

Simulation of Curly Hair

5th Semester Project Laboratory Report

Barnabás Börcsök
borcsok.barnabas@simonyi.bme.hu

Dr. László Szécsi
Advisor




Figure 1. The achieved visual look

Abstract

A self-assessment of the 5th semester Project Laboratory Project is presented in this article. My goal was to achieve hair simulation of acceptable quality both in terms of look and performance. Multiple approaches were considered before arriving at a Position Based Dynamics based solution. The simulation was implemented in C++ with the Open Graphics Library (OpenGL ¹).

CCS Concepts: • **Computing methodologies** → *Physical simulation*.

Keywords: hair simulation, position based dynamics, OpenGL

¹<http://www.opengl.org>

1 Introduction

In the 5th semester of their undergrad studies, BSc students from Budapest University of Technology and Economics embark on their first journey of scientific research. I was always interested in and fascinated by computer graphics, so it came naturally to choose a subject in this area. As I had little hands-on experience in this field, a long time had to be dedicated to research and trying out different simulation methods.

The first part of this paper reflects this, giving an overview of considered methods, and other possible routes that could have been taken to implement hair simulation.

The implementation and all of the code mentioned is available at <http://git.sch.bme.hu/bobarna/brave-2>.

2 Overview of considered methods

There were mainly three methods considered, two of them being substantially different:

2.1 Mass Spring System

The whole idea of doing hair simulation as my project laboratory came after reading the paper by Iben et al. [1]. They outline a method for the artistic simulation of curly hair for use in their film production. The method was featured in the 2012 movie Brave.

This approach models the hair as a chain of particles with given mass, each connected via springs. This results in a mass-spring system, which is then modeled by considering well-known physics formula's such as Newton's second law of motion $F = m * a$, and Hooke's Law, which states that the force F needed to extend or compress a spring by some distance x scales linearly with respect to that distance. The paper by Iben et al. [1] created an elaborate system of springs, with additional core springs making the simulation stable and resulting in the desired look and feel.

The main characteristic of this approach is that a Mass Spring System deals with forces, and tries to stay true to the laws of physics. It accounts for internal and external forces from which accelerations are computed based on Newton's second law of motion. A time integration method is then used to update the velocities and finally the positions of the particles. This will not be the case with Position Based Dynamics.

2.2 Position Based Dynamics (PBD)

The paper by Müller et al. [2] presents an approach that omits the velocity as well as the force layer of the then-present popular approaches for simulation methods of dynamic systems in computer graphics.

Position Based Dynamics (PBD) – as its name suggests – works with the position of particles. A big advantage of such a system lies in its controllability, and the easy reducibility of overshooting problems present in force based systems. Another favorable aspect is that PBD methods are generally easier with the maths and somewhat easier to implement in terms of the needed mathematical and physical background needed to grasp its inner workings. In addition, – as it works directly with the position of particles – collision constraints can be handled easily and penetrations can be resolved completely by projecting points onto the penetrated surface. Although such measures can only be applied on the expanse of physical accuracy, position based dynamics being only a “good enough” approximation of how objects behave in real life.

Chapter 2 of the paper by Umenhoffer et al. [6] builds on the Müller et al. [2] paper and gives a concise overview of the position based dynamics method, and goes on to show the type of constraints that can be applied to control to behaviour of such a system – which is present in both papers.

2.3 Follow-the-Leader (FTL)

The Dynamic Follow-the-Leader (FTL) method outlined by Müller et al. [3] focuses on the fast simulation of hair and fur on animated characters. The sheer number of computation needed for simulating thousands of hair strands, each consisting of numerous particles presents a big challenge. Also, as each strand is inextensible, it is not trivial to come up with an algorithm that does its job in feasible time.

As it will be covered later in this article, the already introduced PBD method needs multiple iterations per frame in order to keep the system from stretching and becoming unstable. The FTL method by Müller et al. [3] presents a method that takes only a single iteration through the particles of each hair strand per frame to achieve the desired results.

As this sounds fascinating, and cuts the computation time needed substantially, in the early stages of my project, I implemented the FTL method alongside PBD. It *works*, and is fast, although due to the iteration count being one, when greater forces are being applied to the system, a substantial amount of stretching was introduced to the system. As it seemed in comparison that the “basic” PBD method yielded the same – or even more accurate – results, and it was feasible to allow myself to use multiple iterations per frame, I did not further investigate if there could have been improvements made to my implementation to eliminate the stretching in the presence of greater forces.

