The Haptic Museum

Margaret L. McLaughlin, Gaurav Sukhatme, Cyrus Shahabi,
Joao Hespanha, Antonio Ortega, and Gerard Medioni
Integrated Media Systems Center, University of Southern California, Los Angeles, CA 90089 USA
mmclaugh@usc.edu

Our IMSC team has used haptics to allow museum visitors to explore three-dimensional works of art by “touching” them, something that is not possible in ordinary museums due to prevailing “hands-off” policies [1, 2]. Haptics involves the modality of touch--the sensation of shape and texture an observer feels when exploring a virtual object, such as a three-dimensional model of a piece of pottery or art glass[3, 4, 5]. The haptic devices used in our research are the Phantom [6] and the CyberGrasp [7]. The Phantom is a desktop device that provides force feedback to the user's fingertip. The image on the left, below, shows a researcher at IMSC exploring the surface of a virtual teapot from USC’s Fisher Gallery using a Phantom. The image on the right shows the researcher calibrating the CyberGrasp, a whole-hand force-feedback glove that can be used to grasp virtual objects. A network of “tendons” transmits grasp forces back to the user’s fingers and palm. Both of these devices can be used by a remote museum visitor retrieving the model of the art object over the Internet or other network. Our mission is to develop seamless, device-independent haptic collaboration such that a museum staff member and a museum-goer or art student at a remote location can jointly examine a vase or bronze figure, note its interesting contours and textures, and consider such questions as "Why did the artist make this side rough but that side smooth?" or "What is this indentation on the bottom for?"

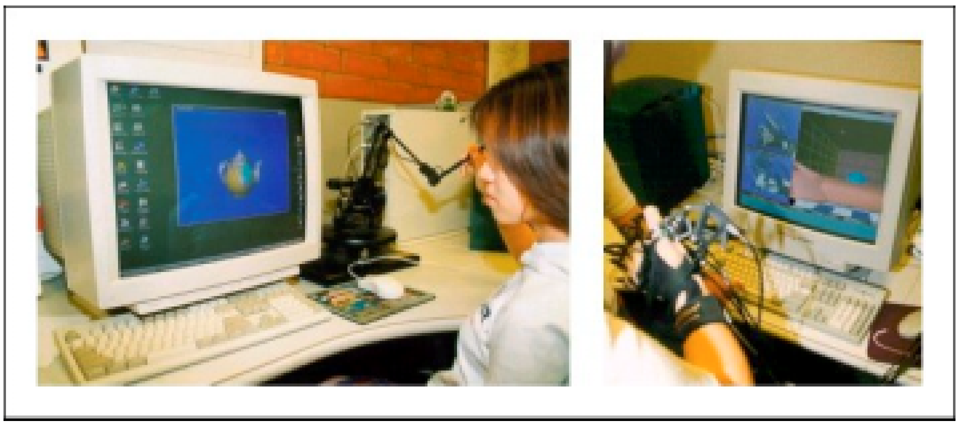

Fig 1 (a). IMSC researcher moving the Phantom over the surface of a
digitized teapot. (b) Researcher calibrating the CyberGrasp force-feedback glove

In support of our mission IMSC researchers are addressing a number of fundamental questions: What are the software tools and models needed to facilitate multi-user exploration of shared virtual environments? What are the spatial coordinates of the intersection between the haptic probe (or rendering device) and the virtual object being probed? How can current haptic displays be improved via better sensing ? How can we optimize low-level force control for haptic devices? What algorithms and techniques are needed to convey the feel of deformable objects? How do we capture users’ exploration with haptic devices? How do we compress haptic exploration data so that it becomes possible to store or transmit long interactive sessions? How do devices like the Phantom and CyberGrasp compare with respect to haptic fidelity? What haptic features facilitate appreciation? What understanding of the work of art is gained by enabling the sense of touch in virtual environments?

Acquisition of models

There are several commercial 3D digitizing cameras available for applications like the museum, such as the ColorScan and the Virtuoso shape cameras. We have chosen the 3Scan system because it provides a turnkey solution to our digitization needs, enabling us to generate fully-textured models in 3DS format compatible with our rendering systems (the Phantom, the World ToolKit environment). The 3Scan system includes a computer-controlled turntable, which can capture images from many views and reconstruct the solid surface of objects. 3Scan application software is equipped with decimation tools, which enables us to control the number of polygons that are created from 100 to 1,000,000.

