

Systematic Literature Review of Data Exchange Strategies for Range-limited Particle Interactions

Theresa Werner¹, Ivo Kabadshow² and Matthias Werner¹

¹*Department of Operating Systems, Chemnitz University of Technology, Chemnitz, Germany*

²*Jülich Supercomputing Centre, Jülich, Germany*

{thewe, werne}@hrz.tu-chemnitz.de, i.kabadshow@fz-juelich.de

Keywords: Communication, Distributed Memory, Molecular Dynamics, Spatial Decomposition, HPC.

Abstract: Molecular dynamics simulations (MDS), no matter in which form, have always spent a lot of effort on the time-consuming part of direct particle-to-particle interactions ($O(N^2)$). Even if the interaction radius of each particle is limited, it remains the most time-critical part, especially when increasing the number of compute nodes to calculate on. This Systematic Literature Review (SLR) focuses on the spatial decomposition approach to MDS and ways to optimize its data exchange. We gathered and compared available concepts related to range-limited interactions and investigated whether they show similarities and how those can be categorized. Based on the findings, we can summarize that all communication schemes are derived from the same basic idea, the so-called shift communication. The concepts differ in which data is communicated and how nodes calculate the forces between particle pairs. Two categories can be distinguished here: home-box-centric and neutral territory methods.

1 INTRODUCTION

Scientists already predicted that in the future the high performance sector would no longer be computation-bound but become communication-bound (Chandramowlishwaran and Vuduc, 2012). HPC networks would grow from several dozens of nodes to several thousands, and in order to make use of that hardware, not only synchronization strategies but also communication strategies must be applied to HPC problems.

One such problem in HPC occurs in molecular dynamics simulations (MDS), or simpler, N -body problems, where a large number of particles are interacting with each other. For the most accurate result in simulation every particle-to-particle (p2p) interaction needs to be calculated. This leads to a complexity of $O(N^2)$ for N particles. Even when Newton's third law is applied to calculate each force only once by considering symmetry, the complexity does not change (note here that the absolute runtime can be reduced though). MD systems typically have hundreds of thousands to billions of particles; a parallel all-to-all force calculation would either require repeated data exchange between nodes or expect each node to hold all the data in every timestep. Both options are undesirable, so scientists developed methods with reduced computational complexity $O(N \log N)$ or $O(N)$ and opened

the field for fast summation methods in MDS, introducing a threshold for less costly force computations. All these methods, in one way or another, make a distinction between near-field and far-field interactions.

A well-known parallelization approach to MDS is spatial decomposition. It splits the simulation space into small rectangular boxes and, when a fast summation method is applied, only exchanges particle data between boxes that are no further apart than a certain cut-off distance.

With the problem of HPC applications becoming communication-bound in mind, this SLR focuses on communication schemes and data exchange strategies between boxes/compute nodes that lie beyond broadcasting and all-to-all communication. It aims to summarize all concepts found and to show connections and similarities and it proposes a categorization for range-limited MDS methods. We found a variety of methods for managing the data exchange which might at first look unrelated, but after closer investigation they reveal a common core concept. All schemes use some kind of shift communication, and there are two approaches to how a node computes certain forces.

Section 2 outlines the range-limited particle interaction problem and its different variants. Section 3 explains the process of writing the SLR and sums up the relevant works. Section 4 analyses and categorizes




Figure 1: Size of the import area (gray) of the black target box with different R_{cut} : (a) $R_{\text{cut}} < b$, (b) $R_{\text{cut}} = b$, and (c) $R_{\text{cut}} > b$.

the found information. And Section 5 provides a summary and gives an outlook to possible future work.

2 CUT-OFF DISTANCE

Considering a spatially decomposed simulation space, the data that needs to be exchanged in order to perform all range-limited (RL) calculations depends on the definition of the RL range, the cut-off radius R_{cut} . Overall there are three ways to characterize RL interactions (for simplicity we assume cubic boxes of edge length b): $R_{\text{cut}} < b$, $R_{\text{cut}} = b$, and $R_{\text{cut}} > b$. Figure 1 depicts those three versions of the RL area. The target box is colored black, the surrounding white boxes which overlap with the gray import area defined by R_{cut} are the boxes of which data is needed for the RL interactions of the black box; the remaining white boxes are already part of the long-range interactions and will not be taken into account for computation or communication, here.

