

A MULTIREOLUTION DISCRETE ELEMENT METHOD FOR TRIANGULATED OBJECTS WITH IMPLICIT TIME STEPPING*

PETER J. NOBLE[†] AND TOBIAS WEINZIERL[†]

Abstract. Simulations of many rigid bodies colliding with each other sometimes yield particularly interesting results if the colliding objects differ significantly in size and are nonspherical. The most expensive part within such a simulation code is the collision detection. We propose a family of novel multiscale collision detection algorithms that can be applied to triangulated objects within explicit and implicit time stepping methods. They are well suited to handle objects that cannot be represented by analytical shapes or assemblies of analytical objects. Inspired by multigrid methods and adaptive mesh refinement, we determine collision points iteratively over a resolution hierarchy and combine a functional minimization plus penalty parameters with the actual comparison-based geometric distance calculation. Coarse surrogate geometry representations identify “no collision” scenarios early on and otherwise yield an educated guess which triangle subsets of the next finer level might yield collisions. They prune the search tree and furthermore feed conservative contact force estimates into the iterative solve behind an implicit time stepping. Implicit time stepping and nonanalytical shapes often yield prohibitive high compute cost for rigid body simulations. Our approach reduces the object-object comparison cost algorithmically by one to two orders of magnitude. It also exhibits high vectorization efficiency due to its iterative nature.

Key words. discrete element method, triangle collision checks, implicit time stepping, multiscale methods, surrogate geometries

MSC codes. 70E55, 70F35, 68U05, 51P05, 37N15

DOI. 10.1137/21M1421842

1. Introduction. The simulation of rigid or incompressible bodies is a challenge that arises in many fields. Notably, it is at the heart of discrete element method (DEM) simulations, where millions of these objects are studied. Progress on simulations with rigid, impenetrable objects depends on whether we can handle high geometric detail accurately: For the analysis of particle flow such as powder, it is mandatory to simulate billions of particles, i.e., incompressible bodies, while the realism of some simulations hinges on the ability to handle particles of different shapes and sizes [2, 13, 14, 15, 16, 22, 23, 24, 25]. It is the support of different shapes and sizes that allows us to simulate complex mixture phenomena, separation of scales, or blockage if many objects, also known as particles, try to squeeze through narrow geometries.

DEM codes spend most of their runtime on collision detection [13, 15, 16, 19]. We tighten the challenge and study a DEM prototype over particles where (i) rigid particles have massively differing size, (ii) rigid particles are discretized by many triangles, and (iii) rigid particles have complicated, nonconvex shapes. Our code supplements each rigid particle with an ϵ -area [2, 22, 24] and considers two particles to be “in contact” if their ϵ -environment overlaps. This yields a weak compressibility model, where the contact points are unique up to an ϵ -displacement. Yet, both the arrangement and the topology of these points still can change significantly between any two time steps, while the collision models using the contact data remain inherently

*Submitted to the journal’s Methods and Algorithms for Scientific Computing section May 21, 2021; accepted for publication (in revised form) March 14, 2022; published electronically July 28, 2022.

<https://doi.org/10.1137/21M1421842>

Funding: This work was supported by the EPSRC through grant 1764342.

[†]Department of Computer Science, Durham University, Durham, DH1 3LE, UK (peter.j.noble@durham.ac.uk, tobias.weinzierl@durham.ac.uk).

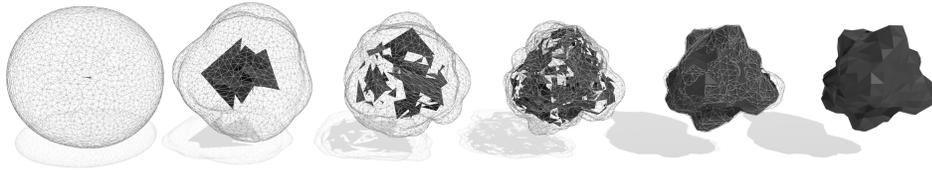


FIG. 1.1. *Surrogate triangle hierarchy for a “bumped sphere.” The actual sphere geometry \mathbb{T}_h is shown on the right. From left to right: Finer and finer surrogate representations including their ϵ -environments. The coarser a representation, the lower the triangle count and the larger the corresponding ϵ . The surrogate triangles are weakly connected. The right most geometry is the actual particle with a closed surface.*

stiff. Large-scale DEM codes require the comparison of many particles per time step. Efficient many-body simulations employ techniques such as neighbor lists and cell or tree metadata structures [2, 10, 15, 22] to narrow down the potential collisions, i.e., to identify particle pairs which might collide as they are spatially close. After this preprocessing or filtering step, they detect collision points between particle pairs. We focus exclusively on the particle-particle comparison challenge, as efficient algorithms for the neighborhood identification—including techniques for challenging shapes and massively differing sizes—are known.

Our work proposes a multiscale contact detection scheme which brings down the compute time for the contact detection aggressively. We can handle complex, nonconvex shapes and even speed up implicit time stepping significantly. The latter is, so far, prohibitively expensive for most codes. To reduce the runtime, our approach phrases the contact search between two particles as an iterative algorithm over multiple resolutions, where coarser particle resolutions act as surrogates (Figure 1.1). This idea enables us to introduce five algorithmic optimizations: First, the surrogates help us identify “no contact” constellations quickly. In this case, we can immediately terminate the search algorithm. The surrogates yield a generalization of classic bounding sphere checks—if two bounding spheres do not overlap, the underlying objects cannot overlap—to highly nonspherical geometries. Second, we exploit that we do not represent surrogate resolutions as a plain level of detail [1] but make them form a tree: If we identify a potential collision, only those sections of the geometry are “up-pixelated” from where a collision point might arise from. We increase the resolution locally. The complexity per iterative step thus does not grow exponentially as we switch to more accurate geometric representations. Instead we tend to have a linear increase of cost as we use finer and finer geometric models. Third, the surrogate resolutions yield conservative estimates of the force that might result from a contact point. Once we employ an implicit time stepping scheme with a Picard iteration, we can permute the iterative solver loop, and the resolution switches such that the Picard iteration forms the outer loop which zaps through resolution levels upon demand. The cheap surrogates provide an educated guess to the iterative force calculation and thus accelerate the convergence. Fourth, we phrase the contact detection as a distance minimization problem [14, 15]. The minimization problem is solved iteratively through an additional, embedded Newton which approximates the Jacobian via a diagonal matrix and runs through a prescribed number of sweeps. Once more, we permute the loop over triangle pairs and the Picard iterations to improve the vectorization suitability. Finally, we acknowledge that an iterative minimization subject to a prescribed iteration count can fail if the underlying geometric problem is ill-posed. In such cases, we eventually postprocess it by falling back to a comparison-based distance calculation.

However, this is only required on the finest mesh, whereas we use nonconvergence within the surrogate tree as sole “refinement criterion” (use a finer mesh) that does not feed forces into the implicit time stepping.

To the best of our knowledge, our rigorous multiscale idea, which can be read as a combination of (i) loop permutation and fusion, (ii) adaptive mesh refinement, (iii) a generalization of volume bounding hierarchies [5, 8, 9, 11, 12], and (iv) an approximate, weak closest-triangle formulation [14, 15], is unprecedented. Its reduction of computational cost plus its excellent vectorization character in combination with the fact that rigid body simulations scale well by construction—the extremely short-range interactions fit well to domain decomposition—brings implicit DEM simulations for triangulated, nonconvex shapes within practical applications into reach. The core algorithmic ideas furthermore have an impact well beyond the realm of DEM. The search for nearest neighbors, i.e., contact within a certain environment, is, for example, also a core challenge behind fluid-structure interaction.

The manuscript is organized as follows: We first introduce our algorithmic challenge and a textbook implementation of both explicit and implicit time stepping for it (section 2). An efficient contact detection between triangulated surfaces of two particles via a minimization problem is introduced next (section 3) before we rewrite the underlying geometric problem as a multiresolution challenge and introduce our notion of a surrogate data structure (section 4). In section 5, we bring both the efficient triangle comparisons and the tree idea together as we plug them into an implicit time stepping code before we allow the nonlinear equation system solvers’ iterations to move up and down within the surrogate tree. This is the core contribution of the manuscript. Following the discussion of some numerical results (section 6), we sketch out future work and close the discussion.

2. Algorithmic framework. We study a system of $|\mathbb{P}|$ rigid bodies (particles). Each particle $p \in \mathbb{P}$ has a velocity $v(p, t)$ and an angular velocity $\omega(p, t)$ which determine its change of position and rotation. Each is described by a triangular tessellation $\mathbb{T}(p, t)$. $t \geq 0$ is the simulation time. We may assume that $\mathbb{T}(p, t)$ spans a well-defined, closed surface represented by a conformal mesh: No two triangles intersect, and we can “run around” a particle infinitely often without falling into a gap. While the triangulation of the object is time-invariant, it moves and rotates over time and therefore depends on t .

Time stepping. A straightforward high-level implementation of an explicit Euler for DEM consists of a time loop hosting a sequence of further loops (Algorithm 2.1): The first inner loop identifies all contact points between the particles. Once we know all contact points per particle, we can determine a velocity and rotation update ($dv(p)$ and $d\omega(p)$) per particle. Before we do so, we update the particles’ positions using their velocity and angular momentum. Finally, we progress in time. Euler–Cromer would result from a permutation of the update sequence.

It is impossible to simulate exact incompressibility with explicit time stepping schemes when no interpenetration is allowed: Everytime we update a particle, we run risk that it slightly penetrates another one due to the finite time step size Δt . At the same time, particles exchange no momentum as long as they are not in direct contact yet. The momentum exchange remains “trivial” until we have violated the rigid body constraint. For these two reasons, we switch to a weak incompressibility model, where each particle is surrounded by an $\epsilon > 0$ area [2, 22, 24]. The area is spanned by the Minkovski sum of the triangles from $\mathbb{T}(p, t)$ and a sphere of radius ϵ minus the actual rigid object. Without loss of generality, we assume a uniform ϵ per particle. Our formalism is equivalent to a soft boundary formulation with an extrusive surface.

Algorithm 2.1 High-level pseudocode for an explicit Euler for rigid particles. The continuous properties $v(p, t)$, $\omega(p, t)$, and $\mathbb{T}(p, t)$ are discretized in time and thus become $v(p)$, $\omega(p)$, and $\mathbb{T}(p)$. t is the (discretized) time, Δt the time step size.

```

1: while  $t < T_{\text{terminal}}$  do                                ▷ We simulate over a time span
2:    $\forall p_i \in \mathbb{P} : \mathbb{C}(p_i) \leftarrow \emptyset$                 ▷ Clear set of collisions for particle  $p_i$ 
3:   for  $p_i, p_j \in \mathbb{P}, p_i \neq p_j$  do                    ▷ Run over all particle pairs
4:      $\mathbb{C}(p_i) \leftarrow \mathbb{C}(p_i) \cup \text{FINDCONTACTS}(\mathbb{T}(p_i), \mathbb{T}(p_j))$ 
5:      $\mathbb{C}(p_j) \leftarrow \mathbb{C}(p_j) \cup \text{FINDCONTACTS}(\mathbb{T}(p_i), \mathbb{T}(p_j))$ 
6:   end for
7:   for  $p_i \in \mathbb{P}$  do
8:      $\mathbb{T}(p_i) \leftarrow \text{UPDATE}(\mathbb{T}(p_i), v(p_i), \omega(p_i), \Delta t)$            ▷ Update geometry
9:   end for                                                ▷ Use velocity, rotation, and time step size
10:  for  $p_i \in \mathbb{P}$  do
11:     $(dv, d\omega) \leftarrow \text{CALCFORCES}(\mathbb{C}(p_i))$ 
12:     $(v, \omega)(p_i) \leftarrow (v, \omega)(p_i) + \Delta t \cdot (dv, d\omega)$      ▷ Update velocity and rotation
13:  end for
14:   $t \leftarrow t + \Delta t$ 
15: end while

```

Algorithm 2.2 Contact identification between two particles p_i and p_j .

```

1: function  $\text{FINDCONTACTS}(\mathbb{T}(p_i), \mathbb{T}(p_j))$ 
2:    $\mathbb{C} = \emptyset$ 
3:   for  $t_i \in \mathbb{T}(p_i), t_j \in \mathbb{T}(p_j), t_i \neq t_j$  do           ▷ Run over all triangles pairs
4:      $c \leftarrow \text{CONTACT}(t_i, t_j)$                           ▷ Find closest point in-between  $t_i$  and  $t_j$ , and compare normal
5:     if  $c \neq \perp$  then                                       ▷  $|n|$  against  $\epsilon$ ; return  $\perp$  if  $|n| > \epsilon$ 
6:        $\mathbb{C} \leftarrow \mathbb{C} \cup \{c\}$ 
7:     end if
8:   end for
9:   return  $\mathbb{C}$ 
10: end function

```

DEFINITION 2.1 (contact point). *Each particle is surrounded by an ϵ -environment. Two particles are in contact if their ϵ -environment overlaps. Overlaps yield contact points which in turn yield forces.*

We parameterize each contact point with its penetration depth: Our contact detection in Algorithm 2.2 identifies the closest path between two triangles. A potential contact point c is located at the center of this line. It is equipped with a normal $n(c)$, which points from the contact point toward either of the closest triangles (Figure 2.1). There is an overlap between the two ϵ -augmented triangles if and only if $|n(c)| \leq \epsilon$, that is, if the ϵ -environments penetrate. In this case c is added to the set of collision points.

