) + (j_2-1)I_1 + \dots + (j_d-1) \prod_{t=1}^{d-1} I_t - \left((j_2-1)I_1 + \dots + (j_d-1) \prod_{t=1}^{d-1} I_t \right) + 1 \tag{133} \\
& = j_1. \tag{134}
\end{aligned}$$

We proceed with all entries except the first and last entry:

$$\forall l \in \{2, \dots, d-1\}: \tag{135}$$

$$\left\langle \bar{I}_l \left(\bar{I}_l^{-1}(j) \right) \right\rangle_l = \left[\frac{\left(1+(j_1-1)+\dots+(j_d-1)\prod_{t=1}^{d-1} I_t\right)-1}{\prod_{t=1}^{l-1} I_t} \right] \tag{136}$$

$$- \left[\frac{\left(1+(j_1-1)+\dots+(j_d-1)\prod_{t=1}^{d-1} I_t\right)-1}{\prod_{t=1}^l I_t} \right] I_l + 1 \tag{137}$$

$$\text{auxiliary calculations:} \tag{138}$$

$$\begin{aligned}
& \left[\begin{aligned}
& = \left[\frac{\left(1+(j_1-1)+\dots+(j_d-1)\prod_{t=1}^{d-1} I_t\right)-1}{\prod_{t=1}^{l-1} I_t} \right] \\
& = \left[\frac{(j_1-1)+\dots+(j_{l-1}-1)\prod_{t=1}^{l-2} I_t}{\prod_{t=1}^{l-1} I_t} + \frac{(j_l-1)\prod_{t=1}^{l-1} I_t + \dots + (j_d-1)\prod_{t=1}^{d-1} I_t}{\prod_{t=1}^{l-1} I_t} \right] \\
& \text{auxiliary calculations:} \\
& \quad \begin{aligned}
& 0 \leq \underbrace{(j_1-1)+\dots+(j_{l-1}-1)}_{1 \leq j_1 \leq I_1} \prod_{t=1}^{l-2} I_t \\
& \leq I_1 - 1 + (I_2-1)I_1 + \dots + (I_{l-1}-1)\prod_{t=1}^{l-2} I_t \\
& = I_1 - 1 + I_2I_1 - I_1 + \dots + \prod_{t=1}^{l-1} I_t - \prod_{t=1}^{l-2} I_t \\
& = \prod_{t=0}^{l-1} I_t - 1 \\
& \prod_{t=0}^{l-1} I_t - 1 < \prod_{t=0}^{l-1} I_t
\end{aligned} \\
& = \left[\underbrace{\frac{(j_1-1)+\dots+(j_{l-1}-1)\prod_{t=1}^{l-2} I_t}{\prod_{t=1}^{l-1} I_t}}_{\geq 0, < 1} + \underbrace{(j_l-1) + \dots + (j_d-1) \prod_{t=l}^{d-1} I_t}_{\in \mathbb{N}_0} \right] \\
& = (j_l-1) + \dots + (j_d-1) \prod_{t=l}^{d-1} I_t \\
& = \left((j_l-1) + \dots + (j_d-1) \prod_{t=l}^{d-1} I_t \right) - \left((j_{l+1}-1) + \dots + (j_d-1) \prod_{t=l+1}^{d-1} I_t \right) I_l + 1 \tag{139}
\end{aligned} \right] \tag{140}
\end{aligned}$$

$$=(j_l - 1) + (j_{l+1} - 1)I_l + \cdots + (j_d - 1) \prod_{t=l}^{d-1} I_t - (j_{l+1} - 1)I_l - \cdots - (j_d - 1) \prod_{t=l}^{d-1} I_t + 1 \quad (141)$$

$$=j_l. \quad (142)$$

We finish with the last entry:

$$\left\langle \bar{I}_l \left(\bar{I}_l^{-1}(\underline{j}) \right) \right\rangle_d = \left\lfloor \frac{\left(1 + (j_1 - 1) + \cdots + (j_d - 1) \prod_{t=1}^{d-1} I_t \right) - 1}{\prod_{t=1}^{d-1} I_t} \right\rfloor + 1 \quad (143)$$

$$= \left\lfloor \underbrace{\frac{1 + (j_1 - 1) + \cdots + (j_{d-1} - 1) \prod_{t=1}^{d-2} I_t}{\prod_{t=1}^{d-1} I_t}}_{<1} + \underbrace{(j_d - 1)}_{\in \mathbb{N}_0} \right\rfloor + 1 \quad (144)$$

$$=j_d. \quad (145)$$

Taking these three results together, we can conclude

$$\bar{I}_l \left(\bar{I}_l^{-1}(\underline{j}) \right) = \begin{bmatrix} j_1 \\ j_2 \\ \vdots \\ j_d \end{bmatrix} \quad (146)$$

$$= \underline{j}. \quad (147)$$

We need to prove that \bar{I}_l and \bar{I}_l^{-1} are bijective, and we need to prove that \bar{I}_l and \bar{I}_l^{-1} are the inverse functions of one another. \square

PROOF OF LEMMA 2. We need to prove that the order independence of the interact function follows the order independence of a series of interactions:

$${}_1 i_g(1 i_g(p_j, p_k), p_{k'}) = {}_1 i_g(1 i_g(p_j, p_{k'}), p_{k''}) \quad (148)$$

$$\rightarrow \tilde{p} *_1 i_g \sigma(p_1, \dots, p_n) = \tilde{p} *_1 i_g (p_1, \dots, p_n). \quad (149)$$

The proof idea is to use the bubble sort strategy to go from an arbitrary permutation σ of the interacting particles to the original order. Therefore, we prove that the necessary swap of two consecutive particles does not change the result:

$${}_1 i_g(1 i_g(p_j, p_k), p_{k'}) = {}_1 i_g(1 i_g(p_j, p_{k'}), p_{k''}) \quad (150)$$

$$\rightarrow \tilde{p} *_1 i_g \sigma(p_1, \dots, p_n) = \underbrace{\tilde{p} *_1 i_g (\sigma(p_1), \dots, \sigma(p_{j-1}), \sigma(p_j), \sigma(p_{j+1}), \dots, \sigma(p_n))}_{=: p'} \quad (151)$$

$$= p' *_1 i_g (\sigma(p_j), \sigma(p_{j+1}), \dots, \sigma(p_n)) \quad (152)$$

$$= {}_1 i_g(1 i_g(p', \sigma(p_j)), \sigma(p_{j+1})) *_1 i_g (\sigma(p_{j+2}), \dots, \sigma(p_n)) \quad (153)$$

$$= {}_1 i_g(1 i_g(p', \sigma(p_{j+1})), \sigma(p_j)) *_1 i_g (\sigma(p_{j+2}), \dots, \sigma(p_n)) \quad (154)$$

$$= \tilde{p} *_1 i_g (\sigma(p_1), \dots, \sigma(p_{j-1}), \sigma(p_{j+1}), \sigma(p_j), \sigma(p_{j+2}), \dots, \sigma(p_n)) \quad (155)$$

$$\text{using bubble sort} \quad (156)$$

$$= \tilde{p} *_1 i_g (p_1, \dots, p_n). \quad (157)$$

\square

PROOF OF LEMMA 3. We need to prove that the function $copy_g^{ALL}$ does not induce overlapping communications.

Overlapping means that two processes communicate either to the same other third process or to each other. The potential overlapping communications are avoided by letting only some processes communicate simultaneously. The communicating processes are distributed in a checkerboard-like structure. Between two communicating processes are always two passive processes. Since each process communicates only with its direct neighbor processes, there cannot be an overlapping communication. To prove that, we need to prove

$$\begin{aligned} \forall k, l', l'' \in \{1, \dots, 3^d\} \quad \forall j', j'' \in \{1, \dots, {}^k N_{cell}^*\} : \\ j' \neq j'' \rightarrow \beta(\gamma(k, j'), l') \neq \beta(\gamma(k, j''), l'') \vee \beta(\gamma(k, j'), l') = \text{undef.}, \end{aligned} \quad (158)$$

where k is the number of the checkerboard-like pattern; $\gamma(k, j)$, Equation (60), is an index of a reading process; and $\beta(\gamma(k, j), l)$, Equation (61), its l th neighbor with which it communicates. Hence, two communicating processes do not have a common process with which they communicate. If β is undefined, there is no process.

Inserting the definition of γ into β results in

$$\beta(\gamma(k, j), l) = \bar{1}_l^{-1} \left(\bar{1}_l \left(\bar{1}_l^{-1} \left({}^k \bar{1}_l^*(j) \cdot 3 + \bar{3}_l(k) - \bar{3} \right) \right) + \bar{3}_l(l) - \bar{2} \right) \quad (159)$$

$$= \bar{1}_l^{-1} \left({}^k \bar{1}_l^*(j) \cdot 3 + \bar{3}_l(k) + \bar{3}_l(l) - \bar{5} \right). \quad (160)$$

From here on, we do a proof by contradiction. Therefore, we negate the statement we want to prove and derive absurdity.