Assuming that each box is handled by its own compute node, RL p2p interactions require a varying amount of data exchange between nodes depending on R_{cut} . Designing a communication-efficient algorithm to make this transfer less costly can be a major time saver since this exchange has to be done for each box in every simulation timestep.

3 SLR METHODOLOGY

This section serves to familiarize the reader with the process of finding relevant material for analysis as defined by (Biolchini et al., 2005). The process consists of three main steps: determining whether an SLR is necessary, the planning stage that asks for the What and How of the filtering of scientific publications, and the execution stage that filters through the results by different conditions and sums up the works that make it through the filtering.

3.1 Need for an SLR

Since communication –apart from avoiding and hiding it– has never been a huge topic in HPC MD problems, an SLR about optimal communication in RL interactions to our knowledge has not been published, yet. The search for possible reviews has been done with the following search string:

("systematic literature review" OR "systematic review" OR "research review" OR SLR) AND ("communication efficient" OR "communication optimal" OR "communication bound") AND ("molecular dynamics" OR "n-body" OR "fast multipole method" OR FMM)

This search string procured five results on *Google Scholar*. All of them are not systematic reviews concerning this topic of data exchange in RL interactions.¹

3.2 Planning

The first step in planning an SLR is to define research questions which help to find the information that is useful to reach the goal of the SLR. For us the goal was to find possible similarities between the data exchange methods for spatial decomposition MDS and, in hindsight to future research, what can be said about the performance of the methods regarding communication. Hence, the research questions formulated without any prior knowledge of the field were:

- R1** Which of the three cut-off radius versions mentioned in Section 2 does the presented scheme belong to?
- R2** What are assumptions of the scheme about the particle system and the simulation space?
- R3** If communication is explicitly described, what is the message complexity/how many messages are received/sent by one box/node for one complete force calculation step?

The second step is to choose the digital library (e.g. *IEEE Explore*). In order to evaluate which platform is most suitable to deliver a diverse result for the review topic, the following two search strings have been used:

S1 ("communication efficient" OR "communication optimal") AND ("molecular dynamics" OR "n-body" OR "fast multipole method" OR FMM) AND NOT "machine learning"

¹The keywords *astrodynamics*, *fluid dynamics*, *weather simulation*, and *aerodynamics* are excluded due to preceding research regarding whether their commonly used algorithms are suitable to solve the problem of RL interactions.

Table 1: Number of Search Results Sorted by Platform.

Digital Library	Results S1	Results S2
Google Scholar	294	≈ 2780
Springer Link	16	491
Science Direct	17	376
ACM Digital Library	7	79
IEEE Explore	5	22

S2 "molecular dynamics" AND ("spatial decomposition" OR "range limited" OR "distance limited")

String S2 hereby serves to cover the halo of works around search string S1 where information might be hidden because of not being mentioned explicitly in the title or abstract. Table 1 shows the number of results per platform for both search strings. *Google Scholar* clearly has the biggest variety and hence shall be used for finding material for this SLR.²

Understandably, not all results are vital to the SLR. Thus, in order to give the SLR a clear scope, include (I) and exclude (E) criteria must be defined: The source **I1** explicitly targets data exchange for the RL interactions, **I2** explicitly targets distributed memory simulations, **E1** does not work with spatial decomposition, **E2** optimizes data exchange by explicitly handling shared memory, **E3** focuses on load balancing, **E4** mentions the use of GPUs, **E5** only uses MD as means to an end, **E6** uses the search string keywords in a different context or with a different meaning (outside the domain).

3.3 Execution

3.3.1 Filtering Sources and Information

The results of the search are first filtered by title and second by abstract (and introduction, and conclusion). The sources matching these two pre-selections are read in full and evaluated once more regarding usefulness while redundancies are removed and promising references are looked up. Table 2 shows how the number of sources is reduced in each step.

Now it needs to be decided which information apart from the explanation of the scheme should be included in the review. The information to include or exclude derived from the research questions is: **in1** the length of the cut-off radius R_{cut} with respect to the box size b , **in2** the number of messages required

²The keyword *cell-based*, although commonly used in molecular dynamics, is excluded because it only leads to results regarding human cells (cancer treatment) and does not help procuring content related to p2p interactions. The keywords *nearest neighbor* and *distributed memory* were checked as well and found to yield no relevant results.

