To appear in the ACM SIGGRAPH conference proceedings

Magnets in Motion

Bernhard Thomaszewski*
WSI/GRIS,
Universität Tübingen

Andreas Gumann†
Institut für Theoretische Physik,
Universität Tübingen

Simon Pabst*
WSI/GRIS,
Universität Tübingen

Wolfgang Straßer*
WSI/GRIS,
Universität Tübingen

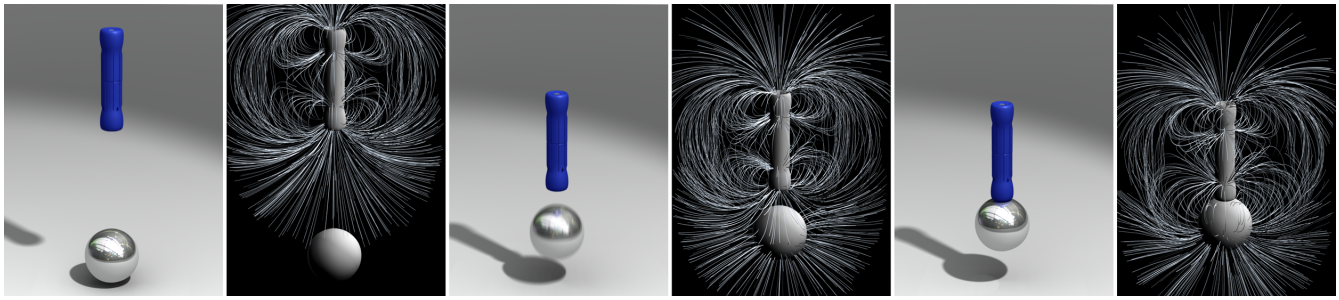

Figure 1: *Interleaved simulation snapshots and magnetic field renderings. A toy magnet, carrying small permanent magnets in its ends, is moved towards a sphere into which it induces a magnetization. As the magnet approaches the sphere, the induced field and thus the attracting forces increase until they exceed the gravitational force and pull the sphere towards the magnet.*

Abstract

We introduce magnetic interaction for rigid body simulation. Our approach is based on an equivalent dipole method and as such it is discrete from the ground up. Our approach is symmetric as we base both field and force computations on dipole interactions. Enriching rigid body simulation with magnetism allows for many new and interesting possibilities in computer animation and special effects. Our method also allows the accurate computation of magnetic fields for arbitrarily shaped objects, which is especially interesting for pedagogy as it allows the user to visually discover properties of magnetism which would otherwise be difficult to grasp. We demonstrate our method on a variety of problems and our results reflect intuitive as well as surprising effects. Our method is fast and can be coupled with any rigid body solver to simulate dozens of magnetic objects at interactive rates.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation;
Keywords: Magnetic fields and forces, rigid body dynamics

1 Introduction

Simulating the dynamic behavior of rigid bodies began to attract the interest of computer graphics researchers more than twenty years ago [Hahn 1988; Baraff 1989]. Given that rigid bodies are often made of metal, a notable lack is that *magnetic forces* have not yet been considered in this context.

*e-mail: {thomaszewski,pabst,strasser}@gris.uni-tuebingen.de
†e-mail: andreas.gumann@uni-tuebingen.de

Integrating magnetic forces into rigid body dynamics (RBD) greatly enriches the range of possible applications: an animator (artist) can create a plethora of new, stunning effects that cannot easily be reproduced using the standard approach of artificial forces and inverse kinematics. Similarly, video games can benefit from magnetics at real time frame rates. Furthermore, the visualization of magnetic field lines, which is customarily done using iron filings, could be greatly improved using computer graphics techniques.

1.1 Goals and Contributions

The purpose of this work is to develop a model for discrete magnetic interaction that is simple and fast while reproducing all macroscopic magnetic effects in a faithful way. Our contributions are:

Symmetric Equivalent Magnetic Dipole Approach We pose magnetic force and torque computation as discrete *magnetic dipole* interaction and provide explicit formulae for all quantities. The essence of our approach is the decomposition of magnetic objects into aggregates of magnetic dipole cells. In the limit of an infinite number of dipoles, the magnetization in the interior of the original object and the magnetic field outside it are *equivalent* to those of the original object. In the same way, the total forces and torques converge to those of the original object with increasing number of cells. Using a finite number of cells means an approximation whose accuracy depends on the ratio of the cells' diameters to the inter-object distance. Hence, the accuracy can be effectively controlled by the number of cells, *i.e.*, the sampling density. Unlike previous approaches our method is *symmetric* since we use dipoles for computing both magnetic force and field. This symmetry automatically ensures preservation of linear and angular momenta.

Adaptive Refinement We present an adaptive algorithm for computing magnetic properties based on a hierarchical multi-resolution sampling. The magnetic field decays very rapidly away from its source and so do force and torque acting on a distant object. In fact, the magnetic field of any magnet resembles a dipole field for large enough distances. We can therefore save computation time using fewer cells in regions where the magnetic field is weak and adapt cell resolution as a function of inter-object distance.

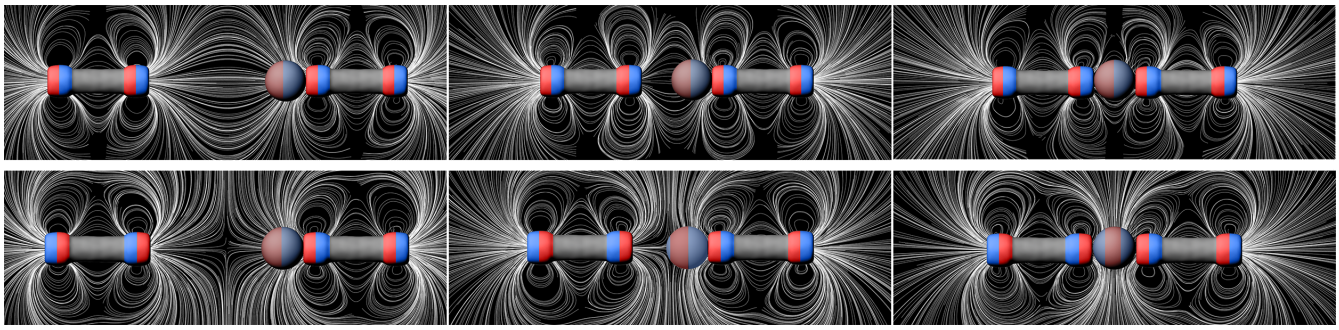

Figure 2: *Inhomogeneous induced magnetization.* Top*: toy magnets have identically oriented homogeneous magnetization (indicated by the coloring) and hence attract each other. The sphere is magnetized homogeneously according to the external field.* Bottom*: magnets have opposite magnetization and therefore repel each other. As the left toy magnet is moved rightwards its magnetic field induces a magnetization in the sphere (*middle*). Continuing this motion, the repelling forces increase but, at the same time, the attracting field induced in the sphere grows up to the point, where the resulting attraction forces exceed the repelling forces and the toy magnet snaps to the sphere (*right*).*

Magnetic Materials For our simulations, we physically describe a variety of different materials with very different magnetic properties and include them in our computation scheme. The materials covered comprise hard ferromagnets or permanent magnets, soft ferromagnets, paramagnets, diamagnets and even superconductors.

