Scalable Graph Networks for Particle Simulations

Karolis Martinkus,[1] **Aurelien Lucchi,** [1] **Nathanaël Perraudin** [2]
[1] ETH Zürich
[2] Swiss Data Science Center
martinkus@ethz.ch, aurelien.lucchi@inf.ethz.ch, nathanael.perraudin@sdsc.ethz.ch

Abstract

Learning system dynamics directly from observations is a promising direction in machine learning due to its potential to significantly enhance our ability to understand physical systems. However, the dynamics of many real-world systems are challenging to learn due to the presence of nonlinear potentials and a number of interactions that scales quadratically with the number of particles N, as in the case of the N-body problem. In this work, we introduce an approach that transforms a fully-connected interaction graph into a hierarchical one which reduces the number of edges to $O(N)$. This results in linear time and space complexity while the precomputation of the hierarchical graph requires $O(N \log(N))$ time and $O(N)$ space. Using our approach, we are able to train models on much larger particle counts, even on a single GPU. We evaluate how the phase space position accuracy and energy conservation depend on the number of simulated particles. Our approach retains high accuracy and efficiency even on large-scale gravitational N-body simulations which are impossible to run on a single machine if a fully-connected graph is used. Similar results are also observed when simulating Coulomb interactions. Furthermore, we make several important observations regarding the performance of this new hierarchical model, including: i) its accuracy tends to improve with the number of particles in the simulation and ii) its generalisation to unseen particle counts is also much better than for models that use all $O(N^2)$ interactions.

1 Introduction

The ability to simulate complex systems is invaluable to many fields of science and engineering. Constructing simulators for such systems by hand can be very labour intensive or even impossible if the underlying processes are not understood or no sufficiently precise and accurate approximations of the interactions are known. To address this, various data-driven methods for learning system dynamics have been investigated (Battaglia et al. 2016; Mrowca et al. 2018; Li et al. 2018; Greydanus, Dzamba, and Yosinski 2019; Sanchez-Gonzalez et al. 2019, 2020; Finzi et al. 2020). However, they either use only local interactions between close-by particles or they simulate systems with just a few particles. Unfortunately, this does not capture many real-world scenarios, where systems are comprised of thousands of particles that interact over long distances. Using only local interactions would cause very high errors in such cases.

In this contribution, we focus in particular on the N-body problem, since it cannot be solved to sufficient accuracy with only local information. The nonlinearity of the interactions and the fact that the system is chaotic if the number of particles $N > 2$ (Roy 2012) makes this problem particularly complicated. This complexity can be seen in the original work on Hamiltonian Neural Networks (Greydanus, Dzamba, and Yosinski 2019) where the performance on the 2-body problem was good, but significantly deteriorated on the 3-body problem.

We propose a hierarchical model which builds on top of existing graph network (GN) architectures such as (Battaglia et al. 2018) and previous work on accurate physical simulations of a few particles (Sanchez-Gonzalez et al. 2019). Our hierarchical architecture is physically motivated by the multipole expansion and inspired by the fast multipole method (FMM) (Greengard 1988). Our method allows us to extend existing models and simulate complex systems that require $O(N^2)$ interactions with thousands of particles which we have observed empirically to be infeasible when a fully connected graph is used. Importantly, models that use our hierarchical approach retain similar accuracy to models working with a fully connected graph and a smaller particle count.

Finally, we note that the fast multipole method and multipole expansion have been applied to various other problems, such as flow simulations (Koumoutsakos and Leonard 1995), acoustics (Gunda 2008), molecular dynamics (Board Jr et al. 1992; Ding, Karasawa, and Goddard III 1992) and even interpolation of scattered data (reconstruction of a 3D object mesh) (Carr et al. 2001). This suggests that our hierarchical method should also facilitate learning on a similarly wide range of problems.

2 Related Work

Recent studies show that neural networks can successfully learn to simulate complex physical processes (Battaglia et al. 2016; Sanchez-Gonzalez et al. 2018; Mrowca et al. 2018; Li et al. 2018; Greydanus, Dzamba, and Yosinski 2019; Sanchez-Gonzalez et al. 2019, 2020). Most of the existing work focuses on introducing better physical biases such as a relational model (Battaglia et al. 2016), conservation law bias (Greydanus, Dzamba, and Yosinski 2019), combining

these with an ODE bias (Sanchez-Gonzalez et al. 2019) and various similar refinements (Chen et al. 2019; Zhong, Dey, and Chakraborty 2019; Saemundsson et al. 2020; Desai and Roberts 2020). Unfortunately, all of these highly accurate and noise resistant models are only able to work with tens of particles in complex spring or gravitational systems. Models which are focused on simulating fluids, rigid and deformable bodies are able to work with thousands of particles (Mrowca et al. 2018; Li et al. 2018; Sanchez-Gonzalez et al. 2020). However, in these scenarios, it is possible to achieve state of the art performance by just using local information (Sanchez-Gonzalez et al. 2020).

There is a long history of attempts to improve the scalability of traditional particle simulations. One of the first efficient methods was the tree code introduced by (Barnes and Hut 1986). It uses a hierarchical spatial tree (i.e. quadtree) to define localised groups of particles. Particles then interact directly with other groups (instead of their member particles) if the distance to the group is sufficiently larger than the radius encompassing all of the particles in that group, to avoid large errors. Interaction list for each particle would be computed by going through the tree top-down and at each level adding all of the cells that are sufficiently far away and haven't been covered by previous interactions to the list. At the lowest level, particles from neighbouring cells would interact directly. This algorithm has $O(N \log(N))$ complexity. Fast multipole method (FMM) (Greengard and Rokhlin 1987) improved this approach by adding cell-cell interactions and then propagating the resulting effects to the children. This algorithm has $O(N)$ complexity for the force estimation. It has been shown by (Dehnen 2014) that by carefully tweaking the implementation of the FMM it is possible to achieve comparable gravitational force errors to a direct summation method that computes all $O(N^2)$ interactions.