In the ideal virtual museum, visitors will be able to “touch” not only rigid body objects such as vases and teapots, but also deformable objects such as tapestries and hangings. One area that IMSC researchers are beginning to explore is the development of reliable vision-based control systems for robotic applications such as the acquisition of images for 3D modeling. Feedback control employing video cameras as sensors was proposed several years back [8, 9], although only recently has the computing power required for robust image processing been available at a low cost. However, most of these applications are still far from ready for use in an industrial setting because of their low reliability. Two topics that have been identified as crucial for the development of reliable vision-based control systems for robotic applications—such as 3D & 4D modeling for haptics—are (1) the development of self-calibrating control algorithms [10] and (2) the use of single-camera image acquisition systems in feedback control [11]. One can use images of an object taken from multiple viewpoints to construct a 3D model of the object to be used for haptics [12]. To automate the procedure of collecting the multiple views one needs to have a camera mounted on a computer-controlled robot arm. This is particularly important to construct 4D models of objects whose shape is evolving (e.g., a work of art as it is being produced). From a controls perspective the research problem is to build algorithms to position the camera. The desired position can be specified directly in terms of its Cartesian coordinates or indirectly in terms of desired locations of parts of the object in the image. The latter falls in the area of vision-based control and is significantly more interesting because the use of vision in the feedback loop allows for great accuracy with not very precise, and therefore relatively inexpensive, robotic manipulators (cf. [12] and [13]).

Contact detection

The basic problem in haptics is to detect contact between the virtual object and the probe. Once this contact is reliably detected, a force corresponding to the interaction physics is generated and rendered using the probe. This process usually runs in a tight servo loop within a haptic rendering system. Instabilities due to a slow loop are well studied in the literature [14]. Lin et al. [15] have proposed an extensible framework for contact detection, which deconstructs the workspace into regions and at runtime identifies the region(s) of potential contacts. Currently, we are focusing on a specific aspect of the problem, namely generating the so-called Surface Contact Point. Due to discretization of both time (finite sampling periods) and space (the model geometry is a collection of polyhedra), the probe will penetrate the surface of the object. Clearly, this is not desirable for physical objects. One technique is to generate a SCP, which is the closest point on the surface to the actual tip of the probe. The force generation can then happen as though the probe were physically at this location rather than within the object. Existing methods in the literature generate the SCP by using the notion of a god-object [16], which forces the SCP to lie on the surface of the virtual object. Our approach to the generation of a SCP is to use information from the current contact detection cycle and *past* information from the contact history to predict the next SCP effectively. The prediction of the SCP will be particularly useful in cases where geometric models of the surface are available (as opposed to a polygonal approximation only). For example, it is possible to model surface texture using a simple periodic model. This information can be used to accurately predict the “next” SCP given a sequence of previous SCPs. It may be noted that we are not trying to solve the problem of detecting and identifying intersections between polyhedra (or other objects). Complementary efforts [17] along those lines already exist and are fairly well studied. We focus instead on incorporating local surface geometry into the prediction of the SCP assuming that a contact detection algorithm is already available. The novelty of our approach is in fusing contact history information with sensed location information at the present time to estimate the next surface contact point. We are experimenting with a well-known linear predictor, the Kalman Filter, as a first step [18].

Fundamentally, we treat SCP generation as an estimation problem. Given no information about the local surface geometry, the Kalman filter will solve the problem of predicting the next SCP by extrapolating the trajectory formed by the previous SCPs. This is especially useful in cases where local surface geometry is available at different resolutions over the surface being explored. In short, this allows the imposition of a local smoothing constraint. This also allows the haptic feedback system to incorporate smoothness constraints in velocity and acceleration, not just position in space. If the local surface geometry is available accurately then the “local” god-object method will produce an SCP estimate that the filter will weigh heavily compared to the predictive component. Thus using our predictive method it should be possible to work with objects at varying levels of detailed modeling.

A further novel aspect of our work is the ability of the SCP generation algorithm to incorporate local physics into the interaction. If the object being explored has varying texture or stiffness, the algorithm can adjust to this by automatically changing the appropriate parameters. Also distinctive is the ability of our algorithm to work with dual representations of objects. For contemporary graphics and haptics rendering, surface polyhedral approximations are commonly used. However, several other analytically "cleaner" representations exist for objects (e.g., generalized cylinders). These representations can be used for the predictive component of the filter and the sensed position of the probe and the "current" position of the SCP form the input to the update step.

