Levels of Detail in Reducing Cost of Haptic Rendering: a Preliminary User Study

Jian Zhang, Shahram Payandeh and John Dill
Experimental Robotics and Graphics Laboratories,
School of Engineering Science
Simon Fraser University
Burnaby, BC V5A 1S6, Canada
{jzhanga/shahram/dill}@cs.sfu.ca

Abstract

Haptic rendering complex objects in virtual environments is computationally intensive. In this paper we start the investigation of a new category of approach to reducing the computation in haptic rendering. Our approach is based on the hypothesis that the accuracy of haptic perception might be limited. Results of the experiments described in this paper suggest that subjects might not be able to distinguish two haptic objects if they are beyond some refinement level. This limitation of haptic perception may be taken advantage of in haptic rendering by replacing a fine object with a coarser object to reduce scene complexity.

1. Introduction

In recent years, haptic feedback has been increasingly used in virtual environment simulations. This modality provides another channel of information presentation in a variety of applications. In order to achieve realistic haptic feedback, the force information should be sent to the haptic device at a frequency as high as 1kHz for the user to perceive the forces as smooth and realistic [24]. This indicates the mathematical equations governing the physics models involved in that need to be iteratively solved at the same high rate. Thus, simulating a complex scene with objects with different physical properties is computationally intensive. A surface model of a deformable object would involve at least hundreds of nodes to present realistic force feedback and surface deformation. For every node, differential equations must be solved. If other computationally demanding tasks such as collision detection and graphic rendering are also present, performing the overall complicated loop in less than 1ms is difficult.

To reduce the amount of computation yet maintain plausible results, effort has been put into finding a better trade-off between simulation accuracy and update rate. For example, [10] used "finite spheres" as an approximation to FEM in a local area; [11], [12] interpolated forces between calculated forces from a low update rate deformation model to feed the demand of high update rate haptic simulation.

Another approach to speed up run-time simulation is to isolate certain procedures in real-time computation, precompute them and later combine precomputed results with states of a simplified run-time model. A boundary element method to compute system response of a linear model offline was proposed in [13]. Precomputation of "elementary deformation" is an important part of [15] for a hepatic surgery simulation system. When the simulated object is represented as polygonal meshes, adaptive refinement can be applied to present the operator more details in the local interaction area, while maintaining simplicity in other parts of the object [9], [16], [17], [18].

This paper is a preliminary investigation of the effect of varying haptic Level of Detail as a way of reducing computational load, using the method developed in [9] for computing force parameters for varying mesh fineness. The operators of virtual environment systems are human beings, who have limited accuracy in visual perception, touch, hearing, smell and taste. We hypothesize that beyond some refinement level, additional details on a haptically rendered object would make no difference to the operator's haptic perception of the object. If two objects with different levels-of-detail are perceived the same haptically, we can adopt the coarser one to reduce the complexity of the haptically rendered scene and reduce the complexity of the simulation.

Some research has been done on human perception of a virtual environment where force feedback devices are used for exploration. It has been shown that force or force-related information can overcome graphical geometry information such as slope when the operator is trying to determine the shape of the object [3]. The abilities of a few haptic textures to enhance the user's performance in a combined graphic and haptic computer user interface are examined [5]. In some situations visual cues can increase human somatosensation acuity [6].

Work closely related to human perception of surface roughness with a rigid link, which can be a link of a force feedback device, includes [1], [7], [8]. West and Cutkosky [7] describes an experiment in which the subjects were asked to feel walls shaped as sinusoidal waves by holding a handle, which is connected to a device that can generate forces. The results showed that the amount of time needed for exploring physical walls and virtual walls of the same

sinusoidal shape is essentially the same. A similar conclusion was also reached in [8]. West and Cutkosky [7] reported that when exploring with a fingertip, less time is needed than with a rigid link. The experiments in [1] showed that a bare finger was more sensitive to the magnitude of roughness than probes. Pooling [1], [7], [8] together we see that exploration of surfaces with a stylus is different than with a bare finger in both time and accuracy, but exploring a real physical surface with a stylus and a virtual surface with a haptic device does not appear to be so different.