In this work, we apply the lessons learned in scaling traditional particle simulations to current approaches of learning simulations from data using graph neural networks.

There have been attempts to incorporate a hierarchical structure to neural simulators of fluids, rigid and deformable bodies (Mrowca et al. 2018; Li et al. 2018). However, in those cases, the focus was to improve the accuracy and effect propagation in simulations with only local interactions between particles. The resulting methods are not as scalable and are ill-suited for simulations of long-range force field interactions. (Li et al. 2018) restricted hierarchy to one level of groups. (Mrowca et al. 2018) included edges between children and all of their grandparents in the hierarchy. This resulted in $O(N \log N)$ complexity when propagating information through their hierarchical graph. To construct the hierarchy, recursive applications of k-means clustering were used, which can be more computationally expensive than using a quadtree or an octree (Arthur and Vassilvitskii 2006). In these hierarchies, nodes only interact with other nodes in the same group. This can result in two particles that are next to each other having no direct interaction, even though close interactions are the most important ones. We also know from physics that cluster-cluster force field interactions can only be approximated if the clusters are sufficiently far away (Dehnen 2014). Nevertheless, neighbouring clusters often interact directly in these hierarchies. These issues make existing hierarchical approaches ill-suited for force field interactions. This is unsurprising, as they were designed for a fundamentally different problem.

Successful non-graph-based neural simulators have also been proposed (Ummenhofer et al. 2019; Finzi et al. 2020). However, they only use local information from the particle's neighbourhood to update its position.

3 Model

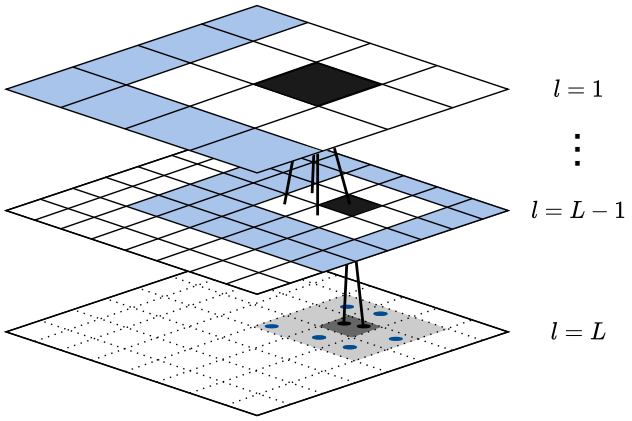

Figure 1: 3 level hierarchy ($L = 3$) from the point of view of the black particles that belong to the black cells. The first level ($l = 1$) has 16 cells, in each subsequent level, the number of cells quadruples. At the lowest level, we have the particles. Each cell is connected to its four child cells at the lower level. Cells at the second-lowest level ($L - 1$) are instead connected to the particles that belong to them. First, during an upward pass cell features are recursively computed from the features of their children. Then, during the downward pass, at each level, the black cell interacts with its near-neighbours (blue). Aggregated interactions are propagated to the cell's children. At the lowest level (L), particles interact directly with particles (blue) from the neighbouring cells (grey).

3.1 Hierarchical Graph

We build the hierarchy by recursively subdividing the space into four parts (2D space is assumed, but our approach naturally generalises to 3D spaces as it is based on quadtrees). When a cell is split we add edges between it and its children. We assume that the particles are roughly uniformly distributed and repeat the splitting $\lfloor \log(N) \rceil$ times. In expectation, this results in having one particle per cell at the lowest level. In the general case, when the data is not uniform, each cell can be split until there are at most k particles in it, where $k \geq 1$ is some constant. For real-world data, this should also result in linear experimental complexity as it does for the FMM (Dehnen 2014; Kabadshow 2012). Empty cells are pruned. We also include another level at the bottom of the hierarchy that holds all of the particles. Each particle has an edge to the cell it belongs to at the level above.

Both, the particles and the cells are represented as nodes in the hierarchical graph. Particle nodes have mass, position and velocity as their features. For the cell nodes, the total mass, centre of mass position and velocity are pre-computed from their children. If Coulomb interactions are simulated, the particle charge and total charge features are added to the particle and cell nodes respectively.

Each cell is also connected to its near-neighbours. Near-neighbours are other cells at the same level that are not directly adjacent to the cell, but whose parents were adjacent to the cell's parent (blue, Figure 1). This way we recursively push down the closest interactions to lower levels, where the interactions are more granular. Ensuring that interacting cells are never next to each other is also necessary in order to avoid potentially large errors (Dehnen 2014). At the lowest level particles are directly connected with other particles that belong to the same cell or the neighbouring cells. Importantly, two particles that are next to each other are never separated. This hierarchy also ensures that the receptive field of each particle is close to being symmetric at every level. The top-level with 4 cells is removed from the hierarchy, as all of the cells are neighbours and do not interact. A visual depiction of the hierarchy can be seen in Figure 1.

It is easy to see that constructing a quadtree with $\log_4 N$ levels takes $O(N \log N)$ time and $O(N)$ space. Next, we demonstrate that the strategy used to construct the hierarchy yields a number of nodes and edges that scale linearly with N (instead of the typical quadratic complexity discussed earlier).

Theorem 1. *There are $O(N)$ nodes in the hierarchy.*

Proof. Based on our assumption that particles are uniformly distributed, after $\log_4(N)$ splits we will have one particle per cell and thus N cells. In each level going from the bottom-up we will have 4 times fewer cells. This gives us the following geometric progression for the total number of cells:

$$N + \frac{N}{4} + ... + 16 = \sum_{l=0}^{\log_4(N)-2} N \left(\frac{1}{4}\right)^l < 2N \ .$$

Considering that in the hierarchy we also include one level with all of the particles, we will have $< 3N = O(N)$ nodes in the hierarchical graph. □

Theorem 2. *There are $O(N)$ edges in the hierarchy and each node has at most a constant number of edges.*

Proof. As we assume that particles are uniformly distributed and that in expectation we have 1 particle per cell in the lowest level, each particle will have $8 + 1$ edges in expectation (since each particle is only connected to particles from its 8 neighbouring cells and their parent cell).