Under the assumption that

$$j' \neq j'' \wedge \beta(\gamma(k, j'), l') = \beta(\gamma(k, j''), l'') \wedge \beta(\gamma(k, j'), l') \neq \text{undef.} \quad (161)$$

$\xleftrightarrow{\text{Equation (159)}}$

$$\begin{aligned} \wedge \bar{1}_l^{-1} \left({}^k \bar{1}_l^*(j') \cdot 3 + \bar{3}_l(k) + \bar{3}_l(l') - \bar{5} \right) = \bar{1}_l^{-1} \left({}^k \bar{1}_l^*(j'') \cdot 3 + \bar{3}_l(k) + \bar{3}_l(l'') - \bar{5} \right) \\ \wedge \bar{1}_l^{-1} \left({}^k \bar{1}_l^*(j') \cdot 3 + \bar{3}_l(k) + \bar{3}_l(l') - \bar{5} \right) \neq \text{undef.} \end{aligned} \quad (162)$$

$\xleftrightarrow{\text{Lemma 1}}$

$$\begin{aligned} \wedge \left(\begin{aligned} & j' \neq j'' \\ & {}^k \bar{1}_l^*(j') \cdot 3 + \bar{3}_l(k) + \bar{3}_l(l') - \bar{5} = {}^k \bar{1}_l^*(j'') \cdot 3 + \bar{3}_l(k) + \bar{3}_l(l'') - \bar{5} \\ & \vee \bar{1}_l^{-1} \left({}^k \bar{1}_l^*(j') \cdot 3 + \bar{3}_l(k) + \bar{3}_l(l') - \bar{5} \right) = \text{undef.} \end{aligned} \right) \\ \wedge \bar{1}_l^{-1} \left({}^k \bar{1}_l^*(j') \cdot 3 + \bar{3}_l(k) + \bar{3}_l(l') - \bar{5} \right) \neq \text{undef.} \end{aligned} \quad (163)$$

\longleftrightarrow

$$\begin{aligned} \wedge {}^k \bar{1}_l^*(j') \cdot 3 + \bar{3}_l(k) + \bar{3}_l(l') - \bar{5} = {}^k \bar{1}_l^*(j'') \cdot 3 + \bar{3}_l(k) + \bar{3}_l(l'') - \bar{5} \\ \wedge \bar{1}_l^{-1} \left({}^k \bar{1}_l^*(j') \cdot 3 + \bar{3}_l(k) + \bar{3}_l(l') - \bar{5} \right) \neq \text{undef.} \end{aligned} \quad (164)$$

Definition 13

$$j' \neq j'' \wedge \underbrace{k_{\bar{1}}^{\bar{1}}(j')}_{\in \mathbb{N}_1^d} \cdot 3 + \underbrace{\bar{3}_l(l')}_{\in [\bar{1}, \bar{3}]} = \underbrace{k_{\bar{1}}^{\bar{1}}(j'') \cdot 3}_{\in \mathbb{N}_1^d} + \underbrace{\bar{3}_l(l'')}_{\in [\bar{1}, \bar{3}]} \wedge k_{\bar{1}}^{\bar{1}}(j') \cdot 3 + \bar{3}_l(k) + \bar{3}_l(l') - \bar{5} \in \mathbb{N}_1^d \cap [\bar{1}, \bar{1}] \quad (165)$$

\longleftrightarrow

$$j' \neq j'' \wedge \underbrace{k_{\bar{1}}^{\bar{1}}(j') - k_{\bar{1}}^{\bar{1}}(j'') \cdot 3}_{\in \mathbb{Z}^d} = \underbrace{\bar{3}_l(l'') - \bar{3}_l(l')}_{\in [-\bar{2}, \bar{2}]} \wedge k_{\bar{1}}^{\bar{1}}(j') \cdot 3 + \bar{3}_l(k) + \bar{3}_l(l') - \bar{5} \in \mathbb{N}_1^d \cap [\bar{1}, \bar{1}] \quad (166)$$

Definition 13

$$j' \neq j'' \wedge \underbrace{k_{\bar{1}}^{\bar{1}}(j') - k_{\bar{1}}^{\bar{1}}(j'') \cdot 3}_{\in \mathbb{Z}^d} = \underbrace{\bar{3}_l(l'') - \bar{3}_l(l')}_{\in [-\bar{2}, \bar{2}]} = \bar{0} \wedge k_{\bar{1}}^{\bar{1}}(j') \cdot 3 + \bar{3}_l(k) + \bar{3}_l(l') - \bar{5} \in \mathbb{N}_1^d \cap [\bar{1}, \bar{1}] \quad (167)$$

Lemma 1

$$j' \neq j'' \wedge j' = j'' \wedge l' = l'' \wedge k_{\bar{1}}^{\bar{1}}(j') \cdot 3 + \bar{3}_l(k) + \bar{3}_l(l') - \bar{5} \in \mathbb{N}_1^d \cap [\bar{1}, \bar{1}] \quad (168)$$

$$\not\vdash (\text{contradiction}) \quad (169)$$

\longrightarrow

$$\forall k, l', l'' \in \{1, \dots, 3^d\} \quad \forall j', j'' \in \{1, \dots, k_{N_{cell}^*}\} : \quad (170)$$

$$j' \neq j'' \rightarrow \beta(\gamma(k, j'), l') \neq \beta(\gamma(k, j''), l'') \vee \beta(\gamma(k, j'), l') = \text{undef.}$$

□

PROOF OF LEMMA 4. We need to prove, after the copy_g^{ALL} function, that the storages of all processes contain the particles of the corresponding cells and all their neighbor particles. Hence, first, we need to prove that each process executes the copy_g^{ALL} function and, second, that on a process after the copy_g^{ALL} function all neighbor particles of all particles of the center storage compartment are in the storage.

First, the idea is to prove that $\gamma(k, j)$ is bijective and the codomain is the set of all indices of the processes $\{1, \dots, N_{cell}\}$:

$$\gamma : \{1, \dots, 3^d\} \times \{1, \dots, k_{N_{cell}^*}\} \rightarrow \{1, \dots, N_{cell}\} \text{ is bijective.} \quad (171)$$

We define a helper function $\underline{\gamma}^*$, prove it is bijective, and transform it into γ .

We declare

$$\underline{\gamma}^* : \{1, \dots, 3\}^d \times \mathbb{N}^d \cap [\bar{1}, k_{\bar{1}}^{\bar{1}}] \rightarrow \mathbb{N}^d \cap [\bar{1}, \bar{1}] \quad (172)$$

and define

$$\underline{\gamma}^*(\underline{k}, \underline{j}) := \underline{k} + (\underline{j} - \bar{1}) \cdot 3. \quad (173)$$

One element of $\underline{\gamma}^*$ is $\langle \underline{\gamma}^*(\underline{k}, \underline{j}) \rangle_w$ and we can rewrite it by

$$\langle \underline{\gamma}^*(\underline{k}, \underline{j}) \rangle_w = 1 + (k_w - 1) + (j_w - 1) \cdot 3 = \underbrace{k_{\bar{1}}^{\bar{1}} \iota^{-1} \left(\begin{pmatrix} k_w \\ j_w \end{pmatrix} \right)}_{\text{is bijective } (\iota \text{ is bijective})}, \quad (174)$$

where $k_{\underline{J}_w} := \binom{3}{k I_w^*}$. The domain of the arguments j and k depend on each other. Hence, it is not obvious that \underline{Y}^* is not leaving the codomain $\mathbb{N}^d \cap [\underline{1}, \bar{1}]$. \underline{Y}^* is monotone. Therefore, it is sufficient to prove that \underline{Y}^* is in the codomain for the smallest and largest arguments. The minimal value of \underline{Y}^* is

$$\min \left(\left\langle \underline{Y}^*(k, j) \right\rangle_w \right) = 1 + (1 - 1) \cdot 3 = \underline{1}. \quad (175)$$

This is in the codomain. The maximal value of \underline{Y}^* is

$$\max \left(\left\langle \underline{Y}^*(k, j) \right\rangle_w \right) = k_w^{\max} + \left(k_w^{\max} I_w^* - 1 \right) \cdot 3. \quad (176)$$

Using the definition of $k_w^{\max} \bar{I}^*$, Equation (58), leads to

$$\max \left(\left\langle \underline{Y}^*(k, j) \right\rangle_w \right) = k_w^{\max} + \left(\left\lfloor \frac{1}{3} (I_w - k_w^{\max} + 3) \right\rfloor - 1 \right) \cdot 3. \quad (177)$$

We substitute

$$I_w - k_w^{\max} + 3 =: \underbrace{T'}_{\in \mathbb{N}} \cdot 3 + \underbrace{T''}_{\in \{0, 1, 2\}} \quad (178)$$

and get

$$\max \left(\left\langle \underline{Y}^*(k, j) \right\rangle_w \right) = k_w^{\max} + \left(\left\lfloor \frac{3T'}{3} + \frac{T''}{3} \right\rfloor - 1 \right) \cdot 3 \quad (179)$$

$$= k_w^{\max} + (T' - 1) \cdot 3. \quad (180)$$

We rearrange the substitution to

$$3T' = I_w - k_w^{\max} + 3 - T'' \quad (181)$$

and plug it in:

$$\max \left(\left\langle \underline{Y}^*(k, j) \right\rangle_w \right) = k_w^{\max} + I_w - k_w^{\max} + 3 - T'' - 3 \quad (182)$$

$$= I_w - \underbrace{T''}_{\in \{0, 1, 2\}}. \quad (183)$$

This means

$$\max \left(\left\langle \underline{Y}^*(k, j) \right\rangle_w \right) \leq I_w. \quad (184)$$

