Virtual Reality International Conferences 2000
Laval 18–21 May 2000

Keynote Address: The Challenges of Large–Volume Haptics

Grigore C. Burdea
CAIP Center, Rutgers University,
96 Frelinghuysen Rd., Piscataway, New Jersey USA.
E-mail: burdea@vr.rutgers.edu

Keywords: haptics, virtual reality, tracking errors, PHANToM arm, safety, haptic suit

Abstract

Haptics is gaining ground as an important sensorial channel that enhances virtual reality interactions. Current commercial interfaces give the user the ability to touch and feel virtual objects. Extending haptics to large volumes is very desirable due to the new generation large-volume displays. However, the work envelope over which such haptic feedback can be produced in VR is limited by the dimensions of the interface, or by the range of 3-D magnetic trackers. Producing large-volume haptics requires significant advances in sensing, actuators, wearable computing, modeling and communication software. Human factor studies should follow in order to quantify the effects of the new technology, as well as address its safety concerns.

1 Introduction

Haptics is gaining ground as an important sensorial channel that enhances virtual reality interactions. Current commercial interfaces give users the ability to touch and feel virtual objects. In order to resist the user's actions these interfaces need to be grounded to either the desk, a wall or ceiling, or on the user's body (either on the back, forearm, or palm) [Burdea, 1996]. Figure 1 is a classification of force feedback interfaces based on their grounding location [Bergamasco, 1993].

Tactile mice, such as the FEELit mouse (Immersion Co., CA) are 2-D interfaces that reproduce the mechanical texture of objects, as well as their 2-D contour. For dextrous tasks, tactile gloves, such as the CyberTouch® (Virtual Technologies, Palo Alto CA) give finger-specific contact information. Force feedback interfaces, such as the PHANToM® (SensAble Technologies, Woburn MA), provide information on object compliance, inertia and weight. For dextrous tasks the CyberGrasp® glove outputs finger-specific resistive forces, but cannot reproduce object weight or inertia.

Extending haptics to large-volume virtual environments is very desirable due to the new generation displays, such as the Baron workbench (Barco Co., Belgium), or the CAVE® (Fakespace Co., Ontario, Canada). Large displays fill a large portion of the user's field of view with graphics, which increases his immersion sensation. The next step would be to allow the user freedom of motion within this large volume, without encumbering interfaces, cables, etc.

Unfortunately, the work envelope over which such haptic feedback can be produced in VR is limited by many factors, including the dimensions of the interface, the range of 3-D magnetic trackers, cost, etc. Producing large-volume haptics requires significant advances in sensing, actuators, wearable computing, modeling and communication. This paper takes a look at the various challenges the VR haptics designer faces today when it comes to large-volume virtual environments.

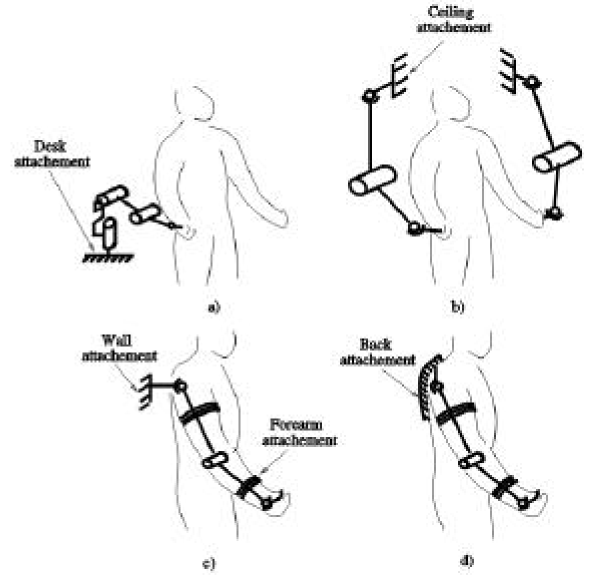

Figure 1. Possible force feedback mechanical grounding locations (adapted from [Bergamasco, 1993]. Reprinted by permission).