Each cell in the hierarchy is connected to its parent (≤ 1), its children (≤ 4) and its near neighbours (≤ 27). There are at most 27 near neighbours because the cell is connected to other cells that belong to its parent or the neighbours of its parent ($1 + 8$ parent cells, each with 4 children - 36 cells), but not itself or the cells that are its immediate neighbours ($36 - 1 - 8 = 27$).

So, in expectation, we will always have at most a constant number of edges per node. As we have $O(N)$ nodes as per Theorem 1, we will also have $O(N)$ edges. □

3.2 Graph Networks

Graph network (GN) models (Battaglia et al. 2018) operate on a graph $G = (V, E, \boldsymbol{u})$, with global features $\boldsymbol{u}$, nodes V and edges E. Graph networks enforce a structure similar to traditional simulation methods, where first interactions (edges) between particles (nodes) are computed. Then all incoming interactions (edges) are aggregated per particle (node) and together with particle features are used to compute updated particle features. Global values such as the Hamiltonian (total energy) of the system can also be computed from the interactions and the particle features. This relational bias has been shown to greatly improve the accuracy of simulators compared to using a single multi-layer perceptron (MLP) (Battaglia et al. 2016; Sanchez-Gonzalez et al. 2018).

In the basic case, a particle system is represented as a fully connected graph, where each node is a particle. The node features we use are mass m, position $\boldsymbol{q}$, velocity $\dot{\boldsymbol{q}}$ and if applicable charge z. The edge matrix E holds sender and receiver IDs for each edge. In our case relative node positions are used in the models. Meaning that one of the edge features used during the forward pass is the distance vector between the sender and the receiver. Node positions are masked everywhere else.

As graph network models perform distinct operations on each edge and each vertex their time and space complexity is linear in the number of edges and the number of vertices. If a fully connected graph is used this results in overall $O(N^2)$ time and space complexity.

Delta Graph Network (DeltaGN). DeltaGN is analogous to previous direct neural simulators (Battaglia et al. 2016; Sanchez-Gonzalez et al. 2018, 2019, 2020). This model directly predicts the change in particle position $\boldsymbol{q}$ and velocity $\dot{\boldsymbol{q}}$:

$$(\boldsymbol{q}, \dot{\boldsymbol{q}})_{t+1} = (\boldsymbol{q}, \dot{\boldsymbol{q}})_t + (\Delta\boldsymbol{q}, \Delta\dot{\boldsymbol{q}}) \ ,$$

where $(\Delta\boldsymbol{q}, \Delta\dot{\boldsymbol{q}}) = \text{GN}_V(V, E, \Delta t)$ are the new vertex features produced by the graph network.

Hamiltonian ODE Graph Network (HOGN). HOGN (Sanchez-Gonzalez et al. 2019) uses a graph network to compute the Hamiltonian of the system (single scalar):

$$H_{\text{GN}}(\boldsymbol{q}, \boldsymbol{p} = m \cdot \dot{\boldsymbol{q}}) = \text{GN}_u(V, E) \ .$$

By deriving this output w.r.t. the inputs of the network (particle position $\boldsymbol{q}$ and momentum $\boldsymbol{p}$) and using Hamilton's equations we can recover the derivatives of particle position and momentum:

$$f^{HOGN}_{\dot{\boldsymbol{q}},\dot{\boldsymbol{p}}} = \left(\frac{\partial H}{\partial p}, -\frac{\partial H}{\partial q}\right) = (\dot{\boldsymbol{q}}, \dot{\boldsymbol{p}}) \ .$$

The position updates for a given Δt are then produced by a differentiable Runge–Kutta 4 (RK4) integrator that repeatedly queries $f^{HOGN}_{\dot{\boldsymbol{q}},\dot{\boldsymbol{p}}}$.

3.3 Adapting Existing Models

The idea behind adapting existing graph network models to use the hierarchy is simple: instead of using a dense representation of the particle interactions as input, we use the sparse particle interaction graph from the lowest level of the hierarchy (Figure 1). When the graph network model constructs updated edge feature matrix E' for this sparse graph, we append to it edges coming to each particle from its parent cell. This augmented edge feature matrix is then used to update vertex features and to calculate any global features in the same manner as done in the original models.

The construction of this special edge – representing all distant interactions for a particle – requires propagating the information through the hierarchy. To this end, the updated lowest level cell feature vectors are built from their children as:

$$\boldsymbol{v}'_c = \boldsymbol{v}_c \oplus \sum_{p \in \text{children}(c)} \phi^{p \to c}(m_c, \dot{\boldsymbol{q}}_c, m_p, \dot{\boldsymbol{q}}_p, (\boldsymbol{q}_c - \boldsymbol{q}_p)) \ ,$$

where v_c is the original feature vector of a cell (total mass and centre of mass velocity), p is the child particle, $\phi^{p \to c}$ is an MLP and $\oplus$ is a concatenation operator. The features of the parent cell c_p in the upper levels are built in a similar way from their children cells c_c:

$$\boldsymbol{v}'_{c_p} = \boldsymbol{v}_{c_p} \oplus \sum_{c_c \in \text{children}(c_p)} \phi^{c_c \to c_p}(m_{c_p}, \dot{\boldsymbol{q}}_{c_p}, \boldsymbol{v}'_{c_c}, (\boldsymbol{q}_{c_p} - \boldsymbol{q}_{c_c})).$$

The $\phi^{c_c \to c_p}$ MLP uses a different set of parameters from $\phi^{p \to c}$, but $\phi^{c_c \to c_p}$ parameters are shared between all of the levels.