With a set of contact points plus their normals and dimensionless masses $M(p_i)$ and $M(p_j)$, we can derive the force acting on a particle.

$$F(c) = \frac{n(c)}{|n(c)|} K_s \left(1 - \frac{|n(c)|}{\epsilon} \right) \sqrt{\frac{1}{\frac{1}{M(p_i)} + \frac{1}{M(p_j)}}}$$

computes the force arising from one contact point by mapping the ϵ -area onto a simple spring with a perpendicular friction force yet without any empirical damping [7]. The force depends on the contact normal $n(c)$ and applies to both colliding particles p_i and p_j subject to a minus sign for one of them. It is calibrated by a spring constant $K_s = 1,000$. A particle's total force is then the sum over the individual contact forces. Taking the center of mass, the total mass, and the mass distribution of a

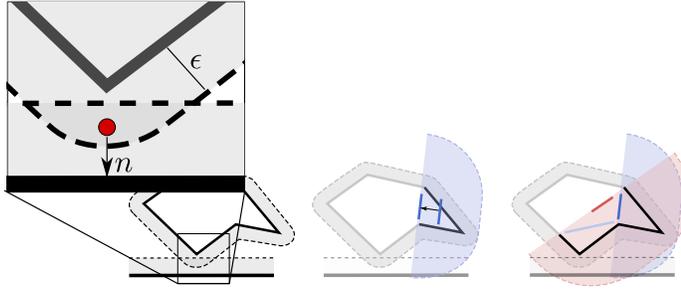


FIG. 2.1. A pair of objects (black, solid) with their ϵ -environment (black, dotted) collide. In the present sketch, one object is a “spherical” particle spanned by six edges, while the other object is a plane at the bottom. Left: Two-dimensional sketch of the contact point concept. The zoom-in shows the contact point which is located in the middle of the overlap region. The contact normal is directed from the middle point toward the closest point on one of the objects involved. Middle: When an object hosts very extruded features, we slightly shrink the surrogate such that the surrogate without ϵ becomes a closer fit around the “unbumped” real geometry. We trade such a surrogate for a bigger ϵ . Empirical evidence suggests that this yields slightly advantageous forces within a multiscale iterative solve. Right: The conservative property of the surrogate triangles states that all fine level geometry (including its ϵ -boundary) must be encompassed by the surrogate’s ϵ . This doesn’t suggest that all surrogate children are included.

particle into account, the total force and total torque determine its acceleration and angular acceleration [4].

This is a simplistic presentation—we ignore, for example, sophisticated interaction functions which distinguish contact points from contact faces—where the discontinuous force that arises at a contact is approximated by a “fade-in” force as two particles approach each other: Over an interval of size ϵ , the force smoothly approximates the target force as the penetration depth $|n(c)|$ increases. The simple physics allow us to focus on the core challenge: how to find contact points efficiently.

Implementation remark 1. As we work with triangulated objects and derive contact points from triangle-triangle comparisons, the algorithm identifies some contact points redundantly: If the closest distance between two objects is spanned by two object vertices x_i and x_j , every triangle combination t_i and t_j , where t_i is adjacent to x_i and t_j to x_j , finds the same contact point and consequently adds it to \mathbb{C} (Algorithm 2.2). The set notation for \mathbb{C} highlights that we do not work with redundant contact points: In the implementation, we run over \mathbb{C} and merge close-by contact points, i.e., points closer than ϵ , into one average point. This filter step prior to any use of the elements from \mathbb{C} is necessary, as we work with floating point arithmetics and an $\epsilon > 0$; i.e., a contact point is not a unique point in space and might temporarily exist multiple times within \mathbb{C} with slightly different coordinates. \square

Implicit time stepping. An implicit Euler for a DEM code has to solve a nontrivial, nonlinear equation system per time step. Nontrivial means that the system’s sparsity pattern depends on the solution of the system itself. It is determined by the contact point search: We obtain entries in the interaction matrix, where the corresponding normal $|n| \leq \epsilon$. Nonlinearity means that the quantities in the interaction matrix (forces) depend on the (guess of the) geometries’ arrangement.

Assumption 1. Our implicit time stepping problem exhibits a contraction property; i.e., Picard iterations can solve the underlying nonlinear equation system.

Algorithm 2.3 High-level pseudocode for an implicit Euler.

```

1: while  $t < T_{\text{terminal}}$  do ▷ We simulate over a time span
2:    $\forall p_i \in \mathbb{P} : \mathbb{T}^{\text{guess}}(p_i) \leftarrow \mathbb{T}(p_i), v^{\text{guess}}(p_i) \leftarrow v(p_i), \omega^{\text{guess}}(p_i) \leftarrow \omega(p_i)$ 
3:   while  $\mathbb{T}^{\text{guess}}, v^{\text{guess}}$  or  $\omega^{\text{guess}}$  change significantly for any  $p$  do
4:      $\forall p_i \in \mathbb{P} : \mathbb{C}(p_i) \leftarrow \emptyset$  ▷ Clear set of collisions for particle  $p_i$ 
5:     for  $p_i, p_j \in \mathbb{P}, p_i \neq p_j$  do ▷ Run over all particle pairs
6:        $\mathbb{C}(p_i) \leftarrow \mathbb{C}(p_i) \cup \text{FINDCONTACTS}(\mathbb{T}^{\text{guess}}(p_i), \mathbb{T}^{\text{guess}}(p_j))$ 
7:        $\mathbb{C}(p_j) \leftarrow \mathbb{C}(p_j) \cup \text{FINDCONTACTS}(\mathbb{T}^{\text{guess}}(p_i), \mathbb{T}^{\text{guess}}(p_j))$ 
8:     end for
9:     for  $p_i \in \mathbb{P}$  do
10:       $(dv, d\omega) \leftarrow \text{CALCFORCES}(\mathbb{C}(p_i))$ 
11:       $(v^{\text{guess}}, \omega^{\text{guess}})(p_i) \leftarrow (v, \omega)(p_i) + \Delta t \cdot (dv, d\omega)$ 
12:       $\mathbb{T}^{\text{guess}}(p_i) \leftarrow \text{UPDATE}(\mathbb{T}(p_i), v^{\text{guess}}(p_i), \omega^{\text{guess}}(p_i), \Delta t)$ 
13:    end for
14:  end while
15:   $\forall p_i \in \mathbb{P} : \mathbb{T}(p_i) \leftarrow \mathbb{T}^{\text{guess}}(p_i), v(p_i) \leftarrow v^{\text{guess}}(p_i), \omega(p_i) \leftarrow \omega^{\text{guess}}(p_i)$ 
16:   $t \leftarrow t + \Delta t$ 
17: end while

```

As our nonlinear system is “well behaved”—we employ reasonably small Δt —we rely on a fixed-point formulation of the implicit time stepping and employ Picard iterations; i.e., we approximate the velocity and angular velocity $(dv, d\omega)(p)(t)$ through a repeated application of the contact detection plus its following force calculation (Algorithm 2.3). We assume that we remain within the Picard iterations’ region of convergence.

Picard avoids the assembly and inversion of an equation system. However, many contact points enter the algorithm temporarily throughout the iterations, which eventually are not identified as contacts. This happens, for example, if we rotate the particles “too far” throughout the iterations. Despite small Δt , we cannot provide an upper bound on the number of Picard iterations required or make assumptions on the (temporarily) identified contact points, i.e., the cost per iteration.

Notation and terminology. Any particle $p \in \mathbb{P}$ has a (closed) volume $V(p)$ which is spanned by its triangular surface $\mathbb{T}(p)$. Since we stick to explicit and implicit single-step, single-shot methods, we omit the parameterization (t) from hereon:

$$p_1 \cap p_2 = \emptyset \quad \forall p_1, p_2 \in \mathbb{P} \text{ with } p_1 \neq p_2,$$

as we have rigid, nonpenetrating objects. Consequently, $\mathbb{T}(p_1) \cap \mathbb{T}(p_2) = \emptyset$. The particles’ triangles do not intersect. Yet, as each particle is surrounded by an ϵ -layer, our particles’ triangles $t \in \mathbb{T}(p)$ unfold into a set of volumetric objects $t^\epsilon \in \mathbb{T}^\epsilon(p)$, and our particles’ volumes are extended $p \subset p^\epsilon$, too. Overlaps between extended volumes do exist and yield contact points

$$(2.1) \quad p_1^\epsilon \cap p_2^\epsilon \neq \emptyset \Rightarrow \text{contact point}$$

which is equivalent to

$$\exists t_1 \in \mathbb{T}(p_1), t_2 \in \mathbb{T}(p_2) : t_1^\epsilon \cap t_2^\epsilon \neq \emptyset \Rightarrow \text{contact point.}$$

A contact point c between two triangles t_i and t_j in (2.1) is located at

$$(2.2) \quad x(c) = \frac{1}{2}(t_1(a_1, b_1) + t_2(a_2, b_2)) \text{ with } \arg \min_{a_1, b_1, a_2, b_2 \in [0,1]} \frac{1}{2} \|t_1(a_1, b_1) - t_2(a_2, b_2)\|^2$$

if t_1 and t_2 are surrounded by the same ϵ -halo. For different ϵ ’s, the weights for $x(c)$ have to be adapted accordingly. a_1, b_2 are Barycentric coordinates within t_1 ;

i.e., $t_1(a_1, b_1)$ returns a position within the triangle. a_2 and b_2 are the counterparts for t_2 . If $x(c)$ is the position of the contact point c , the corresponding normal $n(c) = t_1(a_1, b_1) - x(c)$ or $n(c) = t_2(a_2, b_2) - x(c)$.

3. Iterative contact detection via distance minimization. To find the closest point between two triangles is a classic computational geometry problem [10]. We rely on three different algorithms to solve it.

Direct distance calculation (comparison based). A comparison-based identification of contact points consists of two steps. First, we compute the distance from each vertex of the two triangles to the closest point on the other triangle as well as the distance between each pair of edges between the two triangles [10]. This yields six point-to-triangle distance tests and nine edge-to-edge distance tests. In a second step, we select the minimum distance out of the 15 combinations. This brute force calculation yields an exact solution—agnostic of truncation errors—yet requires up to 30+14 comparisons (if statements) per triangle pair, which we tune via masking, blending, and early termination [20]. Dealing with the cases where intersections between triangles are possible requires an additional six edge-to-plane distance tests, where intersections outside the area of the triangle are discounted.

Iterative distance calculation. As an alternative to a comparison-based approach, we replace the geometric checks with a functional, where we minimize the distance between the two planes spanned by the triangles but add the admissibility conditions over the Barycentric coordinates as Lagrangian parameters [14, 15]. In line with (2.2), let $a, b \in [0, 1]$ describe any point on their respective triangle:

$$\begin{aligned}
 \arg \min_{a_1, b_1, a_2, b_2} J(a_1, b_1, a_2, b_2) := & \arg \min_{a_1, b_1, a_2, b_2} \underbrace{\frac{1}{2} \|t_i(a_1, b_1) - t_j(a_2, b_2)\|^2}_{=: \hat{J}(a_1, b_1, a_2, b_2)} + \alpha_{\text{iterative}} \left(\right. \\
 & \max(0, a_1 - 1) + \min(-a_1, 0) + \max(0, b_1 - 1) + \\
 & \max(-b_1, 0) + \max(0, a_1 + b_1 - 1) + \\
 & \max(0, a_2 - 1) + \min(-a_2, 0) + \max(0, b_2 - 1) + \\
 & \left. \max(-b_2, 0) + \max(0, a_2 + b_2 - 1) \right).
 \end{aligned}
 \tag{3.1}$$

This is a weak formulation of the challenge since any $\alpha_{\text{iterative}} < \infty$ allows the closest distance line between two triangles to be rooted slightly outside the very triangles.

The minimization problem can be solved via Newton iterations. However, the arising Hessian becomes difficult to invert or noninvertible for close-to-parallel or actually parallel triangles. We therefore regularize it by adding a diagonal matrix. After that, we approximate the regularized Hessian and update the minimization and constraints alternatingly:

$$\begin{aligned}
 (a_1, b_1, a_2, b_2)^{(n+0.5)} &= (a_1, b_1, a_2, b_2)^{(n)} - \text{diag}^{-1} \left(\hat{H}(a_1, b_1, a_2, b_2)^{(n)} \right. \\
 & \quad \left. + \alpha_{\text{regularizer}} \text{id} \right) \nabla_{a_1, b_1, a_2, b_2} \hat{J}(a_1, b_1, a_2, b_2)^{(n)} \\
 (a_1, b_1, a_2, b_2)^{(n+1)} &= (a_1, b_1, a_2, b_2)^{(n+0.5)} - \text{diag}^{-1} \left(\tilde{H}(a_1, b_1, a_2, b_2)^{(n+0.5)} \right. \\
 & \quad \left. + \alpha_{\text{regularizer}} \text{id} \right) \nabla_{a_1, b_1, a_2, b_2} \tilde{J}(a_1, b_1, a_2, b_2)^{(n+0.5)}
 \end{aligned}
 \tag{3.2}$$

$\alpha_{\text{regularizer}} > 0$ is small, while (n) is the iteration index. The \hat{J} and its Hessian correspond to the quadratic term of the actual functional in (3.1) also known as