It remains to be shown that I_w can be reached. This is only possible if $T'' = 0$. Hence, we need to prove

$$\exists k_w \in \{1, 2, 3\} : T'' = 0. \quad (185)$$

From $T'' = 0$ it follows that

$$\frac{3T' + T''}{3} = \frac{I_w + 3 - k_w^{\max}}{3} = n \in \mathbb{N} \quad (186)$$

→

$$n + \frac{k_w^{\max} - 1}{3} = \frac{I_w + 2}{3}. \quad (187)$$

Taking the *floor* on both sides leads to

$$\left\lfloor n + \underbrace{\frac{k_w^{\max} - 1}{3}}_{\in \{0, 1, 2\}} \right\rfloor = \left\lfloor \frac{I_w + 2}{3} \right\rfloor \quad (188)$$

→

$$n = \left\lfloor \frac{I_w + 2}{3} \right\rfloor. \quad (189)$$

Inserting this into Equation (187) leads to

$$k_w = ((I_w + 3) - 1) - \left\lfloor \frac{(I_w + 3) - 1}{3} \right\rfloor \cdot 3 + 1 = \left\langle \bar{1}_l(I_w + 3) \right\rangle_1 \in \{1, 2, 3\}, \quad (190)$$

where $\bar{\mathbf{J}} = (3, \dots)^T$. From this, we can conclude

$$\forall w \in \{1, \dots, d\} : \left\langle \underline{\gamma}^* \right\rangle_w \text{ is bijective} \quad (191)$$

→

$$\underline{\gamma}^* \text{ is bijective.} \quad (192)$$

Now we need to transform $\underline{\gamma}^*$ to γ . We substitute \underline{k} and \underline{j} by

$$\underline{k} = \bar{3}_l(k), \quad \underline{j} = {}^{k\bar{1}}_l(j). \quad (193)$$

We know that ι is bijective (Lemma 1, Definition 13) and therefore,

$$k \in \{1, \dots, 3^d\}, \quad j \in N^d \cap \left[\bar{\mathbf{1}}, {}^{k\bar{1}}_l \right]. \quad (194)$$

We also know ι^{-1} is bijective. Hence,

$$\bar{1}_l^{-1} \left(\underline{\gamma}^* \left(\bar{3}_l(k), {}^{k\bar{1}}_l(j) \right) \right) = \bar{1}_l^{-1} \left(\bar{3}_l(k) + \left({}^{k\bar{1}}_l(j) - \bar{\mathbf{1}} \right) \cdot 3 \right) \quad (195)$$

$$= \gamma(k, j). \quad (196)$$

This means γ is bijective, and the $copy_g^{ALL}$ function is executed for all processes.

Second, it remains to be proven that the neighbor particles of all particles of the center storage compartment are in the storage. The function $copy_{(\bar{g}, \mathcal{P}, w)}(\mathbf{p}, l)$, Equation (70), copies the central $\left(\frac{3^d+1}{2}\right)$ storage compartment from the $\beta(w, l)$ -th process to the l th storage compartment on the w th process. We start from the vectorial index view to find the corresponding cell/process where the particles are for the l th storage compartment on the w th process. The vectorial index of the w th process is $\bar{1}_l(w)$ and the vectorial index of the l th storage compartment is $\bar{3}_l(l) \in \{1, 2, 3\}$. The storage center compartment $l = \left(\frac{3^d+1}{2}\right)$ corresponds to the cell belonging to the process. The rest of the storage compartments should represent the surrounding cells. Therefore, we need to shift the storage compartment index by -2 in all dimensions to account for this. Hence, the vectorial index of the corresponding process for the l th storage compartment of the w th process is

$$\bar{1}_l(w) + \bar{3}_l(l) - \bar{\mathbf{2}}. \quad (197)$$

Transforming the vectorial index to a scalar index results in

$$\bar{1}_l^{-1} \left(\bar{1}_l(w) + \bar{3}_l(l) - \bar{\mathbf{2}} \right) = \beta(w, l). \quad (198)$$

Hence, the $\beta(w, l)$ -th process has the corresponding particles in its center storage compartment. The storage gets the particles from the surrounding cells by the $copy$ function from the other processes. With this, we can derive the domain area that the storage covers. The vectorial indices of the cells that belong to the domain are

$$\left\{ \bar{1}_l(\beta(w, l)) : l \in \{1, \dots, 3^d\} \right\} = \left[\bar{1}_l(w) - \bar{\mathbf{1}}, \bar{1}_l(w) + \bar{\mathbf{1}} \right] \cap \left[\bar{\mathbf{1}}, \bar{\mathbf{1}} \right] \cap \mathbb{N}^d. \quad (199)$$

Per Equation (63), the particles that belong to the w th cell are

$$p_j = (\dots, x_j, \dots) \in \mathbf{p}_w^1 : \bar{l}(w) = \left\lfloor \frac{1}{r_c} (x_j - \underline{D}_{\min}) \right\rfloor + \bar{\mathbf{1}} \quad (200)$$

→

$$p_j \in \mathbf{p}_w^1 : \bar{l}(w) - \bar{\mathbf{1}} = \left\lfloor \frac{1}{r_c} (x_j - \underline{D}_{\min}) \right\rfloor. \quad (201)$$

→

$$p_j \in \mathbf{p}_w^1 : \frac{1}{r_c} (x_j - \underline{D}_{\min}) \in \left[\bar{l}(w) - \bar{\mathbf{1}}, \bar{l}(w) - \bar{\mathbf{1}} + \bar{\mathbf{1}} \right) \quad (202)$$

→

$$p_j \in \mathbf{p}_w^1 : x_j \in \left[\left(\bar{l}(w) - \bar{\mathbf{1}} \right) \cdot r_c + \underline{D}_{\min}, \bar{l}(w) \cdot r_c + \underline{D}_{\min} \right). \quad (203)$$

Taking all cells/compartments, Equation (199), of the storage of the w th process leads to

$$p_j \in \bigcirc_{l=1}^{3^d} \langle \text{PROC}_w \underline{\mathbf{p}}^1 \rangle_l : x_j \in \left[\left(\left(\bar{l}(w) - \bar{\mathbf{1}} \right) - \bar{\mathbf{1}} \right) \cdot r_c + \underline{D}_{\min}, \left(\bar{l}(w) + \bar{\mathbf{1}} \right) \cdot r_c + \underline{D}_{\min} \right) \cap \left[\underline{D}_{\min}, \underline{D}_{\max} \right), \quad (204)$$

where $[\bar{\mathbf{1}}, \bar{\mathbf{1}}]$ translates to $[\underline{D}_{\min}, \underline{D}_{\max}]$, Equation (56):

→

$$p_j \in \bigcirc_{l=1}^{3^d} \langle \text{PROC}_w \underline{\mathbf{p}}^1 \rangle_l : x_j \in \left[\left(\bar{l}(w) - \bar{\mathbf{2}} \right) \cdot r_c + \underline{D}_{\min}, \left(\bar{l}(w) + \bar{\mathbf{1}} \right) \cdot r_c + \underline{D}_{\min} \right) \cap \left[\underline{D}_{\min}, \underline{D}_{\max} \right). \quad (205)$$

It remains to be proven that the neighbor particles of all particles of the central storage compartment of the w th process are in the storage. Hence, the domain area covered by the storage includes the area covered by the neighborhood function u . Let $p_k \in \mathbf{p}^1$ and $p_j \in \langle \mathbf{p}^1 \rangle_{u(g, \mathbf{p}^1, k)}$, Equation (48); then

$$|x_k - x_j| \leq r_c \quad (206)$$

→

$$x_j \in [x_k - \bar{\mathbf{1}} \cdot r_c, x_k + \bar{\mathbf{1}} \cdot r_c]. \quad (207)$$

Let now $p_k \in \mathbf{p}_w^1$; then we know, Equation (203):

$$x_k \in \left[\left(\bar{l}(w) - \bar{\mathbf{1}} \right) \cdot r_c + \underline{D}_{\min}, \bar{l}(w) \cdot r_c + \underline{D}_{\min} \right) \quad (208)$$

→

$$x_j \in \left[\left(\bar{l}(w) - \bar{\mathbf{1}} \right) \cdot r_c + \underline{D}_{\min} - \bar{\mathbf{1}} \cdot r_c, \bar{l}(w) \cdot r_c + \underline{D}_{\min} + \bar{\mathbf{1}} \cdot r_c \right) \cap \left[\underline{D}_{\min}, \underline{D}_{\max} \right) \quad (209)$$

→

$$x_j \in \left[\left(\bar{l}(w) - \bar{\mathbf{2}} \right) \cdot r_c + \underline{D}_{\min}, \left(\bar{l}(w) + \bar{\mathbf{1}} \right) \cdot r_c + \underline{D}_{\min} \right) \cap \left[\underline{D}_{\min}, \underline{D}_{\max} \right). \quad (210)$$

