DOI: 10.1111/cgf.15185
ACM SIGGRAPH / Eurographics Symposium on Computer Animation 2024
M. Skouras and H. Wang
(Guest Editors)

COMPUTER GRAPHICS *forum*
Volume 43 (2024), Number 8

Strongly Coupled Simulation of Magnetic Rigid Bodies

L. Westhofen, J. A. Fernández-Fernández, S. R. Jeske & J. Bender

RWTH Aachen University, Germany

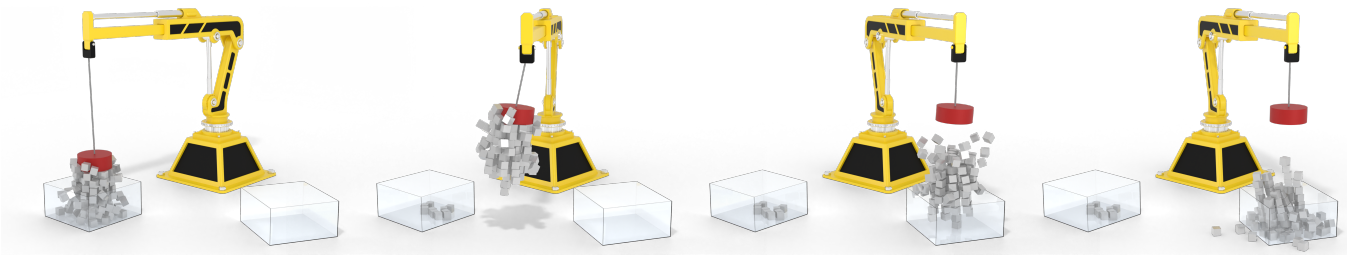

Figure 1: *Simulation sequence of a crane with an electromagnet picking up* 125 *nickel cubes and moving them from one box to another. The crane is simulated as a rigid body system connected with joints and actuated by motors, the electromagnet is attached with a rope to the crane. The whole simulation system is strongly coupled, that is, the movement and handling of rigid bodies, constraints and magnetic effects are handled together, leading to a robust and easy to setup simulation without any bouncing artifacts with an average time step size of* $\Delta t = 1.92\,\text{ms}$.

Abstract

We present a strongly coupled method for the robust simulation of linear magnetic rigid bodies. Our approach describes the magnetic effects as part of an incremental potential function. This potential is inserted into the reformulation of the equations of motion for rigid bodies as an optimization problem. For handling collision and friction, we lean on the Incremental Potential Contact (IPC) method. Furthermore, we provide a novel, hybrid explicit / implicit time integration scheme for the magnetic potential based on a distance criterion. This reduces the fill-in of the energy Hessian in cases where the change in magnetic potential energy is small, leading to a simulation speedup without compromising the stability of the system. The resulting system yields a strongly coupled method for the robust simulation of magnetic effects. We showcase the robustness in theory by analyzing the behavior of the magnetic attraction against the contact resolution. Furthermore, we display stability in practice by simulating exceedingly strong and arbitrarily shaped magnets. The results are free of artifacts like bouncing for time step sizes larger than with the equivalent weakly coupled approach. Finally, we showcase the utility of our method in different scenarios with complex joints and numerous magnets.

CCS Concepts

• ***Computing methodologies*** → ***Physical simulation;***

1. Introduction

Including magnetic forces in rigid-body simulation enables a large spectrum of interesting effects. These range from simulating simple permanent magnets sticking notes on a whiteboard across intricate mechanisms seemingly moving on their own to electromagnetic-powered machinery, to name a few. Studying magnetic effects remains an interesting and challenging topic in computer graphics since they exhibit a highly unique motion given their governing equations. Depending on the orientation of magnets and the material properties, magnetic materials either attract or repel each other. The strength of these forces is strongly dependent on the distance between magnetic objects, that is, the closer two magnetic bodies are, the stronger the force becomes.

In computer graphics, magnetic effects have been studied in the context of rigid-body simulations [TGPS08], the simulation of ferromagnets [KPH18; KH20], magnetic deformables like magnetic silly putty [SNZ*21], thin shells [CNZ*22], or ferrofluids [IYI*12b; HHM19]. While previous approaches have already

shown the intricacy of magnetic phenomena, there are some unique properties, especially of the magnetic force that should be considered when simulating magnetic objects. Primarily, the magnetic forces drastically increase in strength as the distance between two magnetic objects approaches zero. Thus, these forces are very sensitive to small displacements in close proximity to other objects. Additionally, the commonly used magnetic force approximation [TGPS08; KPH18] is unbounded in strength, resulting in infinite forces when the distance between the magnetized objects is zero. Both properties challenge the time discretization of magnetic simulations. Accounting for the magnetic forces explicitly, that is, computing them at the beginning of the time step and keeping them fixed, will lead to inaccuracies at best, and unrecoverable, near-infinite forces at worst. Secondly, the magnetic attraction forces directly oppose the direction of contact forces. Thus, it must be ensured by the contact solver that magnetic objects never reach a configuration leading to near-infinite forces while at best providing an interpenetration-free simulation state.

As a result, providing external magnetic forces to an off-the-shelf rigid body solver may result in the simulation showing unstable behavior like bouncing or explosions. Both can be attributed to the fact that by iterating between external magnetic forces and collision resolution, the solver may fail to find an equilibrium of forces. This approach is known as weakly coupling the magnetic solver with the rigid body solver. Further, this may result in lower computational performance of the system, if we need to increase the temporal resolution to reduce the effects of jittering or to correctly handle strong magnetic attractions. For complex and large simulations featuring many strong magnets and complex collisions, the weakly coupled approach will struggle to remain stable, even for small time steps.

We propose a strongly coupled method for handling linear magnetic effects to solve these issues. To this end, we reformulate the equations of motion as a minimization problem, turning the magnetic effects into an Incremental Potential (IP) formulation. Thus, we present a magnetic potential formulation that is guaranteed to find a stable equilibrium. The resulting method ensures an interpenetration-free and robust simulation of magnetic bodies even in the presence of strong magnets and challenging geometries. In addition, we present a novel, hybrid time integration scheme for magnetic forces based on a hybrid explicit / implicit scheme. Thus, the implicit time integration of the magnetic forces is only used between object pairs close enough to be in the sensitive region of the forces. This allows us to ensure, by construction, that sudden motions, in the form of large correction updates, within the iterative non-linear solution do not disturb the stability of the overall system. Furthermore, magnetic pairs far away from each other are handled explicitly. The forces between these objects are still necessary to capture but are much less sensitive to the corrections that occur during the iterations of the non-linear solver. Handling these forces explicitly results in no contribution to the second-order derivatives, reducing the otherwise 100% fill-in of the global linear system solver. The resulting method allows for a stable and interpenetration-free simulation of strong, linear magnetic rigid bodies with large time step sizes. Thus, our method can successfully simulate a wide range of complex magnetic effects and mechanisms like the electromagnetic crane in Fig. 1.

2. Related work

Outside of the field of computer graphics, the simulation of magnetic fields and their effects have been studied for a long time. For numerous simple scenarios, the magnetic fields can be found analytically by solving Maxwell's equations. While these can be found in the typical physics textbooks [Jac98], finding the respective fields for arbitrary geometries, materials. and scene configuration quickly becomes a challenge. A common approach used e.g. in mechanical engineering [JSS*22], is to use state-of-the-art solvers like COMSOL or ANSYS MAXWELL. These solve Maxwell's equations using different simulation techniques like e.g. the finite element method [KABA11].

In visual computing, we avoid solving the full set of Maxwell's equations for magnetic effects. Thus, we usually solve a simplified problem, restricting ourselves to the simulation of magnetostatics. In the pioneering work of Thomaszewski et al. [TGPS08], the authors simulate rigid bodies by introducing an external magnetic force and torque. Their method can simulate permanent and linear magnets, that is, dia- and paramagnets, and serves as the basis of our approach. Kim et al. [KPH18] extend this approach by introducing a method for simulating ferromagnetic materials. These materials remain magnetized after an initial magnetization and thus feature a non-linear relationship between the magnetization and the magnetic field. Kim and Han [KH20] derive a surface-only approach for simulating ferromagnetic materials called the magnetic boundary method. Their derivation yields a non-diverging magnetic force for the simulation of magnetic objects with, e.g., sharp features. For the volumetric, ferromagnetic model, Kim and Han [KH22] improve the convergence of the magnetization solver by including the changes in magnetization and magnetic fields if the spatial configuration changes. The commonalities of the aforementioned methods lie in the interplay between the handling of magnetic effects and the rigid bodies. These provide the magnetic influences as forces and torques as input for the rigid body solver. While this allows the rigid bodies to be simulated with a multitude of different state-of-the-art methods [BET14], the strong coupling of our method provides for example better robustness for large time steps.