3 Position Based Simulation

3.1 Algorithm Overview

We model an S hair strand as a chain of N particles and a set of M constraints. Each particle $p \in [1, \dots, N]$ has three attributes: mass (m_p), position (\mathbf{x}_p) and velocity (\mathbf{v}_p).

A constraint $c \in [1, \dots, M]$ with cardinality n_c is a function $C_c : \mathbb{R}^{3n_c} \mapsto \mathbb{R}$. It operates on a set of indices $\{i_1, \dots, i_{n_c}\}, i_k \in [1, \dots, N]$. The constraint function also has a stiffness parameter $k_c \in [0..1]$ and a type of either *equality* or *inequality*.

Constraint c with type *equality* is satisfied if $C_c(\mathbf{x}_{p_1}, \dots, \mathbf{x}_{p_{n_c}}) = 0$. If its type is *inequality*, then it is satisfied if $C_c(\mathbf{x}_{p_1}, \dots, \mathbf{x}_{p_{n_c}}) \geq 0$. The stiffness parameter k_c defines the strength of the constraint in a range from zero to one.

Given these notations, the algorithm works in the following way:

Algorithm 1: pseudo code for the PBD simulation

```

(1) foreach  $p$  in particles do
(2)   | initialize  $\mathbf{p}_x = \mathbf{x}_p^0, \mathbf{p}_v = \mathbf{v}_p^0, p_w = 1/m_p$ 
(3) end
(4) loop
(5)   | foreach  $p$  in particles do
(6)     |  $\mathbf{p}_v \leftarrow \mathbf{p}_v + \Delta t \cdot p_w \cdot \mathbf{f}_{external}$ 
(7)   | end
(8)   | dampVelocities( $p_1, \dots, p_N$ )
(9)   | foreach  $p$  in particles do
(10)    |  $\mathbf{p}_{tmp} \leftarrow \mathbf{p}_x + \Delta t \cdot \mathbf{p}_v$ 
(11)   | end
(12)   | foreach  $p$  in particles do
(13)    | generateCollisionConstraints( $p_x \rightarrow$ 
(14)    |  $p_{tmp}$ )
(15)   | end
(16)   | loop numberOfIterations times
(17)    | projectConstraints( $C_1, \dots, C_{M+M_{coll}}, particles$ )
(18)    | /* only the temporary positions of particles are needed */
(19)   | end
(20)   | foreach  $p$  in particles do
(21)    |  $\mathbf{p}_v \leftarrow (\mathbf{p}_{tmp} - \mathbf{p}_x) / \Delta t$ 
(22)    |  $\mathbf{p}_x \leftarrow \mathbf{p}_{tmp}$ 
(23)   | end
(24)   | velocityUpdate( $particles$ )
(25)   | /* velocities of colliding particles are modified according to friction
(26)   | and restitution coefficients. */
(27) end

```

Lines (1)-(3) initialize the particles. We store the inverse masses w in order to be able to store infinitely heavy particles that are stationary during the simulation. ($\frac{1}{\infty} \approx 0$)

The core idea of position based dynamics is shown in lines (9)-(11), (15)-(17) and (18)-(21). In line (10), estimated positions for each particle p are calculated using an explicit forward Euler integration step. In line (16) the iterative solver manipulates these temporary position estimates such that they satisfy the constraints. It does this by repeatedly projecting each constraint in an iterative manner to similar to that of a Gauss-Seidel² solver (see section 3.2). Once the solver finishes with the iterations, in lines (19)-(20) each particle is moved to the calculated (once temporary) positions, and the velocity of each particle is updated accordingly.

Velocities are manipulated in lines (6), (8) and (19).

Line (6) allows to account for external forces in the simulation if some of the forces cannot be converted to positional constraints. (For example the pulling force between subsequent particles are modeled as distance constraints instead of external forces, and so are the forces generated by the particles colliding

with other objects.) We use it to add gravity, and to generate wind effects for demonstration purposes. If only the gravitational force is present, then the line becomes $\mathbf{p}_v \leftarrow \mathbf{p}_v + \Delta t \cdot \mathbf{g}$, where \mathbf{g} is the gravitational acceleration.

In line (8) the velocities can be damped, if necessary. Section 3.5 of the Müller et al. [2] paper gives some more sophisticated methods for damping, although the paper also points out that any form of damping can be used. In our case, simply applying a k_s damping coefficient with value between 98% to 99.9% seemed to be a good enough solution.