Table 2: Number of Search Results After Manual Filtering.

Filter Step	Results S1	Results S2
Google Scholar	294	≈ 2780
Selection by title	15	56
Selection by abstract	6	13
Selection by full text	4	3
Found in references	+0	+2
Redundancies	-1	-1
Total	3	4

for the force calculation and application of one simulation timestep. Communication complexity (per box) if available, **in3** distribution of data (how many boxes per node). **ex1** Specific implementation details, **ex2** physics.

3.3.2 Summary of Sources

Source 1: Molecular Dynamics Simulations on Distributed Memory Machines

Liem, Brown, and Clarke (Liem et al., 1991) made it their goal to avoid redundant force calculation and to minimize inter-processor communication. They achieved it by making use of "proxy" communication and a smart force calculation method. First, they decomposed their rectangular 2D simulation space into small squares and distributed those squares in bigger rectangular tiles over a k -ary 2-cube³. Since they choose their cut-off radius to be equal to the box size b , each node only needs to communicate particle data of boxes lining the "northern" (N) and "eastern" (E) boundaries of their dedicated tile. First, all nodes send their N+E particle data to their neighbor node in the north while simultaneously receiving the particle data from their south neighbor (Fig. 2a). Next, they combine their own particle data with the data from the south neighbor and send both data sets to the east neighbor; simultaneously they receive the west and south-west data from their west neighbor (Fig. 2b).

They also proposed a rule that defines which processor has to calculate which force interactions other than between particles of its own tile. Figure 2c illustrates which processor needs to calculate forces between which particle sets. The white lines between two differently colored patches mean calculating the forces between those particles. The figure is incomplete for all but the yellow processor because the grid is incomplete.

Last, the forces are sent back by first sending all results to the west neighbor and simultaneously re-

³ k -ary n -cubes is a categorization derived from (Tang, 1992), see Source 2.

Table 3: Characteristics of (Liem et al., 1991).

cut-off	$R_{\text{cut}} = b$
data distribution	uniform distribution of tiles of boxes over nodes (multiple boxes per node)




Figure 2: (a) and (b) depict the two steps of particle data communication. The force data is distributed in the reverse pattern. (c) A white line between data patches denotes that the forces between these sets need to be calculated on that node; some pairs are missing since the node grid is incomplete.

ceiving from the east neighbor; the force data of the east neighbor is combined with the own force data and the whole package is sent south, while receiving the forces from N and NE from the north neighbor. Each node can now update its local particles.

For the three-dimensional case, this method is also known by the name *Eighth-Shell (ES) Shift*.

Source 2: Pipelined Global Data Communication on Hypertori

(Tang, 1992) present several effective ways of global communication on hypertori. Since the architecture of 3-cube hypertori resembles the spatial decomposition of MDS, this paper has been included in this review. The methods described can easily be transformed into data exchange methods for RL interactions. The categorization of hypertori by describing them as k -ary n -cubes is also derived from this paper and means hypercubes which have n dimensions with k nodes in each dimension.

Tang first quotes (Saad and Schultz, 1989) on their so-called alternate direction exchange method, which proposes to split a binary n -cube into two binary $(n-1)$ -subcubes, consecutively in each direction. Hence, a binary 3-cube is first split into two binary 2-cubes and then into 4 binary 1-cubes. Tang generalizes this approach to splitting a k -ary n -cube into k k -ary $(n-1)$ -cubes; Figure 3 shows the splitting for a 3-ary 3-cube following that rule.

The communication is now performed in a daisy-chain manner that can be generalized as: first, the data is daisy-chained along dimension one; for each following dimension up to n , the accumulated data of each chain is daisy-chained along the paths of that di-

Table 4: Characteristics of (Tang, 1992).

cut-off	$R_{\text{cut}} \leq b, R_{\text{cut}} > b$
data distr.	—




Figure 3: Splitting the hypertorus: the 3-ary 3-cube in (a) is first split into 3 3-ary 2-cubes like in (b) and then into (3×3) 3-ary 1-cubes like in (c).

mension.