Validation We assess the accuracy of our method using a problem with known analytical solution and compare simulation results to real-world experiments. Additionally, we investigate the influence of the cell sampling density on approximation quality by comparison with exact results and within the scope of several simulations of real-world experiments.

Our method is very fast and, when coupled with a standard RBD solver, allows for interactive simulation of magnetic interactions for dozens of objects. As an example, toy magnets (see Fig, 1) can be used to explore magnetism in its different forms and to experiment with parameters like magnetization. In doing so, the user can recognize many familiar interactions between magnets but also experience magnetic effects, which do not necessarily follow intuition, see Fig. 2.

2 Related Work

Computing magnetic fields and forces on magnetized objects is a common problem in electric and mechanical engineering applications. Existing approaches include the virtual work method, Maxwell's stress tensor method and equivalent source methods. The latter category comprises the equivalent magnetic charge method, the equivalent magnetizing current method and the equivalent magnetic dipole method, which is the most similar to the approach presented in this work. These methods are mainly used for static analysis of magneto-mechanical systems within engineering tasks. Comparative studies of these methods can be found in [Delfino et al. 2001; de Medeiros et al. 1998]. Most of these methods resort in some way to finite element (FE) computations - at least for the computation of the external magnetic field. For interactive simulation, however, such an approach is unfeasible due to the associated numerical costs.

Terzopoulos et al. [1987] pioneered physical simulation in computer graphics. Foundations of rigid body dynamics were laid in [Hahn 1988] and [Baraff 1989] and subsequent developments have led to substantial progress in time stepping algorithms [Stewart and Trinkle 1996] and handling of contact [Baraff 1994; Guendelman et al. 2003; Pauly et al. 2004] and friction [Baraff 1991; Kaufman et al. 2005] to name just a few.

To the best of our knowledge, no attempt to model magnetic interaction between rigid bodies has been reported in computer graphics so far. Remotely related is the work of [Kim and Lin 2004], who model lightning and electrical arcs using a physically based approach. They use a simplified version of the Helmholtz equation for propagating electromagnetical waves. Our method bears some similarity to the Barnes-Hut algorithm [1986], which made the simulation of complex, self-gravitating N-body systems tractable in terms of computation time. Using a hierarchical approach based on spatial decomposition, O($NlogN$) run time complexity is achieved.

Finally, most of the concepts from magnetostatics used in this work are well-known and in-depth explanations can be found, e.g., in the textbooks [Landau et al. 1984; Jackson 1999]

3 Physical Modeling

To simulate the dynamic behavior of permanently as well as non-permanently magnetized rigid bodies, we must calculate the forces and torques acting between a large number of objects. To this end, we have to determine the total magnetic field in every time step.

3.1 Magnetostatics

Our approach is based on an approximate solution for the magnetic field of an object with a given magnetization, and, on a corresponding approximation for the calculation of the forces and torques. These approximations make use of a subdivision of objects into cells and their accuracy is effectively controllable by the number of cells. Both approximations lend themselves readily to an adaptive adjustment of the number of cells in order to assure the required level of accuracy (see 4.2). At the same time, the combination of the approximations is such that linear as well as angular momentum are preserved. For an infinite number of cells, the continuous limit is attained and the way to calculate the forces and torques used in our method coincides with the above-mentioned equivalent magnetic dipole method.

In order to calculate the total magnetic field, taking into account the magnetic properties of the different materials, an appropriate description has to be found for the two groups of materials considered in this work: those with permanent magnetization and those with a magnetization induced by an external magnetic field. The first group mainly comprises ferromagnets with a strong remanence, *i.e.*, residual magnetism. The second group covers diamagnets, paramagnets, ferromagnets with a weak remanence as well as superconductors. Since we assume the magnetized objects to be uncharged,

electric fields can be neglected. Accordingly, the interaction of the magnetized objects poses a magneto-static problem.

The relevant portion of Maxwell's equations for the description of magneto-static problems reads (in SI units)

$$\nabla \cdot \mathbf{B} = 0, \qquad \nabla \times \mathbf{H} = \mathbf{J}\,, \tag{1}$$

$$\mathbf{H} = \frac{1}{\mu_0}\mathbf{B} - \mathbf{M}\,, \tag{2}$$

together with a constitutive relation which links $\mathbf{B}$ and $\mathbf{H}$. In these equations, $\mathbf{B}$ denotes the magnetic flux density or magnetic induction, $\mathbf{H}$ the magnetic field, $\mathbf{J}$ the current density and $\mathbf{M}$ the magnetization or density of the magnetic moment.

For a given current distribution $\mathbf{J}$ we define the associated magnetic moment $\mathbf{m}$ as

$$\mathbf{m} = \frac{1}{2}\int_V d^3s(\mathbf{s}-\mathbf{r}_O) \times \mathbf{J}(\mathbf{s}-\mathbf{r}_O) \tag{3}$$

where $\mathbf{r}_O$ is an arbitrary point inside the volume V occupied by $\mathbf{J}$. We consider $\mathbf{m}$ to be positioned at $\mathbf{r}_O$. The magnetic induction can be obtained by a series expansion of the magnetic vector potential for distances $\mathbf{r}$ large compared to the spatial extent of the magnetized volume V [Jackson 1999]. The leading term of the magnetic induction for a magnetic moment $\mathbf{m}$ located at position $\mathbf{r}_O$ reads

$$\mathbf{B}(\mathbf{r}) = \frac{\mu_0}{4\pi}\left[\frac{3\mathbf{n}(\mathbf{n}\cdot\mathbf{m})-\mathbf{m}}{|\mathbf{r}-\mathbf{r}_O|^3}\right] \tag{4}$$

with $\mathbf{n} = (\mathbf{r}-\mathbf{r}_O)/|\mathbf{r}-\mathbf{r}_O|$. The distribution described by this expression is usually called the field of a dipole. For the special case of a sphere with homogeneous magnetization $\mathbf{M} = \frac{1}{2}[\mathbf{r}\times\mathbf{J}(\mathbf{r})]$, (4) represents the exact solution for all $\mathbf{r}$.

If a magnetized object is considered as an aggregate of magnetic dipoles, its total magnetic induction can be obtained by a linear superposition of the individual dipole fields. For $i = 1 \ldots N$ magnetic moments $\mathbf{m}_i$ located at positions $\mathbf{r}_i$

$$\mathbf{B}(\mathbf{r}) = \frac{\mu_0}{4\pi}\sum_{i=1}^{N}\left[\frac{3\mathbf{n}_i(\mathbf{n}_i\cdot\mathbf{m}_i)-\mathbf{m}_i}{|\mathbf{r}-\mathbf{r}_i|^3}\right] \tag{5}$$

with $\mathbf{n}_i = (\mathbf{r}-\mathbf{r}_i)/|\mathbf{r}-\mathbf{r}_i|$. If we assume the magnetization to be given, the exact solution for the total magnetic induction of the magnetized object is obtained in the continuous limit $N \to \infty$. However, for finite N, this superposition can be used to effectively calculate the magnetic induction with controllable accuracy. Assuming that an object with magnetization $\mathbf{M}$ and volume V is subdivided into N cells with a fraction of the volume V/N, each of the cells carries a magnetic moment $\mathbf{M}V/N$. The total magnetic induction of these N cells, calculated using (5), is valid for distances larger than the spatial extent of the cells, which is V/N. Accordingly, with growing number of cells, the range of validity is increased to smaller and smaller distances from the object. At the same time, the subdivision into cells together with the linear superposition (5) allows for an inhomogeneous magnetization as well as for complex geometries of the objects under consideration.