2. Large–volume Tracking

Figure 2 shows a typical large-volume virtual environment with haptics, which was developed at the Human-Machine Interface Laboratory, Rutgers University. The user stood in front of a large 3-D stereo display and was wearing active glasses. He manipulated virtual balls of varying compliance using a haptic glove (in this case the Rutgers Master II force feedback glove). The user could feel the hardness of the balls and was able to sort them accordingly, into “soft,” “medium,” or “hard” bins [Matossian, 1999].

The above application was designed to also test the possibility of direct interaction, without a mediating graphics hand on the screen. In a direct interaction scheme the user would have reached and grasped the balls that appeared to be floating in space. The key to success would have been an accurate measurement of the user’s hand position in front of the display. A Polhemus Long Ranger® magnetic tracker measured the wrist position and orientation. Unfortunately, the tracker introduced large errors due to metallic interference. This interference, typical of all magnetic trackers, made direct manipulation impossible, and the simulation had to use the classical virtual hand approach.

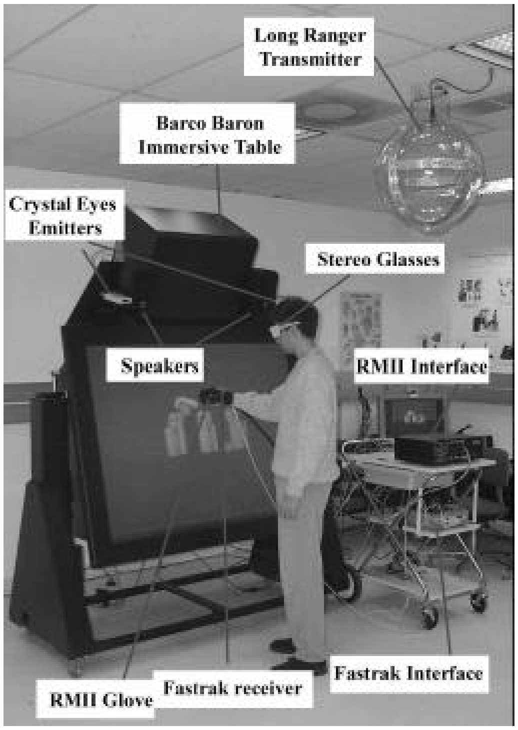

Figure 2. Typical large-volume VR simulation with haptics [Matossian, 1999]. Reprinted by permission.

Calibration measurements were subsequently done to determine the cause of Long Ranger errors [Trefftz and Burdea, 2000]. It was determined that the metal in the Barco Baron display was one such cause. The tracker emitter was subsequently detached from the ceiling, and placed on a wood tripod sold by Polhemus (Colchester, VT). This support allowed adjustment of the emitter height above the floor. Tracker position measurements were then compared with accurate mechanical measurements. It was determined that errors grew with distance from the emitter, as well as with proximity to either the floor or the ceiling, as shown in Figure 3. These errors were due to the metal used in the construction of the building.

In all fairness, the above problems are not limited to Polhemus trackers. Such tracking interference has also been reported for the Ascension MotionStar® full body wireless suit. While a wireless suit did provide more freedom of motion to the user than a tethered tracker, its accuracy suffered due to the influence of the metal in the floor [Marcus, 1997].

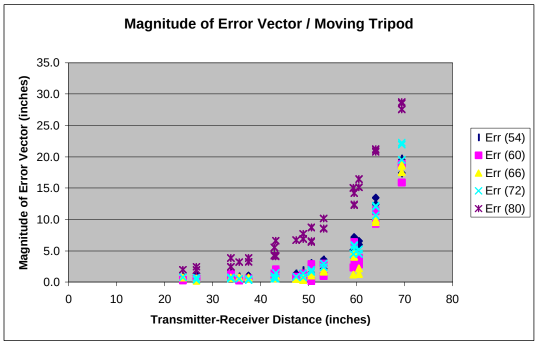

Figure 3. Tracker position errors as a function of distance and proximity to floor or ceiling [Trefftz and Burdea, 2000].