A simulation with a cut-off distance does not require global communication. Daisy-chaining the data all around the cube is not necessary. By introducing a sense of direction, communication can be reduced. Data is transmitted along one dimension j in two phases. Phase one forwards the data in the first direction for as many steps as the cut-off distance is long. For the second phase, the transmission is flipped around and the data is handed down into the opposite direction. This is applied to all j dimensions, and now each node only has a subset of the data of the whole simulation space. The next SLR source does exactly that.

Source 3: Fast Parallel Algorithms for Short-range Molecular Dynamics

(Plimpton, 1995) introduces what will later be called *Full-Shell (FS) Shift* communication. It can be considered an adjusted version of Tang (Source 2) or a version of Liem (Source 1) but without using Newton's third law or force decomposition to further reduce the number of messages.

The communication of one box/node takes place in three phases and is limited to its six immediate neighbors called East (E), West (W), North (N), South (S), Up (U), and Down (D). First, the particle data of the box is sent to the West and particle data from the East is received, then vice versa (see Fig. 4a). Next, the accumulated data from the box and its E and W neighbors is sent to the North, while the SW, S, and SE data is received from the the South. And in return, the data is sent to the South and the NW, N, and NE data is received from the North (see Fig. 4b). Lastly, the data of the whole plane is sent up and down and the data of the U and D planes are received in return (see Fig. 4c). If $R_{\text{cut}} < b$ only the relevant data is sent, and if $R_{\text{cut}} > b$ the neighbors help handing down the data like in Source 2.

In case of $R_{\text{cut}} \leq b$, the communication finishes after six messages and each node can compute all the

Table 5: Characteristics of (Plimpton, 1995).

cut-off	$R_{\text{cut}} \leq b, R_{\text{cut}} > b$
data distr.	one box per node




Figure 4: Full-Shell Shift communication: The black boxes data is first distributed along the axis of red boxes (a), then along the axis of yellow boxes (b), and last among the blue boxes (c). The blue boxes are incomplete for visualization purposes. The lines are the paths along which the data is spread.

forces it needs in order to update its particles. In case of $R_{\text{cut}} > b$, the number of messages sums up to $6k$ with $k = \lceil R_{\text{cut}}/b \rceil$.

Source 4: Zonal Methods for the Parallel Execution of Range-limited N-body Simulations

The work of (Bowers et al., 2007) is based on two works, (Snir, 2004) and (Shaw, 2005). Both earlier works find that the broadcast for the force decomposition has the property that “for any two processors p_1 and p_2 , there is a processor p_3 so that both p_1 and p_2 send all their data to p_3 . Then p_3 can compute all interactions between atoms from p_1 and $[\dots] p_2$ (Snir, 2004).” Based on this property they designed spatial and force decomposition hybrids: (Snir, 2004) worked this idea into his Base-Comb model (advanced version see Fig. 5 upper left) and (Shaw, 2005) into the idea of tower and plate.

(Bowers et al., 2007) realized that the two ideas were two versions of the same basic concept. They compared the two methods’ shared properties and introduced the general class of *neutral territory methods*. Derived from Snir and Shaw they first introduced the *two-zone methods*. Apart from reviewing the Base-Comb and Tower-Plate method, they developed the Cloud, City, and Foam methods depicted in Figure 5. To make sure that these rather complex sets are able to cover the simulation space without missing pairs of interacting boxes, Bowers et al. put the so-called convolution criterion in place which says that the coverage region must include the whole influence region. Influence region here means the region that is defined by the cut-off radius around the box in question; the coverage region is the area from which the box in question imports data (as implied, it might be bigger than the area defined by R_{cut}).

Next, Bowers et al. introduced *k-zone methods*.

Table 6: Characteristics of (Bowers et al., 2007).

cut-off	$R_{\text{cut}} > b$
data distr.	one box per node




Figure 5: Two-zone methods (Bowers et al., 2007): upper left is Snir’s hybrid refined by Shaw, upper right is the cloud method, lower left is the city method, and lower right is the foam method.