The convergence of the total magnetic induction with increasing number of cells is demonstrated in Fig. 3. For this figure, a cuboidal test object with a homogeneous magnetization has been used. The distribution of the magnetic induction rapidly converges to the final configuration, which is reached for a large number of cells. Between the subfigure for 10×5 cells and the one for 20×10 cells, changes are already very small. For more than 40×20 cells, hardly any changes of the magnetic induction are noticeable except for very short distances from the object. For distances in the range of the spatial extent of the object, a few cells are already sufficient for a very good accuracy.

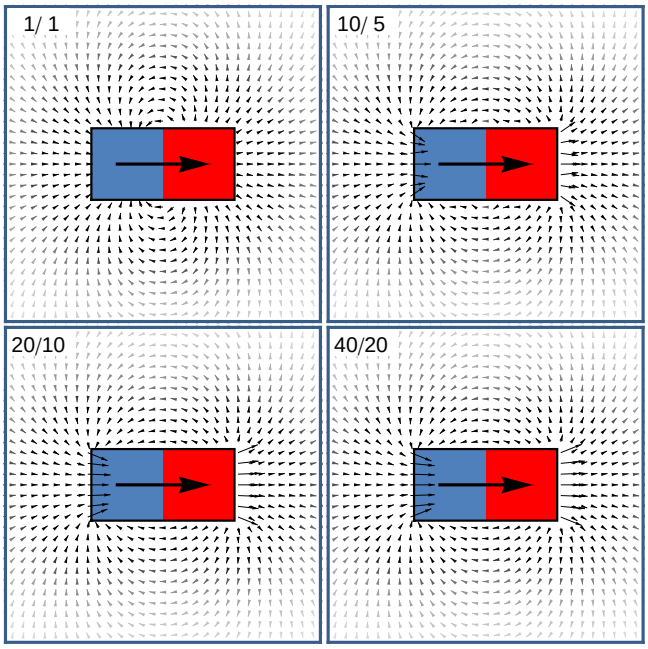

Figure 3: *Magnetic induction $\mathbf{B}(\mathbf{r})$ for an object with homogeneous magnetization. The number of cells in x/y direction used for the calculation is indicated in the upper left corner of each figure. The shade of the vectors is proportional to their magnitude.*

3.2 Forces and Torques

If a current distribution is subject to an external magnetic field, both forces and torques are exerted according to Ampère's law. The leading terms can be obtained by a Taylor series expansion of the magnetic induction around the center of the current distribution. The lowest non-vanishing terms in the expansion of the force and torque on the magnetic moment $\mathbf{m}$ are

$$\mathbf{F} = \nabla(\mathbf{m}\cdot\mathbf{B}) \tag{6}$$

$$\mathbf{T} = \mathbf{m}\times\mathbf{B} \tag{7}$$

The total force and torque acting on a magnetized object, considered as an aggregate of M magnetic dipoles cells, can now be determined: plugging the magnetic induction (5) into (6) and (7), we find explicit expressions for the resulting force $\mathbf{F}_k$ acting on the magnetic moment $\mathbf{m}_k$ of a single cell located at position $\mathbf{r}_k$, which is placed in the field of the N cells of another object with corresponding magnetic moments $\mathbf{m}_i$ located at positions $\mathbf{r}_i$,

$$\mathbf{F}_k = \frac{\mu_0}{4\pi}\sum_{i=1}^{N}\frac{1}{|\mathbf{r}_k-\mathbf{r}_i|^4}\Big[-15\mathbf{n}_{ik}\big((\mathbf{m}_k\cdot\mathbf{n}_{ik})(\mathbf{m}_i\cdot\mathbf{n}_{ik})\big) + 3\mathbf{n}_{ik}(\mathbf{m}_k\cdot\mathbf{m}_i) + 3\big(\mathbf{m}_k(\mathbf{m}_i\cdot\mathbf{n}_{ik}) + \mathbf{m}_i(\mathbf{m}_k\cdot\mathbf{n}_{ik})\big)\Big] \tag{8}$$

with $\mathbf{n}_{ik} = (\mathbf{r}_k-\mathbf{r}_i)/|\mathbf{r}_k-\mathbf{r}_i|$, and accordingly the torque $\mathbf{T}_k$,

$$\mathbf{T}_k = \frac{\mu_0}{4\pi}\sum_{i=1}^{N}\left[\frac{3(\mathbf{m}_k\times\mathbf{n}_{ik})(\mathbf{m}_i\cdot\mathbf{n}_{ik}) - \mathbf{m}_k\times\mathbf{m}_i}{|\mathbf{r}_k-\mathbf{r}_i|^3}\right]. \tag{9}$$

Note that with increasing distance, force decreases more rapidly than torque, by a factor of $1/|\mathbf{r}_k - \mathbf{r}_i|$. Thus, for large distances, the interaction is largely dominated by torque.

In order to obtain the total force $\mathbf{F}$ and torque $\mathbf{T}$ acting on the object, (8) and (9) have to be evaluated for all cells ($k = 1 \dots M$) and summed up appropriately. For the total force, the positions of the cells $\mathbf{r}_k$ relative to the center of mass of the object have to be respected in the summation. In the case of the total torque, we simply find $\mathbf{T} = \sum_{k=1}^{M} \mathbf{T}_k$ for rigid objects. In the continuous limit $M \to \infty$, the force and the torque coincide with the expressions of the equivalent magnetic dipole method and the exact solution is obtained. However, the calculation of the total force and torque gives accurate and reliable results even with a small number of cells per object. This robustness concerning mesh refinement has also been observed in [Delfino et al. 2001] in the context of their FE equivalent magnetic dipole method calculations.

3.3 Interpretation and Proof of Conservation

This section gives a detailed interpretation of all terms involved in (8) and (9) and also shows how to prove preservation of linear and angular momenta. The reader can safely skip this part when reading the paper for the first time.

In order to illustrate the contribution of the different terms of the expressions for the force (8) and torque (9) we investigate six exemplary arrangements of two bar magnets depicted in Fig. 4. For this illustration, the bar magnets are modeled as simple magnetic dipoles with a single cell. Only the force and torque terms acting on the second bar magnet are listed in the following; the ones acting on the first one follow from interchanging $\mathbf{m}_1$ and $\mathbf{m}_2$ and inverting $\mathbf{n}_{12} = -\mathbf{n}_{21}$.

In the first arrangement A1, the two bar magnets are aligned in parallel and positioned side by side. The first term of (8) vanishes since both $(\mathbf{m}_2 \cdot \mathbf{n}_{12})$ and $(\mathbf{m}_1 \cdot \mathbf{n}_{12})$ vanish. The second term of (8) (proportional to $3\mathbf{n}_{12}(\mathbf{m}_2 \cdot \mathbf{m}_1)$) contributes (yellow arrow in the sketch, labeled 2). The third term of (8) vanishes for the same reason as the first one. The first term of (9) vanishes since $(\mathbf{m}_1 \cdot \mathbf{n}_{12})$ vanishes, and the second term of (9) vanishes since $\mathbf{m}_2 \times \mathbf{m}_1 = 0$. Accordingly, the only relevant contribution is given by the second term of the expression for the force and the total force is repulsive.