Algorithm 3.1 A hybrid, batched reformulation of the iterative distance calculation. The iterative sweep over a whole batch of triangles is stripped of a dynamic stopping criterion (top) and therefore yields three types of results per triangle pair: x holds a coordinate if the Barycentric coordinates yield a contact point or \perp if they yield no contact point or \odot if the triangle combination has to be postprocessed with the comparison-based algorithm (bottom).

```

1: for  $t_i \in \mathbb{T}(p_i), t_j \in \mathbb{T}(p_j), t_i \neq t_j$  do
2:    $(a_1, b_1, a_2, b_2, n)(t_i, t_j) \leftarrow 0$  ▷ Variables per triangle pair allow us to vectorize
3:    $J_{\text{old}}(t_i, t_j) \leftarrow \infty$  ▷ Over triangles in  $\mathbb{T}$ 
4:    $J(t_i, t_j) = J(a_1(t_i, t_j), b_1(t_i, t_j), a_2(t_i, t_j), b_2(t_i, t_j))$  ▷ Functional from (3.1)
5:   while  $n \leq N_{\text{iterative}}$  do ▷ Fixed iteration count instead of  $|J - J_{\text{old}}| > C\epsilon$ 
6:      $J_{\text{old}}(t_i, t_j) \leftarrow J(t_i, t_j)$ 
7:     update  $a_1(t), b_1(t), a_2(t), b_2(t)$  ▷ Modified gradient descent from (3.2)
8:      $J(t_i, t_j) = J(a_1(t_i, t_j), b_1(t_i, t_j), a_2(t_i, t_j), b_2(t_i, t_j))$ 
9:      $n \leftarrow n + 1$ 
10:  end while
11:   $x(t_i, t_j) \leftarrow \begin{cases} \frac{1}{2}(t_i(a_1, b_1) - t_j(a_2, b_2))(t_i, t_j) & \text{if } |J(t_i, t_j) - J_{\text{old}}(t_i, t_j)| \leq C\epsilon \wedge \hat{J}(t_i, t_j) \leq 2\epsilon^2 \\ \perp & \text{if } |J(t_i, t_j) - J_{\text{old}}(t_i, t_j)| \leq C\epsilon \wedge \hat{J}(t_i, t_j) > 2\epsilon^2 \\ \odot & \text{otherwise} \end{cases}$ 
12: end for
13:
14: for  $t_i \in \mathbb{T}(p_i), t_j \in \mathbb{T}(p_j), t_i \neq t_j$  do
15:   if  $x(t_i, t_j) = \odot$  then
16:      $\hat{n}(t_i, t_j) \leftarrow$  shortest distance vector between  $t_i$  and  $t_j$  ▷ Use comparison-based algorithm
17:      $\hat{x}(t_i, t_j) \leftarrow$  is central point on line identified by  $\hat{n}(t_i, t_j)$ 
18:      $x(t_i, t_j) \leftarrow \begin{cases} \hat{x}(t_i, t_j) & \text{if } |\hat{n}(t_i, t_j)| \leq 2\epsilon \\ \perp & \text{otherwise} \end{cases}$  ▷ Eliminate  $\odot$  entries in result
19:   end if
20: end for

```

(2.2). \tilde{J} and its Hessian cover the remaining penalty terms, i.e., $J - \hat{J}$. However, we omit the Dirac terms in there; i.e., we explicitly drop terms that enter the formulae for $a_1 = 0, a_1 = 1, b_1 = 0, \dots$. Our solver iterates back and forth between the \tilde{J} -minimization and a fulfillment of the constraints. This modified Newton becomes a gradient descent, where the step size is adaptively chosen by analyzing an approximation of the inverse to the Hessian.

Hybrid distance calculation. If two subsequent iterates $|J(a_1, b_1, a_2, b_2)^{(n+1)} - J(a_1, b_1, a_2, b_2)^{(n)}| \leq C\epsilon$, we have found an actual minimum over functional (3.1) and can terminate the minimization. In this case, we assume that a_1, b_1, a_2, b_2 identify the minimum distance. Without further analysis, it is impossible to make a statement on the upper bound on n . Our hybrid algorithm therefore eliminates the termination criterion and imposes $n \leq N_{\text{iterative}}$. Consequently, it labels a distance calculation as invalid if $|J(a_1, b_1, a_2, b_2)^{(N_{\text{iterative}})} - J(a_1, b_1, a_2, b_2)^{(N_{\text{iterative}}-1)}| > C\epsilon$. The iterative code's result realizes a three-valued logic: "found a contact point," "there is no contact," or "has not terminated" (Algorithm 3.1).

Our hybrid algorithm invokes the modified iterative algorithm. If the result equals "not terminated" (\odot), the hybrid algorithm falls back to the comparison-based distance calculation. It is thus not really a third algorithm to find a contact point but a combination of the iterative scheme with the comparison-based approach serving as a posteriori limiter.

Implementation remark 2. Triangle meshes are typically held as graphs over vertex sets. We flatten this graph prior to the contact detection: A sequence of $|\mathbb{T}|$ triangles is converted into a sequence of $3 \cdot 3 \cdot |\mathbb{T}|$ floating point values; i.e., each

triangle is represented by the coordinates (three components) of its three vertices. Such a flat data structure can be generated once prior to the first contact detection. All topology is lost in this representation, and vertex data is replicated—it is a triangle soup—but the data is well suited to be streamed without indirect memory lookups. Per update of the position and rotation, every coordinate is subject to an affine transformation. \square

4. A multiresolution model. The cost to compare two particles p_i and p_j is determined by their triangle counts $|\mathbb{T}(p_i)|$ and $|\mathbb{T}(p_j)|$. To reduce this cardinality, we construct geometric cascades of triangle models per particle—the *surrogate models* (Figure 1.1)—and plug representations from this cascade into Algorithm 2.2.

DEFINITION 4.1. *A surrogate model $\mathbb{T}_k(p)$, $k \geq 1$, is a triangle-based geometric approximation of a particle described by $\mathbb{T}(p)$. A sequence of surrogate models $\{\mathbb{T}_1(p), \mathbb{T}_2(p), \mathbb{T}_3(p), \dots\}$ for p with its volumetric extensions $\mathbb{T}_k^\epsilon(p)$ hosts efficient (Definition 4.2), conservative (Definition 4.3), and weakly connected (Definition 4.5) abstractions of $\mathbb{T}(p)$.*

Surrogate models are different representations of a particle. The term “cascade” highlights that each particle is assigned a whole sequence of representations. These representations can step in for our real geometry and are a special class of bounding volume techniques [10]. To simplify our notation, $\mathbb{T}_0(p) := \mathbb{T}(p)$; i.e., the $k = 0$ surrogate model is the geometric object itself. The bigger k is, the coarser, i.e., more abstract, the surrogate. We emphasize that the ϵ is a generic symbol; i.e., each surrogate model hosts its own bespoke ϵ -environment.

DEFINITION 4.2. *A surrogate model $\mathbb{T}_{k_i}^\epsilon(p_i)$ is efficient if, for any other model $\mathbb{T}_{k_j}^\epsilon(p_j)$, finding all contact points between $\mathbb{T}_{k_i}^\epsilon(p_i)$ and $\mathbb{T}_{k_j}^\epsilon(p_j)$ is cheaper than finding all contact points between $\mathbb{T}_{\hat{k}_i}^\epsilon(p_i)$ and $\mathbb{T}_{\hat{k}_j}^\epsilon(p_j) \forall 0 \leq \hat{k}_i < k_i$.*

We assume an almost homogeneous cost per triangle-to-triangle comparison—an assumption that is shaky for the hybrid comparison and subject to vectorization efficiency and memory management effects. Hence, the triangle count of surrogate models decreases with increasing k , i.e., $|\mathbb{T}_k| \ll |\mathbb{T}_{k+1}|$.

DEFINITION 4.3. *A surrogate model $\mathbb{T}_{k_i}(p_i)$ inducing $\mathbb{T}_{k_i}^\epsilon(p_i)$ is conservative if*

$$\forall p_i, p_j, k_i, k_j : \mathbb{T}_{k_i}^\epsilon(p_i) \cap \mathbb{T}_{k_j}^\epsilon(p_j) = \emptyset \Rightarrow \mathbb{T}^\epsilon(p_i) \cap \mathbb{T}^\epsilon(p_j) = \emptyset.$$

Conservative means that any two surrogates of two particles are in contact (overlap) if the two particles are in contact. Yet, this does not have to hold the other way round: If their surrogates are in contact, there might still be gaps between the particles; i.e., there might be no contact point.

COROLLARY 4.4. *Let a surrogate model hierarchy $\mathbb{T}_k(p)$, $k \geq 1$, be monotonous if*

$$\forall 1 \leq \hat{k} < k : \forall t_{\hat{k}}^\epsilon \in \mathbb{T}_{\hat{k}}^\epsilon(p) : t_{\hat{k}}^{\epsilon_{\hat{k}}} \subseteq \bigcup_{t \in \mathbb{T}_k^\epsilon(p)} t.$$

Our surrogate hierarchies do not have to be monotonous.

A monotonous cascade of triangles plus ϵ -environments would grow in space as we move up the hierarchy of models. Therefore, we do not impose it, even though monotonicity would imply conservativeness “for free.” Empirical evidence suggests that abandoning monotonicity allows us to work with significantly tighter ϵ -choices

per level. Yet, it also implies that we generally cannot show algorithm correctness through plain induction.

Classic level-of-detail algorithms require a coarsened representation of a triangulated model to preserve certain properties such as connected triangle surfaces or the preservation of certain features such as sharp edges. Our surrogate models, however, are to be used as temporary replacements within our calculations. We thus can ask for weak representations of the geometries.

DEFINITION 4.5. *The triangles of a particle have to span a connected surface. For surrogate models, there is no such constraint. Therefore, the surrogate models can be weakly connected: Their triangles can be disjoint with gaps in-between, or they can intersect each other. A surrogate's ϵ -environment is connected, however, and it covers (overlaps) all connectivity of the original model.*

The connectivity addendum in Definition 4.5 motivates the term weakly connected. It clarifies that we—despite the disjoint, anarchic configuration of a surrogate triangle set (Figure 1.1)—cannot miss out on some geometry extrema such as sharp edges due to tests with surrogates: If there is no contact between two surrogates, there is also no contact between their real discretizations.

Definition 4.3 implies that we can use surrogate models as guards and run through them for coarse to fine: If there are no contact points between two surrogate models, there can be no contact points for the more detailed models. We can stop searching for contact points immediately. It does not, however, hold the other way round. Definition 4.2 implies immediately that the number of triangles that we examine in such an iterative approach is monotonously growing. Definition 4.5 gives us the freedom to construct such triangle hierarchies, as it strips us from many geometric constraints.

DEFINITION 4.6. *Let \mathcal{T} be a directed acyclic graph, where each node represents a set of triangles. The level ℓ of a node is its distance (edge count) from the root node in \mathcal{T} . The resulting graph is a surrogate tree if and only if*

1. *the root node hosts the coarsest surrogate model $\mathbb{T}_{k_{\max}}$;*
2. *the union over all leaf sets yields the particle triangulation \mathbb{T}_0 ;*
3. *any triangle is a surrogate for the union over its children's triangle sets.*

With Definition 4.6, the union over all sets with the same level yields the surrogate model $\mathbb{T}_{k_{\max}-\ell}$. We use $N_{\text{surrogate}}$ to denote the number of children of a surrogate triangle. $N_{\text{surrogate}}$ does not have to be uniform over the tree, i.e., is a generic symbol. If we have a surrogate model, take the nodes within the tree which hold its triangles, and replace the triangles with those triangles stored within children nodes, we obtain the next finer surrogate. A surrogate tree is a generalization of the concept of a cascade of surrogates: The tree formalism allows us also to construct different, hybrid-level surrogate models if we only replace some triangles of a model with their children.

Implementation remark 3. There are multiple ways to construct surrogate trees. We construct our trees through a recursive algorithm. It starts from the triangle set $\mathbb{T} = \mathbb{T}_0$ of the particle and splits this set into $N_{\text{surrogate}}$ subsets of roughly the same size hosting close-by triangles. Per subset, we construct one surrogate and thus obtain a tree of depth one where the root node hosts $|N_{\text{surrogate}}|$ triangles. As long as a node within the tree hosts more triangles than a prescribed threshold, we apply the splitting recursively and thus disentangle the triangle sets further and further: Existing tree levels are pushed down or sieved through the tree hierarchy (Appendix B). A follow-up bottom-up traversal of the tree constructs well-suited surrogate triangles with

appropriate ϵ -choices. As we only require weakly connected surrogates, the steps within this bottom-up traversal are independent of operations on sibling nodes within the tree and can be mapped onto local minimization problems (Appendix A). \square

5. Multiresolution contact detection. With our surrogate tree definition, we are in the position to propose a multiscale algorithm for an explicit Euler which utilizes the tree as an early stopping criterion for the search for contacts. The observations for the explicit time stepping then allow us to derive two implicit time stepping algorithms that exploit the multiscale nature of the geometry to terminate searches early and to supplement the underlying iterative algorithm with educated guesses.

5.1. Explicit Euler. Our explicit Euler exploits the multiscale hierarchy by looping over the resolutions held in \mathcal{T} top-down. The tree is unfolded depth-first, and we implement an early stopping criterion: If a surrogate triangle and a triangle set from another particle do not collide, the children of the surrogate triangle in our triangle tree cannot collide either. The depth-first traversal along this branch of the tree thus can terminate early. The surrogate model unfolds adaptively.

The concept defines a marker (Algorithm 5.1): Let \mathbb{A} identify a set of active nodes from \mathcal{T} . The union of all sets labeled by \mathbb{A} yields all triangles from a particle that participate in collision checks. At the begin of a particle-to-particle comparison, only the particles' roots are active. From there, we work our way down into finer and finer geometric representations as long as the surrogate models suggest that there might be some contacts until we eventually identify real contact points stemming from the finest mesh.

