Enhanced Collision Perception Using Tactile Feedback

Aaron Bloomfield*   Norman I. Badler†

*University of Pennsylvania

†University of Pennsylvania, badler@seas.upenn.edu

Enhanced Collision Perception Using Tactile Feedback

Aaron Bloomfield   Norman I. Badler

Center for Human Modeling and Simulation
University of Pennsylvania
{aaronb, badler} @cis.upenn.edu

Abstract

We used a custom designed tactor suit to provide full body vibrotactile feedback across the human arm for the purpose of enabling users to perceive a physical sense of collisions in a virtual world. We constructed a 3-D virtual environment to test arm reach movements. We present the results of human subject trials that test the benefit of using vibrotactile feedback for this purpose. Our preliminary results presented here show a small, but distinct, advantage with the use of tactors. With additional refinements to the system, improved performance results can be obtained.

1. Introduction

Commercial haptics devices and concomitant research generally focus on new and innovative ways of providing a haptic response to end-effectors, usually the hand. To feel fully immersed in a VR environment, however, one should feel haptic feedback on more than just the hand. Recently a low cost, low weight, and low complexity device has been introduced for VR: the tactor. Formulated with either pneumatic or electronic control, the tactor is essentially a small pressure point or button that can be actuated remotely as a tactile stimulus.

A VR experience of a confined space may be satisfying in a visual but not a haptic sense. By flying through the space one can get excellent visual impressions of its shape and relationships. But if one needs to reach inside the space, say to do complex equipment maintenance or repair, then the flying eye and the disembodied hand are no longer adequate paradigms for virtual equivalents of physical presence in the scene. Besides the end-effector haptic experience of the Phantom and similar devices, as well as haptic gloves, there seem to be no good devices and software interfaces to give a user a sense of confined immersion. Thus, applications which require experiencing and testing the feasibility of physical access for the entire body cannot be supported in realistic (i.e. low) cost VR configurations.

The computations involved in using tactors as a surrogate for physical contact use collision detection between each tactor location on the body and the virtual object geometry. Point-to-object collision checking algorithms are simple and fast [7]. Motion capture is used to find the body pose in 3D space; the 3D tactor locations are computed from the body pose and the known positioning of the tactors on the subject. Intersections between tactor locations and virtual object geometry is checked hierarchically with bounding volumes then with detailed object geometry to optimize performance. Once a collision is detected, its corresponding tactor is energized. Although some latency is unavoidable, the vibrotactile sensation is meant to augment a visual (possibly stereo) view of the scene and aid the user in establishing a collision-free pose in the space.

2. Related Work

Tactor arrays have been the subject of recent research, but none to our knowledge have focused on enhancing collision perception. Tan [5] presents a "wearable haptic display for situational awareness in [an] altered-gravity environment". Using a tactor array worn on the back, subjects tested the system on a NASA reduced gravity aircraft (a aircraft that moves in and out of free-fall). The focus of these experiments were how the subjects felt various haptic sensations in altered gravity [6].

Research at the Naval Aerospace Medical Research Laboratory is focusing on the use of tactor arrays. The Tactile Situation Awareness System (TSAS) has the goal of helping pilots perceive orientation, targeting, and situation awareness information though an array of tactors worn over the torso in a flight jacket. Apparently the experience is natural enough for both quick learning (10 seconds or so) and correct interpretation [2]. Cholewiak, from Princeton's Cutaneous Communication Lab, has focused on using haptics to aid sensory-impaired (blind and deaf) and elderly individuals [1]. Recent research at MIT places tactors in a vest

to convey information to the wearer [4]. Their vest uses Nitinol, a shape memory alloy.

3. The Tactor and Suit

The tactor and the suit, including their design decisions, are described in another article [*reference unavailable at article submission time*], and only the salient points are summarized here.

To induce tactile stimulation, we used small vibrating motors similar to those used in cell phones and pagers. These tactors have a footprint of 28 mm x 34 mm (1 1/8" x 1 3/8"), and are 20 mm (7/8") high. They take up to 1.6V DC, and 100 mA. The tactors are relatively small, and still manage to provide sufficient tactile feedback: they can be felt through moderately thick clothing. Figures 1 and 2 show pictures of the tactors. The tactors have Velcro backing for attachment to the tactor suit.