This is the same as Equation (205), which means that the neighbor particles of all particles of the center storage compartment of the w th process are all in the storage of the w th process, and this is true for all processes. \square

PROOF OF LEMMA 5. We need to prove that the function $\text{dist}_{(\bar{g}, \bar{\mathcal{G}}^1)}^{ALL}$ places for each process the particles from the center storage compartments only to the other storage compartments of the same processes. Hence, it does not try to place them in a not-existing storage compartment. Hence,

$$\begin{aligned} \forall w \in \{1, \dots, N_{\text{cell}}\} \forall p_j^1 \in \langle \langle \mathcal{P}^1 \rangle_w \rangle_{\frac{3^d+1}{2}} : p_j^2 := \left\langle \left\langle \left\langle \text{step}_{\bar{\mathcal{G}}^1}^{ALL} \left(\text{copy}_{\bar{g}}^{ALL}(\mathcal{P}^1) \right) \right\rangle_w \right\rangle_{\frac{3^d+1}{2}} \right\rangle_j \\ \longrightarrow \alpha = \bar{\mathbf{3}}_l^{-1} \left(\left\lfloor \frac{1}{r_c} (x_j^2 - \underline{D}_{\min}) \right\rfloor - \bar{l}(w) + \bar{\mathbf{3}} \right) \in \{1, \dots, 3^d\}. \quad (211) \end{aligned}$$

All particles have a position $\underline{x} \in \mathbb{R}^d$; hence, $p_j^1 = (\dots, \underline{x}_j^1, \dots)$ and $p_j^2 = (\dots, \underline{x}_j^2, \dots)$. The condition in Equation (50) restricts the movement of the particles to be smaller than r_c , meaning for $\underline{\Delta x} := \underline{x}_j^2 - \underline{x}_j^1$:

$$|\underline{\Delta x}| \leq r_c. \quad (212)$$

Definition 5

$$|\underline{\Delta x}| = \left| \begin{pmatrix} \Delta x_1 \\ \vdots \\ \Delta x_d \end{pmatrix} \right| = \sqrt{\Delta x_1^2 + \dots + \Delta x_d^2} \leq r_c \quad (213)$$

→

$$\forall k \in \{1, \dots, d\} : r_c \leq \Delta x_k \leq r_c \quad (214)$$

→

$$\forall k \in \{1, \dots, d\} : \Delta x_k \in [-r_c, r_c]. \quad (215)$$

Using this, we can rewrite α from

$$\alpha = \bar{\mathfrak{I}}_l^{-1} \left(\left\lfloor \frac{1}{r_c} (\underline{x}_j^2 - D_{\min}) \right\rfloor - \bar{\mathfrak{I}}_l(w) + \bar{\mathfrak{I}} \right) \quad (216)$$

to

$$\alpha = \bar{\mathfrak{I}}_l^{-1} \left(\left\lfloor \frac{1}{r_c} (\underline{x}_j^1 + \underline{\Delta x} - D_{\min}) \right\rfloor - \bar{\mathfrak{I}}_l(w) + \bar{\mathfrak{I}} \right). \quad (217)$$

We put the index for the dimension as pre-subscript:

$$\alpha = \bar{\mathfrak{I}}_l^{-1} \left(\begin{pmatrix} \left\lfloor \frac{1}{r_c} ({}_1x_j^1 + {}_1\Delta x - {}_1D_{\min}) \right\rfloor - \left\langle \bar{\mathfrak{I}}_l(w) \right\rangle_1 + 3 \\ \vdots \\ \left\lfloor \frac{1}{r_c} ({}_dx_j^1 + {}_d\Delta x - {}_dD_{\min}) \right\rfloor - \left\langle \bar{\mathfrak{I}}_l(w) \right\rangle_d + 3 \end{pmatrix} \right). \quad (218)$$

We know from Equation (63) that

$$\forall p_j^1 \in \mathbf{p}^1 : p_j^1 \in \mathbf{p}_w^1 \text{ with } w = \bar{\mathfrak{I}}_l^{-1} \left(\left\lfloor \frac{1}{r_c} (\underline{x}_j^1 - D_{\min}) \right\rfloor + \bar{\mathfrak{I}} \right). \quad (219)$$

Taking this into α we get

$$\alpha = \bar{\mathfrak{I}}_l^{-1} \left(\underbrace{\begin{pmatrix} \left\lfloor \frac{{}_1x_j^1 - {}_1D_{\min}}{r_c} + \frac{{}_1\Delta x}{r_c} \right\rfloor - \left\langle \bar{\mathfrak{I}}_l(w) \right\rangle_1 + 3 \\ \vdots \\ \left\lfloor \frac{{}_dx_j^1 - {}_dD_{\min}}{r_c} + \frac{{}_d\Delta x}{r_c} \right\rfloor - \left\langle \bar{\mathfrak{I}}_l(w) \right\rangle_d + 3 \end{pmatrix}}_{\in \mathbb{N} \cap [\bar{\mathfrak{I}}, \bar{\mathfrak{I}}]} \right) \quad (220)$$

→

$$\underline{\alpha} \in \{1, \dots, 3^d\}. \quad (221)$$

Hence, under the condition in Equation (63), no particle leaves the storage compartments of its process through the function $dist_{(\bar{g}, \mathcal{G}^1)}^{ALL}$. \square

PROOF OF LEMMA 6. We need to prove that the function $dist_{(\bar{g}, \mathcal{G}^1)}^{ALL}$ places particles only inside storage compartments that represent cells inside the domain. Hence, processes at the domain's border do not have particles in their outer storage compartments. Let

$$\underline{w} = (w_1, \dots, w_d)^T := \bar{I}_l(\underline{w}) \quad (222)$$

$$\underline{\alpha} = (\alpha_1, \dots, \alpha_d)^T := \left\lfloor \frac{1}{r_c} (x_j^2 - D_{\min}) \right\rfloor - \underline{w} + \bar{\mathbf{3}}. \quad (223)$$

Then

$$\forall w \in \{1, \dots, N_{cell}\} \forall p_j^2 := \left\langle \left\langle \left\langle \text{step}_{\mathcal{G}^1}^{ALL} \left(\text{copy}_{\bar{g}}^{ALL} (\mathcal{P}^1) \right) \right\rangle_w \right\rangle_{\frac{3^d+1}{2}} \right\rangle_j : \\ (k \in \{1, \dots, d\} \wedge w_k = 1) \rightarrow \alpha_k \in \{2, 3\} \wedge (k \in \{1, \dots, d\} \wedge w_k = I_k) \rightarrow \alpha_k \in \{1, 2\}. \quad (224)$$

We do a proof by contradiction. First, we prove the statement with $w_k = 1$. Assuming

$$\exists k \in \{1, \dots, d\} : w_k = 1 \wedge \alpha_k = 1, \quad (225)$$

then

$$\alpha_k = 1 = \left\lfloor \frac{1}{r_c} (kx_j^2 - kD_{\min}) \right\rfloor - 1 + 3. \quad (226)$$

→

$$-1 = \left\lfloor \frac{1}{r_c} (kx_j^2 - kD_{\min}) \right\rfloor \quad (227)$$

→

$$\frac{1}{r_c} (kx_j^2 - kD_{\min}) = [-1, 0) \quad (228)$$

→

$$kx_j^2 = [-r_c + kD_{\min}, 0) \quad \not\leq. \quad (229)$$

This contradicts the condition that no particle leaves the domain, Equation (49). Hence, in this case $\alpha_k \in \{2, 3\}$. Now we prove the second statement with $w_k = I_k$. Assuming

$$\exists k \in \{1, \dots, d\} : w_k = I_k \wedge \alpha_k = 3, \quad (230)$$

then

$$\alpha_k = 3 = \left\lfloor \frac{1}{r_c} (kx_j^2 - kD_{\min}) \right\rfloor - I_k + 3 \quad (231)$$

→

$$I_k = \left\lfloor \frac{1}{r_c} (kx_j^2 - kD_{\min}) \right\rfloor \quad (232)$$

→

$$\frac{1}{r_c} (kx_j^2 - kD_{\min}) = [I_k, I_k + 1) \quad (233)$$

→

$$kx_j^2 = [I_k r_c + kD_{\min}, I_k r_c + kD_{\min} + r_c) \quad \not\leq \quad (234)$$

def. \bar{I} (Equation (56))

→

$$kx_j^2 > kD_{\max} \quad \not\leq. \quad (235)$$

This contradicts the condition that no particle leaves the domain, Equation (49). Hence, in this case $\alpha_k \in \{1, 2\}$. Taking both cases together, the function $dist_{(\bar{g}, \mathcal{G}^1)}^{ALL}$ places particles only inside storage compartments that represent cells inside the domain. \square

PROOF OF LEMMA 7. $\bar{\mathcal{P}}^1 := \text{step}_{\mathcal{G}^1}^{ALL}(\text{copy}_{\bar{g}}^{ALL}(\mathcal{P}^1))$ does not necessarily fulfill the condition, Equation (63), to reiterate $\text{step}_{\mathcal{G}^2}^{ALL}(\text{copy}_{\bar{g}}^{ALL}(\bar{\mathcal{P}}^1))$. The function interact_g could miss interaction due to incomplete neighborhoods. Therefore, the particles in $\bar{\mathcal{P}}^1$ need to be redistributed such that the requirement in Equation (63) is fulfilled. The decomposition of the initial permuted particle tuple is stored in the center storage compartment of the processes. Hence, the condition in Equation (63) can be rewritten for all state transition steps as

$$\forall p_j^t \in \bigcirc_{v=1}^{N_{cell}} \langle \langle \mathcal{P}^t \rangle_v \rangle_{\frac{3^{d+1}}{2}} : p_j^t \in \langle \langle \mathcal{P}^t \rangle_w \rangle_{\frac{3^{d+1}}{2}}, \quad (236)$$

where

$$w = \bar{1}_l^{-1} \left(\left\lfloor \frac{1}{r_c} (x_j^t - D_{\min}) \right\rfloor + \bar{1} \right). \quad (237)$$

We need to prove that the two functions $\text{dist}_{(\bar{g}, \mathcal{G}^1)}^{ALL}$ and $\text{collect}_{\bar{g}}^{ALL}$ achieve this.