When large-volume interactions are desired, the solution is to avoid the use of magnetic trackers altogether. A new generation of trackers produced by InterSense Co. (Burlington, MA) combine inertia and ultrasonic sensing, being aimed at exactly these kind of applications [InterSense, 2000]. Because the IS-900 trackers do not use magnetic fields, they are immune to metal interference. Furthermore, they can be modularly extended to a whole building surface (up to 900 m^2) or can be attached to 3-D displays, as was recently demonstrated at the IEEE Virtual Reality 2000 Conference [Sorid, 2000]. As seen in Figure 4, the user interacted with the stereo image using an InterSense ultrasonic stylus (no force feedback was provided). His hand position was tracked by an IS-900 "Virtual Workbench Tracker" (VWT) placed directly on the Baron projector enclosure.

3. Haptic Interfaces

The ability to correctly track the user's actions is only one of the many necessary conditions for large-volume VR with haptics. Another condition is to have haptic interfaces that have minimal impact on the user's freedom of motion. Ideally, the user should be able to walk freely in a large working volume and have haptic feedback at any location in this volume.

One haptic interface that gives the user more freedom of motion is the CyberPack® produced by Virtual Technologies. As shown in Figure 5, the users wears a backpack containing the electronics and actuators controlling up to two CyberGrasp haptic gloves. This gives the user a much larger work volume than the CyberGrasp alone would allow, the only limitation being the length (and weight) of the tether to the backpack unit and to the glove trackers. With this system the user can walk in front of a large display, or in a CAVE, grasp and feel the compliance of virtual objects. Since the system does not produce force feedback at the wrist, the user cannot feel the weight of the manipulated object, or its inertia.

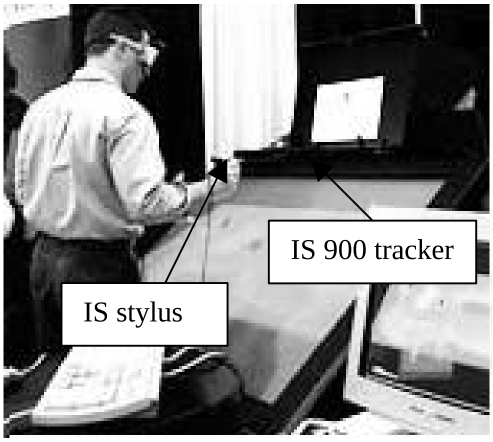

Figure 4. An Intersense inertia/ultrasonic tracker integrated with the Barco Baron 3-D [Sorid, 2000]. © The New York Times. Reprinted by permission.

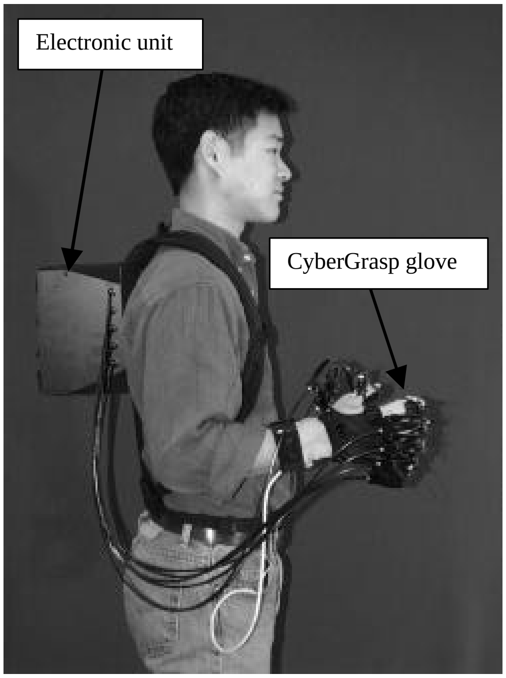

Figure 5. The CyberPack® produced by Virtual Technologies Co. Reprinted by permission.

Virtual Technologies has recently introduced the CyberForce® arm which can be combined with the CyberGrasp glove to provide more realistic haptics simulations [Kramer, 2000]. As shown in Figure 6, the CyberForce attaches to the back of the palm and produces both translating forces and torques to the wrist. As such the user can feel both object mechanical compliance (produced by the haptic glove), and the object weight and inertia (produced by the CyberForce arm). The system is aimed at desktop usage, like CAD/CAM design or service training tasks, thus it is not applicable to large-volume haptics. Of course, there are other force feedback arms that can be attached at the user's wrist, such as the PHANToM Premium, or the Sarcos Master [Hollerbach et al., 1999]. However, these arms need to be grounded on the wall, or floor, which restricts the user's freedom of motion to the workspace of the haptic arm.

