Touch in Immersive Environments

Gaurav Sukhatme
Computer Science
Joao Hespanha
Electrical Engineering-Systems
Margaret L. McLaughlin
Annenberg School for Communication
Cyrus Shahabi
Computer Science

Integrated Media Systems Center
University of Southern California
Los Angeles, CA 90089

In Proc. of the EVA 2000 Scotland Conf. on Electronic Imaging and the Visual Arts, July 2000.

Touch in Immersive Environments

The haptic interface is becoming an increasingly important component of immersive systems. Haptics refers to the **modality of touch** and the sensation of shape and texture an observer feels when exploring a virtual object, such as a three-dimensional model of a vase or teapot. In collaboration with the Fisher Gallery, USC's on-campus art gallery, 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.

Haptic Devices

The haptic devices currently used in our research are the PHANTOM and the CyberGrasp. The PHANTOM, from Sensable, Inc., is a small, desk-grounded robot that permits simulation of single fingertip contact with virtual objects through a thimble or stylus. This device tracks the *x*, *y*, and *z* Cartesian coordinates and pitch, roll, and yaw of the virtual probe as it moves about a three-dimensional workspace. Its actuators also communicate force back to the user's fingertips as it detects collisions with virtual objects, simulating the sense of touch. The CyberGrasp, from Virtual Technologies, is an exoskeletal device that fits over a 22 DOF CyberGlove, providing simulated grasp forces. It is used in conjunction with the Polhemus Fastrak to measure the position and orientation of the forearm in three-dimensional space. Both of these devices can be used by a remote museum visitor who retrieves a 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

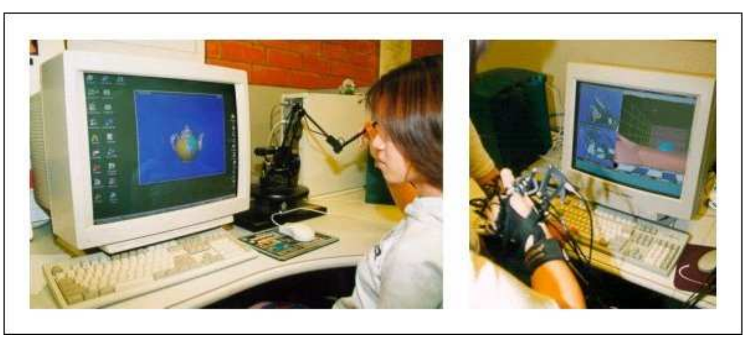

Figure 1. (a) Researcher exploring a digitized teapot from USC's Fisher Gallery with the PHANTOM; (b) (c) Researcher calibrating the CyberGrasp force-feedback glove.

Digitization

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. It saves hours or days of modeling effort for each object. Unlike our previous method of modeling from images, there is no need to manually match points between images; this step is fully automated. Therefore, it saves time. A 3D model takes just 15 minutes to generate.

Figure 2. The 3Scan digitizing system.

The 3Scan enables us to generate fully-textured models in 3D Studio format compatible with our rendering systems (the PHANTOM, the World ToolKit environment). The 3Scan system contains a computer-controlled turntable, which can capture images from many views and reconstruct the solid surface of objects. Our previous method left us with some holes on the object's surface and the bottom. And finally, the 3Scan application software is equipped with decimation tools, which enable us to control the number of polygons that are created from 100 to 1,000,000. Because the system is easy to use we can assign digitization tasks to students and concentrate our energies on some unresolved digitization problems, such as importing object descriptions from data files generated by third party 3D modelers (such as 3DS, VRML format etc.) into a haptic scene graph for the CyberGrasp.