The initial constraints C_1, \dots, C_M are fixed throughout the simulation. In addition to these initial constraints, line (13) generates additional M_{coll} collision constraints. These change from time step to time step. The projection step in line (16) considers both the fixed and the collision constraints. Our implementation does not yet utilize this opportunity to dynamically generate collision constraints. This is one of the possible future works outlined in the subsection **Hair-hair Collisions**.

3.2 The Solver

The input to the solver are $M + M_{coll}$ constraints and the temporary (\mathbf{p}_{tmp}) positions of the particles. The aim of the solver is to modify the estimates such that they satisfy all the constraints. This yields a system of equations. For solving these for particle positions we use a sequence of Gauss-Seidel-type iterations. The original Gauss-Seidel algorithm can only handle linear systems, although we have a non-linear system of equations of hand. Even a simple distance constraint $C(\mathbf{p}_1, \mathbf{p}_2) = |\mathbf{p}_1 - \mathbf{p}_2| - d$ yields a non-linear equation. In addition, the constraints of type *inequality* yield inequalities.

The part we borrow from the Gauss-Seidel method is the idea of solving each constraint independently one after the other. This makes up the essence of the method: we repeatedly iterate through all the constraints and project the particles to valid locations with respect to the given constraint alone. Modifications to point locations immediately get visible to the process. This speeds up convergence significantly, because pressure waves can propagate throughout the chain of particles in a single solver step, and effect which is dependent on the order in which constraints are solved.

In our hair simulation case, we solve the distance constraints by putting the hair particles at the desired distance from each other starting at the scalp, and moving outwards from there.

For projecting the constraints, we utilize the method proposed by Müller et al. [2]. The details of computing the desired delta positions are explained in Appendix A **Constraint Projection**. An example derivation is given in Section 3.3 **Distance Constraint**.

²https://en.wikipedia.org/wiki/Gauss%E2%80%93Seidel_method

There are several ways of incorporating the stiffness parameter. The simplest – and the way we chose to utilize – is simply multiplying the $\Delta \mathbf{p}$ correction by the $k \in [0 \dots 1]$ stiffness parameter. For multiple iteration loops of the solver, the effect of k becomes non-linear. The remaining error after n_s solver iterations is $\Delta \mathbf{p}(1 - k)^{n_s}$. Müller et al. [2] propose to multiply the corrections not by k directly, but by $k' = 1 - (1 - k)^{1/n_s}$. With this transformation, the error becomes $\Delta \mathbf{p}(1 - k')^{n_s} = \Delta \mathbf{p}(1 - k)$, and thus linearly dependent on k and independent of n_s . The resulting stiffness is still dependent on the time step of the simulation. Real time environments typically use fixed time steps in which case this dependency is not problematic.

3.3 Distance Constraint

One of the most basic constraint is the Distance Constraint, forcing two given particles d distance apart. It is an *equality* type constraint, meaning that it is satisfied if $C_{distance} = 0$.

$$C_{distance}(\mathbf{p}_1, \mathbf{p}_2) = |\mathbf{p}_1 - \mathbf{p}_2| - d \quad (1)$$

We are looking for $\Delta \mathbf{p}_1$ and $\Delta \mathbf{p}_2$ corrections such that

$$C_{distance}(\mathbf{p}_1 + \Delta \mathbf{p}_1, \mathbf{p}_2 + \Delta \mathbf{p}_2) = 0.$$

In finding the desired delta values to satisfy the constraint, we follow the method proposed by Müller et al. [2] (see [Appendix Constraint Projection](#)).

The derivatives with respect to the positions are

$$\nabla_{\mathbf{p}_1} C_{distance}(\mathbf{p}_1, \mathbf{p}_2) = \mathbf{n}$$

$$\nabla_{\mathbf{p}_2} C_{distance}(\mathbf{p}_1, \mathbf{p}_2) = -\mathbf{n}$$

$$\mathbf{n} = \frac{\mathbf{p}_1 - \mathbf{p}_2}{|\mathbf{p}_1 - \mathbf{p}_2|}.$$