Force feedback

From a controls perspective the greatest challenges with respect to force feedback are delays (due to computation and communication) and the interaction with an unpredictable disturbance (the human). Any of these easily leads to instability (oscillations) when one tries to increase the stiffness of the objects being rendered. Dealing with uncertainty is the basic subject of control engineering. However, the type of uncertainty present in these systems is particularly challenging. Variable delays, although conceptually simple to model, are difficult to handle because they introduce an infinite dimensional component to the problem [19]. To deal with uncertainty caused by variable delays and the human "on-the-loop" we will employ adaptation and learning. Because of the high nonlinearity of the model, supervisory adaptive control, involving online learning and logic-based switching, will most likely be needed. The research on this type of hybrid controller—characterized by combining continuous dynamics and discrete logic—was triggered by the pioneer work of Martenson [20] and has been pursued ever since [21]. Some of the more successful algorithms in this area go under the name of supervisory control and were proposed by Morse in [22] for linear systems and extended to the nonlinear case by Hespanha in [23].

Haptic recording and playback

An obvious application of recording and playing back haptic exploration data is for training purposes. A student can simply play back the trainer's exploration history to learn how to perform the same actions and apply the same forces and movements. Alternatively, the trainer can replay the student's session in order to evaluate his/her performance. Finally, a system can automatically grade students by computing a distance measure between the data of trainer and student. For such scenarios, analyzing the raw haptic data with no semantic component is sufficient. We have built a system prototype that acquires and plays back data generated from the Phantom [24]. The type of data captured includes movement, rotation, and forces (collisions between virtual objects). The result of this experience is a set of recorded haptic data that can later be played back and analyzed for training purposes.

Once we become aware of the semantics of each subtask performed by a user in a virtual environment, we can describe the entire task more effectively. For example, we can identify if a virtual object is unknown to a user and he/she is trying to identify its characteristics (by examining the hand posture, pattern of texture experience, and kinesthetic abstraction); or if the user is acquainted with the object and is grasping it with more confidence without experiencing its texture or stiffness. Our object-relational database system (OR-DBMS) stores both the semantic data and the original data about the object's characteristics, while our real-time file server stores the raw data with its real-time characteristic. For querying and analysis we need to consult the OR-DBMS first and then playback or do further analysis (or visualization) utilizing the data residing at the real-time file server. The advantage of an object-relational approach is that new data types and their corresponding functions (or methods) can be introduced to the DBMS as first class residents. Consequently, its query language (i.e., SQL3) as well as its other features (e.g., concurrency control, crash recovery) can be immediately utilized for the new data types.

Currently we are beginning work on a data model for storage and retrieval for the CyberGrasp, using two types of haptic data. Grasping data is a set of 26 floating values corresponding to 20 angles for the 20 degrees of freedom of the hand, 3 values for hand coordinates in 3D Cartesian space (x; y; z), and 3 angles for hand orientation (R_x; R_y; R_z). These values can be captured by monitoring the amount of stretch and bending, and the Euclidian distances of the fibers within the Cybergrasp fingers as well as the tracking device located on the center of its palm. Finally, the kinesthetic data is a set of 15 values representing the force applied to each finger, captured through monitoring Cybergrasp tendons, as three floating values F_x, F_y, $F_{z.}$ The frequency of the haptic loop in Cybergrasp is 1KHZ. We can sample the data at the same speed or lower speeds, depending on the application's requirements.

Haptic data compression

In many applications of haptics it will be necessary to store and/or transmit the data generated by these devices (i.e., measurements of position, force, etc). Depending on the specific device these parameters can add up to rates ranging from tens of kilobits to megabits per second. Our goal is to develop tools for compression of the data so that it becomes possible to store or transmit long interactive sessions. We aim at achieving reductions in storage memory needs of greater than an order of magnitude with minimal loss in accuracy. One goal will be to determine the best ways to quantify the loss introduced by compression; as far as we know we are the first to tackle the issue of compression of haptic data. We are exploring compression techniques starting with simple approaches (similar to those used in speech coding), and continuing with methods that will be more specific to the haptic data. We may use one of two lossy methods to compress it: one way is to use the lower sampling rate; the other is to note small changes during movement. For example, for certain grasp motions, not all of the fingers are involved. Further, during the approaching and departing phases the tracker data may be more useful than the CyberGlove data. Vector coding may be more appropriate to encode the time evolution of the multi-featured CyberGrasp data. For cases where the user employs the haptic device to manipulate a static object, we are also considering compression techniques that rely on knowledge of the object, rather than consider coding of an arbitrary trajectory in 3D space.