We prove this by induction. Regarding the base case, we know that the condition in Equations (91) and (92) is true for \mathcal{P}^1 by the definition in Equation (63).

Regarding the induction step, we need to prove that if \mathcal{P}^t fulfills the condition in Equations (91) and (92), then

$$\mathcal{P}^{t+1} = \text{collect}_{\bar{g}}^{ALL} \left(\text{dist}_{(\bar{g}, \mathcal{G}^1)}^{ALL} \left(\text{step}_{\mathcal{G}^1}^{ALL} \left(\text{copy}_{\bar{g}}^{ALL}(\mathcal{P}^t) \right) \right) \right) \quad (238)$$

fulfills also the condition.

$\text{dist}_{(\bar{g}, \mathcal{G}^1)}^{ALL}$, Equation (79), empties all storage compartments except the center storage compartment and then distributes all particles of the center storage compartments according to their position to the other storage compartments of each process. α gives the index of the new storage compartment for each particle. It is trivial that it considers all particles and all processes since it simply iterates over them.

$\text{collect}_{\bar{g}}^{ALL}$, Equation (82), copies for each process its particles from the other processes corresponding to storage compartments. It uses the same strategy as the function $\text{copy}_{\bar{g}}^{ALL}$ to avoid overlapping conditions; hence, Lemma 3 is also valid for the $\text{collect}_{\bar{g}}^{ALL}$ function. The $\text{collect}_{\bar{g}}^{ALL}$ function takes the particles from the $\bar{3}_l^{-1}(\bar{4} - \bar{3}_l(l))$ -th storage compartment of the $\beta(w, l)$ -th process and stores them in the center storage compartment of the w th process.

We need to prove that a particle from each storage compartment of all processes ends up in the suitable process's central storage compartment according to the condition in Equations (91) and (92). We choose without restricting generality a particle p from the $\bar{3}_l^{-1}(\bar{4} - \bar{3}_l(l))$ -th storage compartment of the $\beta(w, l)$ -th process. Hence,

$$p = (\dots, x, \dots) \in \left\langle \left\langle \text{step}_{\mathcal{G}^1}^{ALL} \left(\text{copy}_{\bar{g}}^{ALL}(\mathcal{P}^t) \right) \right\rangle_{\beta(w, l)} \right\rangle_{\bar{3}_l^{-1}(\bar{4} - \bar{3}_l(l))}. \quad (239)$$

Combining this with the distribution of the $\text{dist}_{(\bar{g}, \mathcal{G}^1)}^{ALL}$ function, we get

$$\underbrace{\alpha}_{\bar{3}_l^{-1}(\bar{4} - \bar{3}_l(l))} := \bar{3}_l^{-1} \left(\left\lfloor \frac{1}{r_c} (x - D_{\min}) \right\rfloor - \bar{1}_l(\underbrace{j}_{\beta(w, l)}) + \bar{3} \right), \quad (240)$$

where α is set to $\bar{3}_l^{-1}(\bar{4} - \bar{3}_l(l))$, the index of the storage compartment from which the particles are collected, and j is set to $\beta(w, l)$, the index of the process from which it is collected. Hence, our

particle is in that storage compartment of that process and gets distributed by the mechanism of the $dist_{(\tilde{g}, \mathcal{G}^1)}^{ALL}$ function. We put the indices in and transform

$$\bar{\mathfrak{z}}_{l^{-1}} \left(\bar{\mathfrak{4}} - \bar{\mathfrak{z}}_l(l) \right) = \bar{\mathfrak{z}}_{l^{-1}} \left(\left\lfloor \frac{1}{r_c} (\underline{x} - \underline{D}_{\min}) \right\rfloor - \bar{\mathfrak{I}}_l(\beta(w, l)) + \bar{\mathfrak{z}} \right) \quad (241)$$

Lemma 1

$$\bar{\mathfrak{4}} - \bar{\mathfrak{z}}_l(l) = \left\lfloor \frac{1}{r_c} (\underline{x} - \underline{D}_{\min}) \right\rfloor - \bar{\mathfrak{I}}_l(\beta(w, l)) + \bar{\mathfrak{z}} \quad (242)$$

def. of β (Equation (61))

$$\bar{\mathfrak{4}} - \bar{\mathfrak{z}}_l(l) = \left\lfloor \frac{1}{r_c} (\underline{x} - \underline{D}_{\min}) \right\rfloor - \bar{\mathfrak{I}}_l \left(\bar{\mathfrak{I}}_{l^{-1}} \left(\bar{\mathfrak{I}}_l(w) + \bar{\mathfrak{z}}_l(l) - \bar{\mathfrak{z}} \right) \right) + \bar{\mathfrak{z}} \quad (243)$$

Lemma 1

$$\bar{\mathfrak{4}} - \bar{\mathfrak{z}}_l(l) = \left\lfloor \frac{1}{r_c} (\underline{x} - \underline{D}_{\min}) \right\rfloor - \bar{\mathfrak{I}}_l(w) - \bar{\mathfrak{z}}_l(l) + \bar{\mathfrak{z}} + \bar{\mathfrak{z}} \quad (244)$$

→

$$\bar{\mathfrak{I}}_l(w) = \left\lfloor \frac{1}{r_c} (\underline{x} - \underline{D}_{\min}) \right\rfloor + \bar{\mathfrak{I}} \quad (245)$$

Lemma 1

$$w = \bar{\mathfrak{I}}_{l^{-1}} \left(\left\lfloor \frac{1}{r_c} (\underline{x} - \underline{D}_{\min}) \right\rfloor + \bar{\mathfrak{I}} \right). \quad (246)$$

The functions $dist_{(\tilde{g}, \mathcal{G}^1)}^{ALL}$ and $collect_{\tilde{g}}^{ALL}$ put our particle p in the central storage compartment of the w th process:

$$p \in \left\langle \left\langle collect_{\tilde{g}}^{ALL} \left(dist_{(\tilde{g}, \mathcal{G}^1)}^{ALL} \left(step_{\mathcal{G}^1}^{ALL} \left(copy_{\tilde{g}}^{ALL} (\mathcal{P}^t) \right) \right) \right) \right\rangle_w \right\rangle_{\frac{3^{d+1}}{2}}. \quad (247)$$

Hence,

$$p \in \left\langle \left\langle \mathcal{P}^{t+1} \right\rangle_w \right\rangle_{\frac{3^{d+1}}{2}} \quad (248)$$

fulfills the condition in Equations (91) and (92). \square

B Derivation of the Bounds on Time Complexity and Parallel Scalability

Following [24], an upper bound on the time complexity of the interact function $\tau_{i(g, p', p'')}$ is

$$\forall g \in \{g^1, \dots, g^T\}, p', p'' \in \bigcirc_{w=1}^T \mathbf{p}^w \exists \tilde{g} \in \{g^1, \dots, g^T\}, \tilde{p}', \tilde{p}'' \in \bigcirc_{w=1}^T \mathbf{p}^w : \tau_{i(g, p', p'')} \leq \tau_{i(\tilde{g}, \tilde{p}', \tilde{p}'')} =: \tau_i. \quad (249)$$

Similarly, an upper bound on the time complexity of the evolve function $\tau_{e(g, p)}$ is

$$\forall g \in \{g^1, \dots, g^T\}, p \in \bigcirc_{w=1}^T \mathbf{p}^w \exists \tilde{g} \in \{g^1, \dots, g^T\}, \tilde{p} \in \bigcirc_{w=1}^T \mathbf{p}^w : \tau_{e(g, p)} \leq \tau_{e(\tilde{g}, \tilde{p})} =: \tau_e. \quad (250)$$

An upper bound on the time complexity of the evolve function of the global variable $\tau_{e(g)}$ is

$$\forall g \in \{g^1, \dots, g^T\} \exists \tilde{g} \in \{g^1, \dots, g^T\} : \tau_{e(g)} \leq \tau_{e(\tilde{g})} =: \tau_e. \quad (251)$$