This makes the scaling factor

$$\begin{aligned} s &= \frac{C(\mathbf{p}_1, \mathbf{p}_2)}{w_1 |\nabla_{\mathbf{p}_1} C(\mathbf{p}_1, \mathbf{p}_2)|^2 + w_2 |\nabla_{\mathbf{p}_2} C(\mathbf{p}_1, \mathbf{p}_2)|^2} \\ &= \frac{|\mathbf{p}_1 - \mathbf{p}_2| - d}{w_1 + w_2}. \end{aligned}$$


The final corrections are:

$$\Delta \mathbf{p}_1 = -\frac{w_1}{w_1 + w_2} (|\mathbf{p}_1 - \mathbf{p}_2| - d) \frac{\mathbf{p}_1 - \mathbf{p}_2}{|\mathbf{p}_1 - \mathbf{p}_2|}$$

$$\Delta \mathbf{p}_2 = +\frac{w_2}{w_1 + w_2} (|\mathbf{p}_1 - \mathbf{p}_2| - d) \frac{\mathbf{p}_1 - \mathbf{p}_2}{|\mathbf{p}_1 - \mathbf{p}_2|}$$

Since the type of the Distance Constraint is *equality*, the projection is straight forward: we move both by \mathbf{p}_1 and \mathbf{p}_2 by $\Delta \mathbf{p}_1$ and $\Delta \mathbf{p}_2$ respectively.

Figure 2. Projection of the distance constraint $C(\mathbf{p}_1, \mathbf{p}_2) = |\mathbf{p}_1 - \mathbf{p}_2| - d$. The corrections $\Delta \mathbf{p}_i$ are weighted according to the inverse masses $w_i = 1/m_i$.



3.4 Bending Constraint

Umenhoffer et al. [6] use the PBD technique to simulate cloth materials. They utilize the bending constraint to define the bending resistance of the material, giving it an extra stiffness. Their goal with the Bending Constraint is to keep an initial angle between adjacent triangles formed by particles. They offer the $C(\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3, \mathbf{p}_4) = \arccos(\mathbf{n}_1 \cdot \mathbf{n}_2) - \phi_0$ constraint, where \mathbf{n}_1 and \mathbf{n}_2 are the normal vectors of the triangles, and $\mathbf{p}_2 - \mathbf{p}_1$ forms the common side of the two triangles. The ϕ_0 angle is measured between the normals of the two triangles.

At first, we thought about implementing something similar for our project, but then simplified the “Bending Constraint” for our use-case to essentially distance constraints $C_{distance}(\mathbf{p}_{i-1}, \mathbf{p}_{i+i})$ for $i \in [1 \dots n_{particles} - 1]$ of the same hair strand with the d distance parameter of the being lower than the d parameter of other $C_{distance}$ constraints keeping the particles a given l_{seg} distance apart from each other, and assigning a lower $k_{bending}$ stiffness parameter.

This idea could be further developed when considering other $(\mathbf{p}_j, \mathbf{p}_k)$ pairs and different distances between them. Figure 3 illustrates some example variety in the behaviour of chain of particles under different bending constraints.

3.5 Collision Constraint

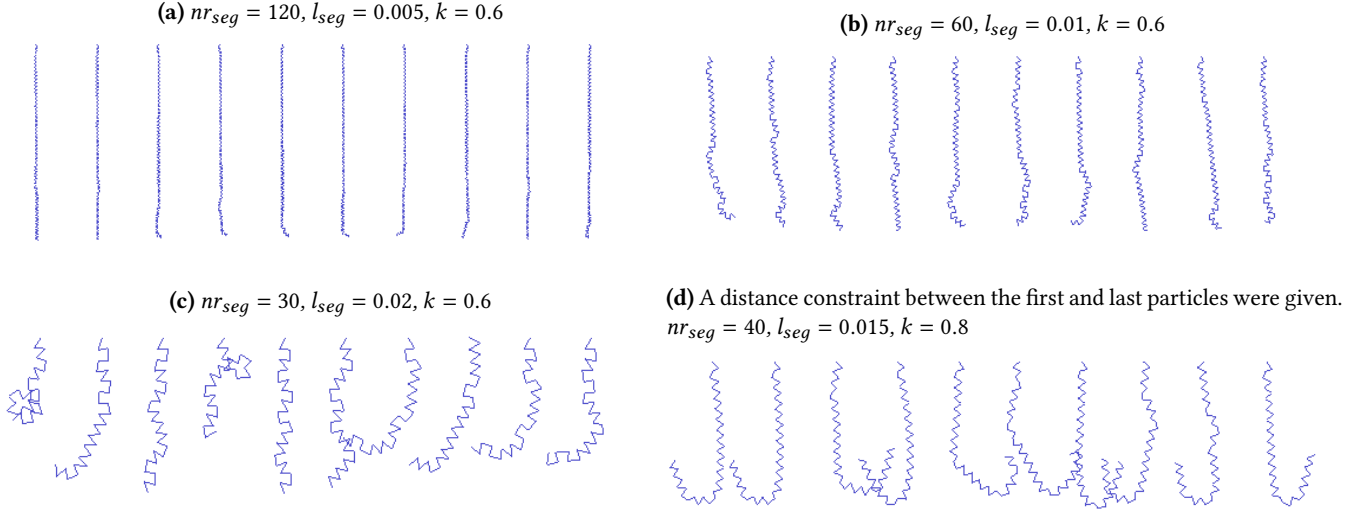
One huge advantage of the position based approach is how simply collision handling can be realized. In each simulation step M_{coll} collision constraints are generated in the (12)-(14) lines of Algorithm 1. For a moving particle \mathbf{p} and points $\mathbf{q}_1, \mathbf{q}_2, \mathbf{q}_3$ forming a triangle with normal \mathbf{n} , the collision constraint is