The previously mentioned research efforts did not show the limited accuracy of human haptic perception of virtual surfaces via a haptic feedback probe. So far, research in this field is still in its early stage. Yamashita, et al [4] conducted an experiment involving users feeling a virtual cylinder via a force feedback device, and found that a "bump" height of 0.1mm is the threshold for human perception of haptic smoothness of the cylindrical surface. However, details of the experiment and analysis were not provided.

In this study we try to determine (for a specific commercially available haptic device [23]) if there are refinement levels of surface resolution beyond which higher levels of detail no longer make a difference to the user. We conducted three experiments on rigid virtual surfaces and deformable virtual surfaces using a haptic feedback device, Phantom [23]. In all three experiments we provided the subjects with graphic representations of the virtual objects in addition to the haptic rendering of the objects.

In the first experiment we asked the subjects to compare rigid objects of the same shape but with different smoothness levels. The smoothness levels are determined by the numbers of polygons on the models. In experiment 2 the subjects were asked to compare deformable objects before and after global refinement. In experiment 3 the deformable objects with different levels of detail were compared. The results of this preliminary study show that replacing fine models with coarser ones in haptic rendering without the operator noticing the difference is possible under certain conditions. Such replacement can reduce computational cost in haptic rendering.

The paper proceeds as follows: System configuration and participants are described first since they are the same for all three experiments. Then the purpose, hypothesis, experiment design and results of each experiment are described. We follow these with discussion and conclusion sections.

2. System configuration

An SGI workstation (Octane) with an R10000 processor and 128 MB of memory and a Phantom 1.5 Premium, a point force feedback device with 6 degrees of freedom in position and rotation and 3 degrees of freedom

in force output, were used in the experiments. The position sensing resolution of the Phantom is 0.03mm. The position of the end point of the Phantom stylus held by the operator was input as the location where the operator is interacting with the virtual environment. The computer calculates the force and position of the contact point for the deformable model. The force data is then sent to the Phantom and the actuators display this force by exerting it to the user's hand.

The software was written in C++, using GHOST (API for the Phantom) for haptic rendering and OpenGL, GLUT and GLUI for graphic rendering. Figure 1 shows the interface used in the experiments. In different experiments different objects are displayed.
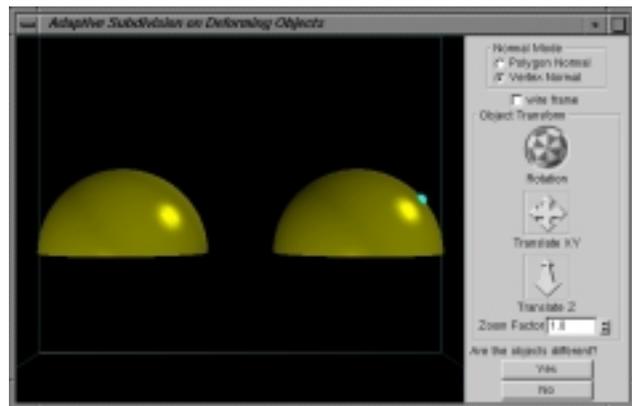

Figure 1: screen shot of the interface used in experiments

3. Participants

The participants of all three experiments were the same 14 males and 14 females between the ages of 18 and 30. All except one were students at Simon Fraser University in Engineering Science, Computing Science, Communications, Kinesiology, Criminology, Physics, and Chemistry. All except one were right handed. All except one had nothing more than trivial previous exposure to the Phantom device.

4. Experiment 1

In experiment 1 we asked if there exists a haptic smoothness level above which people can no longer feel the difference between levels of smoothness in rigid object simulation, given the same visual geometry. Since the objects were modeled as polygonal meshes, perceived smoothness of the objects was related to the number of polygons in the meshes. We presented the subjects with two different levels of smoothness each time, one of which was always the finest, and asked them to answer the question "Are the objects different?". From the frequencies of "yes" and "no" answers for a given pair of smoothness levels, we wish to find a smoothness level, above which further haptic smoothing of the object

surface has no significant effect. An object above such a smoothness level may be replaced by an object just above the level to reduce computation in haptic rendering.

4.1 Experiment design

Eight objects representing different levels of haptic smoothness were used to approximate a hemisphere, 50mm in radius. Each object presented to the subject consisted of a haptic model and a graphic model, both of which were polygonal meshes. The haptic model was responsible for force feedback and was of one of eight levels of smoothness. The graphic model was independent of the haptic model, and always modeled with 2560 polygons. We implemented the graphic model this way so that all the eight objects looked exactly the same, thus eliminating possible visual effects on the subjects' choices. The haptic and graphic models were spatially co-located so that the subjects felt the hemisphere where they saw it. Table 1 gives the number of triangles used in each model.