During the top-down pass at each level cell-cell interactions are computed:

$$\boldsymbol{e}'_{c_j, c_i} = \phi^{c \to c}(\boldsymbol{v}'_{c_i}, \boldsymbol{v}'_{c_j}, (\boldsymbol{q}_{c_i} - \boldsymbol{q}_{c_j})) \ .$$

For each cell, incoming interactions are aggregated (summed) together with interactions propagated from the parent:

$$\boldsymbol{e}'_{\text{parent}_c} = \phi^{c_p \to c_c}(\boldsymbol{v}'_c, \boldsymbol{v}'_{c_p}, (\boldsymbol{q}_c - \boldsymbol{q}_{c_p})) \ .$$

These aggregated interactions $\boldsymbol{e}'_c = \boldsymbol{e}'_{\text{parent}_c} + \sum_{c_j} \boldsymbol{e}'_{c_j, c}$ are used to update cell features:

$$\boldsymbol{v}''_c = \boldsymbol{v}_c \oplus \phi^c(\boldsymbol{v}'_c, \boldsymbol{e}'_c) \ .$$

Finally, the cell-particle edges are computed:

$$\boldsymbol{e}'_{\text{parent}_p} = \phi^{c \to p}(m_p, \dot{\boldsymbol{q}}_p, \boldsymbol{v}''_c, (\boldsymbol{q}_p - \boldsymbol{q}_c)) \ .$$

The $\phi^{c_p \to c_c}$ MLP uses different parameters from $\phi^{c \to p}$. Parameters of $\phi^{c_p \to c_c}$ and $\phi^{c \to c}$ are shared between the hierarchy levels.

If any global parameters $\boldsymbol{u}$ are used by the model they are also used as part of the input to all of the hierarchy MLPs.

It is easy to see that all of these operations have time and space complexity that is linear in the number of nodes and edges.

4 Experiments

4.1 Data

We built our own N-body simulator. It uses a symplectic Leapfrog integrator that preserves the total energy of the system, Plummer force softening that helps avoid the singularity which arises when the distance between two particles goes to zero, individual and dynamic time steps for particles (Dehnen and Read 2011). The time step is set for each particle as a fraction of the base time step based on its acceleration. The particle state (mass, position, velocity) is saved at every base time step. The system we simulate uses periodic boundary conditions. When a particle leaves the unit cell its exact copy enters it on the opposite side.

We simulate 1000 training, 200 validation and 200 test trajectories. Every trajectory is 200 base time steps long. Particle positions are initialised uniformly at random inside the unit cell, mass is set to 1 and each velocity vector component is initialised uniformly at random over $(-1, 1)$. Additionally, if we simulate Coulomb interactions each particle has its charge initialised uniformly at random over $(0.5, 1.5)$ and assigned a random sign. The base time step is set to $\Delta t = 0.01$. Gravitational and Coulomb constants are respectively set to $G = 2$ and $k = 2$. We assume that in our simulated galaxy all of the units are dimensionless.

4.2 Results

We compare the model that uses our hierarchy (Hierarchical DeltaGN) against two baselines: model that uses a fully connected graph (DeltaGN) and a more computationally efficient model that uses 15 nearest neighbour graph (DeltaGN (15 nn)). We also perform an experiment to test if our hierarchical approach is compatible with the HOGN model. Models are trained for 500 thousand steps using a batch size of 100 unless stated otherwise. We exponentially decay the learning rate every 200 thousand steps by 0.1. The initial learning rate for all of the models was set to 0.0003. The models are optimised using ADAM (Kingma and Ba 2014), and the mean square error (MSE) which is computed between the predicted and true phase space coordinates after one time step.

All datasets have the same particle density. When the particle count is increased, the size of the unit cell is increased accordingly.

When the rollout error is reported, it means that 200 test trajectories T are unrolled in an autoregressive manner for a specified number of time steps. The input graph is rebuilt each time using the model's outputs. The error is computed as RMSE between phase-space coordinates of the predictions and the ground truth over the unrolled trajectory.

Gravitational and Coulomb systems are conservative, which means that their total energy should stay constant throughout their evolution. From this arises a common accuracy measure used in N-body simulations - relative energy error between the first and the last system states. We calculate mean relative energy error over all test trajectories:

$$\text{Energy error} = \frac{1}{|T|} \sum_{i=1}^{|T|} \frac{H_{i,0} - \hat{H}_{i,\tau}}{H_{i,0}} \ ,$$