This is similar to idea presented by Liem (Source 1) but with the goal of overlapping communication with computation instead of reducing communication. By defining a rule about which particle set interactions must be calculated when they defined a schedule for the force calculation that is based on the arrival of the remote data packets. The further away a box/node is, the later it appears in the schedule. They applied this method to the Base-Comb, Tower-Plate, and eighth shell ideas.

Two-zone methods and k-zone methods make up the whole of *zonal methods*.

Source 5: Scalable Algorithms for Molecular Dynamics Simulations on Commodity Clusters

In this work (Bowers et al., 2006) aimed to decrease the import region and introduce their MD code named Desmond. It uses the midpoint method to determine which force pairs should be calculated on which node, meaning each node only calculates the forces for particle pairs of which the midpoint of the distance lies within their dedicated area. Figure 6 illustrates the allocation of particle pairs to processors; the midpoints (crosses) mark on which node the force between the pair of particles is calculated; missing particle data needs to be imported.

The pairs handled by one node cannot be further apart than $R_{\text{cut}}/2$ because they do not interact when further apart than R_{cut} , so if one particle is further from the node’s area than $R_{\text{cut}}/2$, the midpoint will not be within its area. This fact reduces the import area per node compared to methods with the “classi-

Table 7: Characteristics of (Bowers et al., 2006).

cut-off	$R_{\text{cut}} < 2b$
data distr.	one box per node




Figure 6: Midpoint method: a particle pairs' midpoint (cross) determines the box in which their interaction force is calculated (Bowers et al., 2006).

cal" use of the cut-off radius. Still, the communication algorithm is the same as described by Plimpton, sending 6k messages to the six nearest neighbors.

Source 6: A Communication-optimal N-body Algorithm for Direct Interactions

Driscoll, Georganas, and Koanantakool (Driscoll et al., 2013) focused on the all-to-all particle interaction by atom decomposition, but they also use their approach on RL interactions with spatial decomposition (although not quite correctly), so they shall be discussed next. Their basic idea is to allocate a group of processors per box to reduce communication.

Their all-to-all atom decomposition approach works as follows: Assuming we have 64 processors $P = 64$ and we dispatch four processors per team $c = 4$ (c standing for copy), we would have $T = P/c = 16$ processor teams. Each processor team leader receives the data for N/T random particles and hands it down to the other processors of the team (step 1 in Figure 7a). Next, that data is shifted askew like displayed in step 2 of Figure 7a, the maximal shift distance being $c - 1 = 3$. After this, the data is shifted in equidistant steps of c along the processor rows until it wraps around the whole simulation space once (step 3 in Figure 7a but for the whole space). In our example this would be given after $T/c = 4$ shifts. Last, by having a reduction step, the data of the whole team is gathered by the team leader. The authors assume that the distribution and reduction steps each have a communication complexity of $O(\log c)$, the skewing shift of $O(1)$, and the equidistant shifts of $O(T/c) = O(P/c^2)$, so we end up with a communication complexity of $O(P)$.

For the case of RL interactions, the method

Table 8: Characteristics of (Driscoll et al., 2013).

cut-off	$R_{\text{cut}} > b$
data distr.	one box per processor group




Figure 7: Processor Teams: (a) the original idea of (Driscoll et al., 2013) for the RL interactions: a wrap-around at the cut-off distance m with fixed import region. (b) The revised idea that considers the import region being unique for each box and chooses c based on m by $c = (m + 1)/k$.

changes slightly. Now, particles are no longer randomly allocated to teams but boxes of the spatial decomposition are assigned to teams. Thus, there is awareness of the particles' spatial coordinates. While distribution, reduction, and initial skewing shift remain unchanged, the equidistant shift now wraps around the cut-off radius m as shown in Figure 7a. However, by splitting the space into four static regions, the authors do not consider that the import region is unique for each box. Their static approach leads to the box in the middle of the region being the only one satisfied while all other boxes do not receive all the data they require or receive data they do not even use. Either this approach must be refined or it is essentially unusable for the cut-off problem.