Haptic collaboration in virtual environments

We are working towards the real-time collection and simultaneous broadcast of haptic information to multiple haptics session participants, so that collaborative exploration of art objects is possible. Our hope is to enable the collection of data from disparate haptic devices, such as the Phantom and the CyberGrasp. We propose to equip two users, for example a museum staff member and a museum visitor (who have no visual contact with each other) with a glove each (or in a variation, with a glove and a Phantom respectively), and allow them to touch the same virtual object in a collaborative task [25]. When the same virtual environment is shared between two distributed sites there may be registration problems [25, 26]. Representations of the virtual object must coincide, but the distributed nature of the communication, especially over the Internet, may introduce considerable latency whose effects may be hard to predict. We hope to make significant progress on the registration of the haptic display systems in collaborative networked environments, and to examine the necessary elements to achieve networked collaboration with disparate haptic devices We plan to address not only integration issues but also questions related to the interaction process itself including feelings of co-presence and performance satisfaction [25], and how these variables are affected by the exploration modality (vision, vision plus haptics, haptics only).

We expect that within a few months the work outlined above will coalesce into a demonstration in which we are able to digitize an art object on the fly, and allow two users to probe it simultaneously by transmitting model information over the network. The contact detection framework outlined above for surface contact point generation via prediction will be used for both haptic devices. In the case of the Cybergrasp, there are several "probes" (the fingers of the user) for which the problem must be solved in parallel. For the first implementation we disregard the case where there might be multiple contacts between a particular finger and the object being manipulated; we assume a single (or no) point of contact between each finger and the object. The demonstration will integrate IMSC's technology for model acquisition, contact point calculation, force generation, and compression.

Summary

The main contributions of our research program will be: (1) development of reliable vision-based control systems for robotic applications such as the acquisition of images for 3D modeling; (2) development of low-level force-control algorithms for haptics rendering, robust with respect to uncertain delays and the interaction with a human "on-the-loop;" (3) application of well-known estimation theory to novel problems in collision detection; (4) techniques for haptic data compression, and methods to evaluate the perceptual impact of lossy compression of haptic data; (5) strategies for the description, storage, and retrieval of a new data type for immersive environments, haptic data; and (6) software and tools for managing real-time, over-the-network exploration of museum objects with disparate haptic devices.

Acknowledgments: This research has been funded by the Integrated Media Systems Center, a National Science Foundation Engineering Research Center, with additional support from the Annenberg Center for Communication at the University of Southern California, the California Trade and Commerce Agency, and Hitachi America, Inc.

References

[1] McLaughlin, M. L., N. Ellison, J. Lucas and S. G. Goldberg, "The Interactive Electronic Exhibition: Restructuring the Boundaries between Museums and their Constituencies," International Communication Association Annual Conference, Montreal, Canada (1997).

[2] Schertz, P. M., J. Jaskowiak and M. L. McLaughlin, "Evaluation of an Interactive Art Museum," SPECTRA, A Publication of the Museum Computer Network, 25 (1), 33-37 (1997).

[3] Srinivasan, M. and C. Basdogan, , "Haptics in Virtual Environments: Taxonomy, Research Status, and Challenges," Computers and Graphics, 21 (4), 393-404 (1997).

[4] Floyd, J., "Haptic Interaction with Three-Dimensional Bitmapped Virtual Environments," in Salisbury, J. K. and Srinivasan, M. A. (Eds), Proc. Fourth PHANTOM Users Group Workshop, AI Lab Technical Report No. 1675 and RLE Technical Report No. 633, MIT (1999).

[5] Wilson, J. P, R. J. Kline-Schoder, M. A. Kenton, and N. Hogan, "Algorithms for Network-Based Force Feedback," in Salisbury, J. K. and Srinivasan, M. A. (Eds.), Proc. PHANTOM Users Group Workshop, AI Lab Technical Report No. 1675 and RLE Technical Report No. 633, MIT (1999).

[6] Massie, T. H. and J. K. Salisbury, "The PHANTOM Haptic Interface: A Device for Probing Virtual Objects," Proc. ASME Winter Annual Meeting, Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems, Chicago, IL (1994)

[7] Virtual Technologies, "CyberGrasp: Groundbreaking Haptic Interface for the Entire Hand," http://www.virtex.com/products/hw_products/cybergrasp.html.