$$C_{collision}(\mathbf{p}, \mathbf{q}_1, \mathbf{q}_2, \mathbf{q}_3) = (\mathbf{p} - \mathbf{q}_1) \cdot \mathbf{n} \geq 0$$

, where we check whether \mathbf{p} lies on the desired side of the given surface. We require that the magnitude of the vector from an arbitrary point on the surface to \mathbf{p} projected onto the normal vector be greater than 0 by utilizing the dot product operation.

The formula for solving the triangle collision constraint is:

Figure 3. Different distance and particle pairing definitions yield significantly different results. Particle masses were slightly randomized for a variance between strands. In each below simulation, nr_{seg} particles of length l_{seg} were generated. Distance constraints acting as bending constraints had a stiffness parameter k . Distance constraints responsible for keeping the particles l_{seg} distance apart always had a $k_{pos} = 1$ stiffness parameter.



$$\Delta \mathbf{p}_{tmp} = -((\mathbf{p}_{tmp} - \mathbf{q}_1) \cdot \mathbf{n}) \cdot \mathbf{n}$$

This formula moves the particle along the triangle normal exactly to the triangle's surface.

3.6 Position Constraint

Maybe the most straight-forward constraint is the Position Constraint, given as:

$$C_{collision}(\mathbf{p}) = |\mathbf{p} - \mathbf{p}_{goal}| = 0$$

$$\Delta \mathbf{p} = \mathbf{p}_{goal} - \mathbf{p}$$

Moving the \mathbf{p} particle to the \mathbf{p}_{goal} position.

We utilize the Position Constraint for keeping the first particle in each strand on the surface of the head.

4 Implementation

After introducing the general position based dynamics simulation method, we dive into the specifics we settled on for our hair simulation.

The code is available at <http://git.sch.bme.hu/bobarna/brave-2>.

4.1 Data Structures

The properties of particles introduced in chapter 3.1 [Algorithm Overview](#) got an additional color property. This decision was made so hair with a gradient color could be simulated.

```
class Particle {
```

```
    vec3 pos;
    vec3 tmp;
    float w;
    vec3 v;
    vec3 color;
}
```

The strands are stored as a vector of vector of Particles. It is practical to keep related data packed right after each other in the memory for improving the performance of operations and function calls.

4.2 Constraints

A Constraint class with an abstract `solve()` method was implemented, storing each derived constraint class (e.g. *PositionConstraint*, *BendingConstraint*) in a heterogeneous collection and calling each overridden `solve()` method when projecting the constraints in the iteration loop (line (16) of Algorithm 1).

This implementation proved to be not fast enough, lagging even with low particle counts. We could not discover the cause of this increased computational time, and decided to settle on another solution. Instead of the above mentioned, we store only the parameters that constraints operate on, and the simulation calls the relevant solve functions with these stored parameters.

4.3 Putting the Hair in the Hair Simulation

The *PBDSimulation* class handles the simulation and rendering of the hair. It takes a *HeadObject* and propagates

hairstrands starting from its surface. Each strands first particle is fixed in place with a *PositionConstraint*, not letting it move away from the scalp.

The geometry of the *HeadObject* is read from an OBJ file. (See the Appendix B.1 OBJ Reader.)