Figure 2. Tactor, top view

Figure 1. Tactor, front view

The tactor assembly consists of a red LED and a motor in a simple parallel circuit. The motor, shown in figure 3 has an off-center weight, which, when rotating, provides the vibration. The motors vibrate at about 60 Hz, depending on the amount of power supplied. Power is

Figure 3: Vibrating Motor

connected to the tactor via the light blue terminal block. The rest of the tactor components are concerned with attaching the motor to the tactor, and the tactor to the

Velcro on the underside. The black rubber housing around the motor is heat-shrinkable tubing, and is intended to keep objects away from the rotating motor. The motor is mounted on a cable mount (similar to the ones used to attach a coaxial cable to a wall or floor) in order to firmly attach it to the rest of the tactor. The tactor assembly is tightly bolted together so that the vibrations are not absorbed by the tactor connections. All the tactor parts were purchased from a electronics supply catalog.

To control the pager motors, we used relay cards that connect to a computer's parallel port. Their output is binary and cannot be modulated in amplitude. Instead we can modulate the frequency of the vibration by rapidly switching the tactor on and off, although the preliminary experiments described here did not utilize this. Customized tactor control boards, such as the TactaBoard described by Lindeman [3], would be a better hardware controller choice for the future. The TactaBoard can control up to 16 tactors, can be chained to other TactaBoards, and plugs into a computer's serial port.

For the tactor suit, we used a long-sleeve, skin-tight shirt designed for athletic use. The shirt, manufactured by Under Armour (Turf Shirt, item 0032), is thin enough to transmit the vibrations and strong enough to endure the sewing necessary to affix the Velcro to the shirt. The skin-tight aspect ensured that the tactors would always lie flat against the subject's skin and yet allow free movement. The elasticity of the suit also admits many different sized people.

The Velcro for the tactors are 2.5 cm (1") squares. Our motion capture system, a ReActor by Ascension Technologies, uses infrared transmitters that are 5 cm (2") circles and

Figure 4. Suit without the tactors

Figure 6. Tactor Placements on the Arm (24 tactors)

thus the Velcro attachments for these are 5 cm (2") squares. Long strips of Velcro are used to secure the ReActor's motion capture wires that run between the infrared transmitters and the battery pack on the user's waist. This suit can be seen in figure 4, and with the tactors attached in figure 5. Although this suit was designed to work with the ReActor, similar Velcro pieces could easily be attached for other commercial motion capture systems.


Figure 5. Suit with the tactors attached

Currently the suit holds up to 24 tactors on the right arm, as shown in figure 6. All 24 tactors were used in this research. The tactors are arranged in six "rings" of four tactors each, spaced uniformly along the forearm and upper arm. The proximal upper arm ring has only two tactors, and the other two tactors of this ring are on the point of the elbow and the back of the hand.

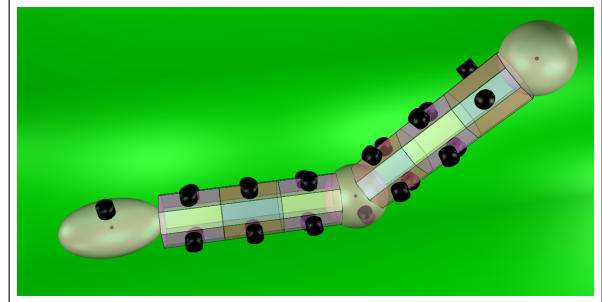For the hands, we used regular glove liners, as they are

inexpensive, thin enough to transmit the vibrations, and also skin-tight. The gloves had to hold a number of motion capture transmitters, as well as one tactor on the back of the hand.

4. Experimental Procedure

Right handed subjects were recruited from the student and staff population at the University. Subjects were paid for their participation. The experiment was performed in front of a projection screen; stereo display was not used.

After signing a consent form, the subjects were shown a video, with narration, of what they needed to do in the experiment. Two questionnaires were given, one after the demo and before the experiment, and the other after the experiment.