Other magnetic effects that have been shown include the field of magnetic fluids including ferrofluids. Ishikawa et al. [IYI*12b; IYI*12a; IYI*13] simulate magnetic fluids by introducing an external magnetic force to a Smoothed Particle Hydrodynamics (SPH) framework. They further allow for the visual simulation of ferrofluids by procedurally generating the characteristic spike shape along the magnetic field lines. Huang et al. [HHM19] provide a physically accurate depiction of ferrofluids by coupling an SPH fluid simulation with a least squares solver of the magnetostatic version of Maxwell's equations. Following their work, Huang and Michels [HM20] provide a surface-only formulation of the previous approach.

Apart from rigid bodies and fluids, the simulation of magnetic fields has also been used for the generation of iron-filing art [YLUH14], for visualizing the magnetic fields in classroom applications [PLH16], and for the simulation of realistic plasma filaments of the sun [PGK*22]. Concerning simulation methods to simulate magnetostatics, Ni et al. [NZWC20] propose a level-

set method showcasing the applicability to magnetic rigid bodies, deformables and fluids and their two-way coupling. Sun et al. [SNZ*21] also demonstrate the applicability of Material Point Method to the aforementioned simulation bodies. They include the magnetic force externally by computing the Maxwell stress tensor which yields a weakly coupled system. The method of Ni et al. [NZWC20] is primarily aimed at the simulation of non-linear magnetic matter using a level-set approach. For magnetoelastic thin shells, Chen et al. [CNZ*22] provide a differentiable simulation framework based on an optimization integrator to among other things deform thin sheets of magnetizable foil to a target state. Similar to our approach, they provide a strongly coupled method. Their method however exclusively simulates magnetic thin shells whereas our method is aimed at rigid bodies. Additionally, no previous method shows the strong coupling of magnetic effects with the collision solver for rigid bodies.

Our method uses an optimization time integrator [KMOW00; MTGG11] to achieve a stable simulation. These types of integrators gained a lot of popularity in computer animation in recent years due to the possibility to robustly run simulations with large time steps [GSS*15]. Using these integrators, Li et al. [LFS*20] introduced the IPC method for the handling of contacts and friction, which we lean on in our method. While we showcase the applicability of a magnetic potential with the contact handling of IPC, other works also include the application of IPC to rigid bodies [FLS*21], affine bodies [LKL*22] as well as multibody dynamics [CLL*22].

3. Method

In this section, we present our method for robustly simulating linear magnetic effects in the context of computer graphics. First, we describe the time integration discretization as an incremental potential (IP) problem. Then, we introduce the magnetic potential we employ to simulate permanent and linear magnets and their relation to the IPC-inspired contact potential. Finally, we present the discretization used to evaluate such a magnetic potential in arbitrarily shaped rigid body objects.

3.1. Incremental Potential

For our goal of a stable rigid body simulation, we pose the underlying equations of motions as a minimization problem as it has been done in current state-of-the-art methods [GSS*15; MEM*20; LFS*20]. This reformulates the equations of motions discretized in time by backward Euler as an incremental potential E, which shall be minimized for the desired degrees of freedom $\mathbf{u}$:

$$\mathbf{u}^{n+1} = \arg\min_{\mathbf{u}} E(\mathbf{u}). \tag{1}$$

Here, the superscript n denotes the n-th time step. Solving the minimization problem can be done by finding the root of the derivative function which equals the equilibrium of forces, that is: $\sum \mathbf{f} = 0$. The root-finding process is usually done by Newton's method, for which we need the Hessian $\mathrm{Hess}(\mathbf{u}) = \frac{\partial^2 E}{\partial \mathbf{u}^2}$, the gradient $\mathrm{grad}(\mathbf{u}) = \frac{\partial E}{\partial \mathbf{u}}$ and the update of our degrees of freedom $\Delta\mathbf{u}$. We iteratively update $\mathbf{u} \leftarrow \mathbf{u} + \Delta\mathbf{u}$ by solving:

$$\mathrm{Hess}(\mathbf{u})\Delta\mathbf{u} = -\mathrm{grad}(\mathbf{u}), \tag{2}$$

for each update step until $\|\mathrm{grad}(\mathbf{u})\|$ approaches 0.

For rigid bodies, we follow the method of Macklin et al. [MEM*20] and formulate the optimization problem in terms of velocities and angular velocities. Let $\underline{\mathbf{v}} = (\mathbf{v}_1, \ldots, \mathbf{v}_N), \underline{\omega} = (\omega_1, \ldots, \omega_N)$ be the stacked vector of velocities and angular velocities of all N rigid bodies respectively, then our problem definition is given as:

$$\left(\underline{\mathbf{v}}^{n+1}, \underline{\omega}^{n+1}\right) = \arg\min_{\underline{\mathbf{v}}, \underline{\omega}} E(\underline{\mathbf{v}}, \underline{\omega}) \tag{3}$$

$$E(\underline{\mathbf{v}}, \underline{\omega}) = E_{\text{inertia}}(\underline{\mathbf{v}}, \underline{\omega}) + E_{\text{other}}(\underline{\mathbf{v}}, \underline{\omega}) \tag{4}$$

$$E_{\text{inertia}}(\underline{\mathbf{v}}, \underline{\omega}) = \sum_a \frac{1}{2}\left[\Delta\mathbf{v}_a^T \mathbf{M}_{\text{mass},a} \Delta\mathbf{v}_a + \Delta\omega_a^T \mathbf{J}_a \Delta\omega_a\right], \tag{5}$$

where a is iterating over the index set of all N rigid bodies, $\Delta\mathbf{v} = \mathbf{v}^{n+1} - \left(\mathbf{v}^n + \Delta t(\mathbf{g} + \frac{\mathbf{F}}{m})\right)$ is the difference between the velocity in the next time step and the predicted velocity, $\Delta\omega = \omega^{n+1} - \left(\omega^n + \Delta t \mathbf{J}^{-1}\tau\right)$ the equivalent for angular velocity, $\mathbf{M}_{\text{mass}} = m\mathbf{I}_3$ the mass matrix and $\mathbf{J}$ the inertia tensor. Additionally, $\mathbf{F}$ denotes external forces excluding gravity $\mathbf{g}$ and τ the external torques.

Oftentimes, a potential is posed in terms of the translation $\mathbf{t}^{n+1}$ and orientation $\mathbf{q}^{n+1}$ of a rigid body in the next time step. In this case, we use the following relations: $\mathbf{t}^{n+1} = \mathbf{t}^n + \Delta t \mathbf{v}^{n+1}$, $\mathbf{q}^{n+1} = \mathbf{q}^n + \frac{1}{2}\Delta t \tilde{\omega}^{n+1}\mathbf{q}^n$, where $\tilde{\omega} = \left(0, \omega_x^{(n+1)}, \omega_y^{(n+1)}, \omega_z^{(n+1)}\right)$ is the extension of the angular velocity to a quaternion and $\tilde{\omega}^{n+1}\mathbf{q}^n$ denotes the quaternion product. For stating the energies in our desired degrees of freedom, we will sometimes refer to the translation and orientations as $\mathbf{t}^{n+1}(\mathbf{v})$ and $\mathbf{q}^{n+1}(\omega)$. Furthermore, we will also use the rotation matrix derived from quaternions, denoted as $R^{n+1} = R(\mathbf{q}^{n+1})$.

The remaining potential E_{other} contains the contact potential E_{contact}, a frictional potential E_{friction}, a potential energy to handle constraints for joints E_{constr}, and the magnetic energy E_{magn}. The contact potential E_{contact} and frictional potential E_{friction} are derived from Li et al. [LFS*20] and are used for the contact and collision resolution. The contact potential increases non-linearly in strength once the distance between two rigid bodies becomes smaller than a given distance threshold $\hat{d}$. We will go into further detail on solving collisions while discussing the implications of combining it with the magnetic energy. Constraints are handled by the potential energy function for constraints [MEM*20; CLL*22]:

$$E_{\text{constr}} = \frac{1}{2}\sum_c k_c C_c^2, \tag{6}$$

where c iterates over the set of all constraints and k denotes the stiffness parameter for each constraint. Thus, we are able to simulate joints and motors.