In the second arrangement A2, we encounter the same situation as in A1, this time with an antiparallel alignment.Consequently, the total force, given by the second term of the expression (8), is attractive.

In the third arrangement B1, all three terms (labeled 1 to 3) contribute to the total force, which is attractive. The first term has the largest prefactor (see (8)) and dominates the total force. The situation in the fourth arrangement B2 is the same as in B1, but with $\mathbf{m}_2$ reversed. Accordingly, all the contributions to the total force change sign and the total force is repulsive. In both B1 and B2, the torque vanishes since $(\mathbf{m}_2 \times \mathbf{n}_{12}) = 0$ and $(\mathbf{m}_2 \times \mathbf{m}_1) = 0$.

The torque comes into play in the last two arrangements C1 and C2. Here, only the third term of the force (8) contributes. The first term of the torque (9), which is proportional to $3(\mathbf{m}_2 \times \mathbf{n}_{12})(\mathbf{m}_1 \cdot \mathbf{n}_{12})$, acts only on $\mathbf{m}_1$ (labeled 4). The second term of (9) contributes for both $\mathbf{m}_1$ and $\mathbf{m}_2$ and is proportional to $-(\mathbf{m}_2 \times \mathbf{m}_1)$ (labeled 5). The reversion of $\mathbf{m}_2$ from C1 to C2 leads to a changing of the sign of all contributions. In C1, the total torque tends to align the two bar magnets as in B1. In C2, the total torque also tends to align the two bar magnets, but in the opposite direction.

From these illustrations, we see that the sum of all forces vanishes in all arrangements shown here. For the sum of all torques, this is less obvious. The general case will be discussed subsequently.

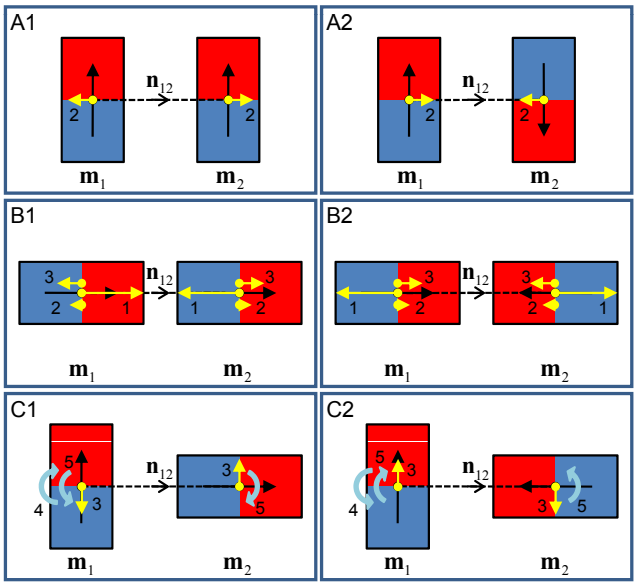

Figure 4: *Six exemplary arrangements of two bar magnets with forces and torques depicted by yellow and blue arrows. Black arrows indicate the magnetic dipole moments $\mathbf{m}_1$ and $\mathbf{m}_2$, "north" and "south" poles of the magnets are shown in red and blue.*

It should be stressed that the approximations for the magnetic induction, forces and torques are formulated in such a way that both linear and angular momentum are preserved. Conservation of momentum is an inherent property of the physical system, but it is not obvious that an approximate description preserves momentum as well. To prove the conservation, it has to be shown that both the total force and the total torque (summed over all cells of all objects) vanish. For the total force, this can easily be seen from (8) by interchanging $\mathbf{m}_k$ and $\mathbf{m}_i$ as well as their positions. Then, the force acting on $\mathbf{m}_k$ is exactly the negative of the force acting on $\mathbf{m}_i$ and the sum of the two vanishes. For the total torque, the derivation involves both forces and torques. The third term of the force (8) as well as the torque (9) contribute to the total torque. A correct combination of the two via $\mathbf{R} \times \mathbf{F} + \mathbf{T}$ with the center of mass $\mathbf{R}$ leads to a vanishing total torque.

3.4 Permanently Magnetized Objects

We will now turn to the physical description of magnetic materials. Permanent magnets, the first group of materials covered in the present work, are usually ferromagnets with a strong remanence or residual magnetism. More precisely, their constitutive relation between $\mathbf{B}$ and $\mathbf{H}$ is given by a hysteresis curve $\mathbf{B}[\mathbf{H}]$ which exhibits a large remanent value of $\mathbf{B}$ for vanishing $\mathbf{H}$. This implies that these materials feature a strong permanent magnetization even after the removal of the (strong) magnetic field used in order to magnetize them. As a second implication, the magnetic field of other magnetic objects, usually much weaker than the field used for the excitation of the permanent magnetization, does not change their magnetization considerably. In the present work, we assume the magnetization of the permanently magnetized objects to be constant (*hard ferromagnets*).

3.5 Induced Magnetization

The second group of magnetic materials exhibits a magnetization only in the presence of an external magnetic field. This group includes diamagnets, paramagnets as well as ferromagnets with a

weak remanence. Superconductors will also be treated within the scope of this group.

In the case of isotropic diamagnets and paramagnets, the linear constitutive relation

$$\mathbf{B} = \mu_0(1+\chi)\mathbf{H} \tag{10}$$

with the magnetic susceptibility of the material χ holds. In the case of a diamagnet, we have $\chi < 0$, whereas in the case of a paramagnet, $\chi > 0$. For the most common diamagnetic and paramagnetic materials, however, the absolute value of the susceptibility is of the order of 10^{-6} and the resulting magnetization as well as the forces and torques are very weak.

As mentioned above, ferromagnets have to be described by a hysteresis curve connecting $\mathbf{B}$ and $\mathbf{H}$. Nevertheless, for the special case of ferromagnets with a weak remanence (*soft ferromagnets*) subject to weak external magnetic fields, a linear approximation in the form of (10) with $\chi > 0$ is adequate. In the case of soft ferromagnets, the magnetic susceptibility varies by several orders of magnitude. Steel, for example, exhibits a value of about $\chi \simeq 700$ whereas $\chi \simeq 20,000$ for mu-metal, which is a is a nickel-iron alloy.

In general, the magneto-static Maxwell equations (1) and (2) have to be solved together with the constitutive relation, *i.e.*, (10) in the present case. For arbitrary geometries, this is a very complicated task because of the boundary conditions that have to be respected. Computationally most efficient are explicit analytic solutions, but they can only be obtained for highly symmetric situations. Fortunately, if we subdivide the objects with an induced magnetization into cells and approximately assume the cells to be spherical, an explicit solution is accessible. Additionally neglecting the coupling between the cells, we obtain an explicit expression for the induced magnetization $\mathbf{M}_i$ of a cell located at position $\mathbf{r}_i$:

$$\mathbf{M}_i = \frac{3}{\mu_0}\frac{\chi}{1+\chi}\mathbf{B}(\mathbf{r}_i)\,. \tag{11}$$

In the present work, we use (5) to compute the magnetic induction due to permanently magnetized objects. Having computed this first contribution, we can then determine the induced magnetization of diamagnets, paramagnets and soft ferromagnets using (11) to obtain the final magnetic induction.