The experiment consisted of a simulation task that was a simple reaching exercise. The users had to reach one's right arm into the puzzles shown in figure 7. The objective was to touch the red sphere with the computer hand in each puzzle without colliding with the rest of the puzzle. Touching the red sphere would cause a successful puzzle completion. The puzzle would then disappear, and the subject would move on to the next one. The puzzles could be completed in any order. If a collision did occur (other than the hand colliding with the red sphere), the puzzle would not complete. This ensured that the subjects made an effort to reduce and eliminate the number of collisions.

The puzzles required the arm be inserted up to about the elbow in order to touch the red sphere. An earlier version of the simulation experiment provided to the subjects was a full arm reaching task where the users had to insert one's entire arm, up to the shoulder, into a puzzle. This was found to be too difficult for the subjects to perform, as many could not complete the experiment. Thus, the puzzles were made simpler by not requiring as much reach.

There were two types of feedback provided in this experiment. The first was the visual feedback, which consisted of
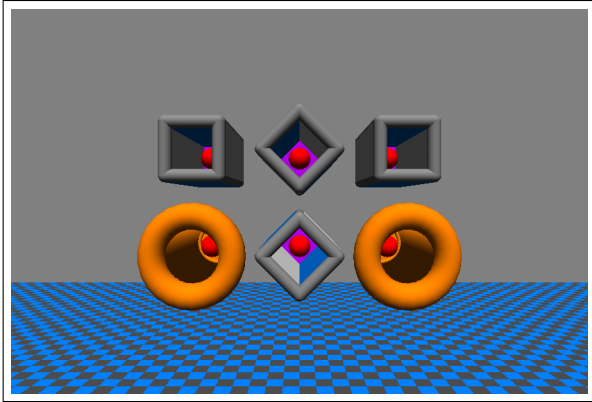
Figure 7. The Simulation Task

visual collision alerts. The part of the arm that was colliding with an object would turn red, as shown in figure 8. The arm was rendered translucent so the subjects could see the occluded sides. The parts of the arm that are not colliding are blue. The collisions are colored magenta in the figure to ensure that the contrast between the colliding areas of the arm and the non-colliding areas of the arm is visible in the grayscale version of the image. The colliding areas (the elbow, the middle of the forearm, and the hand) are a lighter shade of gray in the grayscale version of Figure 8.
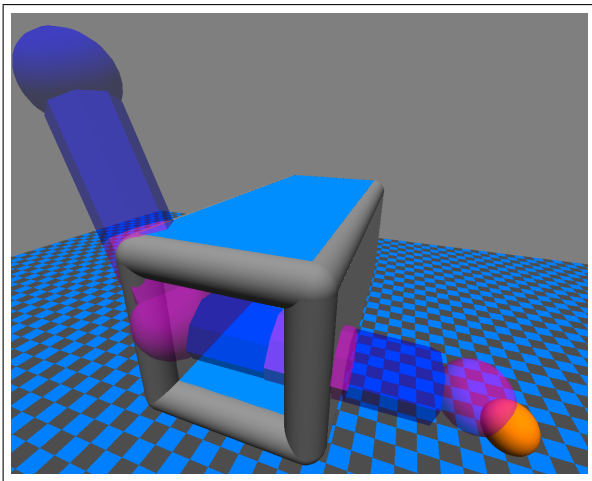

Figure 8. Visual Collision Alerts

Both groups of subjects were provided with the visual collision alerts. The other feedback provided was the haptic feedback, discussed above.

The subjects were randomly split into two groups. The first group, the control or visual group, had only the visual collision alerts. The second group, the tactile or haptic group, had both visual and haptic collision alerts. Both groups wore the tactors, so that arm fatigue would not be a differentiating factor between the two groups. The sub-

jects were given a chance to rest between trials to reduce arm fatigue and tactile saturation. The only difference in the experiment between the two groups was the presence of haptic feedback.

Each group performed six trials of the simulation, where each trial consisted of completing the six puzzles shown in figure 7. The first two trials were not used in the data analysis, as they they were considered learning trials (the results showed that the subjects' performance stabilized after the first two trials).

Furthermore, the group with only visual alerts performed an additional three trials (for a total of nine) with the haptic feedback enabled to see how the performance compared with the first six trials. Thus, all subjects experienced the haptic feedback at some point in the experiment.