Table 1: Numbers of triangles in each haptic mesh and graphic mesh for each object.

objects	haptic model	graphic model
h_1	40	2560
h_2	160	2560
h_3	360	2560
h_4	640	2560
h_5	1000	2560
h_6	1440	2560
h_7	1960	2560
h_8	2560	2560

In each trial the subject saw and felt two hemispheres. One of them was always object h_8; the other was one of the eight objects chosen at random. With eight possible such pairs presented to the subject three times each, there were 24 trials in this experiment. The order of the trials and the relative positions of the two hemispheres to each other were randomized. The graphic display in experiment 1 is shown in Figure 1.

For each trial, starting from the appearance of the objects, subjects were allowed ten seconds to feel the surface of the two objects. When the ten seconds were up, the subject had to click the "yes" or "no" button on the right hand side of the window, under the displayed question "Are the objects different?" (Figure 1). We set a time limit because if there were none, the subject might start taking too many factors into consideration, such as small vibrations of the Phantom device as one presses on a hard surface; or become so suspicious that they started sensing "small differences" that did not actually exist. We set this time limit also because we wanted the experiment to be completed within a reasonable amount of time

before fatigue degraded subjects' performance. Since ten seconds is not a long time, we provided a training session before the experiment where the subject could practice the trials until they could perform the task within ten seconds without difficulty. All of the subjects were able to perform the task proficiently after the training session, although they spent varying amounts of time in training.

Each click on the "yes" or "no" button was recorded by the software. If the subject accidentally clicked on the wrong button, they could subsequently correct their choice. Only the last button click in each trial was taken into account in data analysis. After clicking the button, the subject ended the current trial by closing the window. The window of the next trial then automatically appeared, signifying the start of a new trial.

4.2 Results

Figure 2 shows a histogram of the results for experiment 1. The general trend of the chart indicates that the further the object was from h_8 (i.e. the coarser the object), the more likely subjects were to realize the difference. This agrees with our intuition. We note that even when h_1 was presented with h_8, where the numbers of polygons in the haptic models differed by a factor of 64, subjects occasionally still felt they were the same. We also note that even when h_8 was presented with h_8, about 20% of the times subjects judged them different.
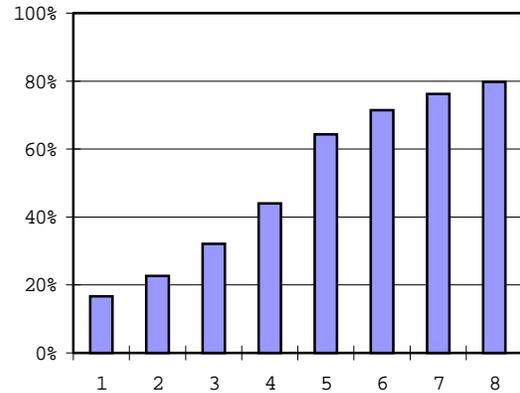

Figure 2: Histogram of experiment 1 data. Abscissa shows object pairs. The ordinate is the number of times subjects confirmed there was no difference between the objects presented. The maximum is 28 subjects multiplied by 3 repetitions for each pair.

In this preliminary study, we undertook a non-parametric approach. Using Friedman's test [20], [21], [22] significant treatment (here treatment refers to the particular pairing, e.g. h_3-h_8) effects were found (S=72.67, p<0.005). Since in each object pair we compare an object with h_8, the pairing of h_8 and h_8 acts as a control treatment. We next compare each treatment with the control treatment to see which objects when paired with h_8 tend to

be identified as different. A simultaneous multiple comparisons test [21] based on Friedman's test was conducted at a significance level of 0.05, yielding the grouping shown in Table 2.

Table 2: Grouping of objects in experiment 1, using Friedman's test.