From (11), it is obvious that the sign of the magnetic susceptibility χ determines the general behavior of the material. In the case of a positive (negative) value, the induced magnetization is oriented parallel (anti-parallel) to the magnetic induction. Accordingly, a paramagnet is attracted by a permanent magnet, whereas a diamagnet is repelled. Soft ferromagnets with $\chi > 0$ behave like paramagnets, but the resulting force and torque are much stronger due to the usually much higher value of χ.

Superconductors expel any external magnetic field from their interior – a phenomenon called Meißner-Ochsenfeld effect. The expulsion of the external magnetic field is achieved by screening currents flowing on the surface of the superconducting material. In the interior of the superconductor, these screening currents generate a magnetic field which exactly balances the external one. Thus, in a very simple physical picture, the magnetic induction vanishes completely inside the superconductor, $\mathbf{B} = 0$. Combining the vanishing magnetic induction with (2), we find $\mathbf{M} = -\mathbf{H}$ inside the superconductor (*perfect diamagnetism*). If we now approximate the magnetic field $\mathbf{H}$ by the vacuum field, we have $\mathbf{H} = \mathbf{B}/\mu_0$ and thus

$$\mathbf{M}_i = -\frac{1}{\mu_0}\mathbf{B}(\mathbf{r}_i) \tag{12}$$

for every cell of the superconductor. Because of the strong induced magnetization which is oriented anti-parallel to the magnetic induction, superconductors are repelled from permanent magnets.

4 Implementation

This section describes some implementational aspects of our method including the computation of magnetic force and torque as well as the adaptive cell hierarchy.

4.1 Algorithm

Algorithm 1 provides an overview of the steps involved in the computation of magnetic force and torque. The magnetization of induced magnets changes according to the magnetic field in their surrounding and has to be recomputed in every time step. This amounts to evaluating the magnetic field of all permanent magnets in the scene at the induced magnets' cell positions (l.3). The following code then computes the magnetic interaction between all magnets in the scene. Note that the bound for the second loop (l.7) is different due to the symmetry of magnetic interaction. For every pair of magnets, we compute force and torque exerted by all cells of one object onto all cells of the other object and vice versa according to (8) and (9) (ll.11-12).

Algorithm 1 Magnetic force and torque computation

```
 1: //Compute magnetization for induced magnets
 2: for all induced magnets do
 3:     computeInducedMagnetization();
 4: end for
 5: //Compute magnetic force and torque
 6: for i = 1 to n_mag do
 7:     for j = i to n_mag do
 8:         for k = 1 to n_cells,i do
 9:             for l = 1 to n_cells,j do
10:                 c_k=cell(k), c_l = cell(l);
11:                 applyForce(i,c_k,c_l), applyTorque(i,c_k, c_l);
12:                 applyForce(j,c_l,c_k), applyTorque(j,c_l, c_k);
13:             end for
14:         end for
15:     end for
16: end for
```

Once forces and torques between two dipoles are computed, they can either be fed directly into the RBD solver (ll.11-12) or be accumulated first. The latter is typically more efficient, especially in a parallel implementation where thread-safety of the RBD code might be an issue. The actual parallelization is simple and efficient since all operations for one iteration are independent of other iterations.

For reasonable magnetization densities, the forces can be safely time stepped using explicit integration, since the distance between two dipole cells cannot fall below one cell diameter and the force is thus limited. However, we cannot rely on the collision code to always guarantee an intersection-free state such that cells may in fact become arbitrarily close. This problem can be dealt with by limiting force and torque to the maximum values corresponding to the closest possible distance between two cells without intersections.

Algorithm 1 has quadratic complexity with respect to the total number of cells in the scene, just as the costs for collision detection scales quadratically with the number of primitives. An adaptive approach is therefore mandatory in order to guarantee sufficient resolution where necessary and yet remain computationally efficient.

Fortunately, we can exploit the controllable accuracy of magnetic induction (5), force (8) and torque (9). For this purpose, we use an adaptive multi-resolution cell hierarchy as described next.

4.2 Adaptive Refinement

As indicated in Section 3, a higher cell density increases the approximation quality of field, force and torque. For complex scenes, however, using a uniformly high cell resolution is prohibitively expensive and we therefore need to decrease the average computation time. The following observation is key to our adaptive strategy: the magnetic field $\mathbf{B}(r)$ of a dipole is proportional to $1/r^3$, r being the distance to the dipole. If we prescribe a sufficient ratio between magnetic field $\mathbf{B}$ and spatial cell density ρ as $\mathbf{B}/\rho = c_0 = const.$ we obtain a means for adapting the number of cells in a region of space to the local magnetic field. The remainder of this section describes the adaptive framework that builds upon this measure.

Cell Hierarchy The basis for the adaptive refinement scheme is formed by a multi-resolution hierarchy of cells, which is similar to the octree stucture used in [Barnes and Hut 1986]. However, instead of partitioning space we construct hierarchies on a per-object basis. The hierarchy for an object is created starting with a single cell at level zero, corresponding to its bounding box. A cell itself has only volume and no shape but we associate with it the bounding box of the geometry that it represents. The cells on level $i+1$ are then obtained through bisection of the cells on level i along their longest axis. We prefer bisection over octasection as it leads to slower growth in the number of cells per level. If we find that a newly created cell lies completely outside the object, it is immediately discarded. The process stops when the cell density of the finest level reaches a prescribed value ρ_0. This value can be determined using, *e.g.*, Fig. 6 which indicates that for an inter-object distance of 0.5 cm it is sufficient to have eight cells per cubic centimeter. Subsequently, we eliminate all those leaf cells whose positions (*i.e.*, their mid-points) lie outside the object or which are closer to the object boundary than half their diameter. This process is applied to each level in order to remove all empty cells without children. Finally, we rescale all cell magnetizations (using the same factor) such that their sum matches the total magnetization of the original object. This step is necessary to guarantee that force and torque computations are consistent among the levels: for sufficiently large distances, a dipole will thus experience the same force and torque from any level of a magnetized object.

Cell Front As mentioned in Section 4.1, force, torque and field are always computed for pairs of magnets. Given such a pair, we have to select a resolution level based on the geometric distance between the two objects. Using a single hierarchy level would mean fixing cell density throughout the entire objects, which is very inefficient if only small regions are close, but large parts are far away from each other (see Fig. 5, bottom). Instead we compute an adapted cell sampling or *cell front* for each of the two objects as follows: starting with the top level cell of one object, we first determine the distance that would be sufficiently large for the current cell density $\rho_c = 1/v_c$, with v_c being the current cell's volume. By construction we have $\rho_0/\rho_c = r_c^3/r_0^3$ and, hence, the sufficiently large distance is $r_c = \sqrt[3]{\frac{\rho_0}{\rho_c}} r_0$. We then check whether the geometric distance d of the cell to the other object is smaller than r_c. To this end, we first check the corresponding bounding spheres for intersection. If they intersect, d is computed using a GJK-test [Gilbert et al. 1988] on the oriented bounding boxes. The current cell density is sufficient if $d >= r$. Otherwise, we apply the same process recursively to all children of the current cell until a sufficient density is reached. The result of this process is a cell front for every object, which reflects the continuous geometric distance to the other object (see Fig. 5). Note that the hierarchy needs only be constructed once at the beginning of the simulation. Additionally, algorithm 1 requires only small changes to ll.8-9 in order to support adaptive refinement.