3.2. Magnetic Potential

In this section, we describe the derivation of the magnetic potential. Our method addresses the simulation of magnetostatic effects, that is the study of magnetic fields, forces, and torques of only bound charges. This restriction allows us to circumvent solving the full set of Maxwell's equations while still covering and well approximating a broad range of magnetic effects like permanent magnets and

slowly changing electromagnets. The correct depiction of magnetostatics for our method relies on the correct handling of the magnetic potential. In contrast to e.g. the gravitational potential, which always exerts an attraction between two points of mass, the magnetic potential is not only dependent on its strength given by the elementary units but also the direction of said units in the surrounding magnetic fields. The origin of the magnetic field is usually depicted as a vector $\mathbf{m}$, called the magnetic dipole moment. The direction of the magnetic dipole moment reflects the orientation of the magnetic field created by it and the magnitude states the resulting strength.

From a physical viewpoint, we know that the magnetic dipole moments tend to align with the magnetic flux density field $\mathbf{B}$, which in 3D is a three-dimensional vector field indicating the direction and strength of the magnetic force. This field is closely related to another vector field, the magnetic field $\mathbf{H}$. In vacuum, the magnetic fields $\mathbf{B}$ and $\mathbf{H}$ are related by the vacuum permeability μ_0:

$$\mathbf{B} = \mu_0 \mathbf{H}. \tag{7}$$

Inside a magnetizable material, the fields change by the magnetization $\mathbf{M}$ stemming from the orientation of the magnetic dipole moments $\mathbf{m}$:

$$\mathbf{B} = \mu_0 (\mathbf{H} + \mathbf{M}), \qquad \mathbf{M} = \frac{d\mathbf{m}}{dV}. \tag{8}$$

Note, that the magnetization is the volume derivative of the magnetic dipole moments inside a magnetic material. For non-magnetic and non-magnetized materials the magnetization is close to zero since either their magnetic response to an external magnetic field is negligibly low or the orientations of the magnetic dipole moments cancel themselves out to a net-zero magnetization. In our work, we will only concern ourselves with permanent and linear magnets. The latter are characterized by a linear relationship between the magnetic field and the magnetization $\mathbf{M} = \chi \mathbf{H}$, which using Eq. (8) yields:

$$\mathbf{B} = \mu \mathbf{H}, \qquad \mu = \mu_r \mu_0, \qquad \mu_r = (1 + \chi). \tag{9}$$

Here, χ denotes the magnetic susceptibility, and μ_r, μ the relative magnetic and magnetic permeability, respectively.

Following the derivations of Jackson [Jac98], the potential energy needed to fully capture the magnetostatic processes includes the work needed to preserve the strength of the magnetic dipoles and their alignment to the magnetic flux density field. Jackson derives the change in work δW needed as:

$$\delta W = \int_V \delta \mathbf{B} \cdot \mathbf{H} dV, \tag{10}$$

where V denotes the occupied volume of all magnetic bodies. Note that Eq. (10) is a general energy description for all magnetic phenomena. Assuming linear materials, the energy potential can be further simplified, yielding the following magnetic potential E_{magn} and energy density function Ψ_{magn}:

$$\Psi_{\text{magn}} = -\frac{1}{2} \mathbf{M} \cdot \mathbf{B}_{\text{ext}}, \qquad E_{\text{magn}} = \int_V \Psi_{\text{magn}} dV, \tag{11}$$

where $\mathbf{B}_{\text{ext}}$ denotes the external magnetic flux density, that is, the magnetic flux density at the integration point without the influences of the body containing said point. Note that Chen et al. [CNZ*22] derive a similar potential. Unlike their approach however, we use a volumetric discretization for the goal of simulating rigid bodies.

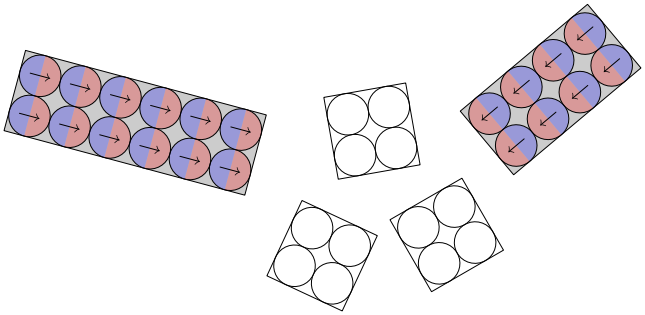

Figure 2: *Magnetic sampling of rigid bodies using a regular spherical sampling. Initially, permanent magnets (grey) are with a magnetic moment from the beginning, while magnetizable objects (white) have a zero magnetic dipole moment. During simulation, the latter receive a non-zero magnetic dipole moment given by the present magnetic fields.*

In the following, we use Eq. (11) as the basis for the simulation of linear magnetizable and permanent magnetic rigid bodies. To be able to integrate our magnetic potential into the optimization-based time integrator, we need to discretize the volume of each rigid body to numerically solve the volume integral.

3.3. Discretization

We denote the stacked vector of the translations and orientations represented as quaternions of the N rigid bodies present in our simulation as $\underline{\mathbf{t}} = (\mathbf{t}_1, \ldots, \mathbf{t}_N)$ and $\underline{\mathbf{q}} = (\mathbf{q}_1, \ldots, \mathbf{q}_N)$. For computing the magnetic quantities, we have to sample the volume of each rigid body to solve Eq. (11). To this end, we approximate the volume of each rigid body using spheres at whose center we place a magnetic dipole moment, following the approach of Thomaszewski et al. [TGPS08]. As the authors denote, using spheres as the discretization primitive has the advantage that the external magnetic flux density induced at any point $\mathbf{B}_{\text{ext}}$ can be expressed by a simple analytical equation and if we increase the resolution of the discretization, we converge to the correct field in the limit. Note that from the discretization, it is assumed that all magnetic fields like $\mathbf{B}$ and $\mathbf{M}$ are constant over the sphere's volume. Depending on the amount of samples we want to take for each rigid body, they form a (not-necessarily tight) sphere packing as can be seen in Fig. 2. The spheres are attached to the rigid body at a fixed position using connectors, which are determined at the beginning of the simulation in local coordinate space. Thus, for the rigid body a with $\|\mathcal{M}_a\|$ samples, where $\mathcal{M}_a$ denotes the indices of samples for body a, the vector of the local sampling positions is given by $(\mathbf{r}_1, \ldots, \mathbf{r}_M)_a$.

The local sampling positions can be transformed into world space, yielding the global position $\mathbf{x}$ of sphere i inside rigid body a as:

$$\mathbf{x}_i(\mathbf{t}_a, R_a) = R_a \mathbf{r}_i + \mathbf{t}_a. \tag{12}$$

Each point's magnetic influence is determined by its magnetic

dipole moment. With our assumption that the fields are constant over the sphere's volume, we can finally discretize the magnetic potential Eq. (11). The magnetic potential between just two samples i and j simplifies to:

$$E_{\text{magn},ij} = -\frac{1}{2} V_i \mathbf{M}_i \cdot \mathbf{B}_{\text{ext},j}(\mathbf{x}_i) = -\frac{1}{2} \mathbf{m}_i \cdot \mathbf{B}_{\text{ext},j}(\mathbf{x}_i), \tag{13}$$

where V_i is the volume of the sphere of sample i and $\mathbf{m} = \int_{V_i} \mathbf{M}_i dV = V_i \mathbf{M}_i$. $\mathbf{B}_{\text{ext},j}$ denotes the external magnetic flux density created by sample j, that is given by:

$$\mathbf{B}_{\text{ext},j}(\mathbf{x}_{ij}, \mathbf{m}_j) = \frac{\mu_0}{4\pi} \frac{3\hat{\mathbf{x}}_{ij} \left(\hat{\mathbf{x}}_{ij} \cdot \mathbf{m}_j\right) - \mathbf{m}_j}{\|\mathbf{x}_{ij}\|^3}, \tag{14}$$

where $\mathbf{x}_{ij} = \mathbf{x}_i - \mathbf{x}_j$ and $\hat{\mathbf{x}} = \mathbf{x}/\|\mathbf{x}\|$. Now, a single sampling point inside a rigid body will create a potential, which should provide an influence towards all samples of the other rigid bodies. Additionally, this point will be influenced by the potentials of all samples of other rigid bodies. Thus, the resulting total energy is the superposition of the magnetic potentials of all pairs (i, j) where i and j are sampling spheres of distinct rigid bodies, which we index by the set $\mathcal{P}$:

$$E_{\text{magn}} = \sum_{(i,j)\in\mathcal{P}} E_{\text{magn},ij}. \tag{15}$$

From this formulation, we observe that every magnetized rigid body influences every other magnetized rigid body. A direct consequence is that the Hessian of the potential function will become a dense matrix.