An upper bound on the time complexity of the stopping condition function $\tau_{f(g)}$ is

$$\forall g \in \{g^1, \dots, g^T\} \exists \tilde{g} \in \{g^1, \dots, g^T\} : \tau_{f(g)} \leq \tau_{f(\tilde{g})} =: \tau_f. \quad (252)$$

An upper bound on the time complexity of the neighborhood function $\tau_{u([g, p], j)}$ is

$$\forall g \in \{g^1, \dots, g^T\}, p \in \{p^1, \dots, p^T\}, j \in \{1, \dots, |p|\} \exists \tilde{g} \in \{g^1, \dots, g^T\}, \tilde{p} \in \{p^1, \dots, p^T\}, \tilde{j} \in \{1, \dots, |\tilde{p}|\} : \tau_{u([g, p], j)} \leq \tau_{u([\tilde{g}, \tilde{p}], \tilde{j})} =: \tau_u. \quad (253)$$

An upper bound on the size of the neighborhood $|u([g, \mathbf{p}], j)|$ is

$$\forall g \in \{g^1, \dots, g^T\}, \mathbf{p} \in \{\mathbf{p}^1, \dots, \mathbf{p}^T\}, j \in \{1, \dots, |\mathbf{p}|\} \exists \tilde{g} \in \{g^1, \dots, g^T\}, \tilde{\mathbf{p}} \in \{\mathbf{p}^1, \dots, \mathbf{p}^T\}, \tilde{j} \in \{1, \dots, |\tilde{\mathbf{p}}|\} : |u([g, \mathbf{p}], j)| \leq |u([\tilde{g}, \tilde{\mathbf{p}}], \tilde{j})| =: \zeta_u. \quad (254)$$

The time complexity of calculating \bar{I}_l and \bar{I}_l^{-1} is $O(d)$. Therefore, the time complexity of the index transformation functions is bound by Cd , where C is a constant. Hence, the time complexity of $\beta, \tau_{\beta(t,l)}$ is bound by

$$\forall \bar{I} \in \mathbb{N}_{>0}^d, \forall t \in \{1, \dots, N_{\text{cell}}\}, l \in \{1, \dots, 3^d\} \exists C_\beta \in \mathbb{R} : \tau_{\beta(t,l)} \leq C_\beta d. \quad (255)$$

The time complexity of $\gamma, \tau_{\gamma(t,l)}$ is bound by

$$\forall \bar{I} \in \mathbb{N}_{>0}^d, \forall k \in \{1, \dots, 3^d\}, j \in \{1, \dots, {}^k N_{\text{cell}}^*\} \exists C_\gamma \in \mathbb{R} : \tau_{\gamma(t,l)} \leq C_\gamma d. \quad (256)$$

The time complexity of computing α is bound by

$$\forall \underline{D}_{\min}, \underline{D}_{\max} \in \mathbb{R}^d, \underline{x} \in [\underline{D}_{\min}, \underline{D}_{\max}], \forall j \in \{1, \dots, N_{\text{cell}}\} \exists C_\alpha \in \mathbb{R} : \tau_{\bar{I}^{-1}} \left(\left[\frac{1}{r_c} (\underline{x} - \underline{D}_{\min}) \right]_{-\bar{I}_l(j) + \bar{3}} \right) \leq C_\alpha d. \quad (257)$$

The time complexity for computing the index transformation in the *collect* function is bound by

$$\forall l \in \{1, \dots, 3^d\} \exists C_c \in \mathbb{R} : \tau_{\bar{I}^{-1}} \left(\frac{1}{4 - \bar{3}_l(l)} \right) \leq C_c d. \quad (258)$$

We use these upper bounds to derive bounds on the time complexity of the present parallelization scheme on a sequential computer and on a distributed-memory parallel machine. In general, an upper bound on the time complexity of the state transition depends on the instance $[g^1, \mathbf{p}^1]$. For each instance, we can bound the number of particles by

$$\forall t \in \{1, \dots, T\} \exists N_{\mathbf{p}}^{\max} \in \mathbb{N} : |\mathbf{p}^t| \leq N_{\mathbf{p}}^{\max}. \quad (259)$$

The time complexity of the sequential state transition τ_S is then bound by

$$\tau_S([g^1, \mathbf{p}^1]) \leq T \left(N_{\mathbf{p}}^{\max} \left(\zeta_u \left(N_{\mathbf{p}}^{\max} \right) \tau_i + \tau_u \left(N_{\mathbf{p}}^{\max} \right) + \tau_e \right) + \tau_f + \tau_\varepsilon \right), \quad (260)$$

where the neighborhood-related terms $\zeta_u(N_{\mathbf{p}}^{\max})$ and $\tau_u(N_{\mathbf{p}}^{\max})$ potentially depend on $N_{\mathbf{p}}^{\max}$. In the presented cell-list-based parallelization scheme, we can further bound the neighborhood function by exploiting that the neighborhood calculation is done separately on each process. Hence, only the particles in that process are taken into account. The number of particles in one cell is bound by

$$\forall t \in \{1, \dots, T\}, w \in \{1, \dots, N_{\text{cell}}\} \exists n_{\max} \in \mathbb{N} : \left| \langle \langle \mathcal{P}^t \rangle_w \rangle_{\frac{3^d+1}{2}} \right| \leq n_{\max}. \quad (261)$$

Then, the number of all particles in one process is bound by

$$\forall t \in \{1, \dots, T\}, w \in \{1, \dots, N_{\text{cell}}\} : \left| \bigcirc_{l=1}^{3^d} \langle \langle \mathcal{P}^t \rangle_w \rangle_l \right| \leq 3^d n_{\max}. \quad (262)$$

The number of particles is bound, and the neighborhood function checks if a distance is smaller than r_c and verifies a function $\Omega(g, p_k, p_j)$. We assume

$$\forall g \in \{g^1, \dots, g^T\}, \mathbf{p} \in \{\mathbf{p}^1, \dots, \mathbf{p}^T\}, j, k \in \{1, \dots, |\mathbf{p}|\} \exists C_u \in \mathbb{R} : \tau_{|\underline{x}_k - \underline{x}_j| \leq r_c} + \tau_{\Omega(g, p_k, p_j)} \leq d C_u. \quad (263)$$

Then, the time complexity and the size of the neighborhood function are bound by

$$\tau_u \leq 3^d n_{\max} d C_u, \quad \zeta_u \leq 3^d n_{\max}. \quad (264)$$

The time complexity of sequentially executing the distributed-memory parallel particle method $\tau_{\hat{S}([g^1, \mathbf{p}^1])}(1)$ is determined by the time complexity of the functions $collect_g^{ALL}$, $dist_{(g, \mathcal{G})}^{ALL}$, $step_g^{ALL}$, and $copy_g^{ALL}$; the evolve function of the global variable for each cell; and the stop function. Then,

$$\begin{aligned} \tau_{\hat{S}([g^1, \mathbf{p}^1])}(1) \leq T \left(\right. & \tau_f + \underbrace{N_{\text{cell}} \tau_e}_{(global\ variable\ evolve)} + \underbrace{\prod_{k=1}^{3^d} k N_{\text{cell}}^* \left(C_\gamma d + 3^d (C_\beta d + C_c d + n_{\text{max}}) \right)}_{(collect)} \\ & + \underbrace{N_{\text{cell}} n_{\text{max}} (C_\alpha d + 1)}_{(dist)} + \underbrace{N_{\text{cell}} n_{\text{max}} \left(\tau_e + 3^d n_{\text{max}} C_u d + 3^d n_{\text{max}} \tau_i \right)}_{(step)} \\ & \left. + \underbrace{\prod_{k=1}^{3^d} k N_{\text{cell}}^* \left(C_\gamma d + 3^d (C_\beta d + n_{\text{max}}) \right)}_{(copy)} \right). \end{aligned} \quad (265)$$

We assume that particles are evenly distributed (i.e., the number of particles in each cell is approximately the same) and that the density of particles remains constant when increasing the number of initial particles \mathbf{p}^1 . Therefore, the domain size has to increase proportionally. Hence, the number of cells N_{cell} increases, while n_{max} stays approximately the same. Under the assumption that all functions are then bound by constants, since they do not depend on N_{cell} , we can simplify

$$\tau_{\hat{S}([g^1, \mathbf{p}^1])}(1) \leq T \left(C_f + N_{\text{cell}} \left(C_e^\circ + C_{dist} + C_{step} \right) + \prod_{k=1}^{3^d} k N_{\text{cell}}^* \left(C_{collect} + C_{copy} \right) \right). \quad (266)$$

For a single processor we can further simplify by using $\prod_{k=1}^{3^d} k N_{\text{cell}}^* = N_{\text{cell}}$ to get

$$\tau_{\hat{S}([g^1, \mathbf{p}^1])}(1) \leq T \left(C_f + N_{\text{cell}} (C_e^\circ + C_{collect} + C_{dist} + C_{step} + C_{copy}) \right). \quad (267)$$