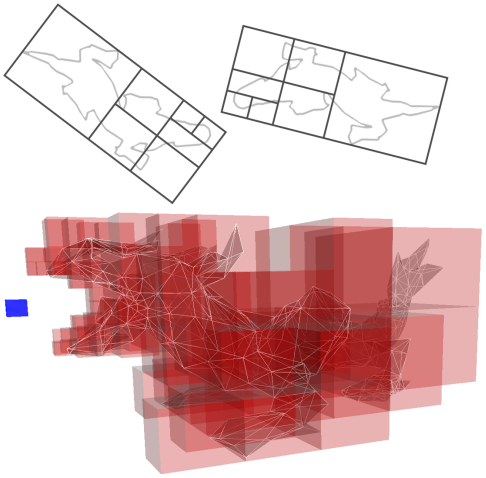

Figure 5: Top*: schematic view of the proximity-based adaptive refinement.* Bottom*: cell front obtained for a dragon model with respect to the blue box.*

5 Validation

Analytical Comparison In order to verify its validity, we apply our method to a test configuration consisting of two cube-shaped ferromagnets with volume $V = (1\text{cm})^3$ exhibiting a homogeneous remanent magnetic induction $\mathbf{B}_r$ of 1 Tesla (see Fig. 6). Using (2), the corresponding magnetization is found to be $\mathbf{M} = \mathbf{B}_r/\mu_0$. This test configuration has been used in [de Medeiros et al. 1998] for the comparison of several of the aforementioned alternative force calculation methods. An analytical solution for this case has been presented in [Akoun and Yonnet 1984]. Since this solution is based on equivalent magnetic charges, which are located on the surface, it is valid only for finite distances $d > 0$. By contrast, our method is also valid for the contact case, *i.e.*, $d = 0$, since the magnetic dipoles cells are always located in the interior of the objects.

The results shown in Fig. 6 were obtained using N^3 uniformly distributed cells for each cube. For the fixed distance of $d = 0.5\,\text{cm}$ used in [de Medeiros et al. 1998], our method reproduces the correct result for the attractive force already with a very low number of cells. For lower distances, the results of our method converge to the exact solution as the number of cells is increased.

The results for the total force as given in Fig. 6 further emphasize the usefulness of adaptivity: for close proximity a high cell resolution is mandatory in order to obtain good accuracy. However, the number of cells necessary for maintaining good approximation quality rapidly decreases with growing inter-object distance.

Momentum Conservation The conservation of linear and angular momentum is an important property of the symmetric equivalent magnetic dipole approach presented in this work. Additionally, it also provides a simple way to check the implementation. For testing purposes, linear and angular momentum are readily computed during the simulation. If the implementation is correct, the sum of both quantities over all objects has to be time-invariant.

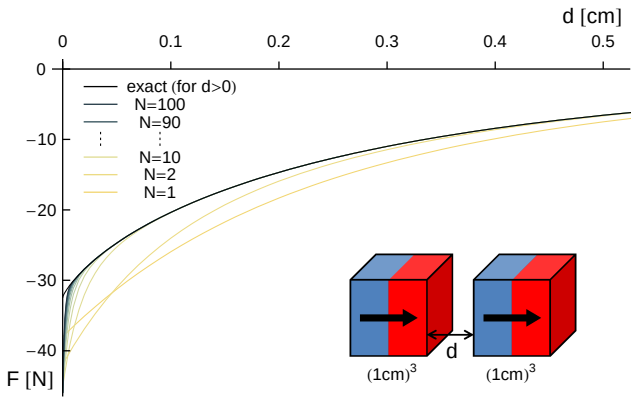

Figure 6: *Comparison of the total force calculated using our method with exact results for a test configuration consisting of two ferromagnetic cubes with homogeneous magnetization. With increasing number of cells per object, (given by N^3 with N as indicated), the results for the force converge to the exact solution.*

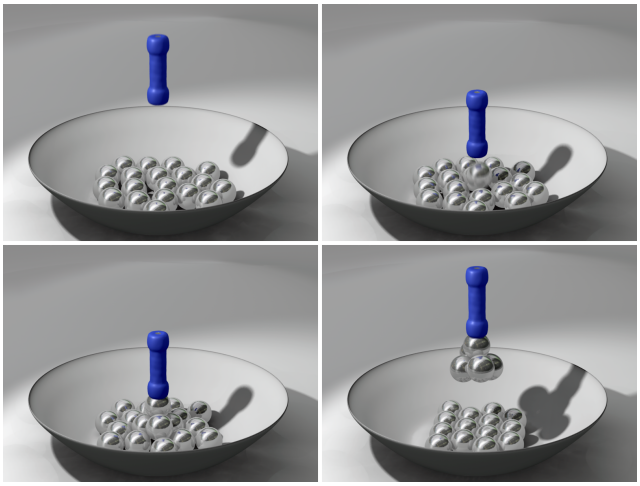

Figure 7: *Snapshots of an interactive simulation in which soft-ferromagnetic spheres are lifted by a permanent magnet.*

We verified this for our code on a number of simulations and obtained very good conservation behavior for the frictionless and non-contact case. During and after contact situations, the conservation deteriorates even for completely elastic collisions. However, this is a peculiarity of the RBD solver and not a defect of our method.

6 Results

This section presents some exemplary results obtained with our method. For our implementation we used the freely available RBD library ODE [Smith 2006].

Many of our examples use *toy magnets*, consisting of a plastic-coated cylinder with two small permanent magnets at the ends (see Figs. 1, 2, 7 and accompanying video). Simulated toy magnets can easily be compared to their real-world counterparts using simple experiments and offer a great potential for fascinating magnetic interaction. The magnetic field of a toy magnet is visualized in Fig. 1.

Examples The example shown in Fig. 2 demonstrates the effect of inhomogeneous induced magnetization. A magnetized sphere is held in contact with a toy magnet, while a second toy magnet is placed at some distance. In the first case, the toy magnets have the same magnetization and are therefore attracted uniformly to each other (Fig. 2, top). Similarly, the sphere is magnetized homogeneously (identical direction) according to the external field and attracts the left toy magnet as well. In the second and more interesting case, the toy magnets have opposite magnetizations leading to a point of zero resulting magnetic field, clearly distinguishable in the left-most figure of the bottom row. Initially, the sphere assumes the magnetization of the right toy magnet, to which it is attached. As the left toy magnet is gradually moved to the right, its magnetic field starts to induce a corresponding magnetization in the sphere, resulting in a directionally inhomogeneous magnetization. However, the left toy magnet is still repelled from the right one. Forcing it to move further rightwards, the repelling forces increase but, at the same time, the attracting field induced in the sphere grows up to the point, where the resulting attraction forces exceed the repelling forces and the toy magnet snaps to the sphere. This fascinating effect, which can be thought of as passing a potential barrier, is clearly perceptible in reality and faithfully reproduced by our method. Note that, although the final geometries of the top and bottom sequences are identical, the resulting fields exhibit subtle local differences.