Algorithm 5.1 Multiresolution contact detection within explicit time stepping. It compares two particles p_i and p_j given by their surrogate trees $\mathcal{T}(p_i)$ and $\mathcal{T}(p_j)$ with each other.

```

1:  $\mathbb{A}_i \leftarrow \text{root}(\mathcal{T}(p_i)), \mathbb{A}_j \leftarrow \text{root}(\mathcal{T}(p_j))$  ▷ Set of active triangles to check
2: while  $\mathbb{A}_i \neq \emptyset \vee \mathbb{A}_j \neq \emptyset$  do
3:    $\mathbb{A}_{i,\text{new}} \leftarrow \emptyset, \mathbb{A}_{j,\text{new}} \leftarrow \emptyset$ 
4:   for  $t_i \in \mathbb{A}_i, t_j \in \mathbb{A}_j$  do
5:      $c \leftarrow \text{CONTACT\_ITERATIVE}(t_i, t_j)$  ▷ Use context-specific  $\epsilon$  depending on  $t_i, t_j$ 
6:     if  $c = \odot \wedge t_i \in \mathbb{T}_0^\epsilon(p_i) \wedge t_j \in \mathbb{T}_0^\epsilon(p_j)$  then ▷ Not converged on nonsurrogate triangles
7:        $c \leftarrow \text{CONTACT\_COMPARISON}(t_i, t_j)$  ▷ Use comparison-based algorithm this time
8:     end if
9:     if  $c \neq \perp$  then
10:      if  $t_i \in \mathbb{T}_0^\epsilon(p_i) \wedge t_j \in \mathbb{T}_0^\epsilon(p_j)$  then ▷ No surrogate triangles
11:         $\mathbb{C}(p_i) \leftarrow \mathbb{C}(p_i) \cup \{c\}, \mathbb{C}(p_j) \leftarrow \mathbb{C}(p_j) \cup \{c\}$  ▷ I.e., proper contact point
12:      else ▷ Unfold
13:        if  $t_i \in \mathbb{T}_0^\epsilon(p_i)$  then
14:           $\mathbb{A}_{i,\text{new}} \leftarrow \mathbb{A}_{i,\text{new}} \cup \{t_i\}$ 
15:        else
16:           $\mathbb{A}_{i,\text{new}} \leftarrow \mathbb{A}_{i,\text{new}} \cup \{\hat{t} : \hat{t} \sqsubseteq_{\text{child}} t_i\}$ 
17:        end if
18:        if  $t_j \in \mathbb{T}_0^\epsilon(p_j)$  then
19:           $\mathbb{A}_{j,\text{new}} \leftarrow \mathbb{A}_{j,\text{new}} \cup \{t_j\}$ 
20:        else
21:           $\mathbb{A}_{j,\text{new}} \leftarrow \mathbb{A}_{j,\text{new}} \cup \{\hat{t} : \hat{t} \sqsubseteq_{\text{child}} t_j\}$ 
22:        end if
23:      end if
24:    end if
25:  end for
26:   $\mathbb{A}_i \leftarrow \mathbb{A}_{i,\text{new}}, \mathbb{A}_j \leftarrow \mathbb{A}_{j,\text{new}}$ 
27: end while

```

LEMMA 5.1. *The hierarchical algorithm yields exactly the same outcome as our baseline code over sets $\mathbb{T}^\epsilon(p_i)$ and $\mathbb{T}^\epsilon(p_j)$. Algorithm 5.1 is correct.*

Proof. The argument relies on three properties:

1. If a contact point is identified for a surrogate triangle, it is not added to the set of contact points. Therefore, a given active set never identifies artificial/too many contact points.
2. A contact point is added if it stems from the comparison of two triangles from the fine grid tessellations which are in the active sets.
3. Let two triangles t_i^ϵ and t_j^ϵ yield a contact point. As surrogates are conservative, they belong to nodes (triangle sets) $\mathbb{T}(p_i)$ and $\mathbb{T}(p_j)$ with $\mathbb{T}(p_i) \sqsubseteq_{\text{child}} \hat{t}_i$ and $\mathbb{T}(p_j) \sqsubseteq_{\text{child}} \hat{t}_j$ in $\mathcal{T}(p_i)$ or $\mathcal{T}(p_j)$, respectively. These surrogates fulfill

$$t_i^\epsilon \cap t_j^\epsilon \neq \emptyset \Rightarrow \hat{t}_i^\epsilon \cap \hat{t}_j^\epsilon \neq \emptyset.$$

and therefore are replaced in the active set by their children in Algorithm 2.1 before the respective algorithm terminates.

The correctness of the algorithm follows from bottom-up induction over the levels of \mathcal{T} : The property holds directly for the finest surrogate levels \mathbb{T}_1 of the tree. Any violation thus has to arise from $\mathbb{T}_k, k \geq 2$, in p_i or p_j . We apply the arguments recursively. \square

We have two triangle-to-triangle comparison strategies on the table (hybrid and comparison based) which are robust, i.e., always yield the correct solution. If we employ the comparison-based approach only, the $c = \odot$ condition never holds, and the corresponding branch is never executed. Otherwise, our algorithmic blueprint implements the hybrid's fallback as it automatically reevaluates the contact search for $c = \odot$. However, it is indeed sufficient to rerun this a posteriori contact search if and only if both triangles stem from the finest triangle discretization.

COROLLARY 5.2. *On the surrogate levels within the tree, it is sufficient to use the (efficient) iterative collision detection algorithm (Algorithm 3.1, bottom), without falling back to the comparison-based variant.*

Proof. Let $\mathbb{T}^\epsilon(p_i) \cap \mathbb{T}^\epsilon(p_j) \neq \emptyset$, i.e., two particles collide. We assume the lemma is wrong; i.e., the tree unfolding terminates prematurely. This assumption formally means

$$\exists t_i \in \mathcal{T}(p_i), t_j \in \mathcal{T}(p_j) : r(p_i, p_j) = \perp,$$

with

$$\exists t_{0,i} \in \mathbb{T}(p_i), t_{0,j} \in \mathbb{T}(p_j) : t_{0,i} \sqsubseteq_{\text{child}} \dots \sqsubseteq_{\text{child}} t_i \wedge t_{0,j} \sqsubseteq_{\text{child}} \dots \sqsubseteq_{\text{child}} t_j \wedge t_i^\epsilon \cap t_j^\epsilon \neq \emptyset.$$

This assumption is a direct violation of the definition of a surrogate model which has to be conservative. \square

5.2. Implicit Euler with multiresolution acceleration. Picard iterations can exploit the multiscale hierarchy by looping over the hierarchy levels top-down: Per iteration of Algorithm 2.3, we have to identify all contact points for the current particle configuration. This search for contact points is the same search as we use it in an explicit Euler. If we replace the contact detection within the inner loop with our multiscale contact detection from section 5.1, we obtain an implicit Euler where the surrogate concept is used *within the Picard loop* as multiresolution acceleration. The surrogate concept enters the algorithm's implementation as a black box.

COROLLARY 5.3. *An implicit Euler using surrogates within the Picard loop body to speed up the search for contact points yields the same output as a flat implicit code with the same number of Picard iterations.*

Proof. This is a direct consequence of Lemma 5.1 and implies the algorithm's correctness. \square

Though we end up with exactly the same number of Picard iterations, the individual iterations are accelerated by the multiresolution technique.

For a localized contact between two particles, the surrogate tree is unfolded along a single or a few branches of the tree. If the nodes within the tree hold roughly the same number of triangles, the number of triangles to be compared grows linearly with the number of Picard steps. We benefit both from a zapping through the resolution levels and the localization of contacts, i.e., the fact that two particles usually collide only in a small area compared to the overall geometric object.

5.3. Implicit multiresolution Euler. A more bespoke implicit multiresolution algorithm arises from ideas inspired by multilevel nonlinear equation system solvers. The multiscale Algorithm 5.2 consists of two nested while loops—the outer loop stems from the Picard iterations, the inner loop realizes the tree unfolding—which we can permute. We obtain an algorithm that runs top-down via the active sets through the surrogate hierarchies and unfolds the trees step by step. Per unfolding step, it uses the Picard loop to converge on the selected hierarchy level. The rationale behind such a permutation is the observation that the efficiency of a nonlinear equation system solver hinges on the availability of a good initial guess. Surrogate resolution levels might be well suited to deliver a good initial guess of what \mathbb{T} looks like in the next time step. This train of thought is similar to the extension of multigrid into full multigrid. The same multigrid analogy suggests that we do not have to converge on a surrogate level, as the level supplements only a guess anyway. In the extreme case, it is thus sufficient to run one Picard iteration per unfolding step only.

Our advanced variant of the implicit Euler is an *outer-loop multiresolution Picard scheme*. Let the Picard loop start from the coarsest surrogate representation per particle (Algorithm 5.2). These representations form our initial active sets. After the Picard step, any surrogate triangle for which the hybrid algorithm has not terminated or for which we identified a contact point is replaced by its next finer representation. In the tradition of value-range analysis, we widen the active set [3]. The Picard loop terminates if the plain algorithm's termination criteria hold; i.e., the outcome of two subsequent iterations does not change dramatically anymore, and no surrogate tree node has unfolded anymore throughout the previous iterate.

The algorithm is completed by a cleanup which removes “obsolete” triangles from the active set: If all children of a surrogate triangle do not certainly contribute a contact point anymore, they are replaced with their parent surrogate triangle. We narrow the active set.

Implementation remark 4. Different to the explicit scheme, we maintain an active set $\mathbb{A}(p_i, p_j)$ per particle-particle combination p_i, p_j : A particle p_i can exhibit a very coarse surrogate representation against one particle while using a very detailed mesh when we compare it to another one. While the number of particle-particle combinations is potentially huge, it is small in practice, as particles are rigid and thus cannot cluster arbitrarily dense. \square

Our genuine multiscale formulation stresses the convergence assumptions: While Assumption 1 guarantees the convergence of the Picard iterations on the finest level,

Algorithm 5.2 Implicit time stepping algorithm, where the Picard and multiresolution loops are intermingled.

```

1:  $\forall p_j \neq p_i \in \mathbb{P} : \mathbb{A}(p_i, p_j) \leftarrow \text{root}(\mathcal{T}(p_i))$   $\triangleright$  Active sets are now parameterized over interactions
2: while  $\mathbb{T}^{\text{guess}}(p_i), v^{\text{guess}}(p_i), \omega^{\text{guess}}(p_i)$  or any  $\mathbb{A}$  change significantly for any  $p_i$  do
3:    $\forall p_j \neq p_i \in \mathbb{P} : \mathbb{A}_{\text{new}}(p_i, p_j) \leftarrow \emptyset, \mathbb{A}_{\text{new}}(p_j, p_i) \leftarrow \emptyset, \mathbb{C}(p_i, p_j) = \emptyset, \mathbb{C}(p_j, p_i) = \emptyset$ 
4:   for  $p_j \neq p_i \in \mathbb{P}$  do
5:     for  $t_i \in \mathbb{A}(p_i, p_j), t_j \in \mathbb{A}(p_j, p_i)$  do
6:        $c \leftarrow \text{CONTACTITERATIVE}(t_i, t_j)$   $\triangleright$  Use context-specific  $\epsilon$  depending on  $t_i, t_j$ 
7:       if  $c = \odot \wedge t_i \in \mathbb{T}_0^\epsilon(p_i) \wedge t_j \in \mathbb{T}_h^\epsilon(p_j)$  then  $\triangleright$  Not converged on nonsurrogate triangles
8:          $c \leftarrow \text{CONTACTCOMPARISONBASED}(t_i, t_j)$   $\triangleright$  Use comparison-based algorithm this time
9:       end if
10:      if  $c \neq \perp \wedge c \neq \odot$  then  $\triangleright$  Implicit guess
11:         $\mathbb{C}(p_i) \leftarrow \mathbb{C}(p_i) \cup \{c\}, \mathbb{C}(p_j) \leftarrow \mathbb{C}(p_j) \cup \{c\}$ 
12:      end if
13:      if  $c = \perp$  then  $\triangleright$  Add only parents
14:         $\mathbb{A}_{\text{new}}(p_i, p_j) \leftarrow \mathbb{A}_{\text{new}}(p_i, p_j) \cup \{\hat{t} : t_i \sqsubseteq_{\text{child}} \hat{t}\}$ 
15:         $\mathbb{A}_{\text{new}}(p_j, p_i) \leftarrow \mathbb{A}_{\text{new}}(p_j, p_i) \cup \{\hat{t} : t_j \sqsubseteq_{\text{child}} \hat{t}\}$ 
16:      else  $\triangleright$  Widen active sets
17:        ...  $\triangleright$  Compare to Algorithm 5.1
18:      end if
19:    end for
20:  end for
21:   $\forall p_j \neq p_i \in \mathbb{P} : \mathbb{A}(p_i, p_j) \leftarrow \mathbb{A}_{\text{new}}(p_i, p_j), \mathbb{A}(p_j, p_i) \leftarrow \mathbb{A}_{\text{new}}(p_j, p_i)$ 
22:  for  $p_i \in \mathbb{P}$  do
23:     $(dv, d\omega) \leftarrow \text{CALCFORCES}(\mathbb{C}(p_i))$   $\triangleright$  Remove redundant contact points first
24:     $(v^{\text{guess}}, \omega^{\text{guess}})(p_i) \leftarrow (v, \omega)(p_i) + \Delta t \cdot (dv, d\omega)$   $\triangleright$  Additional damping might be required
25:     $\mathbb{T}^{\text{guess}}(p_i) \leftarrow \text{UPDATE}(\mathbb{T}(p_i), v^{\text{guess}}(p_i), \omega^{\text{guess}}(p_i), \Delta t)$ 
26:  end for
27:  for  $p_j \neq p_i \in \mathbb{P}$  do  $\triangleright$  Clean up; i.e., add siblings
28:    ...  $\triangleright$  Obsolete if surrogate nodes host only one triangle
29:     $\forall t, \hat{t} \in \mathbb{A}(p_i, p_j)$  with  $t \sqsubseteq_{\text{child}} \hat{t} : \mathbb{A}(p_i, p_j) \leftarrow \mathbb{A}(p_i, p_j) \cup \{t' \in \mathcal{T}(p_i) : t' \sqsubseteq_{\text{child}} \hat{t}\}$ 
30:     $\forall t, \hat{t} \in \mathbb{A}(p_j, p_i)$  with  $t \sqsubseteq_{\text{child}} \hat{t} : \mathbb{A}(p_j, p_i) \leftarrow \mathbb{A}(p_j, p_i) \cup \{t' \in \mathcal{T}(p_j) : t' \sqsubseteq_{\text{child}} \hat{t}\}$ 
31:    ...  $\triangleright$  And remove "redundant" parents
32:     $\forall t \in \mathbb{A}(p_i, p_j) : \mathbb{A}(p_i, p_j) \leftarrow \mathbb{A}(p_i, p_j) \setminus \{\hat{t} \in \mathcal{T}(p_i) : t \sqsubseteq_{\text{child}} \hat{t}\}$ 
33:     $\forall t \in \mathbb{A}(p_j, p_i) : \mathbb{A}(p_j, p_i) \leftarrow \mathbb{A}(p_j, p_i) \setminus \{\hat{t} \in \mathcal{T}(p_j) : t \sqsubseteq_{\text{child}} \hat{t}\}$ 
34:  end for
35: end while
36:  $\forall p_i \in \mathbb{P} : \mathbb{T}(p_i) \leftarrow \mathbb{T}^{\text{guess}}(p_i), v(p_i) \leftarrow v^{\text{guess}}(p_i), \omega(p_i) \leftarrow \omega^{\text{guess}}(p_i)$ 
37:  $t \leftarrow t + \Delta t$ 

```

our multiresolution approach may push the solution into the wrong direction via the surrogate levels and thus make the initial guess on the next finer level leave the single level's convergence domain.