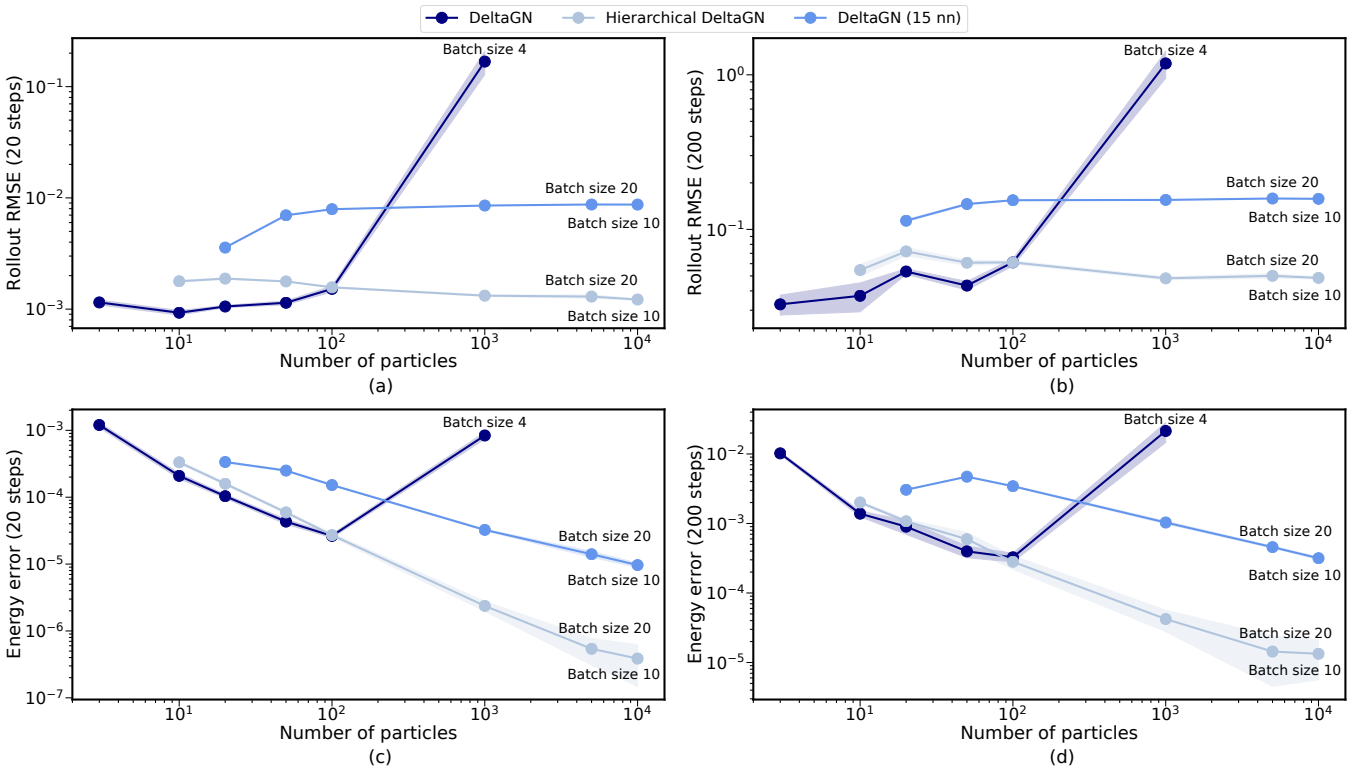

Figure 2: Models trained and evaluated on increasing particle counts. (a) 20 step rollout RMSE, (b) 200 step rollout RMSE, (c) energy error after 20 steps, (d) energy error after 200 steps. Batch sizes smaller than 100 were used if not enough memory was available (shown in the graph).

where τ is the number of time steps the trajectory is unrolled for and $H_{i,t}$ is the Hamiltonian (total energy) of the system at time step t of the trajectory i.

All errors are averaged over 5 independent runs of the model. We report the mean error and the standard deviation.

The experiments were performed on a machine with an Intel Xeon E5-2690 v3 CPU (2.60GHz, 12 cores), 64GB RAM and NVIDIA Tesla P100 GPU (16GB RAM).

Scaling to Larger Particle Counts. From Figure 2 we can see that hierarchical DeltaGN does not perform as well on small particle counts (< 100) where the hierarchy is small. However, it quickly catches up when 100 particles are simulated and is able to simulate large particle counts with even better accuracy. Simulations with more particles are not feasible with DeltaGN that uses the fully-connected graph, as it soon runs out of memory and on the 1000 particle dataset already suffers from terrible performance. This is probably due to the very large edge count per node. The DeltaGN (15 nn) model that uses only local connectivity is able to scale to large particle counts but suffers from poor accuracy due to the ignored long-range interactions. When the same small batch size is used for all of the models the situation stays similar (Figure 3).

The energy error usually decreases with the particle count, because we keep the particle density constant. Thus, when extra particles are added, many long-range interactions are created, but the number of close interactions each particle has stays roughly the same. Close interactions are more error-prone due to much stronger forces at close distances.

Small Batch Size. One easy way to reduce memory usage and to decrease the computation time is to reduce the batch size. However, this potentially comes with a big accuracy penalty. As can be seen from the rollout RMSE (Figure 3a) and the energy error (Figure 3b) plots training models with a batch size of 1 results in a $\sim 10\%$ decrease in accuracy for DeltaGN (15 nn) and a $\sim$ 5 time decrease in accuracy for the other models.

We note that using the same batch size is not entirely fair to models with $O(N)$ complexity. Indeed, they have much fewer edge samples in the batch, compared to the model that uses a fully-connected graph. This lack of edge samples is most likely the reason why some runs of the hierarchical DeltaGN on 20 and 50 particle datasets got stuck in local minima and caused high variance.

Coulomb Interactions. We expect that Coulomb interactions are harder to learn since they can be either attractive or repulsive. We also made this dataset more complex by simulating charges of different magnitudes. In Table 1 we can see that on the 100 particle dataset the hierarchical DeltaGN again performs similarly to the model that uses a fully-connected graph. However, when the particle count increases to 1000 DeltaGN fails, while the hierarchical Delt-