The next test case shows another example for complex behavior obtained with multiple induced magnets. Fig. 7 shows a sequence of four images taken from a simulation of a permanent magnet lifting several soft-ferromagnetic spheres out of a bowl. As the permanent magnet is moved downwards it starts to induce a magnetization in the spheres. At a certain distance, one of the spheres is picked up by the permanent magnet (see upper two images of Fig. 7). Subsequently, the other spheres are further drawn towards the magnet and are lifted as well. Finally, five spheres are attached to the permanent magnet and lifted above the bowl (lower two figures). In the final position, only one of the spheres is directly attached to the permanent magnet whereas the other four spheres are hanging below the first one forming a symmetric configuration. As can be seen in the accompanying video the same behavior can be observed in a real-world experiment.

In two further examples we demonstrate that our method is capable of handling a large number of objects as well as magnets with arbitrary non-convex geometry. Fig. 8 shows a sequence of images from a simulation in which a rectangular permanent magnet lifts four soft-ferromagnetic characters out of a pool of non-magnetic spheres.

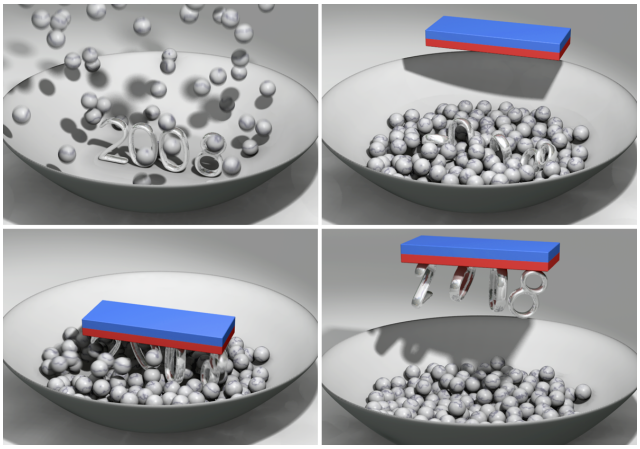

Figure 8: *A strong permanent magnet lifts four soft-ferromagnetic characters out of a pool of non-magnetic spheres.*

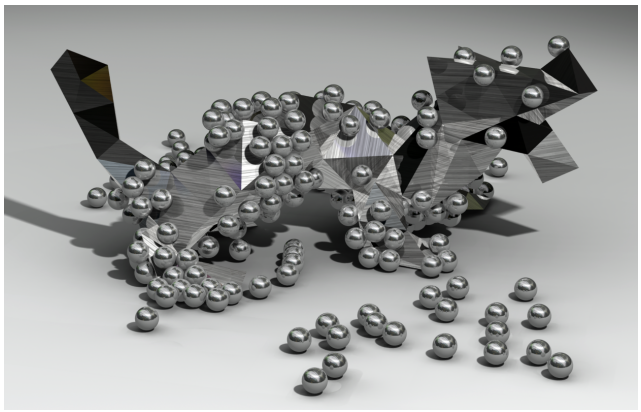

Figure 9: *Simulation scene consisting of a permanently magnetized simple dragon model and a large number of soft-ferromagnetic spheres, 250 in the complete simulation. The spheres are initially dropped into the scene from above the dragon and some of them stick to the surface of the dragon, held by magnetic interaction.*

Fig. 9 shows an image of a scene in which 250 soft-ferromagnetic spheres are dropped onto a simple dragon model. As soon as they approach the dragon, a magnetization is induced in the spheres and they start to interact magnetically with each other and the dragon.

If a superconductor is placed in the vicinity of a hard ferromagnet, the induced magnetization within the superconductor is oriented anti-parallel to the field of the ferromagnet. According to (12), the induced magnetization exactly balances the external field and can thus be of considerable strength. With a sufficiently strong magnetization of the ferromagnet the repelling forces on the superconductor can be large enough to overcome gravity. The simulation shown in Fig. 10 consists of a ferromagnetic ring and a superconducting cube levitating above it. The toroidal shape of the ferromagnet leads to an equilibrium position above the center of the ring for the superconducting cube. As can be seen in the right hand side image of Fig. 10, the ferromagnetic ring is magnetized in the vertical direction. In two points above and below the center of the ring, the magnetic field vanishes. Since the superconductor is repelled from the field, positions close to these two points are energetically favorable. Once put into the upper of the two, the superconducting cube remains in this stable position. From the right hand side image of Fig. 10, one can see that the field lines are repelled from the interior of the superconductor and literally flow around it. This is the Meißner-Ochsenfeld effect mentioned above.

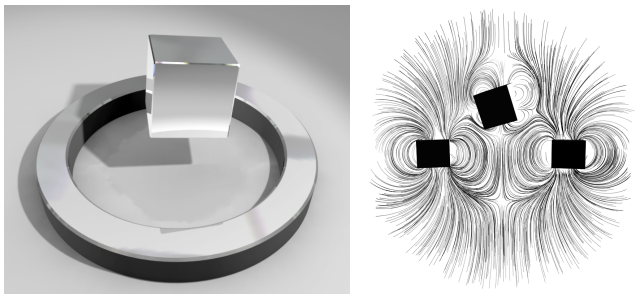

Figure 10: *A superconducting cube levitating above a ferromagnetic ring. Simulation scene on the left and a snapshot of the magnetic field lines on the right.*

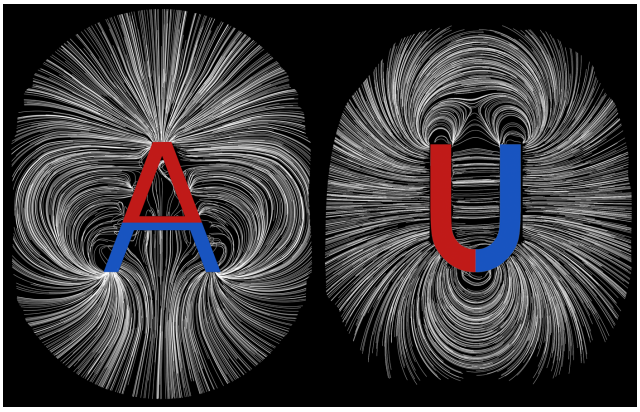

Figure 11: *Our method can be used to compute detailed visualizations of magnetic fields using standard streamline techniques.*

Apart from dynamic simulation, our method can also be used as a tool for computing magnetic fields for arbitrary objects. Using this technique, we can visually discover such fascinating aspects of magnetic fields as, *e.g.*, the relation between field structure and geometric shape (see Fig. 11).

6.1 Performance

The computational impact of our method is determined by the number of cells used to compute magnetic force and torque. Compared to the worst case performance, which scales quadratically with the number of cells, the average run time is greatly reduced using the adaptive sampling described in section 4.2. The benefit of adaptivity depends, however, on the scene under consideration and is generally not constant over time. For the examples described above, the average reduction in computation time due to adaptivity ranges from 1.3 for the scene described in Fig. 2 to 11.1 for the simulation shown in Fig. 9.