Assumption 2. We assume that a Picard iteration on any level of the surrogate trees yields a new solution on the same or a finer resolution which preserves the Picard iteration's contraction property.

LEMMA 5.4. *If Assumption 2 holds and if Algorithm 5.2 terminates, it delivers the correct solution.*

Proof. We have to study two cases over the active sets $\mathbb{A}(p_i, p_j)$ and $\mathbb{A}(p_j, p_i)$. First, assume that $(t_{\mathbb{A}(p_i, p_j)}, t_{\mathbb{A}(p_j, p_i)}) \in \mathbb{A}(p_i, p_j) \times \mathbb{A}(p_j, p_i)$ yields an invalid contact point, i.e., a contact point that does not exist in $\mathbb{T}_0(p_i)$ compared to $\mathbb{T}_0(p_j)$. One of the triangles has to be a surrogate triangle. They are replaced by their children, and the algorithm has not terminated. Instead, we approach the solution further.

In the other case, assume that the algorithm has terminated yet misses a triangle pair $(t(p_i), t(p_j)) \notin \mathbb{T}(p_i) \times \mathbb{T}(p_j)$ which contributes a contact point in the plain model. Due to Definition 4.3 over conservative surrogates,

$$(5.1) \quad \forall t(p_i) \in \mathbb{T}_0(p_j), \exists \hat{t}_{\mathbb{A}(p_i, p_j)} \in \mathbb{A}(p_i) : t(p_i) \sqsubseteq_{\text{child}} \dots \sqsubseteq_{\text{child}} \hat{t}_{\mathbb{A}(p_i, p_j)}$$

such that $\hat{t}_{\mathbb{A}(p_i, p_j)}$ yields a contact point. This point is eventually removed as the corresponding $\hat{t}_{\mathbb{A}(p_i, p_j)}$ is replaced by its children. The algorithm has not terminated yet. \square

COROLLARY 5.5. *The removal of triangles from the active set can cause Algorithm 5.2 to violate Assumption 2.*

Proof. If no triangles are ever removed from the active set, the proof of Lemma 5.4 trivially demonstrates that the algorithm terminates always, as the surrogate tree is of finite depth and width. Even if we overshoot with the Picard iterations, i.e., if we violate the contraction property, we will, in the worst case, get $\mathbb{A}(p_i, p_j) = \mathbb{T}(p_i)$. From hereon, the algorithm converges.

If we, however, remove triangles, it is easy to see that we cannot guarantee that we do not introduce cycles or even amplify oscillations. The contraction property is violated. \square

Implementation remark 5. In practice, Corollary 5.5 implies that we, on the one hand, have to damp the Picard iterations. We artificially reduce the force contributions from coarse surrogate levels to avoid oscillations. On the other hand, we work with a memory set $\mathbb{U}(p_i, p_j)$ in which we hold references to triangles which have been removed from the active set. Once they are readded, we veto any subsequent removal from hereon and, hence, the activation of such triangles' surrogates. \square

5.4. Implementation. There are two reasons why our multiresolution algorithms are expected to yield better performance than a straightforward textbook implementation: First and foremost, we expect the number of triangle-to-triangle comparisons to go down compared to a flat, single level approach. The multiscale algorithm iteratively narrows down the region of a particle where contacts may arise from. These savings on the finer geometric resolutions compensate for additional checks with surrogate triangles. However, any cost amortization has to be studied carefully—in particular for the implicit, nonlinear case, where trees unfold and collapse again—and it hinges upon an efficient realization: In this context, we expect the streaming, comparison-free variants of our algorithm to benefit from vector architectures.

The multiresolution representation of an object can be computed at simulation startup as a preprocessing step. Though we keep the multiresolution hierarchy when particles move and rotate, the flattening of the active sets of triangles from $\mathcal{T}(p)$ into a sequence of coordinates is done on-the-fly, and the flattened data is not held persistently.

On the one hand, this ensures that the memory overhead remains under control. On the other hand, it pays tribute to the fact that the active set changes permanently. To remain fast despite permanently changing active sets, we pick $N_{\text{surrogate}}$ such that the finest nodes within $|\mathbb{T}|$ hold triangle sequences for which streaming instructions such as advanced vector extensions (AVX) already pay off. The tree clusters \mathbb{T} into segments that fit to the architecture, and Algorithm 3.1 hence does not process all triangles from \mathbb{T} in one batch. Instead it runs over subchunks of batches.

We apply this argument recursively and make each nonleaf node within \mathcal{T} hold a set of triangles, too. We make the nodes in the surrogate tree host many triangles and the tree overall shallow such that the per node data cardinality again ensures

that we benefit from vector units. This “tweak $N_{\text{surrogate}}$ ” idea, however, does not fit perfectly to multiscale algorithms where the coarser tree levels typically do not occupy a complete vector length. It would be a coincidence if $|\mathbb{T}|$ and $N_{\text{surrogate}}$ yielded only nodes that fill a vector unit completely on each and every surrogate level. Therefore, we do not run triangle-to-triangle comparisons within the tree directly. Instead, we make the tree/triangle traversal collect all comparisons to be made with a buffer. Once we have identified all triangle collisions to be computed, we stream the whole buffer through the vector units. We merge the triangle representations on-the-fly.

Implementation remark 6. If our surrogate tree hosts more than one triangle per nonleaf node, the algorithm has to be completed by a further cleanup step which ensures that the active set remains consistent with the tree: It runs through the active set of a particle-particle combination once again. If any of a surrogate triangle’s children is part of the active set, the surrogate is removed from the set, but all of its children become active. Without such an additional sweep over the active set, all triangles of a node could be active plus the children of one triangle which implies that we test against the children triangle set plus their surrogate triangle. Restricting the node triangle cardinality for surrogate levels to one renders the additional cleanup unnecessary. \square

6. Runtime results. Our algorithms yield correct results (Lemma 5.1, Corollary 5.3, and Lemma 5.4), but they do not provide an efficiency guarantee. We hence collect runtime results, i.e., gather empirical evidence. Real-world experiments benchmarked against measurements remain out-of-scope for the present paper. We furthermore continue to focus on the actual collision detection and neglect the impact of different time step sizes—in particular comparisons between implicit and explicit schemes facilitating different stable time step choices—the cost of a coarse-grain neighbor search via a grid, e.g., and notably the construction cost for the surrogates which are done offline prior to the simulation run. All experiments are run on Intel Xeon E5-2650V4 (Broadwell) chips in a two socket configuration with 2×12 cores. They run at 2.4 GHz, though TurboBoost can increase this up to 2.9 GHz. However, a core executing AVX2 instructions will fall back to a reduced frequency (minimal 1.8 GHz) to stay within the thermal design power limits [6].

Our node has access to 64 GB TruDDR4 memory, which is connected via a hierarchy of three inclusive caches. They host $12 \times (32 + 32)$ KiB, 12×256 KiB, or 12×2.5 MiB, respectively. We obtain around 109 GB/s in the Stream TRIAD [18] benchmark on the node which translates into 4,556 MB/s per core. The node has a theoretical single precision peak performance between 2.4 (non-AVX mode and baseline speed) and 46.4 Gflop/s per core (AVX 2.0 FMA3 with full turbo boost). All of our calculations are ran in single precision. They are translated with the Intel 19 update 2 compiler and use the flags `-std=c++17 -O3 -qopenmp -march=native -fp-model and fast=2`; i.e., we tailor them to the particular instruction set.

All presented performance counter data is read out through LIKWID [21]. DEM codes are relatively straightforward to parallelize as their particle-particle interaction is strongly localized: We can combine grid-based parallelism (neighbor cells) with an additional parallelization over the particle pairs [15]. The load balancing of these concurrency dimensions, however, remains challenging. As our ideas reduce the comparison cost algorithmically yet do not alter the concurrency character, we stick (logically) to single core experiments to avoid biased measurements due to parallelization or load balancing overheads. Yet, we artificially scale up the setup by

replicating the computations per node over multiple OpenMP threads whenever we present real runtime data or machine characteristics and then break down the data again into cost per replica per core. This avoids that simple problems fit into a particular cache or that memory-bound applications have exclusive access to two memory controllers.

6.1. Experimental setup. We work with two simple benchmark setups before we validate our results for large numbers of particles. In the *particle-particle* setup, we study two spherical objects which are set on direct collision trajectory. They bump into each other and then separate again. The setup yields three computational phases: While the particles approach, there is no collision and no forces act on the particles as we neglect gravity. When they are close enough, the particles exchange forces, and the system becomes very stiff suddenly before the objects repulse each other again and separate. We focus exclusively on the middle phase. Throughout this approach-and-contact situation, the algorithmic complexity of the contact detection is in $\mathcal{O}(|\mathbb{T}|^2)$, as we assume that both particles have the same triangle count.

In the *particle-on-plane scenario*, we drop a spherical object onto a tilted plane. The particle hits the plane, bumps back in a slightly tilted angle, i.e., with a rotation, and thus hops down the plane. This problem yields free-fall phases which take turns with stiff in-contact situations. Furthermore, the area of the free particles which is subject to potential contacts changes all the time as the particle starts to rotate, and the contacts result from a complex geometry consisting of many triangles compared to a simplistic geometry with very few triangles. The underlying computational complexity is in $\mathcal{O}(|\mathbb{T}|)$.

In the *grid scenario*, we finally arrange 24,576 spheres in a Cartesian grid. Each particle slightly overlaps the ϵ -region of its neighbors. As there is no ground plane or gravity, the particles “float” in space. Due to the regular particle layout, the interaction pattern yields a Cartesian topology; i.e., each particle collides with four other particles initially.

Our codes work exclusively with *sphere-like particle shapes*, which result from a randomized parameterization: We decompose the sphere with radius 1 into $|\mathbb{T}|$ triangles. If not stated otherwise, $|\mathbb{T}| = 1,280$. In the first two scenarios the vertices on the sphere which span the triangles are subject to a Perlin noise function, which offsets the vertex along the normal direction of the surface. $\eta_r = 1$ adds no noise and thus yields a perfect, triangulated sphere with radius 1, where all vertices are exactly 1 unit away from the sphere’s origin. Otherwise, the per vertex radius is from $[1, \eta_r]$. As we use a hierarchical noise model, a high η_r yields a degenerated shape which retains a relatively smooth surface.

For the implicit schemes, we consider the result converged when the update to the force and torque applied to every particle underruns a relative threshold of 1%. With this accuracy, single precision is sufficient. The Picard iterations are subject to damping and acceleration: Any update $(dv, d\omega)$ relative to the start configuration of a time step results from the weighted average between the currently computed forces and the forces of the previous step. If forces “pull” into one direction over multiple iterations, the updates behind trials become successively bigger. If the forces oscillate, these oscillations are diminishing. Empirical data suggests that this choice helps us to meet Assumptions 1 and 2. $\epsilon = 10^{-2}$ is uniformly used on the finest mesh level. This is a relative quantity, i.e., chosen relative to the particle diameter. For the plane, we uniformly use $\epsilon = 10^{-2}$.

6.2. Surrogate properties. We first assess our surrogate geometry's properties. Our coarsest surrogate model consists of a single triangle. We compare this triangle's longest edge (diameter) $d_{k_{\max}}$ plus its corresponding $\epsilon_{k_{\max}}$ value to the radius $r_{\text{sphere}} = \frac{\eta_r}{2}$ of the bounding sphere of the fine grid object (Table 6.1). For the surrogate hierarchy, we use $N_{\text{surrogate}} = 8$ as a coarsening factor, a choice we employ throughout the experiments. This implies that an object with $|\mathbb{T}| = 64$ triangles spans three surrogate organized as a tree: They host $|\mathbb{T}| = 64$ (original model), $|\mathbb{T}_{\mu}| = 64/N_{\text{surrogate}} = 8$, and $|\mathbb{T}_{\neq}| = 1$ triangles. In this first test, we approximate low frequency noise by scaling along one axis; i.e., we elongate the sphere along one direction yet do not introduce bumps or extrusions (Figure 6.1). With growing η_r , we obtain increasingly nonspherical objects resembling an ellipsoid. The rationale behind this simplified noise is that we eliminate nondeterministic effects and study the dominant sphere distortion effects.

The combination of $d_{k_{\max}}$ and $\epsilon_{k_{\max}}$ characterizes the shape of our coarsest surrogate model. A large diameter relative to a small halo size describes a disc-like object. A small diameter relative to a large halo size describes a sphere-like object. Different triangle counts for the fine grid model allow us to assess the impact of the level of detail of the fine grid mesh onto the resulting coarsest surrogate geometry.

TABLE 6.1

Different triangle counts $|\mathbb{T}|$ per spherical object scaled along one axis by a factor of μ . Per setup, we study the top level surrogate which contains one triangle and compare the maximum triangle diameter plus its halo size against the bounding sphere radius. Here we only report on the increase in ϵ for the coarsest surrogate, i.e., at the finest level $\epsilon = 0$.