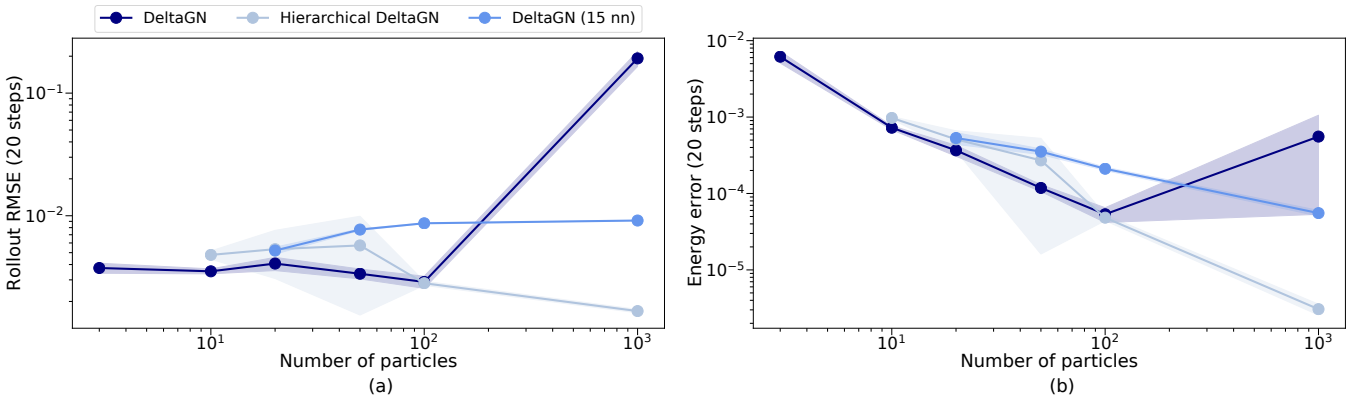

Figure 3: Models trained using a batch size of 1. Their 20 step rollout RMSE (a) and energy error (b).

Model	Number of particles	Rollout RMSE (20 steps)	Energy error (20 steps)	Rollout RMSE (200 steps)	Energy error (200 steps)
DeltaGN	100	$\mathbf{0.0045 \pm 2.7 \cdot 10^{-4}}$	$\mathbf{0.0086 \pm 0.0041}$	$\mathbf{0.0684 \pm 0.0031}$	$\mathbf{0.0584 \pm 0.0202}$
Hierarchical DeltaGN	100	$0.0052 \pm 1.61 \cdot 10^{-4}$	0.0152 ± 0.0036	0.1013 ± 0.0019	0.0931 ± 0.0138
DeltaGN	1000	0.1744 ± 0.0234	0.0592 ± 0.0507	NA	NA
Hierarchical DeltaGN	1000	$\mathbf{0.0058 \pm 1.45 \cdot 10^{-4}}$	$\mathbf{0.0014 \pm 1.90 \cdot 10^{-4}}$	$\mathbf{0.1252 \pm 0.0020}$	$\mathbf{0.0418 \pm 0.0031}$

Table 1: Test accuracy achieved by DeltaGN models on the Coulomb force datasets. DeltaGN used maximum possible batch size of 4 on the 1000 particle dataset, while in all other cases batch size of 100 was used. NA means that during unroll for some of the runs particle velocities grew so large that float32 overflowed and model returned NaN values.

aGN retains similar performance.

The particle charge was supplied to the models alongside the features supplied in the case of gravitational force.

Hierarchical HOGN. Five independent runs of hierarchical HOGN were trained on the 100 and 1000 particle datasets. As can be seen from Table 2 these runs had very high variance. However, one of the hierarchical HOGN runs resulted in the most accurate model we have trained on the 1000 particle dataset. Although, the Hamiltonian and ODE biases did not bring as large of an improvement, as seen when a fully connected graph is used (Sanchez-Gonzalez et al. 2019).

We do believe that the training process of the hierarchical HOGN can be made more stable, but we leave this direction for future work.

Empirical Time Complexity. We timed 10 000 runs of the model forward pass, using different graphs (Figure 4a). We observe, that while GPU parallelism initially counteracts the increased complexity, using a fully-connected graph results in asymptotically quadratic time complexity. While using hierarchical and 15 nearest neighbour graphs results in asymptotically linear scaling. The pre-computation time of the nearest neighbour graphs is asymptotically quadratic, while the pre-computation time of our hierarchical graphs scales as $O(N \log N)$ (Figure 4b). In both cases, pre-computation on the CPU was faster than on the GPU with our implementation. Time scaling is not monotonic for the hierarchical graph pre-computation and the hierarchical DeltaGN, because the depth of the hierarchy is set as $\lfloor \log_4(N) \rceil$.

Generalisation to Unseen Particle Counts. Models that use a fully-connected graph suffer from disastrous performance when they are evaluated on particle counts they were not trained on (Figure 5). The most likely cause is the sharp change in the number of incoming edges for each node. This is corroborated by the fact that DeltaGN (15 nn), which uses a graph with a constant number of neighbours, achieves almost the same accuracy on unseen particle counts as the models trained on that particle count. In our hierarchy, we have a roughly constant number of edges per node (Theorem 2). As a result, the hierarchical DeltaGN generalises to unseen particle counts much better than models that use a fully-connected graph. When more than 100 particles are used, it outperforms DeltaGN (15 nn). Our hierarchical approach can be combined with existing techniques, such as randomly dropping edges during training (Rong et al. 2019) or constraining message size (Cranmer et al. 2019), to further improve the generalisation.

5 Conclusion

We presented a novel hierarchical graph construction technique that can be used to adapt existing graph network models. We show that this hierarchical graph improves the model time and space complexity from $O(N^2)$ to $O(N)$ when applied to particle simulations that require $O(N^2)$ interactions.

