Ph.D. Thesis Evaluation

Thesis Title: Natural GPU-friendly Dynamic Hair Animation
Candidate Name: Petr Kmoch

Summary
This thesis presents methods for hair animation. More precisely, the work uses two observations about the behaviour of real hair and uses that to simplify two existing hair simulation methods. The two observations are:

1. Natural human hair has an elliptical cross-section and is more likely to bend over its major axis. (Section 3.1.2)
2. In may hairstyles, hair strands often stick together to form flat ribbon-like wisps. (Section 3.2.1)

The first observation is used to simplify the Super-Helix hair model, originally presented by Bertails et al. (2006) and the Discrete Elastic Rods model, originally presented by Bergou et al. (2008).

The second observation is used to simplify hair volume representation, during dynamics and collision computations and during rendering.

Finally, the thesis also presents details about a GPU-based implementation of the simulation model, based on nVidia's CUDA platform.

Detailed Evaluation
The thesis claims to have made the following contributions (Section 1.5):

1. Devised new approach to modelling hair bending behaviour and applied it to two existing hair animation methods.
2. Proposed a specific representation of hair volume that captures large number of hairstyles and lends itself to efficient simulation, collision detection and response.
3. GPU implementation of hair simulation.

In light of these, the following points about the thesis are of note:

1. In Chapters 1 and 2, the thesis covers the existing literature in the field of hair simulation thoroughly. These chapters are useful in putting the present work in context.

2. Chapter 3, Sections 3.1 and 3.2 develop the general ideas involved in hair simulation and modelling. Sections 3.3 and 3.5 present, in a lot of detail, the two existing hair animation methods that have been extended in this thesis.
   a) Even though these are already well known and documented in previous literature, the thesis presents them for completeness and that is appreciated.

3. Section 3.4 presents the modification done to the Super Helix model of hair.
   a) The model uses bending over the minor axis as an explicit component, so the change is very simple to make, by setting this bending to zero.
   b) The advantages of this making this change are shown implicitly by the usage of this in haptic interaction with this model (Bonnani and Kmoch 2008). Though smooth haptic response needs a model that can be efficiently evaluated to avoid lag in the interaction experience, exact speed-up advantage offerred by the elimination of bending about the minor axis can only be done by comparing runtime performance of the original super

helix model with the modified one. This is not presented in the thesis.

c) Section 3.4.3 is titled "Evaluation." However, it talks less about evaluation of the implemented model but more about why collision handling is difficult using the super helix model and the computationally expensive reconstruction that has to be done to recover the shape of the recovered hair. This sets up the stage for the second simulation model that this thesis modifies, but it does nothing toward evaluating the first one. So the title of this section should be changed to something more appropriate, and a separate evaluation section should be added that truly evaluates the performance of the modifed super helix model in terms of runtime efficiency and visual quality of the resulting animation.

4. Section 3.6 presents the modification done to the Discrete Elastic Rods Model.
   a) The thesis presents an algorithm for computing twisting of hair strands that replaces the original iterative approach by a direct, explicit solution that minimizes bending over the minor axis.
   b) This depends on the fact that bending stiffness over the minor axis is much larger than the twisting stiffness.
   c) Quantitative comparison with the Newton minimization method is made (Table 3.1) and approximate comparison with the original paper by Bergou et al. is also made (Table 3.2).
   d) Section 3.6.2 talks about the head-hair collision mechanism that is a solved using constraint enforcement during the dynamics simulation. Table 3.3 presents quantitative evaluation of how much additional computation needs to be done when collision enforcement is enabled. It is interesting to see the effect of the detection margin on the computation effort.
   e) Section 3.6.3 decribes the use of "wisps" in representing hair volumes.
      1. The wisps are rendered as a fan of four triangles per quad, where the entire wisp is a strip of quads. These are then texture/relief mapped appropriately. Can geometry shaders be used to dynamically generate individual hair strands at runtime?
      2. The wisps also participate in dynamics via wisp springs between the rim strand and collisions between the wisps. Do the wisp springs and the rendering triangles match up exactly or their dimensions and placements are completely independent of each other? If so, then how does this effect the capability of the rendering proxy in approximating the computed deformation of the wisp?
   f) Section 3.7 describes a prototype implemention of the hair simulation method described in Section 3.6.
      1. The CUDA computation and memory model is presented first, followed by description of the six kernels used to implement major part of the simulation on the GPU.
      2. Rendered output from the GPU-based simulation is shown in Figure 3.22. The wisps appear at very regular intervals in the images shown. Perhaps larger images can be shown to appreciate the quality of output better.

5. Chapter 4 presents the technique used to resolve hair-hair collisions.
   a) The primary acceleration primitive used here are the wisps themselves as collision is computed among wisps and not individual hair strands.
   b) Collisions are classified as aligned and unaligned collisions depending on relative orientation of wisps when they collide.
   c) Collisions are accelerated by buidling a AABB hierarchy over the wisp segments.

6. Chapter 6 describes the simulation results.
   a) The evaluation in Section 5.2 must also specify which specific GPU was used to

generate the results and which version of CUDA was used.

b) Screenshots from all scenarios should be shown, instead of only F1, otherwise the difference between them cannot be fully appreciated.
c) Absolute and relative time taken by each stage of the simulation loop are given in Tables 5.2 and 5.3 respectively.
d) The thesis should include links to video/animation results.
e) It is unfortunate that the collsion accleration structure chosen makes the method non-iteractive. There is a lot of emphasis in the thesis on the fact that the method is amenable for a GPU implemetnation. That this does not lead to an interactive implementation finally, makes the thesis feel incomplete. If the AABB can be replaced by any other acceleration structure, then it should be done and the enchanced results should be included in the thesis.
f) Some comparison with current work in the area like the ones given below would benefit the thesis.
    1. A Reduced Model for Interactive Hairs, Chai et al, SIGGRAPH 2014 (referred in the thesis)
    2. Position-based Elastic Rods, Umetani et al., SCA 2014

7. The thesis must include a future work section where future avenues of research for this problem are discussed. What are the challenges that are left unexplored by this thesis and what are the open problems in the area of hair simulation that the author would like to address in the future?

8. The thesis is very well written. There are a very few typographical errors that a spellchecker should be able to catch.

Final Recommendation

This thesis presents some original and interesting ideas. The ideas are inspired from a non-computer science/mathematics disclipline and used to simplify existing hair simulation models. This amply demonstrates the author's capability for interdisciplinary research in challenging problem areas. The presentation of ideas shows theoretical maturity and the implemented prototype demonstrates an ability to put that theory to practice. The ideas should also be easy to incorporate into existing hair simulation systems, thus making their introduction to computer graphics relevant.

A few suggestions and questions are elucidated in the detailed comments above. If the thesis is modified to take them into account, primarily adding links to video and animation results and implementing a better collision acceleration structure, it would make the thesis more complete and relevant.

Therefore, I recommend that the work presented in this thesis be considered for the award of a doctoral degree to Petr Kmoch.

Parag Chaudhuri
Associate Professor,
Department of Computer Science and Engineering, IIT Bombay.

Date: 10/8/2015