But as we think back to Tang (Source 2) the described algorithm reminds us a lot of daisy-chaining the data but in a multi-layered way. Thus, the same changes as we applied to Tang should be applicable here and make this feasible for simulations with a cut-off. We discard the wrap-around at the cut-off distance and instead perform the shift in two different directions. In order not to waste processors by having some of them idle, we choose c with respect to m . For the 1D problem displayed in Figure 7a, one can choose $c = \lceil (m + 1)/k \rceil$ with $k \in \mathbb{N}$. Why $m + 1$? Because the center box itself must be taken into account. By setting $c = \lceil (5 + 1)/2 \rceil = 3$ it takes one initial skewing shift and one equidistant shift to

Table 9: Characteristics of (Wang et al., 2020).

cut-off	$R_{\text{cut}} < b$
data distr.	one box per node




Figure 8: Ghost communication mode: in case of $R_{\text{cut}} < b$, the box is split into corner (blue), edge (yellow), and face (red) subboxes; the edge length of the corner cubes is equal to R_{cut} . A message comprises of a recombination of the subboxes, thus no superfluous data is sent to the neighbor boxes.

the left and right each by every team in order to distribute the data; Figure 7b shows this adapted version. Obviously, this needs a synchronization step before changing direction, but if Newton’s third law is applied, the right shift happens before the force calculation and the left shift afterwards. For higher dimensions than 1D this method must be refined some more, but it should be safe to assume that the result will be some kind of multi-layered Shift communication. For now, this revised version also has a communication complexity of $O(P)$.

Source 7: Communication Optimization Strategy for Molecular Dynamics Simulation on Sunway TaihuLight

(Wang et al., 2020) are concerned with optimizing the data packages in case of a cut-off radius smaller than the box width. Their goal is to send only the strictly necessary data to the respective recipient.

The idea of Wang et al. is to split each box into corner (blue), edge (yellow), and face (red) subboxes (see Fig. 8). The subboxes act like a padding inside the box with thickness R_{cut} . Each message to one of the six nearest neighbors comprises of a recombination of these subboxes composed of one face, four corner, and four edge subboxes. The communication they use, is the shift communication from Plimpton (Source 3) with $k = 1$, so the total amount of messages for one simulation timestep is six per box.

4 ANALYSIS

The goal of this SLR is to find similarities between the data exchange methods in order to categorize them. We formulated three research questions in Subsection 3.2 the first of which asks for the version of the

Symmetry	N3L Sources: 1	no N3L
Data Selection	Home Box Sources: 3, 6, 7	Neutral Territory Sources: 1, 4, 5
Communication	Shift Comm. Sources: 1, 3, 5, 6, 7	trivial comm. e.g. broadcast

Figure 9: Categorization: We decided on a three-layered categorization. The bottom layer is about the communication scheme chosen, the middle layer about the What and Where of the data, and the top layer about applying Newton’s third law or not.

cut-off radius. Finding similarities based on the cut-off radius yields no definite results. Many schemes can be used on all three versions and one even fits into none of the three categories (Source 5). Regarding assumptions about the system the result is equally inconclusive. An interesting observation, however, is that all authors who propose an explicit communication scheme use the same idea: shift communication. By making different use of received data e.g. calculating forces between particles which are not part of the processor’s home box or making use of Newton’s third law, different versions of the shift came into use e.g. ES, HS, FS Shift. Even the processor team idea of (Driscoll et al., 2013) (Source 6) is an advanced shift communication for the revised cut-off version. Why is it advanced? Because by using processor teams it scales with $O(P)$ whereas all classic shift communication versions scale with $O(P^{1/3})$. Apart from this, one can categorize the sources into home box and neutral territory methods. Sources 3, 8, and 9 are home box methods aiming to satisfy the node’s home box with the data it needs for calculating the forces on its particles. Sources 1 and 4 to 7 are neutral territory methods where each processor also calculates forces between two foreign particles and does not receive enough data to calculate all forces on its own particles by itself. And last, Newton’s third law might be applied to both approaches to reduce redundant force calculation further.

Thus, we categorize this field of spatial decomposition MDS with RL interactions as depicted in Figure 9. We differentiate between three layers. The bottom layer is about the communication scheme which can either be a version of the shift communication or trivial in the form of e.g. a broadcast. The middle layer is concerned with the What and Where of the data and calculations. Here one can choose between home box methods or neutral territory methods. And last, Newton’s third law may or may not be applied to reduce redundant force calculation. In case of choos-

ing shift communication, the middle and upper categorization layer may influence the decision regarding the shift version.