3.3.1. Permanent magnets

We model permanent magnets by having a constant local magnetic dipole moment, that is $\mathbf{m}_{\text{loc},i} = const$. Thus, they produce a different magnetic response based on their direction. Consequently, the global magnetic dipole moment is dependent on the orientation of the associated rigid body, that is, if sample i belongs to rigid body a: $\mathbf{m}_i = R_a \mathbf{m}_{\text{loc},i}$. Thus, our magnetic potential in Eq. (13) is dependent on both the distance of two rigid bodies as well as their orientation, that is:

$$E_{\text{magn},ij} = -\frac{1}{2} R_a \mathbf{m}_{\text{loc},i} \cdot \mathbf{B}_{\text{ext},j}(\mathbf{x}_{ij}, R_b \mathbf{m}_{\text{loc},j}). \tag{16}$$

Note the rotation matrices and global sampling positions dependency on the degrees of freedom $R_a = R(\mathbf{q}(\omega_a))$ and $\mathbf{x}_{ij} = \mathbf{x}(\mathbf{t}(\mathbf{v}_a), R(\mathbf{q}(\omega_a))) - \mathbf{x}(\mathbf{t}(\mathbf{v}_b), R(\mathbf{q}(\omega_b)))$.

3.3.2. Linear magnets

As outlined before, linear magnets possess a linear relationship between their magnetic flux density $\mathbf{B}$ and magnetic fields $\mathbf{H}$. This relationship stems from the ability of the magnetic dipole moments inside the material to align with the external magnetic field, resulting in an intensification or weakening of the external magnetic field. The alignment happens nearly instantaneously for the time scales we are interested in. Thus, also the history of the magnetic moments is irrelevant for us, which would not be the case for ferromagnetic materials.

Now, the direction of the magnetic dipole moments is dependent on the external magnetic flux density at their position $\mathbf{B}_{\text{ext}}(\mathbf{x})$. Since there might be multiple linear magnetizable rigid bodies present inside a scene, their magnetic fields will influence each other, which results in a cyclical dependency on each other's magnetization. One can solve this dependency by including the magnetic dipole moments of linear magnets as degrees of freedom. The resulting optimization problem then becomes:

$$\left(\underline{\mathbf{v}}^{n+1}, \underline{\omega}^{n+1}, \underline{\mathbf{m}}^{n+1}\right) = \underset{\underline{\mathbf{v}},\underline{\omega},\underline{\mathbf{m}}}{\arg\min}\, E(\underline{\mathbf{v}}, \underline{\omega}, \underline{\mathbf{m}}). \tag{17}$$

Since the Hessian matrix for permanent magnets alone is already dense, adding more degrees of freedom, which will also form a dense matrix, will further impair the performance of our method. Thus, we decouple the magnetization process from solving the optimization problem. That means, we compute the strength and direction of the magnetic dipole moments beforehand and only update the direction in response to the rigid bodies orientation in the optimization step. In that sense, we handle linear magnets as permanent magnets from the viewpoint of the optimization solver, which is the same as in Eq. (16). Note that we update the direction of the magnetic dipole moments before they are used to calculate the new values in the next time step.

The magnetic dipole moment $\mathbf{m}_i$ of rigid body a is computed analogously to Thomaszewski et al. [TGPS08], that is:

$$\mathbf{m}_i = V_i \frac{\chi_a}{1 + \chi_a/3} \sum_j \mathbf{B}_{\text{ext},j}\left(\mathbf{x}_{ij}, \mathbf{m}_j\right), \tag{18}$$

where j is iterating over all samples not belonging to rigid body a.

3.4. IPC-based contact

The robustness of our solver can be traced back to the interplay between the magnetic energy and the IPC-based contact and collision model [LFS*20]. Going into detail, collisions in IPC are handled by defining an energy potential for contact. That is, if two objects come too close to each other, a contact potential is activated which continuously and smoothly increases its strength from zero at the contact thickness $\hat{d}$. For our method, we use a cubic penalty function $p(d, \hat{d})$:

$$p(d, \hat{d}) = \begin{cases} k\left(\hat{d} - d\right)^3, & \text{if } d < \hat{d} \\ 0, & \text{otherwise} \end{cases}. \tag{19}$$

Turning the penalty function into a potential for rigid bodies, we firstly find the set of collision primitive pairs $(c_a, c_b) \in \mathcal{C}$, that is, point-triangle and edge-edge pairs, whose distance falls below the collision thickness in the optimization step. For the optimization step, we formulate the distance dist_{c_a,c_b} of the collision primitives c_a, c_b in terms of the translation and orientation of the rigid body:

$$\text{dist}_{c_a,c_b}(\mathbf{t}_a, R_a, \mathbf{t}_b, R_b) = \|R_a \mathbf{r}_{c_a} - R_b \mathbf{r}_{c_b} + \mathbf{t}_a - \mathbf{t}_b\|, \tag{20}$$

where $\mathbf{r}_{c_a}$ and $\mathbf{r}_{c_b}$ are the local vectors from the collision point to the rigid bodies local origin. The contact potential now becomes the sum of the penalty function evaluations for all collision primitive pairs [FLS*21], that is:

$$E_{\text{contact}} = \kappa \sum_{(c_a,c_b)\in\mathcal{C}} p(\text{dist}_{c_a,c_b}(\mathbf{t}_a, R_a, \mathbf{t}_b, R_b), \hat{d}), \tag{21}$$

where κ is the stiffness of the contact potential and dist_{c_a,c_b} describes the distance between the collision primitives c_a and c_b of

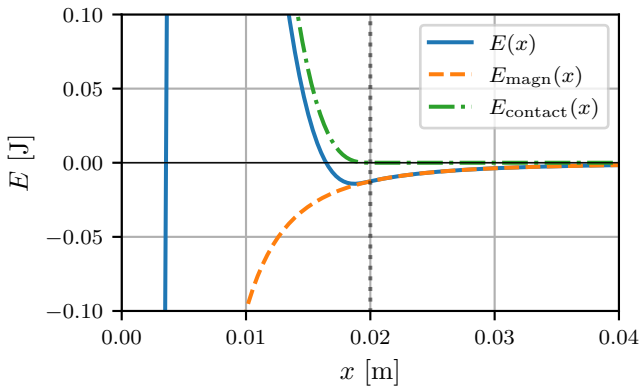

Figure 3: *Plot of potentials against distance. The shape of the total potential E is dependent on the magnetic E_{magn} and contact potential E_{contact}. The contact thickness is indicated by the dotted vertical line. For a well-chosen contact stiffness κ, the total potential exhibits a local minimum near the contact thickness the optimizer will converge to in 1D.*

rigid bodies a and b. Note that the distance is defined between the collision primitives. For the goal of optimization, the distance will be expressed in the desired degrees of freedom using the same connector formula as used for fixing the positions of the magnetic samples in Eq. (12).

Now, the previously defined magnetic potential E_{magn} decreases in value inverse cubically with respect to distance when two magnetic bodies attract each other. The contact potential will increase cubically once the distance falls below the contact thickness $\hat{d}$. We analyze the resulting potential function in one dimension to give us an indication of the behavior of the full-dimensional case. The magnetic potential falls to $-\infty$ for distances approaching zero, while our cubic potential has a well-defined function value for $d = 0$. Thus, the choice of a good stiffness parameter κ is crucial. Note that it is dependent on both the strength of the magnetic attraction as well as the contact thickness. If the stiffness of the penalty function is too low, the strength of the magnetic potential overshadows the total potential energy function. In case this leads to an invalid simulation step like a non-convergent Newton step or an interpenetration, we revert to the current step and double the contact stiffness.

For a properly tuned stiffness parameter κ, we observe that the contact potential rises faster than the magnetic potential falls leading to a local minimum close to the contact thickness, which can be seen in in Fig. 3. The minimum represents the equilibrium of forces between the collision resolution and the magnetic attraction, leading to a robust and intersection-free simulation state.