Group1	Group2
h_1, h_2, h_3, h_4	h_5, h_6, h_7, h_8

The grouping indicates that when any of the first four objects are paired with h_8, the subjects' ability to tell the difference is significantly greater than with a pair of h_8's. When objects in the second group are paired with h_8 the ability for subjects to tell the difference is not significantly different from when two h_8's are presented. This suggests that h_5, h_6 and h_7 can be used as substitutions for h_8 in haptic rendering with most of the users not noticing the difference. We note that h_5 has significantly fewer polygons (1000) than h_8 (2560), suggesting that significant saving in collision detection and haptic rendering computation is possible, enabling a higher update rate for the scene.

5. Experiment 2

Experiment 1 suggests a coarser model may be substituted for a finer one without the user noticing the difference in rigid object simulations. However, whether this is also the case for deformable objects remains to be seen. There are two choices for refining a deformable model: global subdivision and local subdivision. Global subdivision means every polygon is subdivided whereas local subdivision means that the subdivision occurs only where significant deformation takes place. The subdivision technique we used is described in [9]. In this experiment 2 we ask if global subdivision affects subjects' ability to discriminate among deformable models. In this experiment the subjects grasp a point on the object's surface and deform the object (stretching or pushing). With no sliding motion on object surfaces as in experiment 1, "smoothness" is hard to define, so in this experiment and experiment 3 we use the term "level-of-detail" (LOD). Subjects were asked to compare two deformable objects with different haptic LODs. If subjects did not distinguish between two objects of different haptic LODs, we can adopt the coarser haptic model to reduce computational cost.

5.1 Experiment design

In this experiment we used the hemisphere of experiment 1, 50mm in radius. Three variations (h_1, h_2, h_3) represented by polygonal meshes were presented with three different haptic LODs. The second LOD was the result of applying global subdivision to the first one

(lowest LOD), and the third LOD was obtained similarly from the second. As in experiment 1, when an object was presented, a haptic model and a graphic model provided haptic and graphic feedback and were co-located. The haptic model was one of the three LODs. The graphic models of h_1 and h_2 were obtained by applying two or one subdivision step(s) to their corresponding haptic model. The graphic model of h_3 has the same geometry as its haptic model. Thus graphic models of all three objects had the same number of polygons and were at the highest LOD. We implemented them this way because we wanted the visual representation to be the same for all objects. The number of triangles increases by a factor of four after each subdivision step, so the number of triangles on the haptic model increases rapidly (almost exponentially) with the number of subdivision steps. This was part of the reason why we had only three levels of haptic model. The numbers of triangles in haptic and graphic models for each LOD are summarized in Table 3.

Table 3: Number of triangles in haptic and graphic models for each object in Experiment 2.

object	haptic model	graphic model
h_1	40	640
h_2	160	640
h_3	640	640

A technical challenge is the difference in the numbers of triangles on haptic and graphic models for h_1 and h_2. The subject felt the haptic model, while seeing the graphic model. The graphic model could not be precomputed because it was continuously being deformed by the subject. Assigning all node positions to be the same as in the haptic model was not a solution, since that would leave some nodes on the graphic model undefined, given the fact that there were more nodes in the graphic models than in haptic models for h_1 and h_2. We addressed this problem by applying a subdivision algorithm to the haptic model to obtain the graphic model for each display frame. Surface subdivision being expensive is the second reason why we had only three levels of haptic model in this experiment. However, with the advancement of computer hardware, more levels of haptic model in this experiment should not be a problem in the near future.

For the experiment two deformable hemispheres were rendered in a window. One of them was always h_3, the highest haptic LOD, acting as a reference object. However, the subject was not aware of this. The subject could move the tip of Phantom device to touch the objects. Once contact was made, the contacted triangle on the haptic model of that object was attached to Phantom tip. The subject then moved the tip around, stretching or depressing the object. The subject clicked a button on the Phantom stylus to detach the object from Phantom tip before making contact with the other object. The subject

was allowed to deform each object as many times as they wished. During deformation the subject observed and felt the objects. When the objects were deforming, the change of shape inevitably played an important role in helping subjects build an impression of the objects.

A pair of objects consists of h_n and h_3. Each of the three possible pairings in this experiment was repeated three times for each subject, for a total of nine trials. The order of the trials and the relative positions of the two objects were randomized. Figure 3 shows a screen shot of experiment 2.
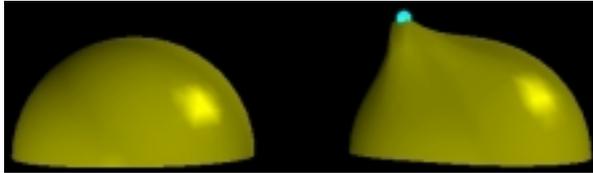

Figure 3: Graphic display in experiment 2. The subject is deforming the right-hand object.