The experimental procedure described above was approved by the Institutional Review Board of the Office of Regulatory Affairs of the University of Pennsylvania (protocol # 708331).

5. Statistics and Experimental Measures

The main data point used to compare the subjects' results was the number of collisions that occurred during the simulation. One collision was defined as a single tactor activation for a single frame of the motion capture system, which ran at 30 frames per second. This ensured that a longer collision was weighed more than a shorter collision. For the control group, the collisions were determined as if they had activating tactors. The collision counts were then averaged across the last four simulations. The first two simulations were excluded from analysis because they served as training runs

The subjects were told that completion time, while important, was not nearly as important as reducing the number of collisions. Using completion time as a data point did not give any useful results.

Subjects were given questionnaires before and after the experiment, which provided both demographic and subjective experimental data. The statistical results of the experiment itself, including the number of collisions, provided the objective experimental data.

6. Experimental Results

18 right handed subjects, 11 male, participated in the experiment. The subjects were randomly assigned to groups of 9 each. The tactile group performed the trials with 79% of the collisions per trial compared to the control group, as shown in table 1. Although the group size was relatively small ($n = 9$), the increase in performance for the tactile group consistently stayed around 80% after the first three subjects in each group.

Group	Average collisions	Collisions st. dev.	Time (seconds)
Control	1743	859	121
Tactile	1375	653	119

Table 1. Experimental results (averaged per trial)

Although the trials were timed, the subjects were told that it was more important to reduce the number of collisions than improve (i.e. decrease) the simulation time. There was no significant difference in the completion times between the groups.

There was no statistically significant difference in average age (27.9) or education between the two groups. There was a small difference of the amount of 3-D experience between the two groups. The tactile group, which performed the tests with fewer collisions, had a slightly higher amount of 3-D experience than the control group (average response of 3.4 versus 3.0, on a self-rated scale of 1 ("none") to 5 ("a lot")). 3-D experience was defined on the questionnaire as experience with "games, animation programs, graphical programming, etc."

These results, while encouraging, were not what we were expecting. We had expected that the tactile group would show a more significant increase in performance (a decrease in the number of collisions). With the refinements discussed in the future work section, below, we feel that this increased performance can be achieved. Reasons for the lack of more definitive results are also discussed in the future work section.

When the control group performed the three additional trials with the haptic feedback, the performance increased slightly by 12% (average collisions decreased from 1743 to 1526). The first of the three additional trails was considered a training run, and was not counted in this analysis. A number of subjects commented that they were accustomed to the experiment without any tactile feedback, and the additional tactile feedback was an unneeded and confusing sensation. However, the performance results dispute this, as the performance did improve. The subjects' collision counts had plateaued during the first six trials, so the learning curve was not a factor in the increased performance of the last three additional trials.

The subjects were asked a number of questions on the questionnaires, and the responses were limited to integers in the range from 1 (very low or very little) to 5 (very high or very much).

The subjects felt strongly that the visual collision alerts contributed to the perception of collisions (average response of 4.2). This value did not vary significantly between groups. The subjects felt slightly more strongly that the haptic alerts alerts contributed to the perception of collisions (average response of 4.3). Interestingly, the tactile group felt *less* strongly that the haptic feedback contributed to the perception of collisions (average response of 3.9) compared to the control group (average response of 4.7). This is surprising, as we expected that the control group, which was forced to perceive the 3-D world without haptic feedback in the first six trials, would feel less strongly that the haptic alerts contributed to the perception of collisions. The reason may be because the control group was unable to feel the collisions until the final three trials, and thus they felt it contributed more. Currently we cannot explain reason for this, and we would like to investigate it in future research.