3.5. Optimizations

Approaching $\|\mathbf{x}_{ij}\| \to \infty$ the change in magnetic energy for a small perturbation approaches zero. Thus, in the far distance, we can argue, that the contribution of the magnetic energy hardly changes with respect to distance, that is $E_{\text{magn}}(\underline{\mathbf{x}}) \approx E_{\text{magn}}(\underline{\mathbf{x}} + \Delta\underline{\mathbf{x}})$. As a consequence, we can model the magnetic influence as a constant force from the current state of the simulation and add its influence to the inertia term.

$$\mathbf{F}_{\text{magn,exp}} = \sum_{(i,j)\in\mathcal{P}} \mathbf{F}_{\text{magn},ij} \tag{22}$$

$$\mathbf{F}_{\text{magn},ij} = \frac{\mu_0}{8\pi\|\mathbf{x}_{ij}\|^4}\left[-15\hat{\mathbf{x}}_{ij}\left(\left(\mathbf{m}_i\cdot\hat{\mathbf{x}}_{ij}\right)\left(\hat{\mathbf{x}}_{ij}\cdot\mathbf{m}_j\right)\right) + 3\hat{\mathbf{x}}_{ij}\left(\mathbf{m}_i\cdot\mathbf{m}_j\right) + 3\mathbf{m}_i\left(\mathbf{m}_j\cdot\hat{\mathbf{x}}_{ij}\right) + 3\mathbf{m}_j\left(\mathbf{m}_i\cdot\hat{\mathbf{x}}_{ij}\right)\right]. \tag{23}$$

For the torque, we need to differentiate between the mechanical torque $\tau_{\text{mech},ij}$, stemming from an application of the magnetic force away from the center of gravity, and the magnetic torque $\tau_{\text{magn},ij}$, which aligns the magnetic moments towards the external magnetic flux density:

$$\tau_{\text{totalmagn,exp}} = \sum_{(i,j)\in\mathcal{P}} \left[\tau_{\text{magn},ij} + \tau_{\text{mech},ij}\right] \tag{24}$$

$$\tau_{\text{mech},ij} = R_a\mathbf{r}_i \times \mathbf{F}_{\text{magn},ij} \tag{25}$$

$$\tau_{\text{magn},ij} = \frac{\mu_0}{8\pi\|\mathbf{x}_{ij}\|^3}\left[3\left(\hat{\mathbf{x}}_{ij}\cdot\mathbf{m}_j\right)\mathbf{m}_i\times\hat{\mathbf{x}}_{ij} - \mathbf{m}_i\times\mathbf{m}_j\right]. \tag{26}$$

As a criterion for switching between the explicit and implicit formulation, we use a distance threshold ϑ. Thus, if the distance between two magnetic samples is greater than ϑ, we assume the change in magnetic energy over the state path to be insignificantly small and use the force formulation. In the other case, that is, if $\|\mathbf{x}_{ij}\| < \vartheta$ we register the pair (i, j) for the optimization integrator.

The advantage of switching between the explicit and implicit formulation lies in the connectivity of the Hessian matrix. Since the explicit formulation is handled as a constant force, it does not create a non-zero contribution to the energy Hessian. Thus, we only spend time computing the magnetic forces and torques at the beginning of the time step, while for the implicit formulation, the entries of the energy Hessian and gradient must be computed for each Newton iteration.

Finally, a full algorithmic outline of our method can be found in Alg. 1.

4. Results

Our results are simulated on a machine with an AMD Ryzen Threadripper PRO 5975WX CPU with 32 cores and 3.60 GHz base core frequency and 256 GB of RAM. We use OpenMP for parallelization using all threads if possible.

For minimizing the incremental potential, the Newton-type solver requires the gradient and the Hessians of all potentials we have outlined. We integrate our method into STARK by Fernández-Fernández et al. [FLL*24] which includes a symbolic differentiation framework [FLW*23] to derive the expressions for all potentials. Furthermore, we project our Hessian to the cone of positive definite matrices. Thus, we can safely use Conjugate Gradient (CG) with a block-diagonal preconditioner to solve the linear system. If not stated otherwise, we used a residual for the conjugate gradient solver of 10^{-12} and stopped the Newton iteration once the acceleration derived from the gradient, i.e: $\|\frac{1}{\Delta t}M^{-1}\text{grad}(\mathbf{v},\omega)\|_\infty$ was smaller than $1m/s^2$ In our experiments, the contact stiffness κ

Algorithm 1: Magnetic Strong Coupling Algorithm

for *all magnetizable rigid bodies a* **do**
  for *all samples* $i \in \mathcal{M}_a$ **do**
    Compute magnetic dipole moment $\mathbf{m}_i$.
for *all magnetic rigid bodies a* **do**
  for *all magnetic rigid bodies* $b \neq a$ **do**
    for *all samples* $i \in \mathcal{M}_a$ **do**
      for *all samples* $j \in \mathcal{M}_b$ **do**
        if $\|\mathbf{x}_{ij}\| < \vartheta$ **then**
          Register the pair (i, j) for the optimization integration.
        else
          Calculate the magnetic force between i and j and add the force $\mathbf{F}_{ij}$ and torque τ_{ij} as external contributions.
while $\|\frac{1}{\Delta t} M^{-1} \text{grad}(\mathbf{v}, \omega)\|_\infty >$ *threshold* **do**
  Assemble Hessian $\text{Hess}(\mathbf{u})$
  Solve $\text{Hess}(\mathbf{u})\Delta\mathbf{u} = -\text{grad}(\mathbf{u})$
  Perform line search $\alpha \leftarrow \text{line\_search}(\Delta\mathbf{u})$
  if *step is valid* **then**
    $\mathbf{u} \leftarrow \mathbf{u} + \alpha\Delta\mathbf{u}$
  else
    $\kappa \leftarrow 2\kappa$
    $\Delta t \leftarrow \frac{\Delta t}{2}$
Update simulation state $\underline{\mathbf{t}}, \underline{\mathbf{q}}, \underline{\mathbf{v}}, \underline{\omega}$

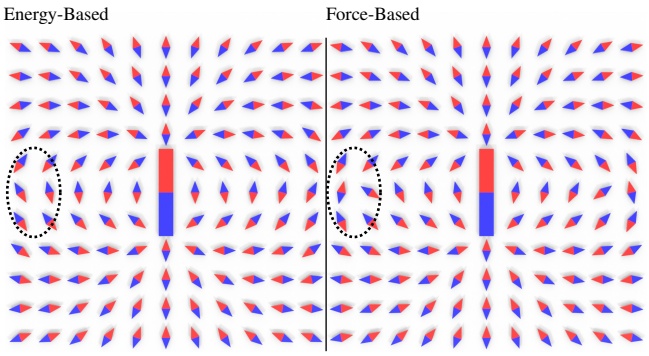

Figure 4: *Magnetic Fields: We compare our potential-based approach with a force-based one [TGPS08]. Our method converges faster to a static equilibrium due to the implicit handling of the magnetic effects. These differences can be noticed for example after* 6.6 s *of simulated time in the outlined regions.*

varies between 10^7 and 10^{14}, and in the case where the line search does not yield a valid configuration, we also half the time step size. An overview of the scenes and their simulation statistics can be seen in Tab. 1.

4.1. Magnetic Fields

Firstly, we check whether our magnetic potential correctly handles the depiction of the underlying magnetic fields. Thus, we compare the differences of our implementation with a strictly force-based implementation [TGPS08]. To this end, we set up a strong permanent magnet surrounded by 118 tiny magnetic compass needles. The former is sampled by 988 spheres ($\|\mathbf{m}\| = 25$ J/T) while each needle is sampled by 4 ($\|\mathbf{m}\| = 0.05$ J/T). The needles are initially aligned in a single direction and shall indicate the direction of the magnetic field created by the permanent magnet. Both the force-based and the energy-based implementation converge to the correct solution, but there are differences in the dynamic behavior between both. Primarily, the energy-based implementation reaches the static equilibrium of forces in a shorter time since the compass needles spin less back and forth during the simulation. That fact can be attributed to the implicit integration of the magnetic effects, since we respect how these are changing throughout the time step with our method.

4.2. Optimizations

We evaluate our hybrid explicit / implicit magnetic energy scheme by looking at the computational performance. Concerning the former, we set up a synthetic scenario with two permanent magnets sampled by a total of 10000 magnetic samples. The samples are arranged in a single dimension, such that decreasing the distance threshold directly and linearly corresponds to the number of magnetic pairs that are handled explicitly instead of implicitly. In order to record the difference in computational time, we vary the distance threshold ϑ from handling all sample pairs fully implicitly to fully explicitly. From the experiment, we deduce that the maximum speedup factor gained in our implementation is around 31. Note that this speedup is just a theoretical upper bound for our implementation, since the fully explicit approach may fail to find a stable configuration. In practice, we choose a distance parameter between two and five times the maximum diameter of any moving object in the scene giving us a good trade-off between stability and efficiency.