Our Research Approach

During the last several years, we have been addressing a number of fundamental questions about haptics: What are the software tools and models needed to facilitate multi-user exploration of shared virtual environments? How can current state-of-the-art 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 art objects such as textiles and hangings? 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? What haptic features facilitate immersion in virtual environments? (McLaughlin *et al*, 1999; McLaughlin *et al*, 2000). Some of the basic areas of our research program are the following:

- Novel low-level force-control algorithms for haptics rendering. These algorithms are robust with respect to uncertain delays (due to computation and communication) and the interaction with a human "on-the-loop." With such algorithms, it will be possible to improve the quality of haptic rendering significantly.
- System identification techniques for haptic probes. The novelty of these algorithms is the nature of their application. These algorithms have not been previously used in this context.
- An estimation theory based method for collision detection. Estimation theory based algorithms are used extensively in robotics and other applications. Our approach in using them for collision detection is novel.
- Software and tools for managing real-time collaboration over the Internet, including realistic mutual touch, with disparate haptic devices.
- Development of reliable vision-based control systems for robotic applications such as the acquisition of images for 3D modeling.

Our collaborations with other USC colleagues have also resulted in

- Techniques for haptic data compression, and methods to evaluate the perceptual impact of lossy compression of haptic data.
- Strategies for the description, storage, and retrieval of haptic data.

Below we report on our current approach to the areas of contact detection, force feedback, and haptic collaboration.

Contact detection

What are the best techniques for detecting the collision of the virtual hand with the art object? 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. Kalman filtering is an established technique for state estimation. To our knowledge, predictive techniques for estimating SCPs have not been studied elsewhere.

The Kalman filter is implemented in two steps. The propagation step uses the system matrix to generate the next state from the previous state:

$$x_{k+1/k} = \Phi x_{k/k} \qquad P_{k+1/k} = \Phi P_{k/k} \Phi^T + Q$$

The update step computes the Kalman gain as follows:

$$K = P_{k+1/k}\left(P_{k+1/k} + R\right)^{-1}$$

The correction to the state estimate is now computed by scaling the residual by the Kalman gain:

$$x_{k+1/k+1} = x_{k+1/k} + K(res) \qquad P_{k+1/k+1} = (I - K)P_{k+1/k}$$

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.

The second 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 (this is a well- known technique where the parameters of the system are estimated online using the same filter which is performing state update).

The third novel aspect of the work 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 (technically, the propagation step within the Kalman filter) and the sensed position of the probe and the “current” position of the SCP form the input to the update step.

Force control

Haptic devices fall roughly into two categories: *impedance displays* that measure positions and exert forces, and *admittance displays* that measure forces and set positions (Yoshikawa *et al.*, 1994). Typically, impedance displays have low inertia and are torque actuated, whereas admittance displays have high inertia and are driven by a position or velocity control loop. Our work focuses mostly on impedance displays, such as the PHANTOM (Massie & Salisbury, 1994), due to their ease of operation, low cost, and versatility. We argue below that, employing advanced control techniques, it is possible to obtain very high fidelity haptic rendering with this type of device. We proceed to describe the state-of-the-art and limitations of the current haptic rendering systems.

Consider a human operator holding the PHANTOM's stylus. The functionality expected from the haptic device is to make the user believe that the stylus she is holding moves in a virtual environment and should therefore not penetrate the objects but rather react to them according to Newtonian mechanics. Suppose then that we desire to simulate the existence in the virtual world of a single planar object whose interior is defined by

$$\{x \in \mathbf{R}^3 : n_{obj}'x < 1\},$$

where $n_{obj} \in \mathbf{R}^3$ denotes the outer normal to the plane that defines the object. This prototypical problem is known in the literature as the rendering of a "virtual wall" (Colgate *et al.*, 1993).

Denoting by $x_{sty} \in \mathbf{R}^3$ the Cartesian coordinates of the stylus' tip, as the human exerts a force f_{hum} on the stylus with mass m_{sty}, it should move according to

$$m_{sty} \ddot{x}_{sty} = f_{hum} + f_{rec}, \tag{1}$$

where f_{rec} denotes the reaction force that the virtual object exerts on the stylus. Clearly, if the stylus is in a position that corresponds to virtual free space (i.e., if $n_{obj}'x > 1$), then $f_{rec} = 0$. However, when the stylus attempts to penetrate the virtual object, f_{rec} becomes nonzero. For a non-deformable object, f_{rec} is completely defined by the constraint