The hair strands are propagated throughout the surface of the *HeadObject* by proposing random $(\mathbf{u}, \mathbf{v}) \in [0 \dots 1] \times [0 \dots 1]$ coordinates and settling on them if they satisfy requirements for the position of the hair – like the corresponding face in world coordinates has an upward pointing normal vector. This propagation of strands on the surface of the head could be improved, as pointed out in sections [More and better customization](#) and [Sequence for randomizing hair positions](#).

5 Results

The implementation is available at git.sch.bme.hu/bobarna/brave-2. The program is able to handle up to around 500 individual hair strands, each consisting of around 30 particles in real time. With the implemented OBJ reader, it is possible to achieve diverse visual results in no time.

An example result shown in Figure 1, and can be seen in the supplementary videos [available here](#).³ An 8-minute video was also made to supplement this writing. The oral assessment in Hungarian is available at <https://drive.google.com/file/d/1sKdWr3YFZV9OKiugLX7eyXjixu4DY0Yq/view?usp=sharing>.

6 Conclusion

After investigating and trying out multiple methods to simulate hair, we arrived at the Position Based Dynamics (PBD) method. An overview of the PBD method, and an introduction to the constraint types used in our implementation was given. An implementation was made in C++ and OpenGL, which easily simulates the hair strands in real time, although the limit of the real-time method can easily be surpassed when adding around a thousand strands each with more than 30 particles.

7 Future Work

As this was part only of a semester-long Project Laboratory course, a great deal of time was dedicated in the early stages of the project to investigate and explore the subject area. This left only a sub-optimal time for implementing all desired aspects of the simulation.

7.1 Hair-hair Collisions

As it is computationally expensive to simulate each particle colliding with all of the other particles, a particle density

field could be used. One such method of grouping hair particles and their corresponding velocities into a 3D voxel grid is outlined in the paper by Petrovic et al. [4], which was also utilized in the Müller et al. [3] paper for handling hair-hair collisions. Such a method could be used to dynamically generate collision constraints in line (13) of Algorithm 1.

7.2 Bending Constraint Improvements

As already mentioned in Section 3.4 [Bending Constraint](#). A huge improvement could be made by improving the bending constraint either by further exploring the possibilities of defining different distance constraints between pairs of particles.

An alternative route could be to explore the possibility of a different type of bending constraint, for example by keeping a given angle between chunks of subsequent particles.

7.3 More and better customization

A drawable texture map is considered for the ability to tell the simulation where to put the hair strands on the surface of the head.

The OBJ reader could also be further developed into reading material, texture and/or color data for the objects.

A common interface for importing and exporting different hairstyles could be developed.

7.4 Better customizability

In the current implementation there is no way to properly customize the look and feel of the hair without accessing the source code. A possible improvement could be to read material and texture data for the hair from the OBJ file. Another huge improvement could be to set the constraints in real time, giving the user much better control of the hair. Other properties, such as length of the hair segments and color of the hair could be set as well.

7.5 Sequence for randomizing hair positions

In the current implementation, a standard C++ random number generator is utilized. Utilizing a low-discrepancy number sequence such as the Halton sequence for randomizing hair strand positions would be an improvement.

7.6 Rendering

This semester left no time to try out different methods for rendering and visualizing the hair. The current implementation connects the individual particles with `GL_LINES`, resulting in a sub-optimal visual experience. Offline (e.g. ray-tracing) methods could utilize the current simulation method. The paper by Rapp [5] discusses the real-time aspect of hair rendering in great detail, supplemented with a broad discussion

³https://drive.google.com/drive/folders/1n79xSokrBe2lKQjxX_DacF7wByXkbpZP?usp=sharing

on different parts of the (OpenGL) rendering pipeline that could be utilized.

7.7 Moving to the GPU

The current implementation runs solely on the CPU. A big leap forward would be the utilization of the GPU. Compute shaders for computations are considered. The paper by Umenhoffer et al. [6] discusses a parallelization of the different constraint solves, although as they work on the same dataset, some computation problems are present during the parallelization of the problem.

Tessellation and geometry shaders could be utilized for the rendering and display of the particles, resulting in a far better image with the same simulation method and properties.

Acknowledgments

To Dr. László Szécsi, associate professor at the Computer Graphics Group (Department of Control Engineering and Information Technology), whom I had the chance to consult with during the semester.