Figure 6. The CyberForce® arm coupled with the CyberGrasp glove produced by Virtual Technologies Co. Reprinted by permission.

Another way to create large-volume simulations with haptics is to use a treadmill that tilts. If the user walks or runs on this equipment, and the computer tracks his motion, then a very large walk-through simulation can be created. Researchers at Sarcos Co. (Salt Lake City, UT) and University of Utah added a force feedback arm attached to a harness worn by the user [Hollerbach et al., 1999]. As shown in Figure 7, the treadmill was placed in front of a three-screen display showing a mountain slope. The resistive force produced by the force feedback arm replicated the gravity component that is associated to walking uphill. A safety rope was tied to the ceiling, in order to prevent injury in case of a fall. The system limitation refers to its lack of force feedback to the hands, thus its potential applications do not include manipulation, assembly or other dextrous tasks.

Figure 7. The treadmill with force feedback arm system developed at the University of Utah. Reprinted by permission.

Recently, Japanese researchers have proposed the replacement of the treadmill with an active floor. As shown in Figure 8, the floor is composed of modular actuator tiles which can change slope to replicate a moderately curved 3-D surface [Noma et al., 2000]. User motion is tracked by a vision-based system, which transmits information to the computer running the simulation.

The advantage of the active floor approach, compared to the walking-in-place treadmill paradigm, is that the active floor allows natural walking over the whole surface of the floor. Thus there is no need for a force feedback arm attached to the user's back, and the simulation volume, and naturalness of interaction are enhanced. The limitation in this case is the amount and size of slope that can be replicated, which depend on floor actuators. Furthermore, there is a lack of force feedback to the hands, similarly to the previous example.

If haptic feedback to more than the hands or the legs is desired, then the only solution is a haptic suit. A haptic suit (whether for tactile or for force feedback) needs to meet certain requirements in terms of weight, bulkiness, energy consumption, sensing, etc. The weight needs to be small in order to reduce user fatigue. The imbedded sensors need to be compact and rugged, while the actuators need to have low energy consumption. The wearable computer integrated with such as suit needs to be powerful enough, such that most of the computation is done locally. Ideally, the suit should be similar to a jumpsuit, such as those used by aviators, but provide

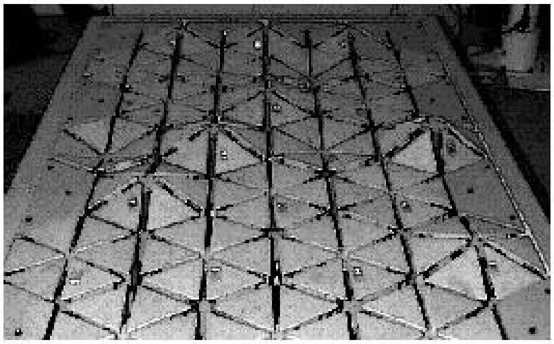

Figure 8. The ground surface simulator
[Noma et al., 2000].
© IEEE 2000. Reprinted by permission.

all the sensing, feedback, computing and communication functions described above.

Figure 9 shows a MotionStar wireless tracking suit produced by Ascension. The suit incorporates 20 magnetic sensors placed at critical locations on the body. These are wired to a backpack where the signal conditioning and communication electronics are located. The suit can work independently for up to three hours, providing up to 100 readings every second, from a range of three meters. There is no hand gesture information, nor is there any haptic feedback.

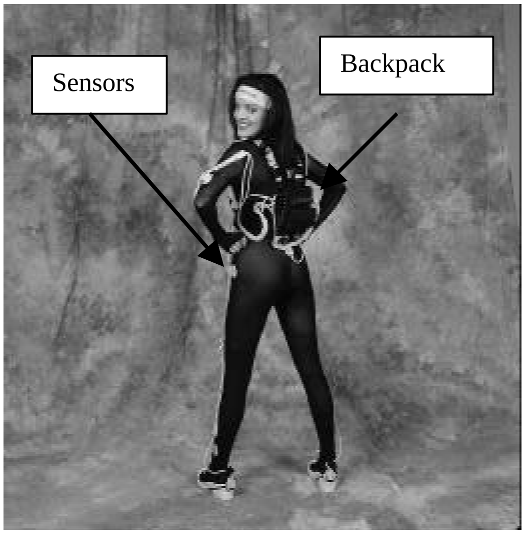

Figure 9. The Ascension MotionStar® tracking suit.
Reprinted by permission.