$$n_{obj}' x_{sty} \geq 1, \tag{2}$$

whereas, for a deformable object, f_{rec} is a function of the amount of penetration and the velocity. Typically, for frictionless deformable objects,

$$f_{rec} = - k\,(n_{obj}' x_{sty} - 1) - b\, n_{obj}' \dot{x}_{sty}, \tag{3}$$

where $k > 0$ denotes the stiffness of the virtual object and b its damping coefficient. Since the satisfaction of the constraint (2) for a non-deformable object may require impulsive forces, it is common to model all virtual objects as deformable. Those that are essentially non-deformable are assumed to have very large stiffness coefficients k. Equation (3) is often represented as

$$f_{rec} = H(s)(n_{obj}' x_{sty} - 1),$$

where the operator

$$H(s) := - k - bs, \qquad s := \frac{d}{dt}, \tag{4}$$

is called the *impedance* of the virtual object (Hogan, 1989). The key problem in the control of haptic devices is to emulate the motion defined by equation (1), as the user interacts with the device. We shall see below that this is particularly challenging when the impedance has a large stiffness coefficient.

The equations of motion of an impedance haptic display such as the PHANTOM can be expressed as

$$M(q)\, \ddot{x}_{sty} + C(q, \dot{q})\, \dot{x}_{sty} + N(q, \dot{q}) = J^{-T}(q)\, \tau + f_{hum}, \tag{5}$$

where q denotes a vector of generalized coordinates, τ the vector of joint torques, and f_{hum} the force applied on the tip of the device by the human operator. The functions M, C, N, J describe the dynamic parameters of the system. In particular, M and C denote the effective mass and Coriolis matrices, respectively, N the vector of effective external forces, and J^{-T} the transpose inverse of the manipulator Jacobian matrix J (Murray *et al.*, 1994). Although we focus on a 3-DOF device, the concepts generalize to any haptic device that can be modeled by a generic manipulator equation like (5). The 3-DOF PHANTOM (Massie and Salisbury, 1994), the 6-DOF PHANTOM (Chen, 1999), and the CyberGrasp are examples of such devices.

To emulate (1) using (5), one would need to set

$$\tau = J^{T} C\, \dot{x}_{sty} + J^{T} N + J^{T} \left(\frac{1}{m_{sty}} M - I\right) f_{hum} + \frac{1}{m_{sty}} J^{T} M\, f_{rec},$$

where I denotes the 3×3 identity matrix. However, since impedance haptic displays typically do not possess sensors capable of measuring the force f_{hum}, the control law

$$\tau = J^{T} C \dot{x}_{\text{sty}} + J^{T} N + \frac{1}{m_{\text{sty}}} J^{T} M f_{\text{rec}} \tag{6}$$

is often used instead. This results in a closed-loop system of the form

$$m_{\text{sty}} \ddot{x}_{\text{sty}} = m_{\text{sty}} M^{-1} f_{\text{hum}} + f_{\text{rec}}.$$

In off-the-shelf haptic devices, the following simplified version of (6) is often used (Carignan & Cleary, 2000)

$$\tau = \frac{1}{m_{\text{sty}}} J^{T} f_{\text{rec}}, \tag{7}$$

resulting in a closed-loop of the form

$$m_{\text{sty}} \ddot{x}_{\text{sty}} = m_{\text{sty}} M^{-1} f_{\text{hum}} + f_{\text{rec}} - m_{\text{sty}} M^{-1} C \dot{x}_{\text{sty}} - m_{\text{sty}} M^{-1} N. \tag{8}$$

For (8) to provide a good approximation to (1), two conditions would have to be verified:

1. $M = m_{\text{sty}} I$ and therefore $m_{\text{sty}} M^{-1}$ is approximately equal to the identity. For this to be true the mass of the stylus should be large when compared to the mass of the manipulator's links.
2. $M^{-1} C \dot{x}_{\text{sty}}$ and $M^{-1} N$ are negligible.

The above requirements are rarely met for small impedance displays and, even for a device as well engineered as the PHANTOM, the unbalanced weight can be as high as 0.2 Newton and the user perceives a mass at the tip of about 100 grams (significantly higher than the mass of the stylus) (Massie & Salisbury, 1994). In practice, this results in noticeable *haptic distortion*, also known as non-transparency of the haptic interface (Adams & Hannaford, 1999).