A subject's task was similar to that in experiment 1 except that they could deform the objects and the time limit was increased to 20 seconds.

5.2 Results

Figure 4 shows experiment 2 results. Using Friedman's test, significant treatment effects were found (S=34.621, p<0.005). Since in each object pair we compare an object with h_3, the pairing of h_3 and h_3 acts as a control treatment. We then compare each treatment with the control treatment to see which objects, when paired with h_3, tend to be identified as different by subjects. A simultaneous multiple comparisons test based on Friedman's test was conducted at a significance level of 0.05. Table 4 shows grouping results from the test.
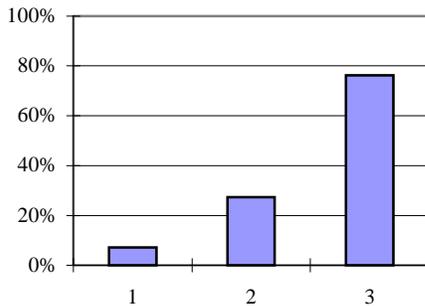

Figure 4: Experiment 2 data histogram. Abscissa is the object pairs. Ordinate is number of times subjects confirmed no difference between objects presented. The maximum count is 28 subjects multiplied by three repetitions per pair.

The grouping in Table 4 indicates that when global subdivision was applied to deformable objects, even though the numbers of polygons were the same for graphic models of the two objects, subjects were able to differentiate them. This also indicates that in a haptically and graphically rendered simulation, substituting a coarse object for its globally subdivided counterpart as a means to save computation is not plausible because users would notice the difference. We also notice that the height of the third column is only about 80%, again showing that subjects occasionally judged identical objects as being different.

Table 4: grouping of objects in experiment 2.

Group1	Group2
h_1, h_2	h_3

6. Experiment 3

Local surface subdivision has recently been explored by some researchers as a trade-off between simulation accuracy and object complexity [16][17][18]. The main idea of local surface subdivision is that only the part of the object being interacted with by the user is refined during the deformation process, while the rest of the object remains unrefined. In experiment 2 we investigated global subdivision as a trade-off between simulation rate and level-of-detail (LOD). To complete the discussion of reducing computational cost by substituting a less refined deformable object for a more refined one, experiment 3 investigated the effect of local subdivision. In this experiment we asked subjects to compare a locally subdivided object with a high (global) LOD object. The local subdivision method we used is described in [9]. If subjects cannot differentiate a locally subdivided object and the high LOD object, in haptic rendering we can substitute a less computationally demanding object to reduce computation.

6.1 Experiment design

We used the hemisphere of the earlier experiments, 50mm in radius, with 4 LODs. Since the hemisphere is approximated by polygonal meshes, each LOD represents a different number of polygons constituting those meshes. The fourth (highest) LOD has four times as many polygons as the first (coarsest) LOD. The second LOD and the third LOD are the result of applying one-ring and two-ring subdivision to the first LOD, respectively [9]. One-ring is defined as the area where all vertices have a direct connection to the vertex being manipulated. Two-ring is defined as the area where all vertices are connected to the manipulated vertex via at most one additional vertex. When a hemisphere was presented to the subject, the geometry of the haptic model was used for the graphic model. At each display frame after the haptic model was deformed, the positions of the nodes were copied into the

graphic model for graphic rendering. Table 5 gives the number of polygons on each model.

Two deformable hemispheres were rendered in a window in the experiment. One of the hemispheres was always h_4, which was the highest LOD, acting as a reference object. The other was of one of the four objects in Table 5. The way the subject manipulated the objects was similar to that in experiment 2. In this experiment the comparison between objects was also based on both haptic and visual cues.