η_r	$ \mathbb{T} = 80$		$ \mathbb{T} = 320$		$ \mathbb{T} = 1,280$		r_{sphere}
	$d_{k_{\max}}$	$\epsilon_{k_{\max}}$	$d_{k_{\max}}$	$\epsilon_{k_{\max}}$	$d_{k_{\max}}$	$\epsilon_{k_{\max}}$	
1.0	0.09	0.49	0.09	0.50	0.08	0.52	0.50
1.2	0.10	0.55	0.10	0.56	0.10	0.56	0.60
1.4	0.34	0.54	0.12	0.65	0.11	0.66	0.70
1.8	0.14	0.84	0.89	0.53	1.35	0.51	0.90
2.6	1.54	0.58	2.24	0.51	2.36	0.52	1.30

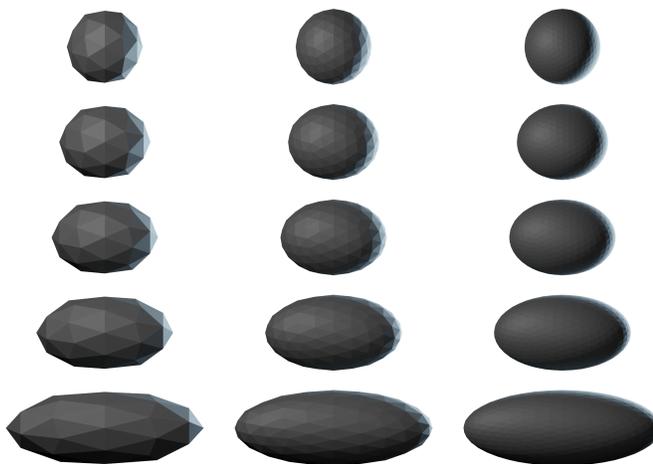


FIG. 6.1. A series of objects created from unit spheres with increasing level of detail ($80 \leq |\mathbb{T}| \leq 1,280$ from left to right) and a scale factor ($1.0 \leq \eta_r \leq 2.6$ from top to bottom).

Our surrogate model of choice on the coarsest level (see the penalty term in (A.1) of the appendix) almost degenerates to a point if the underlying triangulated geometry approximates a sphere. It approaches a bounding sphere. The triangle count approximating a spherical object does not have a significant qualitative or quantitative impact on this characterization of the coarsest surrogate triangle. Once the triangulated mesh becomes less spherical, the surrogate triangle starts to align with the maximum extension of the fine mesh. It spreads out within the geometry along the geometry's longest diameter, an effect that is the more distinct the higher the fine geometry's triangle count. The halo layer $\epsilon_{k_{\max}}$ around the surrogate triangle, which is analogous to a sphere's radius if the surrogate triangle approaches a point, remains in the order of $r = 0.5$. This is the radius of the original unit sphere ($\eta_r = 1$).

For a close-to-spherical geometry, our volumetric surrogate model never exceeds 135% of the bounding sphere volume ($|\mathbb{T}| = 1, 280$). For the highly nonspherical cases ($\eta_r = 2.6$) our surrogate volume can be as little as 37% ($|\mathbb{T}| = 80$) of the simple bounding sphere volume. This advantageous property results from the observation that a growth of $d_{k_{\max}}$ anticipates any extension of the geometry, while the $\epsilon_{k_{\max}}$ ensures that the minimal geometry diameter, which is at least as large as the original sphere, remains covered by the surrogate triangle plus its halo environment.

Observation 1. For highly nonspherical sets of triangles, our surrogate formalism yields advantageous representations. For spherical observations, it resembles the bounding sphere. This holds for all levels of the surrogate cascade.

In a surrogate tree, fine resolution tree nodes (surrogate triangles) are characterized by the low triangle count measurements in Table 6.1, where localized patches are highly nonspherical (large η_r). Surrogate triangles belonging to coarser levels inherit characteristics corresponding to larger $|\mathbb{T}|$. We conclude that our triangle-based multiresolution approach is particularly advantageous as an early termination criterion (“there is certainly no collision”) on the rather fine surrogate resolution levels within the surrogate tree or is overall tighter fitting than bounding sphere formalisms for nonspherical geometries.

6.3. Hybrid single level contact detection. Even though our multiscale approach intends to reduce the number of distance calculations, a high throughput of the overall algorithm continues to hinge on the efficiency of the core distance calculation. We hence continue with studies around the explicit Euler, where we omit the multiscale hierarchy. We work with the finest particle mesh representation only.

The assessment of the core comparison efficiency relies on our sphere-on-plane setup and the particle-particle setup. They represent two extreme cases of geometric comparisons: With the plane, a complex particle with many triangles hits very few triangles. As long as the triangles spanning the plane are large relative to the particle diameter, it is irrelevant how the plane is modeled, i.e., if it consists of solely one or two triangles or an arrangement of multiple triangles. The multiscale algorithms asymptotically approach a $|\mathbb{T}| : 1$ comparison with \mathbb{T} given by the particle. When we collide two particles—we use the same triangle count for both—we can, in the worst case, run into a $|\mathbb{T}| : |\mathbb{T}|$ constellation. For both setups, we use strictly spherical particles ($\eta_r = 1$).

For the triangle-triangle comparisons, two code variants are on the table: We can run the comparison-based (baseline) algorithm, or we can use the hybrid code variant which runs four steps of the iterative scheme before it checks if the two last iterates of the contact point differ by more than $C\epsilon_h$. We use $C \approx 1$. If the difference exceeds the

threshold, our algorithm assumes that the code has not converged and hence reruns the comparison-based code to obtain a valid contact assessment. The comparison-based code variant is 4-way vectorized and relies on Intel intrinsics. The hybrid variant is vectorized over batches of eight packed triangle pairs using an OpenMP `simd` annotation.

Our hybrid approach outperforms a sole comparison-based approach robustly for the $|\mathbb{T}| : |\mathbb{T}|$ setups (Table 6.2). For strongly ill-balanced triangle counts, the insulated comparison-based approach is superior (Table 6.3). The comparison-based code variant is not able to benefit from streaming SIMD extensions (SIMD) or AVX at all (not shown), while the hybrid AVX usage increases with increasing triangle counts. We end up with up to 40–45% “turbo-mode” peak performance which we have to calibrate with the AVX frequency reduction [6]. The relative number of fallbacks, i.e., situations where the iterative scheme does not converge within four iterations, decreases with growing geometry detail, while the same effect is not as predominant for the particle-on-plane scenario.

Our data confirms the superiority of the hybrid approach for the particle-particle comparisons [14, 15]. They confirm that the approximation of the Hessian does not significantly harm the robustness, even though the number of fallbacks becomes nonnegligible. Our arithmetic intensity determines how much improvement results from the hybrid strategy. It is significantly higher for the particle-particle setup as

TABLE 6.2

Particle-particle scenario. We compare a comparison-based realization (top) against a hybrid realization (bottom). Per setup, we present the time-to-solution ($[t]=s$) per Euler step, i.e., one run through all possible triangle combinations, and we augment this data with MFlop/s rates split up into scalar and vectorized contributions. Vector calculations are categorized as 128 bit packed (SIMD) or 256 bit packed (AVX) for four and eight simultaneous 32 bit floating point operations, respectively. For the hybrid setup, we finally quantify how many triangle pairs had to be checked a posteriori, i.e., as fallback, by the comparison-based algorithm. This runtime is included in the data. All measurements are given as the average per core.

	$ \mathbb{T} $	Runtime	Scalar	Packed 128B	Packed 256B	Fallback
Comp.	12	$6.59 \cdot 10^{-2}$	$1.52 \cdot 10^{-2}$	$3.22 \cdot 10^3$		
	36	$5.17 \cdot 10^{-2}$	$1.90 \cdot 10^{-3}$	$3.46 \cdot 10^3$		
	140	$3.80 \cdot 10^{-1}$	$3.00 \cdot 10^{-4}$	$3.18 \cdot 10^3$		
	1,224	$4.35 \cdot 10^1$	$0.00 \cdot 10^0$	$3.04 \cdot 10^3$		
Hybrid	12	$4.99 \cdot 10^{-2}$	$6.27 \cdot 10^1$	$1.11 \cdot 10^3$	$1.07 \cdot 10^4$	7.7%
	36	$2.76 \cdot 10^{-2}$	$1.27 \cdot 10^2$	$9.69 \cdot 10^2$	$1.45 \cdot 10^4$	4.5%
	140	$1.83 \cdot 10^{-1}$	$1.97 \cdot 10^2$	$3.25 \cdot 10^2$	$1.84 \cdot 10^4$	1.2%
	1,224	$1.99 \cdot 10^1$	$2.42 \cdot 10^2$	$5.30 \cdot 10^1$	$2.10 \cdot 10^4$	0.028%

TABLE 6.3

Experiments from Table 6.2 for the particle-on-plane setup.

	$ \mathbb{T} $	Runtime	Scalar	Packed 128B	Packed 256B	Fallback
Comp.	12	$2.60 \cdot 10^{-2}$	$3.87 \cdot 10^{-2}$	$3.33 \cdot 10^3$		
	36	$6.50 \cdot 10^{-2}$	$1.96 \cdot 10^{-2}$	$3.85 \cdot 10^3$		
	140	$1.84 \cdot 10^{-1}$	$1.28 \cdot 10^{-2}$	$4.09 \cdot 10^3$		
	1,224	$1.98 \cdot 10^0$	$2.90 \cdot 10^{-3}$	$4.27 \cdot 10^3$		
Hybrid	12	$2.80 \cdot 10^{-2}$	$1.06 \cdot 10^1$	$1.25 \cdot 10^3$	$9.34 \cdot 10^3$	6.3%
	36	$6.70 \cdot 10^{-2}$	$1.56 \cdot 10^1$	$1.34 \cdot 10^3$	$1.19 \cdot 10^4$	5.1%
	140	$1.79 \cdot 10^{-1}$	$2.15 \cdot 10^1$	$1.34 \cdot 10^3$	$1.35 \cdot 10^4$	4.8%
	1,224	$1.85 \cdot 10^0$	$3.03 \cdot 10^1$	$9.32 \cdot 10^2$	$1.48 \cdot 10^4$	3.6%

opposed to the particle-on-plane setup. Therefore, only the former prospers through vectorization. We see an increased fallback for decreased geometric detail due to the larger relative ϵ , which is used to identify fallback conditions.

Observation 2. As long as we do not compare extreme cases (single triangle versus a lot of triangles), the hybrid approach is faster. It is thus reasonable to employ it on all levels of the surrogate tree, even though it might be reasonable to skip iterative comparisons a priori if the coarsest surrogate level is involved. The latter observation does not result from a mathematical “nonrobustness” but is a sole machine effect.

6.4. Multiresolution comparisons for explicit time stepping. Within an explicit time stepping code, our multiresolution approach promises to eliminate unnecessary comparisons since it identifies “no collision” constellations quickly through the surrogates: Whenever it compares two geometries, the algorithm runs through the resolution levels top down (from coarse to fine). The monotonicity of the surrogate definition implies that we can stop immediately if there is no overlap between two surrogates. The code either employs the pure geometry-based approach or the hybrid strategy on all levels. We focus on spherical particles ($\eta_r=1$) discretized by 1,224 triangles and average over 100 time steps such that we run into no-collision phases for the particle-particle setup and see the particle roll down the plane for the particle-on-plane setup.

Our measurements confirm the superiority of the hybrid scheme in the surrogate context (Table 6.4): In line with section 6.3, no multiresolution setup with comparison-based contact detections on surrogate levels is able to outperform the configurations where all levels are tackled through the hybrid approach. Further studies where different variants are used on different (surrogate) levels are beyond scope.

In our two-particles scenario, the particles yield $1,224^2$ comparisons per time step if no surrogate helper data structure is used. Our data reflects the quadratic (particle-particle) or linear (particle-on-plane) computational complexity. In both cases, the number of triangle comparisons is reduced by more than an order of magnitude through the surrogate hierarchy, and the surrogate version outperforms its single level baseline robustly. The hierarchical scheme’s additional computational cost (overhead) is negligible, though it does not significantly alter the ratio of iterative checks to fallback comparisons in the hybrid scheme.

TABLE 6.4

Time-to-solution ($[t]=s$) and number of triangle distance checks for an explicit Euler step for the particle-particle collision (top) and the particle-on-plane setup (bottom). We compare a comparison-based setup to a hybrid approach on a single level versus a surrogate hierarchy which is traversed from coarse to fine. For the hybrid configuration, we present the number of comparison-based fallbacks versus the number of iterative comparisons. Each iterative comparison of two triangles consists of four Newton steps. Only if these four steps fail to converge, the algorithm issues the comparison-based postprocessing. Both the iterative comparisons plus the (fewer) comparison-based postprocessing steps determine the runtime (right column).

Method	Comparison-based		Hybrid		
	#tri. comp.	Runtime	#tri. comp.	#iterative	Runtime
Single level	$1.50 \cdot 10^6$	$3.96 \cdot 10^0$	$1.66 \cdot 10^3$	$1.50 \cdot 10^6$	$1.87 \cdot 10^0$
Surrogate hierarchy	$8.20 \cdot 10^3$	$2.60 \cdot 10^{-2}$	$1.44 \cdot 10^2$	$7.64 \cdot 10^3$	$1.80 \cdot 10^{-2}$
Single level	$6.27 \cdot 10^5$	$1.68 \cdot 10^0$	$1.57 \cdot 10^4$	$6.27 \cdot 10^5$	$8.00 \cdot 10^0$
Surrogate hierarchy	$5.27 \cdot 10^3$	$7.68 \cdot 10^{-2}$	$4.96 \cdot 10^2$	$5.03 \cdot 10^3$	$4.00 \cdot 10^{-2}$

Observation 3. Our surrogate technique efficiently reduces the number of comparisons between two geometries, as “no-collision” setups are identified with low computational cost.

6.5. Multiresolution comparisons for implicit time stepping. Implicit methods are significantly more stable than their explicit counterpart. The price to pay for this is an increased computational complexity: The Picard iterations imply that we have to run the core contact point detection more often per time step. The Picard iterations’ update of collision point detection furthermore implies that the surrogate tree does not unfold linearly anymore. While the explicit time stepping algorithm runs through the tree from coarse to fine, the implicit scheme might, through the iterative updates of the rotation and position, walk up and down within the tree. The increased stability of implicit time stepping allows users to pick larger time step sizes which, in practice, compensate to some degree of the increased compute load. This problem-specific cost amortization is beyond scope here; i.e., we compare implicit and explicit schemes with the same time step size. All data results from spherical particles and the default triangle counts. We average all measurements over 100 time steps and include the “roll down the plane” situation for the particle-on-plane setup.