The speedup can be attributed to the decrease in non-zero entries in the Hessian matrix while using the explicit / implicit integration scheme. We investigate the extent of our schemes by analyzing the amount of entries in the simulation of a single frame of a non-magnetic bowl filled with 64 weakly permanently magnetized spheres. The scene setup and the associated fill-in patterns of the energies' Hessians can be seen in Fig. 5. The left-most pattern depicts the non-zero Hessian entries related to the contact handling. These amount to 246 non-zero entries which will not change while varying the magnetic threshold. Resolving magnetic interactions implicitly only while they are near a close contact, yields the left-most of the magnetic Hessian matrices. Here, ϑ is slightly larger than $2R$, where R is the radius of the magnetic balls. This Hessian contains 312 non-zero entries, which are only 66 additional entries in comparison to the contact Hessian. While this is close to the lowest bound of non-zero entries that are achievable, it is not advisable to choose the threshold this small since the change from the explicit to implicit integration here leads to a rather harsh jump in the potentials. A more reasonable value we have found experimentally is

Scene	#Magnetic Samples	#Bodies	Avg. Newton Iter.	Avg. CG Iter.	Δt [ms]	T_{step} [ms]	T_{total} [s]
Magnetic Fields	1460	119	1.34	2.7	10*	2135.87	3206.011
Collision Resolution	2	2	1.008	3.85	1	1.19	1.784
Strong Magnetic Forces	2	2	1.003	12.19	1	1.87	18.670
Perpetuum Motion Simulator	401	4	1.049	1.39	1*	26.36	790.938
Magnetic Crane	167	132	26.28	7364.2	1.92*	1943.98	22307.2
Sphere of Magnetic Objects	31	28	10.19	538.82	1*	174.02	4351.453

Table 1: *Simulation statistics of our method. This includes the number of magnetic samples, amount of rigid bodies, the average number of Newton iterations, the average number of CG iterations, the time step size Δt, the time per step T_{step} and the total time the simulation took T_{total}. The asterisk denotes that for the simulations, we allow the time step size to vary. Thus, the entry denotes the average time step size of the simulation.*

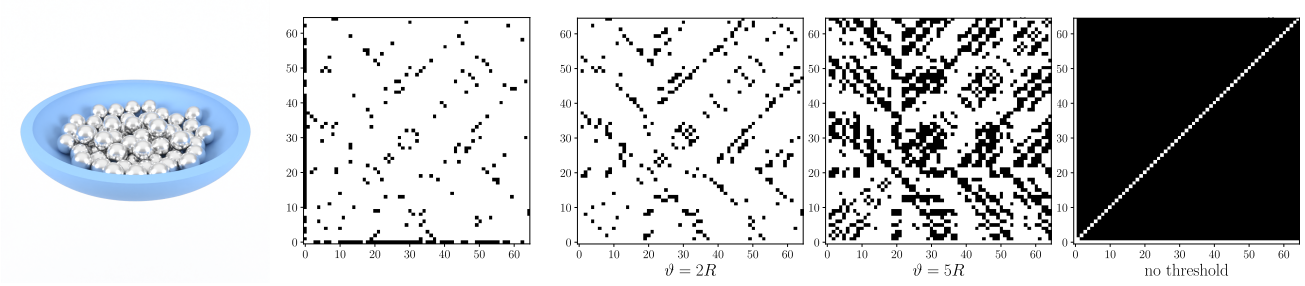

(a) **(b)** **(c)**

Figure 5: *Evaluation of hybrid explicit / implicit optimization in terms of Hessian fill-in: (a) Static scenario including a non-magnetic bowl and permanently magnetized ball magnets. (b) The non-zero entries of the Hessian matrix resulting from the contacts. Each black square represents a non-zero entry between the rigid bodies associated with the pixel coordinates. (c) The non-zero entries of the Hessian matrix from the magnetic potential with different thresholds. The thresholds ϑ are expressed in the radius R of the magnetic spheres. The left-most pattern represents the smallest threshold while still maintaining a strongly coupled simulation, the middle one a good approximation and the right-most one the fully implicit simulation.*

to set $\vartheta = 5R$. This results in 1302 non-zero entries and still yields a speedup of 1.24 in this small scale example in comparison to the fully implicit scenario. The associated Hessian is depicted as the right-most one and contains 4032 entries.

4.3. Weak vs. Strong Coupling

4.3.1. Collision Resolution

As we have previously outlined, weak coupling may lead to jittering or explosions when strong magnetic forces and large time step sizes are present. We showcase this issue by simulating the following synthetic scenario. Two cube-shaped permanent magnets ($\|\mathbf{m}\| = 5$ J/T) with a volume of $(0.1\,\text{m})^3$ sampled by a single magnetic sampling point centered in the cube are placed $0.25m$ apart on the x-axis. Their magnetic dipole moment is aligned with the x-axis leading to a strong attraction between the magnets. For this setup, we compare our implementation with the weakly coupled approach from Thomaszewski et al. [TGPS08] implemented in Bullet3 [Cou15]. While their method was originally implemented into the Open Dynamics Engine (ODE) [Smi06], Bullet3 also implements the same MLCP solver for collision resolution as ODE, which we use for our experiments. We provide the magnetic effects as external forces and torques which are calculated at the beginning of a time step. Note that for our implementation in Bullet3, we have set contact damping to zero and increased the contact stiffness to the largest possible value. For both approaches, we have no restitution. The remaining values like positions, masses, and inertia tensors are chosen to be identical between both implementations.

We compare both approaches for a fixed time step size of $\Delta t = 1\,\text{ms}$ and a moderate attraction strength. Evaluating this scene, we keep track of the distance between both magnetic cubes over time, for which the result is shown in Fig. 6. It shows the magnetic force and collision resolution working against each other in the weakly coupled approach. Since the magnetic force increases cubically in strength in relation to the inverse distance, the collision resolution was not able to find the force equilibrium right away at $t = 0.79$s. This results in an overly strong collision response leading to a small bounce between the magnetic bodies even if the restitution is set to zero. Note that this behavior stems only from the weak coupling between both systems. We further tried exchanging the MLCP solver of Bullet3 with the sequential impulse and NNCG solver, which yielded similar results. In any case, the weakly coupled approach however finds the balance of forces at around $t = 1$s. The strongly coupled approach finds the balance of forces right away at the first impact since the energy-based solver finds the energetically minimal configuration between the decreasing magnetic energy and the

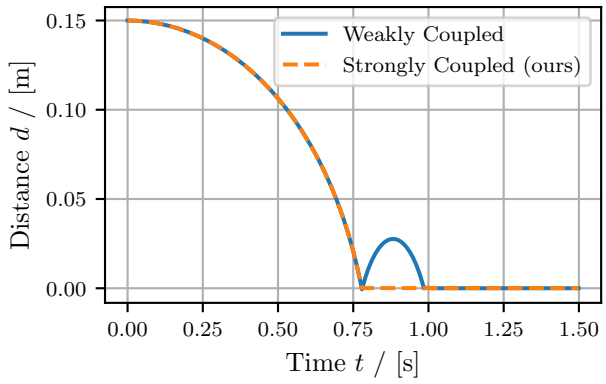

Figure 6: *Collision Resolution: Distance between two cuboidal permanent magnets with a side length of* $0.1m$ *attracting each other. Comparing the weakly coupled and strongly coupled approach, we observe that the former experiences some bouncing before coming to rest. The strongly coupled approach converges safely to the correct configuration without bouncing or jittering.*

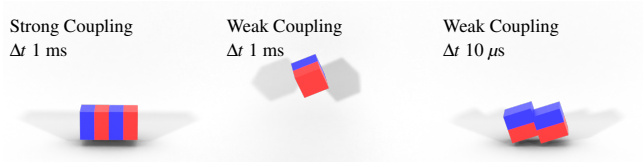

Figure 7: *Strong magnetic forces: In the presence of strong magnetic forces, the weakly-coupled approach fails to exhibit stable contact behavior. Our strongly-coupled approach remains stable throughout.*

increasing barrier energy. While it may also be possible to find a satisfactory result with a weakly coupled approach, this will require additional effort in terms of stabilization and finding the correct equilibrium of forces right away, which comes naturally with the strongly coupled method.