An embedded computer used in a haptic suit should take care of low-level control and computing (including graphics), such that the communication with the host can be done over lower-bandwidth wireless lines. The Xybernaut Co. (Fairfax VA) is producing a Mobile Assistant IV wearable computer, which is shown in Figure 10. The central processing unit is worn on the belt and contains a Pentium processor. Its output is transmitted to a head mounted display, and remotely to a host computer through a wireless connection. In such a configuration the Mobile Assistant IV could display graphics to the user, based on position information received from the remote host. Else new sensors could be developed, which may be read directly by the wearable computer, without the need for remote position data. This is especially critical for VR, due to the real-time requirement, which is affected by time delays in the communication with the remote host.

Figure 10. The Mobile Assistant IV wearable computer. Courtesy of CAIP Center.
Reprinted by permission.

Burdea and colleagues proposed a concept haptic suit using bellow-type pneumatic actuators, as illustrated in Figure 11 [Burdea et al., 1991][Burdea and Coiffet, 1994]. Pneumatics was preferred to other kind of actuators due to its large power/weight ratio, as well as its cleanliness. Each major joint had a pair of actuators working in opposition. In this way the known lower mechanical bandwidth of pneumatic actuators could be improved. All the actuator-regulating valves were located in a belt, together with the connection to the power/control source. A “honeycomb” plastic back plate was used for rigidity, as well as a distribution structure for the pressurized air used by the suit actuators. Such as suit could provide force feedback to arms, legs, torso, and work both as a force feedback suit and as a force-amplifying one. In force amplification mode, the suit could help its user lift heavy (real) weights. Weather for force feedback, or for force amplification the suit would need to lower the center of gravity of the body in order to increase safety/equilibrium. One solution could be to add heavy boots that “anchor” the suit. Otherwise the boots could have powerful magnets that could lock the feet on the metallic floor of a simulation room. Since the suit would have to have its own position sensing, the interference from the metallic floor should be minimal.

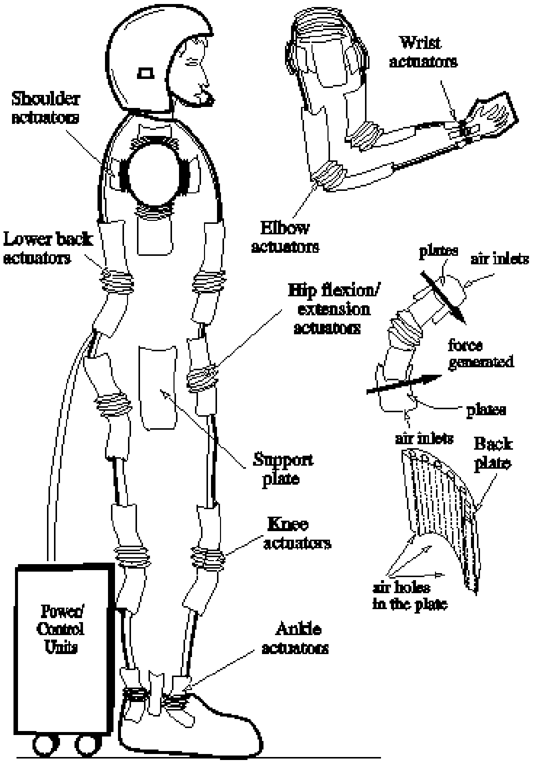

Figure 11. The Jedi force amplification/feedback suit Adapted from Burdea et al., [1991]. © Editions Hermes. Reprinted by permission.

4. Physical Modeling and Programming

Developing a haptic suit, or similar interfaces is only one of the challenges of large-volume haptics. Another important aspect is modeling and programming of the simulation, including distributed, multi-user interactions in the large virtual world.