Another difficulty that arises with the control law (6) is that there is always some delay associated with computing f_{rec}. This delay is mainly due to three factors: time-sampling; computation delay in the contact detection algorithms that determine if the tip of the stylus is penetrating a virtual object; and communication delay between the computer that is executing the contact detection algorithm and the haptic device. The delay due to time-sampling is unavoidable and dominates in current haptic applications. Computation delays become significant in very complex scenes for which the contact detection algorithms require more than one sampling period (usually 0.5-1.0 milliseconds). Communication delays dominate when the virtual environment is distributed between several computers at distinct geographical locations (cf. Haptic collaboration below).

When one takes the delay into consideration, equation (8) actually becomes

$$m_{\text{sty}} \ddot{x}_{\text{sty}} + b\, n_{\text{obj}}' \dot{x}_{\text{sty}} (t - t_{\text{delay}}) + k\, (n_{\text{obj}}' x_{\text{sty}}(t - t_{\text{delay}}) - 1) = f_{\text{hum}}, \tag{9}$$

where t_{delay} denotes the total delay. We have used here the model (3) and, for simplicity, assumed that the device does not introduce haptic distortion, i.e., $M = m_{\text{sty}} I$, $M^{-1} C \dot{x}_{\text{sty}} = 0$, and $M^{-1} N = 0$. Although (9) is stable when $t_{\text{delay}} = 0$, it rapidly becomes unstable for $t_{\text{delay}} > 0$, especially when the stiffness coefficient k is large. The dynamic range of impedance

$H(s)$ that a haptic interface is capable of rendering in a stable fashion is called the *z-width* and is an important measure of performance.

The methods proposed in the literature to deal with instability caused by delay can be roughly classified as model-based or nonmodel-based. *Nonmodel-based* methods design the haptic interface so as to make it passive from the perspective of the human operator (Adams & Hannaford, 1999; Colgate & Schenkel, 1997; Miller *et al.*, 1999). Nonmodel-based methods provide a high degree of robustness with respect to the behavior of the human. In fact, they guarantee stability as long as the human behaves passively (and there is some energy dissipation in the interface). However, and because of this, these methods are very conservative and lead to small z-width. *Model-based* approaches use a more precise characterization for the behavior of the human operator to reduce conservativeness and increase z-width (Carignan & Cleary, 2000; Gillespie & Cutkosky, 1996; Tsai & Colgate, 1995). In (Carignan & Cleary, 2000; Gillespie & Cutkosky, 1996) the force exerted by a human limb is modeled as a second order system of the form

$$f_{\text{hum}} = - k_{\text{hum}}\, n_{\text{obj}}'\, x_{\text{sty}} - b_{\text{hum}}\, n_{\text{obj}}'\, \dot{x}_{\text{sty}} + g_{\text{hum}}, \tag{10}$$

where $k_{\text{hum}} > 0$ denotes the stiffness of the limb, b_{hum} its damping coefficient, and g_{hum} a bias force. This model is inspired by empirical studies of limb dynamics such as (Hajian & Howe, 1997). In (Tsai & Colgate, 1995), the human is modeled as a general linear time-invariant system with input $n_{\text{obj}}'\, x_{\text{sty}}$ and output f_{hum}. The main difficulty in model-based approaches is their low tolerance to errors in the modeling of the human operator. In (Tsai & Colgate, 1995), robust control techniques are proposed to deal with these errors at the expense of added conservativeness.