When parallelizing it, we need to consider that a process is the smallest computational unit. The processes are distributed on a number of processors n_{CPU} . We define a ‘‘processor’’ (CPU) as running concurrently and operating on its own separate memory address space. For convenience we assume $\frac{N_{\text{cell}}}{3^d} \in \mathbb{N}$. Further, to avoid communication conflicts, the processes need to be distributed to the processors according to which checkerboard pattern they belong. All processes of the k th checkerboard pattern are the $\gamma(k, j)$ -th processes with k fixed and $j \in \{1, \dots, k N_{\text{cell}}^*\}$. The processes on one processor must be either from one checkerboard pattern or the entire checkerboard pattern. We formulate the time complexity of the parallelized distributed-memory parallel particle method as

$$\begin{aligned} \tau_{\hat{S}([g^1, \mathbf{p}^1])}(n_{CPU}) \leq T \left(C_f + \Xi_{calc}(N_{\text{cell}}, n_{CPU}, d) (C_e^\circ + C_{dist} + C_{step}) \right. \\ \left. + \Xi_{com}(N_{\text{cell}}, n_{CPU}, d) (C_{collect} + C_{copy}) \right), \end{aligned} \quad (268)$$

where

$$\Xi_{calc}(N_{\text{cell}}, n_{CPU}, d) := \begin{cases} \left\lceil \frac{3^d}{n_{CPU}} \right\rceil \frac{N_{\text{cell}}}{3^d} & \text{if } n_{CPU} \in \{1, \dots, 3^d\} \\ M_2 & \text{if } n_{CPU} \in \{3^d, \dots, N_{\text{cell}}\} \\ 1 & \text{if } n_{CPU} > N_{\text{cell}}, \end{cases} \quad (269)$$

and

$$\Xi_{com}(N_{cell}, n_{CPU}, d) := \begin{cases} N_{cell} & \text{if } n_{CPU} \in \{1, \dots, 3^d\} \\ n_1 M_1 + n_2 M_2 & \text{if } n_{CPU} \in \{3^d, \dots, N_{cell}\} \\ 3^d & \text{if } n_{CPU} > N_{cell} \end{cases}, \quad (270)$$

where the number of checkerboard patterns that have more processors assigned to them is

$$n_1 := n_{CPU} \bmod 3^d, \quad (271)$$

the number of checkerboard patterns that have fewer processors assigned to them is

$$n_2 := 3^d - n_1, \quad (272)$$

the maximum number of processes per processor is

$$M_1 := \left\lceil \frac{N_{cell}}{3^d \lfloor \frac{n_{CPU}}{3^d} \rfloor} \right\rceil, \quad (273)$$

and for checkerboard patterns with fewer processors assigned to them, the maximum number of processes per processor is

$$M_2 := \left\lceil \frac{N_{cell}}{3^d \lfloor \frac{n_{CPU}}{3^d} \rfloor} \right\rceil. \quad (274)$$

For $n_{CPU} \in \{1, \dots, 3^d\}$, the number of processors did not reach the number of checkerboard patterns. In this case, each checkerboard pattern is completely on a processor to avoid communication conflicts. The limiting factor for the calculation (Ξ_{calc}) is then the maximum number of entire checkerboard patterns on one processor. This is $\lceil \frac{3^d}{n_{CPU}} \rceil \frac{N_{cell}}{3^d}$, where $\lceil \frac{3^d}{n_{CPU}} \rceil$ is the maximum number of checkerboard patterns on one processor and $\frac{N_{cell}}{3^d}$ the number of processes in one checkerboard pattern. The communication (Ξ_{com}) is then sequential since a processor does the communication of each process sequentially, and the communication for each checkerboard pattern needs to be done sequentially.

For $n_{CPU} \in \{3^d, \dots, N_{cell}\}$, a checkerboard pattern can be distributed on more than one processor. Therefore, the limiting factor for the calculation (Ξ_{calc}) is the maximum number of processes per processor, which is M_2 . In M_2 , the term $\frac{N_{cell}}{3^d}$ is again the number of processes in one checkerboard pattern and $\lfloor \frac{n_{CPU}}{3^d} \rfloor$ is the minimum number of processors per checkerboard pattern. The limiting factor for the communication (Ξ_{com}) is more complex since the communication is carried out for each checkerboard pattern separately, one after the other. Hence, there is a number (n_1) of checkerboard patterns that have one processor more assigned ($\lceil \frac{n_{CPU}}{3^d} \rceil$) to them and a number (n_2) of checkerboard pattern with less ($\lfloor \frac{n_{CPU}}{3^d} \rfloor$). Processes of the same color can communicate in parallel, but the algorithm needs to go sequentially through all colors. Hence, to sum up the maximum number of communications per processor for all checkerboard patterns, we calculate $n_1 M_1 + n_2 M_2$.

Each processor has exactly one process or no process for n_{CPU} , reaching or exceeding N_{cell} . The calculation (Ξ_{calc}) is saturated, and the limiting factor is 1. The communication (Ξ_{com}) is also saturated, and the limiting factor is the number of checkerboard patterns 3^d .

These upper bounds on the time complexities allow us to derive closed-form expressions for the bounds on the speedups for both strong scaling according to Amdahl's law [3] and weak scaling according to Gustafson's law [12].

First, the speedup of the cell-list scheme on one processor over the sequential state transition is

$$speedup_{\text{cell}}(N_{\text{p}}^{\max}) := \frac{\tau_{S([g^1, p^1])}}{\tau_{\tilde{S}([g^1, p^1])}(1)} \quad (275)$$

$$\approx \frac{N_{\text{p}}^{\max 2} C_{ud} + N_{\text{p}}^{\max} (3^d n_{\max} \tau_i + \tau_e) + \tau_f + \tau_e}{N_{\text{p}}^{\max} \left(3^d n_{\max} C_{ud} + 3^d n_{\max} \tau_i + \tau_e + \frac{\tau_e + C_{\text{collect}} + C_{\text{dist}} + C_{\text{copy}}}{n_{\max}} \right) + \tau_f} \quad (276)$$

$$\in \frac{O(N_{\text{p}}^{\max 2})}{O(N_{\text{p}}^{\max})} = O(N_{\text{p}}^{\max}). \quad (277)$$

Second, Amdahl's law [3] provides an upper bound on the speedup of the cell-list scheme on multiple processors when the problem size is fixed for increasing processors n_{CPU} (strong scaling):

$$speedup_{\text{Amdahl}}(n_{\text{CPU}}) = \frac{\tau_{S([g^1, p^1])}(1)}{\tau_{\tilde{S}([g^1, p^1])}(n_{\text{CPU}})} \quad (278)$$

$$\approx \frac{C_f + N_{\text{cell}} (C_e + C_{\text{collect}} + C_{\text{dist}} + C_{\text{step}} + C_{\text{copy}})}{C_f + \Xi_{\text{calc}}(N_{\text{cell}}, n_{\text{CPU}}, d)(C_e + C_{\text{dist}} + C_{\text{step}}) + \Xi_{\text{com}}(N_{\text{cell}}, n_{\text{CPU}}, d)(C_{\text{collect}} + C_{\text{copy}})}. \quad (279)$$

Third, Gustafson's law [12] provides an upper bound on the speedup of the cell-list scheme on multiple processors when the ratio of problem size to process number is constant while increasing the number of processors (weak scaling); $\frac{N_{\text{cell}}}{n_{\text{CPU}}} = \text{const}$. We achieve this by setting $N_{\text{cell}} = n_{\text{CPU}} \cdot N'_{\text{cell}}$, where N'_{cell} is constant. For a perfectly fitting processor interconnect network topology, we predict a linear speedup on average with steps as visualized in Figure 7(c):

$$speedup_{\text{Gustafson}}(n_{\text{CPU}}) = \frac{\tau_{\tilde{S}([g^1, p^1(n_{\text{CPU}})])}(1)}{\tau_{\tilde{S}([g^1, p^1(n_{\text{CPU}})])}(n_{\text{CPU}})} \approx \frac{C_f + n_{\text{CPU}} \cdot N_{\text{cell}} (C_e + C_{\text{collect}} + C_{\text{dist}} + C_{\text{step}} + C_{\text{copy}})}{C_f + \Xi_{\text{calc}}(n_{\text{CPU}} \cdot N_{\text{cell}}, n_{\text{CPU}}, d)(C_e + C_{\text{dist}} + C_{\text{step}}) + \Xi_{\text{com}}(n_{\text{CPU}} \cdot N_{\text{cell}}, n_{\text{CPU}}, d)(C_{\text{collect}} + C_{\text{copy}})}. \quad (280)$$

Acknowledgments

We thank Pietro Incardona (University of Bonn) for discussions on practically used distributed-memory parallelization schemes.