A haptic suit should provide sufficient data to allow an avatar to be controlled in real time. Boston Dynamics (Cambridge MA) has developed the “DI-Guy” library, which allows programmers to control the actions of a full-body human character. The library was first developed for the military, as illustrated in Figure 12 [Boston Dynamics, 2000]. The character moves realistically, by receiving simple pose commands from the user, and then smoothly interpolating between such frames. The ability of controlling DI-Guy in real time is due to level-of-detail optimization, task-level control and other techniques. At the present time commands are inputted through either off-line programming, or interactively through a simple point-and-click interface. Furthermore, collision detection is performed, but no force feedback in computed.

Figure 12. The DI-Guy library being applied to military tasks. © Boston Dynamics. Reprinted by permission.

The GI-Guy API library would have to be extended to allow input from a suit. It is clear that collision detection needs to be precise and fast, and be extended to more than the hand. This calls for a multiple body collision algorithm similar to that developed by Cohen and colleagues at the University of North Carolina in Chapel Hill [Cohen et al., 1995]. Such an algorithm uses an approximate bounding box collision detection followed by pair-wise exact collision detection.

The next step is to calculate a collision response, which in the case of force feedback means calculating the forces/torques that result from the interaction in the virtual environment and which need to be felt by the user. Since many forces need to be computed in real time, simplified models should be used, in order to avoid time delays. Furthermore, the usual point-contact model used by GHOST and other haptic libraries will not suffice. Thus there is a need to extend these libraries to accommodate the programming needs of large-volume haptics.

Large virtual environments are ideal for multi-user interactions, thus modeling needs to be extended to accommodate force feedback in shared virtual simulations [Buttolo et al., 1996]. Such interactions may be classified as collaborative tasks and cooperative ones. In collaborative virtual environments with haptics, users take turns interacting with a given virtual object. As such the force feedback loop closes locally, and it is much easier to be realized. The effects of transmission time delays are minimal, and thus the requirements on the transmission line linking the remote users are also less stringent.

In a cooperative task, several (remote) users interact with a shared virtual object at the same time. Users feel forces resulting from their own actions, as well as forces applied by the remote user(s). As such the modeling is more complex, and the limits on acceptable time delays are much more stringent. This is due to the degrading effect time delays have on the force feedback quality, and on the stability of the haptic interface. It is expected that newer Internet2-mediated environments will have smaller time delays in the force feedback loop, however such delays could never be totally eliminated.

Alternate schemes will have to be implemented, such as sensorial transposition from force to tactile feedback, and also sensorial redundancy of the feedback modalities. Iwata [2000] reports on the use of artificial viscous forces added to slow down the user's actions, and thus reduce errors in the face of large time delays (up to 3 seconds). Others have proposed "haptic dead reckoning" as a way to have large simulations with haptics [Wilson et al., 1999]. This method is an extension to the haptic domain of the well-known approach used in distributed military simulations. In haptic dead reckoning, all force feedback is computed locally at fast rates, and the local client sends kinematics data to a centralized server. The server maintains a complex model of the shared virtual environment and transmits object property information (such as compliance, damping, etc.) back to the client. This scheme is aimed at reducing the time delay by reducing the amount of data that is sent over the network.

5. Safety and Human Factors

Assuming that all the above challenges have been met, there still remain at least two more to be overcome. The first is user safety. All the approaches to the problem of large-volume haptics, reviewed in this paper, have intrinsic risks to the user. The first risk is the potential for accidents due to the cables/wires that the user needs to drag while walking in large-volume simulations. Another risk factor is the haptic interfaces themselves. Since these devices need to be sufficiently stiff and required to output large forces, they may injure the user inadvertently. Finally, if shared simulations are the case, with collaborative tasks between remote users, then such users may injure each other. Computer and communication malfunctions add to the potential troubling factors.

A typical approach to user's safety used in tele-robotics is to add a "dead man's switch," which powers down the robot if the user so desires. It is not clear that such a technique will work here, due to he more complex interaction, and the close proximity of the user to the haptic interface. Of course, in the case of a haptic suit the user is inside the haptic interface! Thus the suit has to have built-in mechanisms to safeguard the user trapped inside, by limiting, for example, range of motion, or force levels. Limiting force levels and range of motion in turn degrades the quality of the haptic feedback produced, so compromises need to be made.