Table 5: Number of polygons on models for each object in experiment 3.

| object | haptic model | | graphic model |
	Number of rings of local subdivision	number of polygons	number of polygons
h_1	0	360	360
h_2	1	384	384
h_3	2	444	444
h_4	complete subdivision	1440	1440

A pair of objects consists of h_n and h_4. Each of the four possible pairings is repeated three times for each subject, for a total of 12 trials. The order of the trials and the relative positions of the two objects presented were randomized. Figure 5 shows a screen shot for experiment 3.
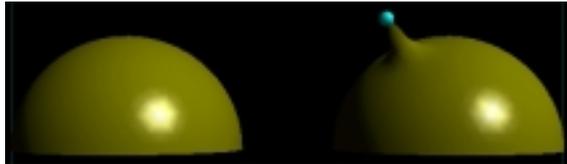

Figure 5: Graphic display in experiment 3. The right object (h_4) is deformed.

Subjects' tasks were the same as for experiment 2.

6.2 Results

Figure 6 shows experiment 3 results. We note that the height of the fourth column is less than 80%.

Friedman's test [20, 21, 22] shows significant treatment effects (S=29.54, p<0.005). Since all objects are compared with h4, the pair (h4, h4) acts as a control treatment. A simultaneous multiple comparison test [21] on columns of data against the control at a significance level of 0.05 results in the grouping in Table 6.

The grouping indicates that local subdivision technique is an effective approximation of the highest LOD object. We note in Table 5 the number of triangles in h4 is three times of that of h3. The grouping in Table 6 shows that in haptic rendering we may be used to substitute h_3 for h_4 to

reduce the complexity by a factor of three without most users noticing the difference, or h_2 for h_4 for a factor close to four.
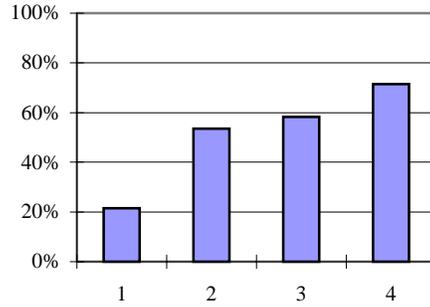

Figure 6: Experiment 3 histogram. Abscissa shows the object pairs. Ordinate is number of times subjects confirmed no difference between objects presented. The maximum count is 28 subjects multiplied by three repetitions per pair.

Table 6: Grouping of objects in experiment 3.

Group 1	Group 2
h1	h2, h3, h4

7. Discussion

In experiment 1 statistical analysis divided the eight objects into two groups. This division indicates that replacing a highly refined rigid surface with a less refined rigid surface may be possible in haptic rendering, given the specific haptic device used. However, the analysis does not point out the cause of this grouping. The factors that cause this division, if identified, may be used to find a less refined object that can replace a highly refined one. Intuitively, the size of polygons relative to the curvature and the height of bumps across edges may be among the factors. By curvature, we mean the angle between two polygons sharing an edge. The bump height is defined as the length between the black dots in Figure 7 (this distance is similar to chordal deviation in smooth surface approximation). Table 10 shows the values of possible factors on each haptic model of objects in experiment 1.
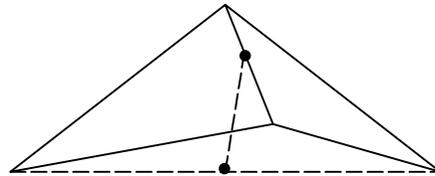

Figure 7: Bump height is defined as the length between the black dots.

Experiments described in [4] showed that the threshold of perceiving haptically rendered surface roughness is at a

"bump" height of about 0.1mm. The haptic device used in [4] is similar to that used in this study. The bump heights of h_5 and h_6 in Table 7 appear consistent with the threshold obtained in [4].

Since bump height and ratio between polygon size and object size in experiment 1 both decrease as the object becomes smoother (Table 7), it is difficult to examine their effects individually with our experiment configuration. More specific experiments should be designed in the future to test these effects separately.

Table 7: Edge lengths, angles between neighbor polygons and ratios between polygon sizes and object sizes in experiment 1

Object	Average edge length (mm)	Average edge length / hemisphere radius	Average angle across edges (degree)	Average bump height (mm)
h_1	29.6	0.592	160	4.55
h_2	15.0	0.300	170	1.14
h_3	10.0	0.201	173	0.51
h_4	7.54	0.151	175	0.29
h_5	6.04	0.121	176	0.18
h_6	5.03	0.101	177	0.13
h_7	4.31	0.086	177	0.09
h_8	3.78	0.076	178	0.07

Experiment 2 shows the subjects successfully identified the difference between the two objects most of the time. In all three experiments the subjects were encouraged to make any comments they wished. Some subjects commented that they made judgments based on shape difference deforming the objects. From such comments it seems that visual judgment might play a role in affecting the result of experiment 2. When the haptic model consists of only a small number of triangles, displacing one triangle from its original position would cause movements visible over a large area. On the other hand, when the haptic model consists of a large number of triangles and with the same number of triangles displaced because of the user's interaction with one triangle, the area of deformation is smaller, or more localized.