We expect to improve upon the existing algorithms for force control in two fundamental ways:

1. Employ nonlinear compensation to cancel the nonlinear terms present in (8) to improve transparency of the haptic interface. Because of device modeling errors, this compensation may need to be robustified using inverse optimality (Freeman & Kokotovic, 1996) and nonlinear robust control (Aloliwi & Khalil, 1998; Arcak & Kokotovic, 1999; Jiang *et al.*, 1996; Jiang & Hill, 1999; Krstic *et al.*, 1996).
2. Utilize Multiple-model Supervisory Control (MMSC) (Anderson *et al.*, 1999; Hespanha *et al.*, 1999; Hespanha & Morse, 1999 *a*, *b*; Hespanha *et al.*, 2000) to estimate online a model for the human behavior and use it for control. This effectively results in an adaptive version of (Carignan & Cleary, 2000; Gillespie & Cutkosky, 1996; Tsai & Colgate, 1995). This application of MMSC is closely related to the noise suppression problem in (Fujii *et al.*, 1998). MMSC can also be used to online estimate and precisely cancel the nonlinear terms that appear in (8), for improved haptic transparency.

MMSC is a form of adaptive control that combines online learning with logic-based switching. The research on this type of *hybrid controller*—characterized by combining continuous dynamics and discrete logic—was triggered by the pioneering work of Mårtenson (1985) and has been pursued ever since (cf., Morse, 1995 and references therein). Some of the more successful algorithms in this area were proposed by Morse (1996) for linear systems and extended to the nonlinear case in Hespanha (1998) and in Hespanha and Morse (1999a, 1999b). In supervisory control, one starts by designing off-line an array of “candidate controllers.” These controllers should be chosen so that, for each possible value of the unknown process model, there is at least one controller that achieves the desired performance. A specially designed supervisor is then used to orchestrate the switching among the candidate controllers. To decide which controller should be put in the loop, the supervisor continuously monitors the so called “performance signals.” Each performance signal measures the likelihood that the corresponding controller will yield good performance at controlling the process.

MMSC is characterized by very fast adaptation to sudden changes in the process or the control objectives. Such changes are expected in haptics due to the unpredictability of the human-in-the-loop. In fact, it is shown in Hajian and Howe (1997) that changes in the parameters of the human limb dynamics become noticeable over periods of time larger than 20 milliseconds. The use of online adaptation was recently suggested as a promising direction to improve haptic rendering (Carignan & Cleary, 2000; Gillespie & Cutkosky, 1996).

Haptic Collaboration

We plan to design and test a system for haptic collaboration across a network in support of “communication by touch”. The basic idea is to provide a framework for multiple users (each with his or her own haptic device connected to a computer) to share a common experience of touch. This will allow users to exert forces on each other through the network as well as exert forces on common objects. We have already implemented a prototype system (see Figure 3 (a) (b)), which allows users with different device types to touch each other over a network.

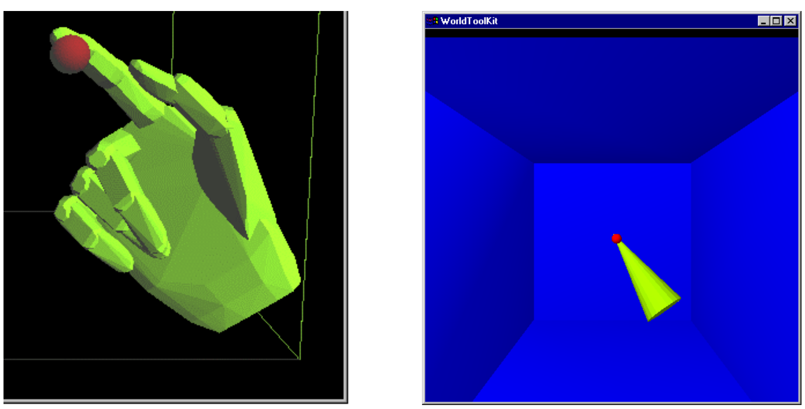

Figure 3. Mutual touch in a virtual environment. (a). A CyberGrasp user (the green hand) experiences the touch of a PHANTOM user at a remote location (the red ball represents the PHANTOM tip). (b) A PHANTOM user feels force against her fingertip as it collides with a single surface contact point on the virtual hand.(represented by the cone).

Several basic research issues need to be addressed to ensure that our collaboration system works reliably and scales to large numbers of users.