An optimum balance of safety and performance can only be achieved through diligent human factor trials and design iterations. Such human factor studies need also be conducted in order to assess the efficacy of the haptics technology in general, and for large-volume simulations in particular. Since the technology is new such studies are yet to be done, which in turn is a challenge to the interface designer and to the application developer.

6. Conclusions

The present review is by necessity limited. It is clear that many challenges exist today to the desire to extend haptics to large-volume virtual reality interactions. These are summarized below:

- Non-intrusive, accurate, long-range tracking of multiple user body segments;
- Light, high-power, wearable haptic interfaces, with on-suit computing and wireless communication;
- Physical modeling of collisions and forces over the whole body;
- Reduction of time delay effects on distributed virtual environments with cooperative haptic interactions between remote participants;
- User's safety;
- Human factors studies of haptics for large-volume virtual environments.

Acknowledgements

The author's research reported here was supported in part by grants from the National Science Foundation (BES-9708020), from Office of Naval Research (N14-98-1-0264) and from Rutgers University (CAIP and SROA grants).

References

Bergamasco, M., "Theoretical Study and Experiments on Internal and External Force Replication," *IEEE Workshop on Force Display on Virtual Environments and Its Application to Robotic Teleoperation," IEEE*, New York, 1993, pp. 45–58.

Boston Dynamics, "DI-Guy – Populate your Virtual World with Life-Like Humans," www.bdi.com.

Burdea, G., N. Langrana, and W. You, "Supersuit Project Report II: Concepts," Rutgers University, Piscataway NJ, August 1991.

Burdea, G., and P. Coiffet, *Virtual Reality Technology*, John Wiley & Sons, New York, 1994.

Burdea, G., *Force and Touch Feedback for Virtual Reality*, John Wiley & Sons, New York, 1996.

Buttolo, P., R. Oboe, B. Hannaford, and B. McNeely, "Force Feedback in Shared Virtual Simulations," Proceedings of MICAD, France, 6. pp.

Cohen, J., M. Lin, D. Manocha, and M. Ponamgi, "I-COLLIDE: An Interactive and Exact Collision Detection System for Large-Scale Environments," *Proceedings of ACM Interactive 3D Graphics Conference*, pp. 189–196, 1995.

Hollerbach, J., W. Thompson, and P. Shirley, "The Convergence of Robotics, Vision, and Computer Graphics for User Interaction," *The International Journal of Robotics Research*, vol. 18, no. 11, pp. 1088–1100, November 1999.

InterSense Co., "InterSense IS-900 Precision Motion Tracker," Company brochure, Burlington, MA, 2000. Also at www.isense.com.

Iwata, H., "Challenges in Networked Haptics," Haptics in Virtual Environments Workshop, IEEE Virtual Reality 2000, New Brunswick NJ, March 2000.

Kramer, J., "The Haptic Interfaces of the Next Decade," Panel Session, *IEEE Virtual Reality 2000 Conference*, March 2000.

Marcus, M., "Practical Aspects of Motion Capture Technology for the Entertainment Industries," Mirage Virtual News Bulletin, www.itc.co.uk/mirage, 1997.

Matossian, V., ``Integration of a Force Feedback Glove in Large Volume Partly Immersive VR Simulations," Master Thesis, University of Paris, June 1999.

Noma, H., T. Sughihara and T. Miyasato, "Development of Ground Surface Simulator for Tel-E-Merge System", *Proceedings of IEEE Virtual Reality 2000*, IEEE, 2000, pp. 217–224.

Sorid, D., "Giving Computers a Sense of Touch," *The New York Times*, page G 11, March 23, 2000.

Trefftz, H. and G. Burdea, "Calibration Errors in Large-Volume Virtual Environments," CAIP TR-243, Rutgers University, 2000.

Wilson, J., R. Kline-Schoder, M. Kenton, and N. Hogan, "Algorithms for Networked-Based Force Feedback," Creare Inc. Hanover NH, 1999.