Note also that phantom is a single point contact device, which makes it difficult to sense the overall shape of an object by feeling the force. Subjects reported that shape differences between two objects are much easier to detect visually than haptically with the Phantom. In experiment 1 we ensured the same number of polygons appeared in each graphic model. How this fact affected the subjects haptic perception might be further investigated.

Experiment 3 shows that local subdivision is an efficient tool to approximate a finer object with a coarser one. The algorithm to determine the physical parameters is described in [9], which ensures the force the subject feels before and after subdivision is the same, making the two objects feel the same. Furthermore, local subdivision produces triangles of the same size as the finer object. This makes the two objects look alike visually. Therefore local subdivision makes the two objects appear haptically and visually similar. The reason that columns two and three in Figure 6 are not as high as column four may be related to the difference in the visual appearances of the two objects where they were not affected much by deformation. In these parts the fine object looks smoother than the coarser object, although the difference is subtle. However, the fact that the heights of column 2 and 3 in Figure 6 are very close to that of column 4 suggests the subjects focused on the deforming parts of the two objects, which would justify our effort to subdivide locally.

The numbers of rings of subdivision [9] we chose in experiment 3 were one and two. In this study we did not investigate what criteria should be applied to determine the number of rings to subdivide. In future studies criteria should be established. In an application, an algorithm could also be developed to decide adaptively which polygons to subdivide in the deforming area.

In experiment 2 and 3 the subjects could not slide on the deformable surface. This limitation might make the haptic experience not as realistic as in the real world. Further experiments that allow subjects to slide on deformable surfaces should be conducted.

The heights of the last columns in Figures 2, 4 and 6 are less than 80% of the maximum height possible suggesting that humans are intrinsically suspicious about what they see and what they feel. How this affects application design needs further investigation.

8. Conclusion

In this paper we present the results of a preliminary study of perception of touch in a virtual environment via a haptic device. We designed three experiments to determine if there were smoothness levels beyond which the difference is no longer significant. The purpose of this work was to explore possibilities for reducing the computation needed for haptic rendering. If the difference between two objects is indiscernible to the subject, the object with less complexity should be used in haptic rendering to reduce computation.

Recognizing the shape of a rigid object via a haptic feedback device is a process of temporal integration [8]. However, this integration does not apply to deformable objects since the object's shape is being changed, hence rigid objects are fundamentally different from deformable objects. To take this difference into account, experiment 1 was designed for rigid object surface exploration and

experiments 2 and 3 for deformable object manipulation. Experiment 1 showed that a rigid surface may be replaced by a less refined one in haptic rendering, without most users not noticing the difference. Experiments 2 and 3 showed that local surface subdivision is an efficient tool to approximate a globally detailed surface with a locally detailed surface, provided local refinement occurs in the interaction area. Further study is needed to determine the weighting of factors affecting perception of deformable objects when both haptic feedback and visual feedback are provided. Overall, this study shows that substituting a relatively coarser object for a finer object to reduce computational cost is possible in haptically rendered virtual environments, regardless whether the object is rigid or deformable.

Acknowledgement

References

[1] Klatzky, R.L. and Lederman, S.J., "Tactile roughness perception with a rigid link interposed between skin and surface", *Perception & Psychophysics*, 61(4), 1999, pp591-607.

[2] Lederman, S.J., Klatzky, R.L. and Hamilton, C.L., "Perceiving roughness via rigid probe: psychophysical effects of exploration speed and mode of touch", *Haptics-e The Electronic Journal of Haptic Research, 1(1)*, available at www.haptics-e.org.

[3] Robles-De-La-Torre, G. and Hayward, V., "Force can overcome object geometry in the perception of shape through active touch", *Nature*, vol.412, July 26, 2001, pp445-448.