[8] Hutchinson, S., G. Hager, and P. I. Corke, "A Tutorial on Visual Servo Control," *IEEE Trans. on Robotics and Automation*, 12(5), 651-670 (1996).

[9] Corke, P. I., "Visual Control of Robot Manipulators—A Review," in K. Hashimoto (Ed.), *Visual Servoing*, pp. 1-32, World Scientific (1994).

[10] Hager, G. D., W.-C. Chang, and A. S. Morse, "Robot Hand-eye Coordination Based on Stereo Vision," *IEEE Contr. Syst. Mag.*, vol. 15, pp. 30-39 (1995).

[11] Hespanha, J. P., "Single-camera Visual Servoing." Submitted to the 39th Conf. on Decision and Contr., Jan. 2000.

[12] Tang, C.-K., and G. Medioni, "Robust Estimation of Curvature from Noisy 3D Data for Shape Description," in *Proc. of the IEEE Conference on ICCV*, pp. 426-433, 1999.

[13] Hespanha, J. P., Z. Dodds, G. Hager, and A. S. Morse, "What Tasks can be Performed with an Uncalibrated Stereo Vision System?" *The Int. J. of Computer Vision*, 35(1), pp. 65-85 (1999).

[14] Colgate, J. E., M. C. Stanley, and J. M. Brown, "Issues in the Haptic Display of Tool Use," Proc. IEEE/RSJ

[15] Lin, M. C., A. Gregory, S. Ehmann, S. Gottschalk, and R. Taylor, "Contact Determination for Real-Time Haptic Interaction in 3D Modeling, Editing, and Painting," in Salisbury, J. K. and Srinivasan, M. A. (Eds.), Proc. Fourth PHANTOM Users Group Workshop, AI Lab Technical Report No. 1675 and RLE Technical Report No. 633, MIT (1999)

[16] Zilles, C. B., and J.K. Salisbury, "A Constraint-based God-object Method for Haptic Display," Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems, Pittsburgh, PA, pp. 146-151 (1995).

[17] Lin, M. C., and S. Gottschalk , “Collision Detection between Geometric Models: A Survey,” Proceedings of IMA Conference on Mathematics of Surfaces (1998).

[18] S. I. Roumeliotis, G. S. Sukhatme, and G. A. Bekey. “Smoother Based 3-d Attitude Estimation for Mobile Robot Localization. In *IEEE International Conference on Robotics and Automatio*n (1999).

[19] Kolmanovski, V., S.-I. Niculescu, and K. Gu, “Delay Effects on Stability: A Survey,” in *Proc. of the 38th Conf. on Decision and Control*, pp. 1993-1998 (1999).

[20] Martenson, B., “The Order of any Stabilizing Regulator is Sufficient A Priori Information for Adaptive Stabilization,” *Syst. & Contr. Lett.*, vol. 6, pp. 87-91 (1985).

[21] Morse, A. S., “Control Using Logic-Based Switching,” in A. Isidori, ed., *Trends in Control: A European Perspective*. London: Springer-Verlag, pp. 69-113 (1995).

[22] Morse, A. S., “Supervisory Control of Families of Linear Set-point Controllers—Part 1: Exact Matching," *IEEE Trans. Automat. Contr.*, vol. 41, pp. 1413-1431 (1996).

[23] Hespanha, J. P., *Logic-Based Switching Algorithms in Control*. PhD thesis, Yale University, New Haven, CT, (1998).

[24] C. Shahabi, G. Barish, B. Ellenberger, N. Jiang, M. Kolahdouzan, A. Nam, and R. Zimmermann, “Immer-sidata Management: Challenges in Management of Data Generated within an Immersive Environment,” *Multimedia Information System*s, October (1999).

[25] Basdogan, C., C.-H. Ho, M. Slater, and M. A. Srinavasan, "The Role of Haptic Communication in Shared Virtual Environments," Proc. PHANTOM Users Group, PUG98 (1998) http://www.sensable.com/community/98papers/19%20basdogan-pug98.pdf

[26] Buttolo, P. J., R. Hewitt, R. Oboe, and B. Hannaford, "Force Feedback in Virtual and Shared Environments," online, available HTTP: WWW URL brl.ee.washington.edu/BRL/publications/Rep100.ps (1997).