References

- [1] Yaser Afshar and Ivo F. Sbalzarini. 2016. A parallel distributed-memory particle method enables acquisition-rate segmentation of large fluorescence microscopy images. *PLoS One* 11, 4 (2016), e0152528. <https://doi.org/10.1371/journal.pone.0152528>
- [2] B. J. Alder and T. E. Wainwright. 1957. Molecular dynamics simulation of hard sphere system. *J. Chem. Phys.* 27 (1957), 1208–1218.
- [3] Gene M. Amdahl. 1967. Validity of the single processor approach to achieving large scale computing capabilities. In *Proceedings of the April 18–20, 1967, Spring Joint Computer Conference*. 483–485.
- [4] Michael Bergdorf, Ivo F. Sbalzarini, and Petros Koumoutsakos. 2010. A Lagrangian particle method for reaction-diffusion systems on deforming surfaces. *J. Math. Biol.* 61 (2010), 649–663. <https://doi.org/10.1007/s00285-009-0315-2>
- [5] George C. Bourantas, Bevan L. Cheeseman, Rajesh Ramaswamy, and Ivo F. Sbalzarini. 2016. Using DC PSE operator discretization in Eulerian meshless collocation methods improves their robustness in complex geometries. *Comput. Fluids* 136 (2016), 285–300.
- [6] Janick Cardinale, Grégory Paul, and Ivo F. Sbalzarini. 2012. Discrete region competition for unknown numbers of connected regions. *IEEE Trans. Image Process.* 21, 8 (2012), 3531–3545.
- [7] G. H. Cottet and S. Mas-Gallic. 1990. A particle method to solve the Navier-Stokes system. *Numer. Math.* 57 (1990), 805–827.

- [8] P. Degond and S. Mas-Gallic. 1989. The weighted particle method for convection-diffusion equations. Part 1: The case of an isotropic viscosity. *Math. Comput.* 53, 188 (1989), 485–507.
- [9] Jeff D. Eldredge, Anthony Leonard, and Tim Colonius. 2002. A general deterministic treatment of derivatives in particle methods. *J. Comput. Phys.* 180 (2002), 686–709.
- [10] R. A. Gingold and J. J. Monaghan. 1977. Smoothed particle hydrodynamics—Theory and application to non-spherical stars. *Roy. Astron. Soc. Mon. Not.* 181 (1977), 375–378.
- [11] R. Greenlaw, H. J. Hoover, and W. L. Ruzzo. 1995. *Limits to Parallel Computation: P-completeness Theory*. Oxford University Press. <https://books.google.de/books?id=YZHnCwAAQBAJ>
- [12] John Gustafson. 1988. Reevaluating Amdahl’s Law. *Commun. ACM* 31 (May 1988), 532–533. <https://doi.org/10.1145/42411.42415>
- [13] Alexis Hérault, Annamaria Vicari, and Ciro Del Negro. 2008. A SPH thermal model for the cooling of a lava lake. In *Proceedings of the 3th SPHERIC Workshop*.
- [14] R. W. Hockney and C. R. Jesshope. 2019. *Parallel Computers 2: Architecture, Programming and Algorithms*. CRC Press. <https://books.google.de/books?id=c2ytDwAAQBAJ>
- [15] R. W. Hockney and J. W. Eastwood. 1988. *Computer Simulation Using Particles*. Institute of Physics Publishing.
- [16] Pietro Incardona, Antonio Leo, Yaroslav Zaluzhnyi, Rajesh Ramaswamy, and Ivo F. Sbalzarini. 2019. OpenFPM: A scalable open framework for particle and particle-mesh codes on parallel computers. *Comput. Phys. Commun.* 241 (2019), 155–177.
- [17] Masaki Iwasawa, Ataru Tanikawa, Natsuki Hosono, Keigo Nitadori, Takayuki Muranushi, and Junichiro Makino. 2016. Implementation and performance of FDPS: A framework for developing parallel particle simulation codes. *PASJ* 68, 4 (2016), 54.
- [18] Sven Karol, Tobias Nett, Jeronimo Castrillon, and Ivo F. Sbalzarini. 2018. A domain-specific language and editor for parallel particle methods. *ACM Trans. Math. Softw.* 44, 3 (2018), 34.
- [19] Jonathan R. Karr, Jayodita C. Sanghvi, Derek N. Macklin, Miriam V. Gutschow, Jared M. Jacobs, Benjamin Bolival, Nancyra Assad-Garcia, John I. Glass, and Markus W. Covert. 2012. A whole-cell computational model predicts phenotype from genotype. *Cell* 150 (2012), 389–401.
- [20] Wing Kam Liu, Sukky Jun, and Yi Fei Zhang. 1995. Reproducing kernel particle methods. *Int. J. Numer. Meth. Fluids* 20 (1995), 1081–1106.
- [21] Susan M. Mniszewski, James Belak, Jean-Luc Fattebert, Christian F. A. Negre, Stuart R. Slattery, Adetokunbo A. Ade-doyin, Robert F. Bird, Choongseok Chang, Guangye Chen, Stéphane Ethier, Shane Fogerty, Salman Habib, Christoph Junghans, Damien Lebrun-Grandié, Jamaludin Mohd-Yusof, Stan G. Moore, Daniel Osei-Kuffuor, Steven J. Plimpton, Adrian Pope, Samuel Temple Reeve, Lee Ricketson, Aaron Scheinberg, Amil Y. Sharma, and Michael E. Wall. 2021. Enabling particle applications for exascale computing platforms. *Int. J. High Perform. Comput. Appl.* 35, 6 (2021), 572–597. <https://doi.org/10.1177/109434202111022829>
- [22] J. J. Monaghan. 2005. Smoothed particle hydrodynamics. *Rep. Prog. Phys.* 68 (2005), 1703–1759.
- [23] Mark T. Nelson, William Humphrey, Attila Gursoy, Andrew Dalke, Laxmikant V. Kalé, Robert D. Skeel, and Klaus Schulten. 1996. NAMD: A parallel, object-oriented molecular dynamics program. *Int. J. Supercomput. Appl. High Perform. Comput.* 10, 4 (1996), 251–268. <https://doi.org/10.1177/109434209601000401>
- [24] Johannes Pahlke and Ivo F. Sbalzarini. 2023. A unifying mathematical definition of particle methods. *IEEE Open J. Comput. Soc.* 4 (2023), 97–108. <https://doi.org/10.1109/OJCS.2023.3254466>
- [25] Szilárd Páll and Berk Hess. 2013. A flexible algorithm for calculating pair interactions on SIMD architectures. *Comput. Phys. Commun.* 184, 12 (2013), 2641–2650. <https://doi.org/10.1016/j.cpc.2013.06.003>
- [26] J. V. W. Reynders, J. C. Cummings, M. Tholburn, P. J. Hinker, S. R. Atlas, S. Banerjee, M. Srikant, W. F. Humphrey, S. R. Karmesin, and K. Keahey. 1996. POOMA: A framework for scientific simulation on parallel architectures. In *Proceedings of the 1st International Workshop on High-level Programming Models and Supportive Environments*, A. Bode, M. Gerndt, R. G. Hackenberg, and H. Hellwagner (Eds.). Technical University of Munchen; Res. Centre Julich; Central Inst. Appl. Math.; 10th IEEE Int. Parallel Process. Symposium; IEEE Comput. Soc. Tech. Committee on Parallel Process.; ACM SIGARCH, IEEE Comput. Soc. Press, Los Alamitos, CA, USA, 41–49.
- [27] Ivo F. Sbalzarini. 2010. Abstractions and middleware for petascale computing and beyond. *International Journal of Distributed Systems and Technologies* 1 (April 2010), 40–56. <https://doi.org/10.4018/jdst.2010040103>
- [28] Ivo F. Sbalzarini, Anna Mezzacasa, Ari Helenius, and Petros Koumoutsakos. 2005. Effects of organelle shape on fluorescence recovery after photobleaching. *Biophys. J.* 89, 3 (2005), 1482–1492.
- [29] I. F. Sbalzarini, J. H. Walther, M. Bergdorf, S. E. Hieber, E. M. Kotsalis, and P. Koumoutsakos. 2006. PPM—A highly efficient parallel particle-mesh library for the simulation of continuum systems. *J. Comput. Phys.* 215, 2 (2006), 566–588.
- [30] Birte Schrader, Sylvain Reboux, and Ivo F. Sbalzarini. 2010. Discretization correction of general integral PSE operators in particle methods. *J. Comput. Phys.* 229 (2010), 4159–4182.

- [31] Stuart Slattery, Samuel Temple Reeve, Christoph Junghans, Damien Lebrun-Grandié, Robert Bird, Guangye Chen, Shane Fogerty, Yuxing Qiu, Stephan Schulz, Aaron Scheinberg, Austin Isner, Kwitae Chong, Stan Moore, Timothy Germann, James Belak, and Susan Mniszewski. 2022. Cabana: A performance portable library for particle-based simulations. *J. Open Source Softw.* 7, 72 (2022), 4115. <https://doi.org/10.21105/joss.04115>
- [32] Siddhartha Verma, Guido Novati, and Petros Koumoutsakos. 2018. Efficient collective swimming by harnessing vortices through deep reinforcement learning. *PNAS* 115, 23 (2018), 5849–5854. <https://doi.org/10.1073/pnas.1800923115>
- [33] Jens H. Walther and Ivo F. Sbalzarini. 2009. Large-scale parallel discrete element simulations of granular flow. *Eng. Computat.* 26, 6 (2009), 688–697.

Received 24 January 2023; revised 20 August 2024; accepted 11 September 2024