4.3.2. Strong Magnetic Forces

Next, we test both methods for exceedingly strong magnetic forces. Again, we use two cuboidal permanent magnets ($\|\mathbf{m}\| = 500$ J/T), but this time spaced 2m apart. Furthermore, we increase the magnetic dipole moment's strength by a factor of 100 in comparison to the previous scene. For this scene, we compare the weakly coupled approach in Bullet3 using $\Delta t = 1\,\text{ms}$ and $\Delta t = 10\,\mu\text{s}$ with our strongly coupled approach with a maximum time step size of $\Delta t = 1\,\text{ms}$. The results can be seen in Fig. 7. The weakly coupled approach fails to correctly handle the collision with the larger time step size since the magnetic attraction is too strong for the collision resolution. Thus, we need to lower the time step size. But also for the extremely low time step size of $10\,\mu\text{s}$, the collision resolution may lead to issues like interpenetration. While these interpenetrations may be attributed to a possibly asymmetrical collision handling of Bullet3 [CLW24], the simulation of strong magnets with large time steps remains infeasible for the weakly coupled approach. By strongly coupling the magnetic effects to the collision solver, we can correctly simulate the collision without any interpenetration or other visual artifacts.

4.4. Complex scenarios

4.4.1. Perpetual Motion Simulator

This scene depicts a metal ball track which gives the illusion of running indefinitely without any outer influences. The mechanism consists of a nickel ball ($\mu_{\text{nickel}} = 1.26 \cdot 10^{-4}\text{H}\,\text{m}^{-1}$) which drops from a funnel into a track of practically non-magnetizable metal which loops the ball back into the funnel. Since the metal ball could not loop back into the funnel through the gravitational potential alone, it is being accelerated by an electromagnet ($\|\mathbf{m}\| = 13.75$ J/T) which switches off the moment the ball passes by. This gives the ball enough momentum to follow the track and fall into the funnel repeating its motion indefinitely. The mechanism is a finely tuned system between gravitational energy, magnetic energy, and friction. Simulating it is a challenging task with previous approaches, since we not only need a well-tuned solver for each effect stemming from the previous potentials but also a robust collision solver to handle the nickel ball running along the track. As shown in Fig. 8, our method successfully simulates this scenario leading to an accurate depiction of the real-life equivalent. Thus, our method not only handles the interplay between strong magnetic forces, collisions, and friction but can also be used for slowly switching electromagnets where the resulting electromagnetic effects are negligibly small.

4.4.2. Magnetic Crane

In this experiment, we showcase the the applicability of our method to larger scenarios as well as the interaction with constraints and subsequently joints. The scene in Fig. 1 depicts an electromagnet attached to a crane by a rope, moving 125 magnetizable nickel cubes from one box to another. The electromagnet contains 42 magnetic samples with a magnetic dipole moment of 250J/T pointing in the upwards direction. The magnetizable cubes are sampled with one sphere each. After the crane lowers the electromagnet towards the first box, it is turned on leading to a rapid magnetization and attraction of all cubes. The following collision is handled robustly leading to a fast orientation of the cubes onto the magnet without excessive bouncing. In the end, the crane moves over to a second box, where we turn off the electromagnet, dropping the cubes.

4.5. Sphere of Magnetic Objects

Finally, we show that our method can robustly handle arbitrarily shaped magnetic objects and the complex interaction between them. Thus, we place a mix of 27 magnetizable nickel armadillos and bunnies into a closed glass sphere and let them drop into a pile. Next, we activate a magnet placed below the pile, letting all objects attract each other. Afterwards, we drag the magnet across the sphere's surface until all objects stick to the top of the sphere. Lastly, we deactivate the magnet, letting the objects pile up at the bottom again. This sequence is depicted in Fig. 9. In this example, we show that our solver stably handles this scenario leading to no inter-penetrations.

Figure 8: *Perpetual Motion Simulator: A magnetizable ball is running indefinitely from the funnel into the tracks and loops back into the funnel. The additional momentum for the ball to loop back into the funnel is provided by turning on and off an electromagnet at the correct time.*

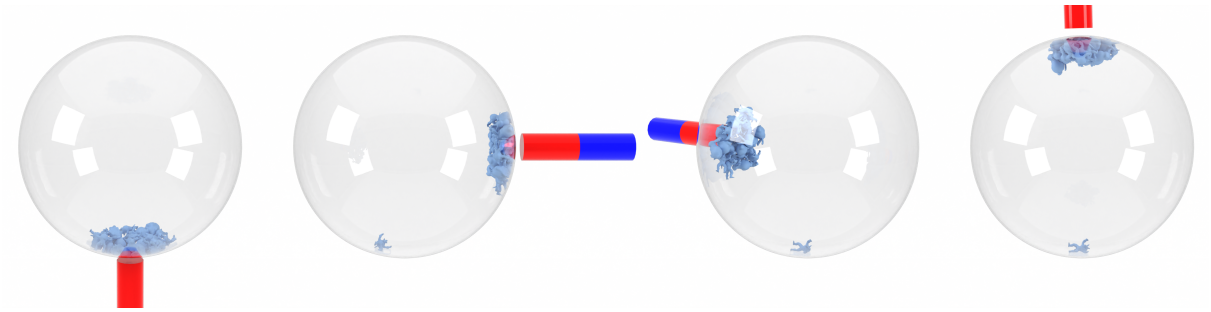

Figure 9: *Sphere of Magnetic Objects: A glass sphere contains a multitude of magnetizable objects of arbitrary shapes. A magnet in close proximity around the sphere, attracting objects and moving them around until the objects appear to stick to the ceiling of the sphere. Then the magnet is turned off and the objects fall back to the ground.*

5. Conclusion

We presented a strongly coupled method for the simulation of linear magnets with rigid bodies. The basis of our method relies on the magnetic potential discretized by spherical sampling points. This potential is added to the ones gained by restating the backward Euler discretized equations of motion for rigid bodies as an optimization problem. For contact, we used an IPC-like collision resolution. Additionally, we provided insights into why this coupling works remarkably well from a theoretical standpoint. Furthermore, we optimized the system with a distance-based, hybrid explicit / implicit magnetic potential approach to alleviate the computational effort by reducing the non-zero Hessian entries. Finally, we have shown the stability and robustness of our solver in multiple challenging scenarios ranging from strong magnetic forces to arbitrarily shaped collision geometries. Thus, by strongly coupling magnetic effects, we obtain a method, which allows for the interpenetration-free and robust simulation of magnets while keeping large time step sizes.

5.1. Limitations & Future Work

In our work, we have shown our method in conjunction with rigid bodies. However, the magnetic potential should also be usable for magnetic deformables like volumetric deformable magnets, magnetic cloth, and thin shells. Using IPC to simulate the aforementioned objects as magnets should be entirely possible given our exploration of the interaction of the magnetic potential and contact potential in Sec. 3.4.

One stepping stone to achieve this goal is also to refine the optimization we provided. While the distance-based criterion works great for speeding up many scenes with distant magnetic effects, our heuristic is purely distance-based. A better alternative would be to incorporate the information of change in the magnetic energy to argue when we can safely use an explicit or similarly fast approximation of the magnetic influences. Further integrating optimization methods like the Fast Multipole Method [Rok85; GR87] are also possible improvements.

Finally, adapting our magnetic potential to include ferromagnetic and electromagnetic effects remains an interesting avenue for future research.

Acknowledgements

The presented investigations were carried out at RWTH Aachen University within the framework of the Collaborative Research Centre SFB1120-236616214 "Bauteilpräzision durch Beherrschung von Schmelze und Erstarrung in Produktionsprozessen" and funded by the Deutsche Forschungsgemeinschaft e.V. (DFG, German Research Foundation). The sponsorship and support is gratefully acknowledged. Additionally, we thank the Stanford Computer Graphics Laboratory for providing the model of the

Stanford Bunny. Open Access funding enabled and organized by Projekt DEAL.

References

[BET14] BENDER, JAN, ERLEBEN, KENNY, and TRINKLE, JEFF. "Interactive Simulation of Rigid Body Dynamics in Computer Graphics". *Computer Graphics Forum* 33.1 (2014), 246–270. ISSN: 1467-8659. DOI: 10.1111/cgf.12272. URL: http://dx.doi.org/10.1111/cgf.12272 2.

[CLL*22] CHEN, YUNUO, LI, MINCHEN, LAN, LEI, et al. "A unified newton barrier method for multibody dynamics". *ACM Transactions on Graphics (TOG)* 41.4 (2022), 1–14 3.