5 CONCLUSION

5.1 Summary

This SLR gathered scientific works on data exchange strategies for range-limited interactions in MDS and aimed to find similarities between these works in order to propose a categorization. Its target were spatial decomposition approaches, which split the simulation space into small rectangular boxes that are then distributed over the compute nodes.

As it turns out, all sources that introduce an explicit communication scheme use the same idea called shift communication (see Fig. 4). Apart from that, one can distinguish between two categories of data selection strategies: home-box-centered methods that aim to satisfy the home box or home node with all the data required to calculate the forces on its particles, and neutral territory methods that have nodes calculate forces between particles which do not reside in the node's home box. Additionally Newton's third law (N3L) may be applied to reduce redundant force calculations. Thus, we propose the three-layered categorization displayed in Figure 9 where the bottom layer is concerned with the selection of the communication scheme, the middle layer with what data should be moved where, and the top layer with whether N3L is applied or not.

5.2 Future Work

Future work could be the design of a 2D and 3D processor team algorithm as proposed by (Driscoll et al., 2013) and applying communication schemes to higher levels in MDS, for example for the interaction of multipoles on the same tree depth in the Fast Multipole Method (FMM).

ACKNOWLEDGEMENTS

This research was funded by DFG project FMHub, project Nr. 443189148.

REFERENCES

Biolchini, J., Gomes Mian, P., Cruz Natali, A. C., and Horta Travassos, G. (2005). Systematic review in soft-

ware engineering. In *Technical Report RT-ES 679-05, COPPE/UFRJ PESC*.

- Bowers, K. J., Chow, D. E., Xu, H., Dror, R. O., Eastwood, M. P., Gregersen, B. A., Klepeis, J. L., Kolosvary, I., Moraes, M. A., Sacerdoti, F. D., Salmon, J. K., Shan, Y., and Shaw, D. E. (2006). Scalable Algorithms for Molecular Dynamics Simulations on Commodity Clusters. In *SC '06: Proceedings of the 2006 ACM/IEEE Conference on Supercomputing*, pages 43–55.
- Bowers, K. J., Dror, R. O., and Shaw, D. E. (2007). Zonal methods for the parallel execution of range-limited N-body simulations. *Journal of Computational Physics*, 221(1):303–329.
- Chandramowlishwaran, A. and Vuduc, R. W. (2012). Communication-Optimal Parallel N-body Solvers. In *2012 IEEE 26th International Parallel and Distributed Processing Symposium Workshops & PhD Forum*, pages 2462–2465.
- Driscoll, M., Georganas, E., Koanantakool, P., Solomonik, E., and Yelick, K. (2013). A Communication-Optimal N-Body Algorithm for Direct Interactions. In *2013 IEEE 27th International Symposium on Parallel and Distributed Processing*, pages 1075–1084.
- Liem, S. Y., Brown, D., and Clarke, J. H. R. (1991). Molecular dynamics simulations on distributed memory machines. *Computer Physics Communications*, 67(2):261–267.
- Plimpton, S. (1995). Fast Parallel Algorithms for Short-Range Molecular Dynamics. *Journal of Computational Physics*, 117(1):1–19.
- Saad, Y. and Schultz, M. H. (1989). Data communication in hypercubes. *Journal of Parallel and Distributed Computing*, 6(1):115–135.
- Shaw, D. E. (2005). A fast, scalable method for the parallel evaluation of distance-limited pairwise particle interactions. *Journal of Computational Chemistry*, 26(13):1318–1328.
- Snir, M. (2004). A Note on N-Body Computations with Cutoffs. *Theory of Computing Systems*, 37(2):295–318.
- Tang, Z. (1992). Pipelined Global Data Communication on Hypertoruses. *Journal of Computer Science and Technology*, 7(3):247–256.
- Wang, B., Chen, Y., and Hou, C. (2020). Communication Optimization Strategy for Molecular Dynamics Simulation on Sunway TaihuLight. In *2020 IEEE 22nd International Conference on High Performance Computing and Communications; IEEE 18th International Conference on Smart City; IEEE 6th International Conference on Data Science and Systems (HPCC/SmartCity/DSS)*, pages 571–578.