References

- [1] Hayley Iben, Mark Meyer, Lena Petrovic, Olivier Soares, John Anderson, and Andrew Witkin. 2013. Artistic Simulation of Curly Hair. In *Proceedings of the 12th ACM SIGGRAPH/Eurographics Symposium on Computer Animation (Anaheim, California) (SCA '13)*. Association for Computing Machinery, New York, NY, USA, 63–71. <https://doi.org/10.1145/2485895.2485913> <http://graphics.pixar.com/library/CurlyHairB/paper.pdf>.
- [2] Matthias Müller, Bruno Heidelberger, Marcus Hennix, and John Ratcliff. 2007. Position Based Dynamics. *J. Vis. Comun. Image Represent.* 18, 2 (April 2007), 109–118. <https://doi.org/10.1016/j.jvcir.2007.01.005>
- [3] Matthias Müller, Tae Kim, and Nuttapon Chentanez. 2012. Fast Simulation of Inextensible Hair and Fur. *VRIPHYS 2012 - 9th Workshop on Virtual Reality Interactions and Physical Simulations*. <https://doi.org/10.2312/PE/vriphys/vriphys12/039-044>
- [4] L. Petrovic, M. Henne, and John Anderson Pixar. 2006. Volumetric Methods for Simulation and Rendering of Hair. <https://graphics.pixar.com/library/Hair/paper.pdf>
- [5] M. Rapp. 2014. Real-Time Hair Rendering. <http://markusrapp.de/wordpress/wp-content/uploads/hair/MarkusRapp-MasterThesis-RealTimeHairRendering.pdf>
- [6] Tamás Umenhoffer, Artúr M. Marschal, and Péter Suti. 2016. Simulation methods for elastic and fluid materials. http://cg.iit.bme.hu/~umitomi/publications/GRAFGE02016_PhysicsSimulators.pdf.

A Constraint Projection

Projecting a set of points according to a constraint means moving the points such that they satisfy the constraint. As we move the points in the simulation loop directly, the most important issue is the conservation of linear and angular momentum. Let $\Delta \mathbf{p}_i$ be the displacement of particle i by the projection. Linear momentum is conserved if

$$\sum_i m_i \Delta \mathbf{p}_i = \mathbf{0}, \quad (2)$$

which amounts to conserving the center of mass. Angular momentum is conserved if

$$\sum_i \mathbf{r}_i \times m_i \Delta \mathbf{p}_i = \mathbf{0}, \quad (3)$$

where the \mathbf{r}_i are the distances of the \mathbf{p}_i to an arbitrary common rotation center. If a projection violates either the preservation of linear or angular momentum, *ghost forces* get introduced into the system, which act like external forces dragging and rotating the object. However, only internal constraints need to conserve the momenta. As we reduce the problem of collision handling to simply moving the vertex to a valid point, it is somewhat trivial that for example collision constraints are allowed to have global effects on the object.

The method that Müller et al. [2] propose conserves both momenta for internal constraints. Let us look at a constraint C with cardinality n on the points $\mathbf{p}_1, \dots, \mathbf{p}_n$ with stiffness k . We let \mathbf{p} be the concatenation $[\mathbf{p}_1^T, \dots, \mathbf{p}_n^T]^T$. For internal constraints, C is independent of translation and rotation, which means that rotating or translating the points does not change the value of the constraint function. Therefore, the gradient $\nabla_{\mathbf{p}} C$ is the direction of maximal change. If the correction $\Delta \mathbf{p}$ is chosen to be along $\nabla_{\mathbf{p}} C$ both momenta are automatically conserved if all masses are equal (we handle different masses later). Given \mathbf{p} we want to find a correction $\Delta \mathbf{p}$ such that $C(\mathbf{p} + \Delta \mathbf{p}) = 0$. This equation can be approximated by

$$C(\mathbf{p} + \Delta \mathbf{p}) \approx C(\mathbf{p}) + \nabla_{\mathbf{p}} C(\mathbf{p}) \cdot \Delta \mathbf{p} = 0. \quad (4)$$