[CLW24] CHEN, YI-LU, LY, MICKAËL, and WOJTAN, CHRIS. "Primal–Dual Non-Smooth Friction for Rigid Body Animation". (Sept. 2024) 9.

[CNZ*22] CHEN, XUWEN, NI, XINGYU, ZHU, BO, et al. "Simulation and optimization of magnetoelastic thin shells". *ACM Trans. Graph.* 41.4 (July 2022). ISSN: 0730-0301. DOI: 10.1145/3528223.3530142. URL: https://doi.org/10.1145/3528223.3530142 1, 3, 4.

[Cou15] COUMANS, ERWIN. "Bullet physics simulation". *ACM SIGGRAPH 2015 Courses*. 2015, 1 8.

[FLL*24] FERNÁNDEZ-FERNÁNDEZ, JOSÉ ANTONIO, LANGE, RALPH, LAIBLE, STEFAN, et al. "STARK: A Unified Framework for Strongly Coupled Simulation of Rigid and Deformable Bodies with Frictional Contact". *2024 IEEE International Conference on Robotics and Automation (ICRA)*. 2024 6.

[FLS*21] FERGUSON, ZACHARY, LI, MINCHEN, SCHNEIDER, TESEO, et al. "Intersection-free rigid body dynamics". *ACM Transactions on Graphics* 40.4 (2021) 3, 5.

[FLW*23] FERNÁNDEZ-FERNÁNDEZ, JOSÉ ANTONIO, LÖSCHNER, FABIAN, WESTHOFEN, LUKAS, et al. "Symx: Energy-based simulation from symbolic expressions". *arXiv preprint arXiv:2303.02156* (2023) 6.

[GR87] GREENGARD, LESLIE and ROKHLIN, VLADIMIR. "A fast algorithm for particle simulations". *Journal of computational physics* 73.2 (1987), 325–348 10.

[GSS*15] GAST, THEODORE F, SCHROEDER, CRAIG, STOMAKHIN, ALEXEY, et al. "Optimization integrator for large time steps". *IEEE transactions on visualization and computer graphics* 21.10 (2015), 1103–1115 3.

[HHM19] HUANG, LIBO, HÄDRICH, TORSTEN, and MICHELS, DOMINIK L. "On the accurate large-scale simulation of ferrofluids". *ACM Transactions on Graphics (TOG)* 38.4 (2019), 1–15 1, 2.

[HM20] HUANG, LIBO and MICHELS, DOMINIK L. "Surface-only ferrofluids". *ACM Transactions on Graphics (TOG)* 39.6 (2020), 1–17 2.

[IYI*12a] ISHIKAWA, TOMOKAZU, YUE, YONGHAO, IWASAKI, KEI, et al. "Visual simulation of magnetic fluid taking into account dynamic deformation in spikes". *Image Electronics and Visual Computing Workshop*. 2012 2.

[IYI*12b] ISHIKAWA, TOMOKAZU, YUE, YONGHAO, IWASAKI, KEI, et al. "Visual Simulation of Magnetic Fluids." *GRAPP/IVAPP*. 2012, 319–327 1, 2.

[IYI*13] ISHIKAWA, TOMOKAZU, YUE, YONGHAO, IWASAKI, KEI, et al. "Visual Simulation of Magnetic Fluids Using Dynamic Displacement Mapping for Spike Shapes". *IIEEJ Transactions on Image Electronics and Visual Computing* 1.1 (2013), 51–57 2.

[Jac98] JACKSON, JOHN DAVID. *Classical Electrodynamics*. 3rd. Wiley, 1998 2, 4.

[JSS*22] JESKE, STEFAN RHYS, SIMON, MAREK SEBASTIAN, SEMENOV, OLEKSII, et al. "Quantitative evaluation of SPH in TIG spot welding". *Computational Particle Mechanics* (Apr. 2022). DOI: 10.1007/s40571-022-00465-x 2.

[KABA11] KENNEDY, MARK W, AKHTAR, SHAHID, BAKKEN, JON ARNE, and AUNE, RAGNHILD E. "Analytical and experimental validation of electromagnetic simulations using COMSOL®, re inductance, induction heating and magnetic fields". *COMSOL users conference, Stuttgart Germany*. 2011, 1–9 2.

[KH20] KIM, SEUNG-WOOK and HAN, JUNGHYUN. "Simulation of Arbitrarily-shaped Magnetic Objects". *Computer Graphics Forum*. Vol. 39. 7. Wiley Online Library. 2020, 119–130 1, 2.

[KH22] KIM, SEUNG-WOOK and HAN, JUNGHYUN. "Fast stabilization of inducible magnet simulation". *SIGGRAPH Asia 2022 Conference Papers*. 2022, 1–8 2.

[KMOW00] KANE, COURO, MARSDEN, JERROLD E, ORTIZ, MICHAEL, and WEST, MATTHEW. "Variational integrators and the Newmark algorithm for conservative and dissipative mechanical systems". *International Journal for numerical methods in engineering* 49.10 (2000), 1295–1325 3.

[KPH18] KIM, SEUNG-WOOK, PARK, SUN YOUNG, and HAN, JUNGHYUN. "Magnetization dynamics for magnetic object interactions". *ACM Transactions on Graphics (TOG)* 37.4 (2018), 1–13 1, 2.

[LFS*20] LI, MINCHEN, FERGUSON, ZACHARY, SCHNEIDER, TESEO, et al. "Incremental potential contact: intersection-and inversion-free, large-deformation dynamics." *ACM Trans. Graph.* 39.4 (2020), 49 3, 5.

[LKL*22] LAN, LEI, KAUFMAN, DANNY M, LI, MINCHEN, et al. "Affine body dynamics: Fast, stable & intersection-free simulation of stiff materials". *arXiv preprint arXiv:2201.10022* (2022) 3.

[MEM*20] MACKLIN, MILES, ERLEBEN, KENNY, MÜLLER, MATTHIAS, et al. "Primal/dual descent methods for dynamics". *Computer Graphics Forum*. Vol. 39. 8. Wiley Online Library. 2020, 89–100 3.

[MTGG11] MARTIN, SEBASTIAN, THOMASZEWSKI, BERNHARD, GRINSPUN, EITAN, and GROSS, MARKUS. "Example-based elastic materials". *ACM Trans. Graph.* 30.4 (July 2011). ISSN: 0730-0301. DOI: 10.1145/2010324.1964967. URL: https://doi.org/10.1145/2010324.1964967 3.

[NZWC20] NI, XINGYU, ZHU, BO, WANG, BIN, and CHEN, BAOQUAN. "A level-set method for magnetic substance simulation". *ACM Trans. Graph.* 39.4 (Aug. 2020). ISSN: 0730-0301. DOI: 10.1145/3386569.3392445. URL: https://doi.org/10.1145/3386569.3392445 2, 3.

[PGK*22] PADILLA, MARCEL, GROSS, OLIVER, KNÖPPEL, FELIX, et al. "Filament based plasma". *ACM Transactions on Graphics (TOG)* 41.4 (2022), 1–14 2.

[PLH16] PARK, JIYOUNG, LEE, KYUNGOK, and HAN, JUNGHYUN. "Interactive visualization of magnetic field for virtual science experiments". *Journal of Visualization* 19.1 (2016), 129–139 2.

[Rok85] ROKHLIN, VLADIMIR. "Rapid solution of integral equations of classical potential theory". *Journal of Computational Physics* 60.2 (1985), 187–207. ISSN: 0021-9991 10.

[Smi06] SMITH, RUSS. *Open Dynamics Engine*. 2006. URL: https://www.ode.org/ 8.

[SNZ*21] SUN, YUCHEN, NI, XINGYU, ZHU, BO, et al. "A material point method for nonlinearly magnetized materials". *ACM Trans. Graph.* 40.6 (Dec. 2021). ISSN: 0730-0301. DOI: 10.1145/3478513.3480541. URL: https://doi.org/10.1145/3478513.3480541 1, 3.

[TGPS08] THOMASZEWSKI, BERNHARD, GUMANN, ANDREAS, PABST, SIMON, and STRASSER, WOLFGANG. "Magnets in motion". *ACM Transactions on Graphics (TOG)* 27.5 (2008), 1–9 1, 2, 4, 5, 7, 8.

[YLUH14] YOON, WONBAE, LEE, NAMIL, UM, KIWON, and HAN, JUNGHYUN. "Computer-generated iron filing art". *The Visual Computer* 30.6 (2014), 889–895 2.