Averaged over 100 time steps, we observe that the number of Picard iterations is small and bounded (Table 6.5). We study the impact of a switch to the iterative scheme, with the hybrid fallback on the finest level, and observe that it slightly increases the Picard iteration count. The usage of a multiscale method merged into the Picard iterations increases the iteration count, too. Both modifications yield flawed contact point guesses and thus require us to run more Picard iteration steps overall. The wrong guesses have to be compensated later on.

Observation 4. Both the iterative approximation of contact points and the “one Picard step before we widen the active set” strategy increase the total number of required Picard iterations.

The surrogate hierarchy yields an efficient early termination criterion for our collision detection. If there is no collision, the code does not step down into the fine grid resolutions. This property carries over from the explicit to the implicit algorithm (Figure 6.2). An increase of the computational cost by a factor of 4.6 is acceptable in return for an implicit scheme. We, however, observe that this increase holds for brief point contacts only. It raises to a factor of 13.1 if contacts persist. In our example, this happens once the spherical object starts to roll and slide down the tilted plane.

Within our multiresolution framework, the cost per Picard iteration is not uniform and constant but depends heavily on the surrogate tree fragments that are used. The costs are in particular nonuniform for nonsimplistic setups. The growth in Picard iterations (on average) per time step (Table 6.5) increases the number of triangle-to-triangle checks, compared to the explicit schemes (Table 6.4), by exactly this factor

TABLE 6.5
Average number of Picard iterations per time step for our first two scenarios.

Method	Particle-particle		Particle-plane	
	Comparison-based	Iterative	Comparison-based	Iterative
Single level or surrogate within Picard	4.6	4.9	7.1	7.1
Multiscale Picard	6.2	6.2	13.0	13.1

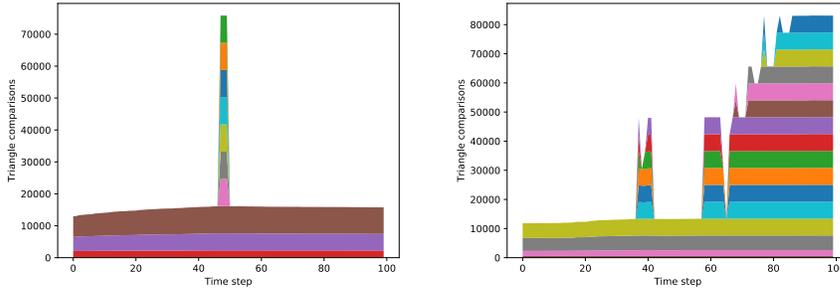


FIG. 6.2. Number of triangle-to-triangle comparisons over time per surrogate representation level. The data stems from the particle-particle collision (left) and the particle-on-plane setup (right) subject to the implicit time stepping. The lowest layer illustrates the triangle comparison count for the first iteration, the next layer for the second iteration, and so forth. The number of layers corresponds to the total number of iterations.

if we stick to a plain geometry model. Yet, it does not manifest in an explosion of the runtime (Table 6.6) if we employ the surrogate trees. They help to reduce the compute cost dramatically, as we study only those parts of the surrogate tree which might induce a collision. We prune the tree per Picard iteration.

Observation 5. Our multiscale Picard approach is particularly beneficial for strongly instationary setups where the topology of particle interactions changes quickly.

Permuting and fusing the Picard iteration loops and the traversal over the surrogate tree reduce the number of triangle-to-triangle comparisons further. This observation holds for the particle-particle setup. It does not hold for the particle-on-plane setup. The advanced version benefits from the fact that we memorize the active set in-between two Picard iterations: While the implicit version from section 5.2 runs through the whole tree starting from the root in every iteration, our advanced version starts from a certain resolution and unfolds at most one level per Picard step. We save the progress through coarser resolutions, and we do not step all the way down in early iterations. This state-based approach works as our narrowing is effective: If we step down into a “wrong” part of the tree and find out that these fine resolutions do not contribute toward the final force, we successively remove these fine resolutions from the (active) comparison sets again. For a sphere rolling or hopping down a plane, the active set remains almost invariant throughout the Picard iterations, and we do

TABLE 6.6

Time per time step ($[t]=s$) and triangle distance comparisons for our implicit schemes for the particle-particle collision (top) and the particle-on-plane setup (bottom).

Method	Comparison-based		Hybrid		
	#tri. comp.	Runtime	#tri. comp.	#iterative	Runtime
Single level	$6.89 \cdot 10^8$	$1.69 \cdot 10^2$	$8.78 \cdot 10^5$	$7.34 \cdot 10^8$	$7.44 \cdot 10^1$
Surrogate within Picard	$3.97 \cdot 10^6$	$1.08 \cdot 10^0$	$7.14 \cdot 10^4$	$3.70 \cdot 10^6$	$7.20 \cdot 10^{-1}$
Multiscale Picard	$1.82 \cdot 10^6$	$7.70 \cdot 10^{-1}$	$2.25 \cdot 10^4$	$1.70 \cdot 10^6$	$4.70 \cdot 10^{-1}$
Single level	$4.43 \cdot 10^8$	$9.61 \cdot 10^1$	$1.11 \cdot 10^7$	$4.43 \cdot 10^8$	$4.41 \cdot 10^1$
Surrogate hierarchy	$3.59 \cdot 10^6$	$4.80 \cdot 10^{-1}$	$3.44 \cdot 10^5$	$3.42 \cdot 10^6$	$2.20 \cdot 10^{-1}$
Multiscale Picard	$3.50 \cdot 10^6$	$4.90 \cdot 10^{-1}$	$3.11 \cdot 10^5$	$3.35 \cdot 10^6$	$1.90 \cdot 10^{-1}$

not benefit from the narrowing or an early termination. We do, however, benefit here from the adaptive localization of the contact detection within the tree.

Observation 6. The multiscale Picard approach in combination with a hybrid contact detection keeps the cost of the implicit time stepping bounded by a factor of four compared to an explicit scheme.

6.6. Multiresolution comparisons for implicit time stepping with large scale differences. To assess collisions between objects of different scales, we first instantiate the particle-particle scenario with both particles being spheres approximated by $|T| = 80$ triangles ($\eta_r = 1$). We fix one of the spheres with a unit diameter and scale the other one up. All data studies situations where the two particles are close and yield contact points.

There is no clear trend relating the number of Picard iterations to the scale of the involved particles (Table 6.7). As the scale difference grows, the number of triangle-triangle comparisons reduces slightly (cmp. how often the iterative scheme is invoked), while more iterative updates run into the fallback, i.e., require a posteriori, if-based triangle comparisons. The cost to compare a large with a small particle thus remains in the same order as the comparison of two equally sized particles. There is no clear runtime trend.

The bigger the size difference between two triangles, the more likely it is that the hybrid scheme runs into the fallback (cmp. data in Table 6.3, where the slope triangles are large compared to the particle triangles). The frequent fallbacks team up with the effect that coarse (surrogate) triangles within the larger particle overlap large numbers of surrogate triangles on the smaller collision partner. Relatively large fractions of the small particle's surrogate tree have to be unfolded.

To distinguish the impact of the particle sizes from the impact of the triangle, we next subdivide the upscaled colliding particle such that the triangle sizes of both particles remain almost invariant and comparable. Doubling the particle size thus corresponds to four times more triangles. Within the surrogate tree, we keep the number of triangles per leaf roughly constant.

Despite the increase in the number of triangles per upscaling, the number of comparisons and the runtime quickly plateau (Table 6.8). The additional cost due to the increase in the triangle count is bounded by a factor of four, even if we continue

TABLE 6.7

Time ($[t]=s$) plus triangle distance comparisons per implicit time step for a particle-particle setup (80 triangles), where one particle is scaled relative to the other one.

Scale	Triangles	Iterations	#tri. comp.	#iterative	Time
1	80	13	$1.59 \cdot 10^1$	$1.45 \cdot 10^2$	$8.72 \cdot 10^{-2}$
2	80	13	$1.93 \cdot 10^1$	$1.27 \cdot 10^2$	$6.80 \cdot 10^{-2}$
4	80	18	$4.31 \cdot 10^1$	$1.78 \cdot 10^2$	$1.00 \cdot 10^{-1}$
8	80	14	$4.10 \cdot 10^1$	$9.82 \cdot 10^1$	$6.82 \cdot 10^{-2}$

TABLE 6.8

Time ($[t]=s$) and triangle distance per time step for two colliding particles of different size with comparable triangle sizes.

Scale	Triangles	Iterations	#tri. comp.	#iterative	Runtime
1	80	13	$1.59 \cdot 10^1$	$1.45 \cdot 10^2$	$8.72 \cdot 10^{-2}$
2	320	19	$4.18 \cdot 10^1$	$5.05 \cdot 10^2$	$2.25 \cdot 10^{-1}$
4	1,280	20	$4.66 \cdot 10^1$	$5.88 \cdot 10^2$	$2.60 \cdot 10^{-1}$
8	5,120	17	$3.80 \cdot 10^1$	$5.44 \cdot 10^2$	$2.40 \cdot 10^{-1}$

to increase the geometric detail beyond a factor of four. The cost per triangle reduces as one particle grows relative to the other.

As the size of the larger object grows, its surrogate tree becomes deeper. The algorithm unfolds the larger tree, but this unfolding is a localized process. We end up with roughly scale-invariant compute cost which is dominated by the number of triangles stored within one node of the finest tree level.

Observation 7. The multiscale Picard approach is particularly effective if we compare objects of massively different size yet with the same level of detail. In this case, the cost per triangle comparison remains almost invariant.

The observation is particular encouraging for simulations with objects of massively different size, where the large objects exhibit a high level of detail. Such situations are found in fluid-structure interaction, where a collision induces a force on a bending object. Such objects often are modeled via adaptive meshes which have to refine around the contact point.

6.7. Many particle experiments. Our multiscale Picard method shows an improvement in time-to-solution for individual particles which collide with other particles of different size or simple geometric object such as a plane. To be useful in real DEM codes, the algorithms must remain performant when we run simulations with many particles.

As long as the neighborhood search is efficient, our achievements for individual particle-particle interactions carry over to dynamic simulations with large particle counts. If particles move around and bounce into each other, yet each particle only interacts with few other objects per time step, our multiscale concepts simply are to be generalized from two objects to a small object count. The situation changes if particles enter a stationary regime, i.e., if the particles are densely clustered and almost at rest.

We therefore finally simulate 24,576 particles that barely overlap and are initially at rest. This imitates a basin of granulates; e.g., each particle is made up of 1,224 triangles; i.e., we handle a total memory footprint of approximately 3.87 GiB which is substantially larger than the 30 MiB total L3 cache.

The hierarchical scheme yields a massive reduction of (iterative) triangle-triangle comparisons, and it notably succeeds to avoid too many comparisons which subsequently have to be postprocessed (Table 6.9). The runtime per particle-particle comparison is significantly lower than in previous dynamic setups. Consequently, the cost per particle is lower, too, even though we have around six collision partners per particle.

TABLE 6.9

Measurements for 24,576 particles which are densely clustered yet almost at rest. The runtime is the runtime per time step per core. We present the total time and the time per particle-particle comparison.

Method	Comparison-based Time			Hybrid Time			
	#tri. comp.	Total	Per comp	#tri. comp.	#iterative	Total	Per comp
Single level	$6.50 \cdot 10^9$	$1.92 \cdot 10^6$	$9.67 \cdot 10^2$	$1.29 \cdot 10^8$	$6.50 \cdot 10^9$	$8.36 \cdot 10^5$	$4.20 \cdot 10^2$
Surrogate within Picard	$2.49 \cdot 10^6$	$4.70 \cdot 10^2$	$2.37 \cdot 10^{-1}$	$2.92 \cdot 10^2$	$2.49 \cdot 10^6$	$3.77 \cdot 10^2$	$1.90 \cdot 10^{-1}$
Multiscale Picard	$9.08 \cdot 10^5$	$2.30 \cdot 10^2$	$1.16 \cdot 10^{-1}$	$1.09 \cdot 10^2$	$1.06 \cdot 10^6$	$2.13 \cdot 10^2$	$1.07 \cdot 10^{-1}$

Since the particles are almost at rest and only have tiny overlaps, the arising forces are very small. Consequently, few Picard iterations are sufficient to converge for the overall setup. The surrogate trees unfold only around the localized contacts, and this unfolding is very efficient; i.e., we barely remove triangles from the active set while we run the Picard iterations.

Observation 8. The multiscale approach remains advantageous for larger particle assemblies where the particles are densely clustered.

7. Conclusion. We present a family of multiresolution contact detection algorithms that exhibit low computational cost and high vectorization efficiency. Few core ideas guide the derivation of these algorithms: We rigorously phrase the underlying mathematics in a multiresolution and multimodel language, where low-cost resolutions (surrogates) or algorithms (iterative contact search) precede an expensive follow-up step which becomes cheaper through good initial guesses or can be skipped in many cases. We replace dynamic termination criteria behind iterative algorithms with fixed iteration counts. While this might mean that we terminate prematurely, a fixed iteration count allowed us to unroll loops and to permute them. The permutation of loops is our last ingredient which we apply on multiple levels: We switch the traversal of triangles with Newton iterations, and we switch the Picard iterations with the tree unfolding.

The present work is solely algorithmic and has theoretical character. A natural next step is its application to large-scale, massively parallel, real-world simulations. This introduces at least three further challenges: First, the parallelization between multiple compute nodes makes the algorithm more complex and introduces significant upscaling and load balancing challenges. Second, we only use a simple contact model to compute interaction forces which might be inappropriate for actual challenges from sciences and engineering. Finally, our studies solely rely on one, fixed time step size. They do not even employ the fact that implicit schemes allow for larger time step size, a fact which has to be employed and studied for real-world runs. Furthermore, we rely—so far—on a naive assumption that the Picard iterations converge. A more robust code variant would either identify nonconvergence via force, rotation, and movement deltas that do not decrease over the Picard iterations or exploit the fact that we know how accurate our surrogate models are via their ϵ -value. In both cases, surrogate levels could be skipped automatically.