The vibrotactile feedback could act as either a signal, a spatial indicator, or both. Acting as a signal would cause the subject to know there was a collision somewhere, but not where. This would be similar to an auditory alarm sounding when a collision occurred. Acting as a spatial indicator would impart, via the vibrations, exactly where the collision occurred. The subjects gave an average response of 2.9 (where 1 meant it acted solely as a signal, and 5 meant it acted completely as a spatial indicator). This number did not vary significantly between the two groups. These results were again encouraging, but less than what we had expected. We thought that the results would be higher. With the refinements discussed in the future work section, we feel that an improved response on this question can be achieved. Three of the subjects were very petite, and the suits provided were too big for them (future research will ensure that smaller sized suits are available). This caused the tactor locations to be significantly less accurate than with the rest of the subjects. Excluding these three subjects, the average response for this question was 3.1.

Tactile saturation, where the skin becomes less sensitive to a repeated stimulus, was not perceived by the subjects to be a problem (average response of 1.5). The subjects felt the tactor activations were highly accurate (average response of 4.3), indicating that they felt that the right tactor was activating for any given collision. Lastly, the tactors were moderately comfortable on the subject's arm (average response of 3.2). From subjects' comments, the moderate comfort level was largely due to the weight of the tactors.

7. Conclusions

In designing the research, we expected to see a more dramatic decrease in the number of collisions in the tactile group. What we found was a small but definite improvement with the use of haptic feedback for full body collision perception. With various refinements, discussed below, we feel greater performance results can be obtained.

Full body haptic feedback is an eventual necessity for a fully immersive experience in a virtual world. One cannot feel fully "in" a virtual world when only the hand is receiving haptic feedback. There are many virtual situations that may present a reason to use haptic feedback; collisions are one of them. Refinements on our work suggest promising options for future research in full body haptic feedback for improved collision perception.

8. Future Work

A common feedback from the subjects was the lack of tactile response on the hand. Because of the current tactor design, and as the glove needed to hold six motion capture transmitters, there was little room for the tactors on the glove. Whenever a subject's hand collided with an object, they always felt the haptic sensation on the back of one's hand. As the hand collisions accounted for about 80% of the total collisions, this led to less than ideal situation of haptic feedback. Our original design, which included full arm reaching puzzles, as opposed to the forearm reaching tasks used in this experiment, would have focused a lower percentage of the collisions on the hand. These puzzles were too difficult to complete for many subjects. A different tactor design that could fit under the motion capture sensors would allow for more feedback on the hand and solve this problem. We feel that this would cause a significant increase in the haptic feedback group's performance. This is the next area of work for this project.

The use of a stereo display (HMD or stereo projection screen) would increase a user's perception of the 3-D scene. This is an aspect we intend to add in future experiments.

With the two types of feedback discussed, there are a total of four possibilities of groups. The other two not studied in this experiment are haptic only collision alerts, and no collision alerts of any type. In both of these groups the user would see one's arm moving, but would not see a red highlight if a part of the arm collided with an object. Combining these with auditory alerts (an auditory alarm if a collision occurs), the number of groups would increase to eight. We would like to run more experiments that compare different combinations of input modalities.

References

[1] R. Cholewiak and A. Collins. The generation of vibrotactile patterns on a linear array: Influences of body site, space, and time. *Perception & Psychophysics*, 62 (6):1220–1235, 2000.

[2] L. Knapp. Lost? go vest, young man. http://www.wired.com/news/gizmos/0,1452,44376-2,00.html, June 2001.

[3] R. Lindeman and J. Cutler. Controller design for a wearable, near-field haptic display. In *Proceedings of the 11th International Symposium on Haptic Interfaces for Virtual Environments and Teleoperator Systems*, pages 397–403, 2003.

[4] M. Nakamura and L. Jones. An actuator for the tactile vest – a torso-based haptic device. In *Proceedings of the 11th International Symposium on Haptic Interfaces for Virtual Environments and Teleoperator Systems*, pages 333–339, 2003.

[5] H. Z. Tan. Perceptual user interfaces: Haptic interfaces. *Communications of the ACM*, 43(3):40–41, Mar. 2000.

[6] R. Traylor and H. Tan. Development of a wearable haptic display for situation awareness in altered-gravity environment: Some initial findings. In *Proceedings of the 10th International Symposium on Haptic Interfaces for Virtual Environments and Teleoperator Systems*, pages 159–164, 2002.

[7] X. Zhao and N. Badler. Near real-time body awareness. *Computer Aided Design*, 26 (12):861–868, Dec 1994.