Model	Number of particles	Batch size	Rollout RMSE (20 steps)	Energy error (20 steps)	Rollout RMSE (200 steps)	Energy error (200 steps)
Hierarchical DeltaGN	100	100	$0.0016 \pm 1.74 \cdot 10^{-5}$	$(2.71 \pm 0.18) \cdot 10^{-5}$	0.0612 ± 0.0015	$(2.82 \pm 0.63) \cdot 10^{-4}$
HOGN	100	50	$\mathbf{(2.09 \pm 0.26) \cdot 10^{-4}}$	$\mathbf{(6.30 \pm 2.14) \cdot 10^{-6}}$	$\mathbf{0.0139 \pm 0.0023}$	$\mathbf{(4.41 \pm 1.03) \cdot 10^{-5}}$
Hierarchical HOGN	100	100	0.0046 ± 0.0042	$(8.57 \pm 8.48) \cdot 10^{-5}$	0.1145 ± 0.0814	0.0016 ± 0.0015
Hierarchical HOGN (best)	100	100	0.0017	$3.06 \cdot 10^{-5}$	0.0579	$3.55 \cdot 10^{-4}$
Hierarchical DeltaGN	1000	100	$0.0013 \pm 1.99 \cdot 10^{-5}$	$(2.37 \pm 0.40) \cdot 10^{-6}$	0.0482 ± 0.0011	$(4.22 \pm 1.44) \cdot 10^{-5}$
Hierarchical HOGN	1000	20	0.0023 ± 0.0011	$(4.87 \pm 1.99) \cdot 10^{-6}$	0.0691 ± 0.0299	$(4.27 \pm 1.42) \cdot 10^{-5}$
Hierarchical HOGN (best)	1000	20	**0.0012**	$\mathbf{2.03 \cdot 10^{-6}}$	**0.0377**	$\mathbf{2.87 \cdot 10^{-5}}$

Table 2: Hierarchical HOGN. Due to the very high variance of the hierarchical HOGN runs we also provide the results for the best run. All models were trained with a batch size of 100 when possible and largest batch size otherwise. HOGN ran out of memory on the 1000 particle dataset even with a batch size of 1.

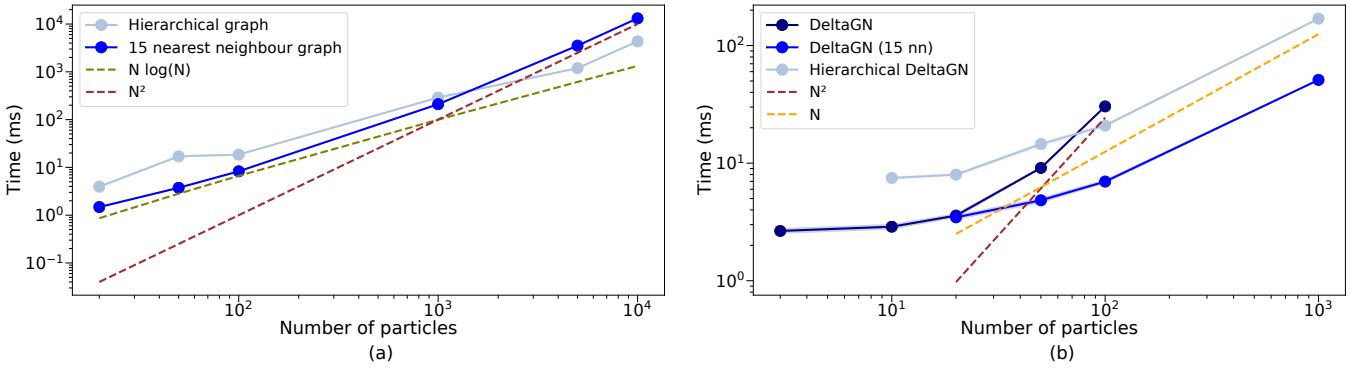

Figure 4: (a) Time scaling of the model forward pass using a batch size of 100. (b) Time scaling of graph pre-computation on a CPU. Averaged over 10 000 (a) and 1000 (b) runs respectively.

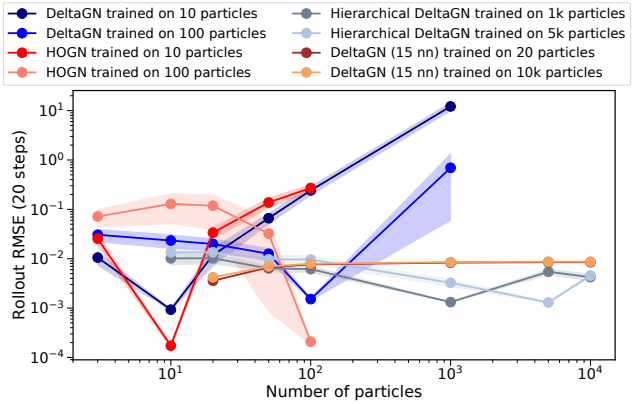

Figure 5: Models trained on one particle count evaluated on different particle counts. 20 step rollout RMSE. The trained models from Figure 2 were used.

A pre-computation step also requires $O(N \log N)$ time and $O(N)$ space. This theoretical improvement is also observed in practice and allows us to train models on much larger datasets than previously possible. We were able to achieve good accuracy on a dataset with 10 000 particles, while prior models that use a fully-connected graph failed on 1000 particles. Our approach also displayed much better accuracy than a standard baseline based on the nearest neighbour graph. Furthermore, we observed improved generalisation to different particle counts when the hierarchy is used. We hypothesise this is due to the regularisation enforced by the hierarchy (the number of edges per particle tends to be more constant), although this requires more investigation.

While we mostly focused on the faster model that directly predicts the system's state change, we have also shown that this hierarchical approach is compatible with Hamiltonian and ODE biases. However, the inclusion of these biases does not seem to have the desired impact on the model's accuracy and further work is required to reduce the variance of the resulting model.

Finally, this work takes an important step towards learning to simulate larger and more realistic dynamical systems in a way that is compatible with many existing state-of-the-art approaches. Developing neural simulators is important because it can even lead to the discovery of novel physics formulas (Cranmer et al. 2020). In general, this method could be used to extend graph networks that are applied to other problems where global information is needed.

Acknowledgements

The authors would like to thank Roberta Huang, Thomas Hofmann, Janis Fluri, Tomasz Kacprzak and Alexandre Refregier for their involvement in the early phase of this project.

References

Arthur, D.; and Vassilvitskii, S. 2006. How slow is the k-means method? In *Proceedings of the twenty-second annual symposium on Computational geometry*, 144–153.