Restricting $\Delta \mathbf{p}$ to be in the direction of $\nabla_{\mathbf{p}} C$ means choosing a scalar λ such that

$$\Delta \mathbf{p} = \lambda \nabla_{\mathbf{p}} C(\mathbf{p}). \quad (5)$$

Substituting Equation 5 into Equation 4, solving for λ and substituting back into Equation 5 gives us the final formula for $\Delta \mathbf{p}$:

$$\Delta \mathbf{p} = -\frac{C(\mathbf{p})}{|\nabla_{\mathbf{p}} C(\mathbf{p})|^2} \nabla_{\mathbf{p}} C(\mathbf{p}) \quad (6)$$

which is a regular Newton-Raphson step for the iterative solution of the non-linear equation given by a single constraint. For the correction of an individual point \mathbf{p}_i we have

$$\Delta \mathbf{p}_i = -s \nabla_{\mathbf{p}_i} C(\mathbf{p}_1, \dots, \mathbf{p}_n), \quad (7)$$

where the scaling factor

$$s = \frac{C(\mathbf{p}_1, \dots, \mathbf{p}_n)}{\sum_j |\nabla_{\mathbf{p}_j} C(\mathbf{p}_1, \dots, \mathbf{p}_n)|^2} \quad (8)$$

is the same for all points. If the points have individual masses, we weight the corrections $\Delta \mathbf{p}_i$ by the inverse masses

$w_i = 1/m_i$. In this case a point with infinite mass, i.e. $w_i = 0$ for example does not move, just as expected. Now Equation 5 is replaced by

$$\Delta \mathbf{p}_i = \lambda w_i \nabla_{\mathbf{p}_i} C(\mathbf{p})$$

yielding

$$s = \frac{C(\mathbf{p}_1, \dots, \mathbf{p}_n)}{\sum_j w_j |\nabla_{\mathbf{p}_j} C(\mathbf{p}_1, \dots, \mathbf{p}_n)|^2} \quad (9)$$

for the scaling factor and for the final correction

$$\Delta \mathbf{p}_i = -s w_i \nabla_{\mathbf{p}_i} C(\mathbf{p}). \quad (10)$$

Section 3.3 gives an example of calculating $\Delta \mathbf{p}$ by considering the distance constraint function $C(\mathbf{p}_1, \mathbf{p}_2) = |\mathbf{p}_1 - \mathbf{p}_2| - d$.

B Supplementary development

B.1 OBJ Reader

An OBJ reader utility was implemented as part of the project. As the OBJ file format⁴ describes a wide range of properties for objects, our OBJ reader handles only the subset of these description options – those that are needed for the project. Unsupported line types were simply ignored, still resulting in a successful file read if possible.

Table 1. Supported OBJ Data Types

Type of data	Format
Comment	# comment
Geometric vertex data	v x y z
Texture coordinates	vt u v
Vertex normals	vn i j k
Triangular faces	f v1/vt1/vn1 v2/vt2/vn2 ... v3/vt3/vn3

- “v x y z” defines a vertex with position (x, y, z).⁵
- “vt u v” defines a 2D texture coordinate with position (u, v).⁶
- “vn i j k” specifies a normal vector with components i, j and k.
- “f v1/vt1/vn1 v2/vt2/vn2 v3/vt3/vn3” Specifies a face with the given indices. For the first vertex, the *v1st* previously defined vertex position is used, the *vt1st* texture coordinate and *vn1st* normal vector is used.⁷

⁴<http://paulbourke.net/dataformats/obj/>

⁵The obj file format allows for a *w* coordinate (also called the weight) for describing rational curves and surfaces. The default value of *w* is 1.0. My implementation handles only the assignment of *x*, *y* and *z* values.

⁶The obj file format allows for 3D texture coordinates as well, but this implementation handles only 2D texture coordinates.

⁷The obj file format allows for face definitions far beyond only triangles, although this implementation handles only face definitions of the above format.

B.2 Recording the simulation on-the-fly

A recording utility was made to work in tandem with the application. In essence, the rendered images are written into bmp or png files if the capturing is enabled. In the implementation⁸, toggling the capturing is mapped to the C key.

For handling the export of the rendered image into files, the `stb_image_write.h`⁹ single-file public domain library was used.

The program outputs the image files in the `renders` folder at 24 FPS, with the naming convention `renderXXXX[.bmp/.png]`, where XXXX is the number of the frame with left-padded zeroes.

After the image files are written, the user can start the `make_video.sh` script to assemble the output `.mp4` and delete all the rendered frames. This script uses the free and open-source `ffmpeg` command line utility¹⁰

⁸<https://git.sch.bme.hu/bobarna/brave-2>

⁹https://github.com/nothings/stb/blob/master/stb_image_write.h

¹⁰<https://ffmpeg.org/>