[4] Yamashita, J., Fukui, Y., Morikawa, O. and Sato, S., "Spatial Resolution in Haptic Rendering", *SIGGRAPH 2001 Conference Abstracts and Applications*, pp264.

[5] Oakley, I., McGee, M.R., Brewster, S. and Gray, P., "Putting the feel in 'look and feel'", *CHI 2000*, April 1-6, 2000, pp415-421.

[6] Jackson, S.R., "'Action binding': dynamic interactions between vision and touch", *TRENDS in Cognitive Sciences*, vol.5, No.12, December 2001, pp505-506.

[7] West, A.M., and Cutkosky, M.R., "Detection of real and virtual fine surface features with a haptic interface and stylus", *Proceedings of the ASME Dynamic Systems and Control Division (1997)*, DSC-vol.61, pp159-165.

[8] Kirkpatrick, A.E., and Douglas, S.A., "Application based evaluation of haptic interfaces", *Proceedings of the 10th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, Orlando, FL, March, 2002, pp32-39.

[9] Zhang, J., Payandeh, S. and Dill, J.,"Haptic subdivision: an Approach to Defining Levels of Detail in Haptic Rendering", *Proceedings of the 10th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, Orlando, FL, March, 2002, pp201-208.

[10] De, S., Kim, J. and Srinivasan, M.A., "A Meshless Numerical Technique for Physically Based Real Time Medical Simulations", *Proceedings of Medicine Meets Virtual Reality*, 2001, IOS Press, Amsterdam, Jan. 2001, pp. 113-118.

[11] Cavusoglu, M.C., and Tendick, F., "Multirate Simulation for High Fidelity Haptic Interaction with Deformable Objects in Virtual Environments", *Proceedings of the 2000 IEEE International Conference on Robotics & Automation*, April 2000, pp. 2458-2465.

[12] Zhuang, Y. and Canny, J., "Haptic Interaction with Global Deformations", *Proceedings of the 2000 IEEE International Conference on Robotics & Automation*, pp. 2428-2433.

[13] James, D.L., and Pai, D.K., "ArtDefo: Accurate Real Time Deformable Objects", *Proceedings of the SIGGRAPH 1999 Annual Conference on Computer graphics,* Los Angeles, CA USA, August 8 - 13, 1999, pp. 65-72.

[14] Braraff, D. and Witkin, A., "Large Steps in Cloth Simulation", *Proceedings of the 25th Annual Conference on Computer Graphics*, Orlando, FL USA, July 19-24, 1998, pp. 43-54.

[15] Cotin, S., Delingette, H. and Ayache, N., "Real Time Elastic Deformations of Soft Tissues for Surgery Simulation", *IEEE Transactions on Visualization and Computer Graphics*, Vol. 5, No. 1, January-March, 1999, pp. 62-73.

[16] Dehunne, G., Desbrun, M., Cani, M. and Barr, A.H., "Dynamic Real-time Deformations using Space & time Adaptive Sampling", *SIGGRAPH 2001*, 12-17 August 2001, Los Angeles, CA, USA.

[17] Hutchinson, D., Preston, M., Hewitt, T., "Adaptive Refinement for Mass/Spring Simulations", *7th Eurographics Workshop on Animation & Simulation*, Poitiers, August 31-September 1, 1996, pp31-45.

[18] Kawai, H., Matsumiya, M., and Takemura, H., "Elastic Object Manipulation Using Coarse-to-Fine Representation of Mass-Spring Models", *SIGGRAPH 2001*, *Conference Abstracts and Applications*, pp176.

[19] Loop, C., "Smooth Subdivision Surfaces Based on Triangles", *Master's thesis*, University of Utah, Department of Mathematics, 1987.

[20] Devore, J.L., *Probability and Statistics for Engineering and the Sciences*, published by Duxbury Press.

[21] Hollander, M. and Wolfe, D.A., *Nonparametric Statistical Methods*, published by John Wiley & Sons.

[22] Larsen, R.J. and Mark, M.L., *An Introduction to Mathematical Statistics and its Applications*, published by Prentice-Hall.

[23] Product literature, www.sensable.com

[24] Minsky, M., Ouh-young, M. and Steele, O., "Feeling and Seeing: Issues in Force Display", *Proceedings of Symposium on Interactive 3D Graphics*, 1990, pp235-243.