Barnes, J.; and Hut, P. 1986. A hierarchical O (N log N) force-calculation algorithm. *nature* 324(6096): 446–449.

Battaglia, P.; Pascanu, R.; Lai, M.; Rezende, D. J.; et al. 2016. Interaction networks for learning about objects, relations and physics. In *Advances in neural information processing systems*, 4502–4510.

Battaglia, P. W.; Hamrick, J. B.; Bapst, V.; Sanchez-Gonzalez, A.; Zambaldi, V.; Malinowski, M.; Tacchetti, A.; Raposo, D.; Santoro, A.; Faulkner, R.; et al. 2018. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261* .

Board Jr, J. A.; Causey, J. W.; Leathrum Jr, J. F.; Windemuth, A.; and Schulten, K. 1992. Accelerated molecular dynamics simulation with the parallel fast multipole algorithm. *Chemical Physics Letters* 198(1-2): 89–94.

Carr, J. C.; Beatson, R. K.; Cherrie, J. B.; Mitchell, T. J.; Fright, W. R.; McCallum, B. C.; and Evans, T. R. 2001. Reconstruction and representation of 3D objects with radial basis functions. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, 67–76.

Chen, Z.; Zhang, J.; Arjovsky, M.; and Bottou, L. 2019. Symplectic recurrent neural networks. *arXiv preprint arXiv:1909.13334* .

Cranmer, M.; Sanchez-Gonzalez, A.; Battaglia, P.; Xu, R.; Cranmer, K.; Spergel, D.; and Ho, S. 2020. Discovering Symbolic Models from Deep Learning with Inductive Biases. *arXiv preprint arXiv:2006.11287* .

Cranmer, M. D.; Xu, R.; Battaglia, P.; and Ho, S. 2019. Learning Symbolic Physics with Graph Networks. *arXiv preprint arXiv:1909.05862* .

Dehnen, W. 2014. A fast multipole method for stellar dynamics. *Computational Astrophysics and Cosmology* 1(1): 1.

Dehnen, W.; and Read, J. I. 2011. N-body simulations of gravitational dynamics. *The European Physical Journal Plus* 126(5): 55.

Desai, S.; and Roberts, S. 2020. VIGN: Variational Integrator Graph Networks. *arXiv preprint arXiv:2004.13688* .

Ding, H.-Q.; Karasawa, N.; and Goddard III, W. A. 1992. Atomic level simulations on a million particles: The cell multipole method for Coulomb and London nonbond interactions. *The Journal of chemical physics* 97(6): 4309–4315.

Finzi, M.; Stanton, S.; Izmailov, P.; and Wilson, A. G. 2020. Generalizing convolutional neural networks for equivariance to lie groups on arbitrary continuous data. *arXiv preprint arXiv:2002.12880* .

Greengard, L. 1988. *The rapid evaluation of potential fields in particle systems*. MIT press.

Greengard, L.; and Rokhlin, V. 1987. A fast algorithm for particle simulations. *Journal of computational physics* 73(2): 325–348.

Greydanus, S.; Dzamba, M.; and Yosinski, J. 2019. Hamiltonian neural networks. In *Advances in Neural Information Processing Systems*, 15379–15389.

Gunda, R. 2008. Boundary element acoustics and the fast multipole method (FMM). *Sound and Vibration* 42(3): 12.

Kabadshow, I. 2012. *Periodic boundary conditions and the error-controlled fast multipole method*. Forschungszentrum Jülich.

Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* .

Koumoutsakos, P.; and Leonard, A. 1995. High-resolution simulations of the flow around an impulsively started cylinder using vortex methods. *Journal of Fluid Mechanics* 296: 1–38.

Li, Y.; Wu, J.; Tedrake, R.; Tenenbaum, J. B.; and Torralba, A. 2018. Learning Particle Dynamics for Manipulating Rigid Bodies, Deformable Objects, and Fluids. In *International Conference on Learning Representations*.

Mrowca, D.; Zhuang, C.; Wang, E.; Haber, N.; Fei-Fei, L. F.; Tenenbaum, J.; and Yamins, D. L. 2018. Flexible neural representation for physics prediction. In *Advances in neural information processing systems*, 8799–8810.

Rong, Y.; Huang, W.; Xu, T.; and Huang, J. 2019. Dropedge: Towards deep graph convolutional networks on node classification. In *International Conference on Learning Representations*.

Roy, A. E. 2012. *Predictability, stability, and chaos in N-body dynamical systems*, volume 272. Springer Science & Business Media.

Saemundsson, S.; Terenin, A.; Hofmann, K.; and Deisenroth, M. 2020. Variational Integrator Networks for Physically Structured Embeddings. In *International Conference on Artificial Intelligence and Statistics*, 3078–3087.

Sanchez-Gonzalez, A.; Bapst, V.; Cranmer, K.; and Battaglia, P. 2019. Hamiltonian graph networks with ode integrators. *arXiv preprint arXiv:1909.12790* .

Sanchez-Gonzalez, A.; Godwin, J.; Pfaff, T.; Ying, R.; Leskovec, J.; and Battaglia, P. W. 2020. Learning to simulate complex physics with graph networks. *arXiv preprint arXiv:2002.09405* .

Sanchez-Gonzalez, A.; Heess, N.; Springenberg, J. T.; Merel, J.; Riedmiller, M.; Hadsell, R.; and Battaglia, P. 2018. Graph networks as learnable physics engines for inference and control. *arXiv preprint arXiv:1806.01242* .

Ummenhofer, B.; Prantl, L.; Thuerey, N.; and Koltun, V. 2019. Lagrangian fluid simulation with continuous convolutions. In *International Conference on Learning Representations*.

Zhong, Y. D.; Dey, B.; and Chakraborty, A. 2019. Symplectic ode-net: Learning hamiltonian dynamics with control. *arXiv preprint arXiv:1909.12077* .