On the methodological side, there are three natural extensions of our work: First, our surrogate mechanism always kicks off from the surrogate tree's root when it searches for contact points. For time stepping codes, this is not sophisticated. It might be advantageous to memorize the tree configurations in-between two subsequent time steps and thus to exploit the fact that many particle configurations change only smoothly in time. Scenarios such as our particle-on-plane setup may particularly benefit from this, as coarse level surrogates, where the force estimates are just as likely to push the solution in the wrong direction, can be skipped based on hints from the previous time step.

Second, we work with multiple spatial representations, i.e., accuracies, but we stick to single precision all the way through. It is a natural extension to make our iterative algorithm use a reduced precision on coarse surrogate models, i.e., early throughout the algorithm. Any machine imprecision can be recovered in our case through a slight increase of ϵ . Such a mixed precision strategy is particularly attractive in an era where more and more compute devices are equipped with special-purpose, reduced-precision linear algebra components. An additional flavor of mixed

precision arises once we use more complex contact models. It is not clear if such complex models are required on the surrogate levels, too; i.e., it might be appropriate to implement a multiphysics model where surrogate contact detection complements accurate interaction models on the finest mesh level.

Finally, we next will have to tackle large-scale systems implicitly: DEM models are notoriously stiff, yet the stiffness is localized, as typically not all particles in a setup interact with all other particles. It is a natural extension to investigate local time stepping, where each particle advances with its own Δt , and to make the surrogate representations naturally follow and inform these local time step choices.

Appendix A. Surrogate triangles. The construction of good surrogate models is a challenge of its own, as there are infinitely many surrogate models for a given particle p . We rely a functional minimization with a fixed coarsening factor to construct the surrogate hierarchy bottom-up. Let one surrogate triangle for a set of triangles \mathbb{T}_{k-1} be spanned by its three vertices $x_1, x_2, x_3 \in \mathbb{R}^3$ which follow

$$(A.1) \quad \arg \min_{x_1, x_2, x_3} \frac{1}{\beta_{\text{size}}} \sum_{x \in \{x_1, x_2, x_3\}, t \in \mathbb{T}_{k-1}} \|x - t\|^{\beta_{\text{size}}} + \frac{\alpha_{\text{area}}}{2} |(x_2 - x_1) \times (x_3 - x_1)|^{-2} + \frac{\alpha_{\text{inside}}}{\beta_{\text{normal}}} \sum_{x \in \{x_1, x_2, x_3\}, x_t \in t \ \forall t \in \mathbb{T}_{k-1}} (\max(0, (x - x_t) \cdot N(x_1, x_2, x_3)))^{\beta_{\text{normal}}}.$$

$\beta_{\text{size}} \geq 2$ is a fixed integer parameter which we usually pick very high, and the term $\|x - t\|$ denotes the distance between a point x and a triangle t . The sum thus minimizes the maximum distance between the vertices of the surrogate triangle and the triangles from \mathbb{T}_{k-1} . We try to make the surrogate triangle as small as possible. The second term acts as a regularizer that avoids that the surrogate triangle degenerates and becomes a single point or a line. Without it, the first term would yield a single point, i.e., $x_1 = x_2 = x_3$.

The third term exploits the fact that each triangle t of p has a unique outer normal $N(t)$. Even though our surrogate models can be weakly connected, it is thus possible to assign each surrogate triangle an outer normal, too. The penalty term over the scalar product drops out due to the max function if the surrogate's normal points into the same direction as the triangles' normals. Effectively, this term ensures that the surrogate triangle nestles closely around a particle and that spikes do not induce a blown-up surrogate (Figure 2.1). Once (A.1) yields a surrogate triangle, ϵ is chosen such that the triangle is conservative for \mathbb{T}_{k-1} .

Appendix B. Surrogate trees. Let $N_{\text{surrogate}} > 1$ be the surrogate coarsening factor. We construct a surrogate tree top-down (Algorithm B.1):

- The first triangle that we insert into \mathcal{T} is the coarsest surrogate model, i.e., a degenerated object description consisting of one triangle. In line with Definition 4.6, this is the root node of our surrogate tree.
- Recursively dividing creates a tree over sets where all nonleaves have cardinality one. The leaf sets have a cardinality of roughly $N_{\text{surrogate}}$. The number of children per tree node is bounded and typically around $N_{\text{surrogate}}$.
- We construct the surrogate triangles by copying one triangle out of the underlying triangle set. Then, we iteratively minimize the functional (A.1).
- To obtain surrogates with reasonably small ϵ , we cluster the triangle sets through a tailored k -means algorithm [17]. The subsets $\mathbb{T}_{\text{local},i}$ thus are reasonable compact.

Algorithm B.1 Top-down algorithm to construct a cascade of surrogate models for a given triangulation \mathbb{T} . The algorithm yields a tree defined through $\sqsubseteq_{\text{child}}$ and hence allows us to derive a vast set of different, locally adaptive surrogate models.

```

1: function CONSTRUCTSURROGATE( $\mathbb{T}$ )
2:   Construct surrogate triangle  $t$  for  $\mathbb{T}$  ▷ Solve (A.1)
3:   Assign  $t$  smallest  $\epsilon$  such that  $t^\epsilon$  is conservative surrogate
4:   Create trivial graph  $\mathcal{T}$  with single node  $\{t^\epsilon\}$  and no edges
5:   CONSTRUCTSURROGATERECURSIVELY( $t^\epsilon, \mathbb{T}$ )
6: end function
7: function CONSTRUCTSURROGATERECURSIVELY( $t_{\text{local}}^\epsilon, \mathbb{T}_{\text{local}}$ )
8:   if  $\mathbb{T}_{\text{local}} \leq N_{\text{surrogate}}$  then
9:     Add node  $\mathbb{T}_{\text{local}}$  to  $\mathcal{T}$ 
10:    Add edge  $\mathbb{T}_{\text{local}} \sqsubseteq_{\text{child}} \{t_{\text{local}}^\epsilon\}$  to  $\mathcal{T}$ 
11:   else
12:    Split  $\mathbb{T}_{\text{local}}$  into  $N_{\text{surrogate}}$  sets  $\mathbb{T}_{\text{local},0}, \mathbb{T}_{\text{local},1}, \mathbb{T}_{\text{local},2}, \dots$  of roughly same size
13:    for  $i$  do
14:      Construct surrogate triangle  $t_{\text{new}}$  for  $\mathbb{T}_{\text{local},i}$  ▷ Solve (A.1)
15:      Assign  $t_{\text{new}}$  smallest  $\epsilon$  such that  $t_{\text{new}}^\epsilon$  is conservative surrogate over  $\mathbb{T}_{\text{local},i}$ 
16:      Add node  $\{t_{\text{new}}^\epsilon\}$  to  $\mathcal{T}$ 
17:      Add edge  $\{t_{\text{new}}^\epsilon\} \sqsubseteq_{\text{child}} \{t_{\text{local}}^\epsilon\}$  to  $\mathcal{T}$ 
18:      CONSTRUCTSURROGATERECURSIVELY( $t_{\text{new}}^\epsilon, \mathbb{T}_{\text{local},i}$ )
19:    end for
20:   end if
21: end function

```

There are many alternative paradigms to construct surrogate trees: Bounding sphere hierarchies would be an alternative. Our approach is relatively slow, yet is exclusively used as preprocessing.

Acknowledgments. This work made use of the facilities of the Hamilton HPC Service of Durham University.

REFERENCES

- [1] T. AKENINE-MILLER, E. HAINES, AND N. HOFFMAN, *Real-Time Rendering*, 4th ed., CRC Press, Boca Raton, FL, 2018.
- [2] F. ALONSO-MARROQUIN AND Y. WANG, *An Efficient Algorithm for Granular Dynamics Simulation with Complex-Shaped Objects*, preprint, arXiv:0804.0474, 2008, <https://arxiv.org/abs/0804.0474>.
- [3] K. APINIS, H. SEIDL, AND V. VOJDANI, *Enhancing top-down solving with widening and narrowing*, in *Semantics, Logics, and Calculi*, Springer, Cham, 2016, pp. 272–288, https://doi.org/10.1007/978-3-319-27810-0_14.
- [4] D. BARAFF, *An Introduction to Physically Based Modeling: Rigid Body Simulation I—Unconstrained Rigid Body Dynamics*, SIGGRAPH '97 course notes, Carnegie Mellon University, 1997.
- [5] G. BAREQUET, B. CHAZELLE, L. J. GUIBAS, J. S. MITCHELL, AND A. TAL, *BOXTREE: A hierarchical representation for surfaces in 3D*, *Comput. Graph. Forum*, 15 (1996), pp. 387–396, <https://doi.org/10.1111/1467-8659.1530387>.
- [6] D. CHARRIER, B. HAZELWOOD, E. TUTLYAEVA, M. BADER, M. DUMBSER, A. KUDRYAVTSEV, A. MOSKOVSKY, AND T. WEINZIERL, *Studies on the energy and deep memory behaviour of a cache-oblivious, task-based hyperbolic PDE solver*, *Int. J. High Perform. Comput. Appl.*, 33 (2019), pp. 973–986.
- [7] P. A. CUNDALL AND O. D. L. STRACK, *A discrete numerical model for granular assemblies*, *Géotechnique*, 29 (1979), pp. 47–65, <https://doi.org/10.1680/geot.1979.29.1.47>.
- [8] H. DAMMERTZ, J. HANIKA, AND A. KELLER, *Shallow bounding volume hierarchies for fast SIMD ray tracing of incoherent rays*, *Comput. Graph. Forum*, 27 (2008), pp. 1225–1233, <https://doi.org/10.1111/j.1467-8659.2008.01261.x>.
- [9] C. EISENACHER, G. NICHOLS, A. SELLE, AND B. BURLEY, *Sorted deferred shading for production path tracing*, *Comput. Graph. Forum*, 32 (2013), <https://doi.org/10.1111/cgf.12158>.
- [10] C. ERICSON, *Real-Time Collision Detection*.

- [11] S. GOTTSCHALK, M. LIN, AND D. MANOCHA, *OBSTree: A hierarchical structure for rapid interference detection*, in Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques, ACM, 1996, <https://doi.org/10.1145/237170.237244>.
- [12] M. HELD, J. KLOSOWSKI, AND J. MITCHELL, *Real-time collision detection for motion simulation within complex environments*, in Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques, ACM, 1996.
- [13] K. IGLBERGER AND U. RÜDE, *Massively parallel granular flow simulations with non-spherical particles*, *Comput. Sci. Res. Develop.*, 25 (2010), pp. 105–113, <https://doi.org/10.1007/s00450-010-0114-4>.
- [14] K. KRESTENITIS AND T. KOZIARA, *Calculating the minimum distance between two triangles on SIMD hardware*, in Proceedings of the 23rd Conference on Computational Mechanics, Swansea University, 2015.
- [15] K. KRESTENITIS, T. WEINZIERL, AND T. KOZIARA, *Fast DEM collision checks on multicore nodes*, in Proceedings of the 12th International Conference on Parallel Processing and Applied Mathematics, Lublin, Poland, R. Wyrzykowski, J. J. Dongarra, E. Deelman, and K. Karczewski, eds., *Lecture Notes in Comput. Sci.* 10777, Springer, Cham, 2018, pp. 123–132.
- [16] T. Y. LI AND J. S. CHEN, *Incremental 3D collision detection with hierarchical data structures*, in Proceedings of the ACM Symposium on Virtual Reality Software and Technology 1998, 1998, pp. 139–144, <https://doi.org/10.1145/293701.293719>.
- [17] J. B. MACQUEEN, *Some Methods for Classification and Analysis of Multivariate Observations*, 1967.
- [18] J. D. MCCALPIN, *Memory bandwidth and machine balance in current high performance computers*, *IEEE Computer Society Technical Committee on Computer Architecture Newsletter*, 2 (1995), pp. 19–25.
- [19] A. D. RAKOTONIRINA AND A. WACHS, *Grains3D, a flexible DEM approach for particles of arbitrary convex shape - Part II: Parallel implementation and scalable performance*, *Powder Technol.*, 324 (2018), pp. 18–35, <https://doi.org/10.1016/j.powtec.2017.10.033>.
- [20] E. SHELLSHEAR AND R. YTTERLID, *Fast distance queries for triangles, lines, and points using SSE instructions*, *J. Comput. Graph. Tech.*, 3 (2014), pp. 86–110.
- [21] J. TREIBIG, G. HAGER, AND G. WELLEIN, *LIKWID: A lightweight performance-oriented tool suite for x86 multicore environments*, in Proceedings of the 2010 39th International Conference on Parallel Processing Workshops, IEEE Computer Society, 2010, pp. 207–216.
- [22] A. WACHS, L. GIROLAMI, G. VINAY, AND G. FERRER, *Grains3D, a flexible DEM approach for particles of arbitrary convex shape - Part I: Numerical model and validations*, *Powder Technol.*, 224 (2012), pp. 374–389, <https://doi.org/10.1016/j.powtec.2012.03.023>.
- [23] T. WEINHART, C. LABRA, S. LUDING, AND J. Y. OOI, *Influence of coarse-graining parameters on the analysis of DEM simulations of silo flow*, *Powder Technol.*, 293 (2016), pp. 138–148.
- [24] S. ZHAO AND J. ZHAO, *A poly-superellipsoid-based approach on particle morphology for DEM modeling of granular media*, *Int. J. Numer. Anal. Methods Geomech.*, 43 (2019), pp. 2147–2169.
- [25] W. ZHONG, A. YU, X. LIU, Z. TONG, AND H. ZHANG, *DEM/CFD-DEM modelling of non-spherical particulate systems: Theoretical developments and applications*, *Powder Technol.*, 302 (2016), pp. 108–152, <https://doi.org/10.1016/j.powtec.2016.07.010>.