Table 1 summarizes computation times for the examples used in this work, showing total computation times and the time spent for computing magnetic forces and torques. To allow easier interpretation we report the computation times without any parallelization, although many parts of the algorithm lend themselves to easy parallelization strategies and our first results showed nearly optimal speedup on a workstation with four CPUs. As expected, computing the magnetic interaction takes up most of the simulation time. Depending on the nature of the individual scenes, however, the rigid body dynamics and especially the collision handling can also require a considerable amount of the total computation time. Still, simulations with dozens of magnets and up to several hundreds of cells run at interactive rates.

7 Conclusion

Limitations Our treatment of induced magnetization is approximate as we do not account for the fact that induced fields act back on the external field. The correct treatment of the problem requires a solution of the Maxwell equations including boundary conditions. A solution could be to reiterate the magnetic field computation, including induced fields, until convergence. Our first experiments in this direction indicated that additional effort is necessary to ensure convergence, which seems very sensitive to the strength of the magnetic field as well as the magnetic susceptibility of the material. Also, the dipole approximation is a possible point of improvement when considering the mutual influence of the induced magnets.

Example	#cells	T_{mag}	T_{rbd}	T_{tot}	Gain	fps
TST	48	0.29	0.09	0.38	1.30	65
BowlSpheres	256	3.9	0.25	4.16	3.76	6
Bowl2008	128	1.65	2.39	4.05	1.97	6.2
SuperCube	142	0.33	0.02	0.35	5.97	71.4
DragonLow	1824	53.1	2.76	55.85	6.72	0.45
DragonHigh	3628	57.6	2.92	60.5	11.1	0.41

Table 1: *Average computation times in milliseconds for one time-step of $1ms$ on an Intel Core2Duo 2.4GHz CPU (using only one core). Computation times for magnetic forces (T_{mag}) and time spent in the RBD code (T_{rbd}) are listed separately.* #cells *refers to the total number of cells in the scene and* Gain *denotes the speed-up for magnetic force computations due to adaptivity, which was used in all examples.* TST *is the scene described in Fig. 2 with two toy magnets and a soft-ferromagnetic sphere.* BowlSpheres *is the simulation shown in Fig. 7.* Bowl2008 *is shown in Fig. 8,* SuperCube *refers to the simulation of a superconducting cube shown in Fig. 10.* DragonLow *and* DragonHigh *are simulations of a magnetic dragon model and 250 soft-ferromagnetic spheres (Fig. 9) using different cell resolutions.*

Our method preserves linear and angular momenta, but this does not imply energy conservation for which a theoretical analysis is difficult because of the energy accumulated in the magnetic field. The rigid body solver used in our implementation does not preserve energy, which is why we did not try to assess this behavior experimentally. It would therefore be interesting to carry out further tests using, *e.g.*, symplectic solvers [Hairer et al. 2006], which are known for their excellent energy preservation behavior.

Future Work Since our approach exhibits good approximation quality and robustness already with a low cell density, it could be used as a fast preview stage to heavy-weight FEM simulations in engineering applications. In this regard, it could be a very interesting extension to include DC or even AC currents as sources of magnetic field and to include force and torque acting on the same currents. This would for example allow to model a virtual electric motor and simulate the interplay of electric, mechanical and magnetic parts.

Finally, with hindsight to educational applications, the integration of a haptic interface would further increase the usefulness of our method.

Acknowledgements

We would like to thank Markus Huber for helping us with the examples and Christian Grünzweig for providing us with literature. Furthermore, we thank Eitan Grinspun for valuable comments and for helping us to revise and restructure the exposition. Finally, we thank the reviewers for many helpful comments and ideas.

References

AKOUN, G., AND YONNET, J.-P. 1984. 3d analytical calculation of the forces exerted between two cuboidal magnets. *IEEE Trans. Magn. 20*, 5, 1962–1964.

BARAFF, D. 1989. Analytical methods for dynamic simulation of non-penetrating rigid bodies. In *SIGGRAPH '89*, ACM, 223–232.

BARAFF, D. 1991. Coping with friction for non-penetrating rigid body simulation. In *SIGGRAPH '91*, ACM, 31–41.

BARAFF, D. 1994. Fast contact force computation for nonpenetrating rigid bodies. In *SIGGRAPH '94*, ACM, 23–34.

BARNES, J., AND HUT, P. 1986. A hierarchical O(N log N) force-calculation algorithm. *Nature 324*, 4, 446–449.

DE MEDEIROS, L. H., REYNE, G., AND MEUNIER, G. 1998. Comparison of global force calculations on permanent magnets. *IEEE Trans. Magn. 34*, 5, 3560–3563.

DELFINO, F., MANELLA, A., MOLFINO, P., AND ROSSI, M. 2001. Numerical calculation of total force upon permanent magnets using equivalent source methods. *COMPEL: The International Journal for Computation and Mathematics in Electrical and Electronic Engineering 20*, 2, 431–447.

GILBERT, E. G., JOHNSON, D. W., AND KEERTHI, S. S. 1988. A fast procedure for computing the distance between complex objects in three-dimensional space. *IEEE Journal of Robotics and Automation 4*, 2, 193–203.

GUENDELMAN, E., BRIDSON, R., AND FEDKIW, R. 2003. Nonconvex rigid bodies with stacking. In *SIGGRAPH '03*, ACM, 871–878.

HAHN, J. K. 1988. Realistic animation of rigid bodies. In *SIGGRAPH '88*, ACM, 299–308.

HAIRER, E., LUBICH, C., AND WANNER, G. 2006. *Geometric numerical integration*, 5th ed. Springer-Verlag, Berlin.

JACKSON, J. D. 1999. *Classical Electrodynamics*, 3rd ed. Wiley, New York.

KAUFMAN, D. M., EDMUNDS, T., AND PAI, D. K. 2005. Fast frictional dynamics for rigid bodies. In *SIGGRAPH '05*, ACM, 946–956.

KIM, T., AND LIN, M. C. 2004. Physically based animation and rendering of lightning. In *PG '04: Proceedings of the Computer Graphics and Applications, 12th Pacific Conference*, IEEE Computer Society, 267–275.

LANDAU, L. D., LIFSHITZ, E. M., AND PITAEVSKII, L. P. . 1984. *Electrodynamics of continuous media*, 2nd ed. Pergamon Press, Oxford.

PAULY, M., PAI, D. K., AND GUIBAS, L. J. 2004. Quasi-rigid objects in contact. In *SCA '04: Proceedings of the 2004 ACM SIGGRAPH/Eurographics symposium on Computer animation*, Eurographics Association, 109–119.

SMITH, R., 2006. Open dynamics engine (ODE). `http://www.ode.org`.

STEWART, D., AND TRINKLE, J. 1996. An implicit time-stepping scheme for rigid body dynamics with inelastic collisions and coulomb friction. *International Journal of Numerical Methods in Engineering 39*, 2673–2691.

TERZOPOULOS, D., PLATT, J., BARR, A., AND FLEISCHER, K. 1987. Elastically deformable models. In *SIGGRAPH '87*, ACM, 205–214.