1. **Model consistency**: A key issue is maintaining model consistency. When multiple users share a common world model, one approach is to incrementally send updates of the world model between them so that the entire model (which can be substantial) does not have to be sent back and forth. This leads to model inconsistencies. If care is not taken then over time the rendering experience is haptically distorted and not useful. We are particularly interested in the scaling properties of this problem.
2. **Latency:** A second key issue is latency. When a user feel forces due to his/her interaction with the virtual environment, these forces need to be rapidly computed as functions of the actions of the other users in the system. The latency of this computation introduces a lag in the control system computing the resultant force that is to be rendered, causing instability and poor rendering (cf. previous section).
3. **Mutual exclusion:** When multiple users have "modify" access to the same environment the problem of mutual exclusion comes into the picture. When one user is “writing” to an object (moving it, deforming it, etc.) one might want to “lock” it so that the other users cannot modify its properties.

Consider n users collaborating in a virtual environment populated with m objects where $n>m$. Assuming a simplified world where the only interaction possible between a user and an object is a point contact, there are potentially n (user-object) contact force vectors at each timestep. In addition if we assume that contact between objects is also restricted to single points then there are $m(m-1)/2$ such contacts. Algorithms that deal with maintaining consistency and latency therefore have to deal with input which is $O(n + m\text{^}2)$. We are interested in algorithms that scale well, i.e., which efficiently maintain consistency and address latency for large m and n. The same applies to mutual exclusion as well.

One approach to this has been “dead-reckoning”. Creare Inc. is building a haptic display toolkit to stabilize network-based force-feedback applications when communication or processing time delays are present (Wilson *et al*., 1999). This toolkit and in particular their use of haptic “dead reckoning” is potentially applicable to the haptic collaboration component of our work.

Dead-reckoning, in its simplest form, refers to the process of estimating the next state of the system based on the current state and the first derivative. It is a process of estimating (tracking) the state of a dynamical system by local linearization. Dead-reckoning has received considerable attention in the robotics literature (see Sukhatme *et al*., 1998 and references therein). The key problem with dead-reckoning is drift. Due to local linearization, the estimate of the state of the system and the actual state diverge rapidly. Stabilizing drift in dead-reckoned systems is done primarily by feedback from sensing. From time to time, the actual state of the system needs to be measured with (high) certainty and the state estimate needs to be updated. The classical formulation of this is given in basic filter theory where the underlying model is assumed to be linear. The Kalman filter is an example of such an approach. In our previous work we have used Kalman filtering techniques for stabilizing drift in mobile robots (see Roumeliotis *et al.* 1999a, 1999b). The canonical formulation of the problem is to track the locus of a dynamical system, in our case, the position of the haptic probes and the objects in the environments.

The Kalman filter is implemented in two steps. The propagation step uses the system matrix to generate the next state from the previous state:

$$x_{k+1/k} = \Phi x_{k/k} \qquad P_{k+1/k} = \Phi P_{k/k} \Phi^T + Q$$

The update step computes the Kalman gain as follows:

$$K = P_{k+1/k}\left(P_{k+1/k} + R\right)^{-1}$$

The correction to the state estimate is now computed by scaling the residual by the Kalman gain:

$$x_{k+1/k+1} = x_{k+1/k} + K(res) \qquad P_{k+1/k+1} = (I - K)P_{k+1/k}$$

Fundamentally, we treat tracking as an estimation problem. The Kalman filter solves the problem of predicting the next state of an object by extrapolating the trajectory formed by its previous motion. This is especially useful in cases where local state information is available at different resolutions over the environment being explored. Further, if *a priori* knowledge about dynamics of the various objects in the scene is available, then this can be reflected in the system matrix. An algorithm based on linear tracking can also incorporate local physics into the model. If the object being explored has varying texture or stiffness, the algorithm can adjust to this by automatically changing the appropriate parameters (this is a well- known technique where the parameters of the system are estimated online using the same filter which is performing state update), see (Roumeliotis 1998). The third novel aspect of the work 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 (technically, the propagation step within the Kalman filter).

Thus the problem of dealing with a large number of regular state updates can be decomposed into local state estimation based on extrapolation and occasional full-state update when the error in the estimate is high enough. It may be noted that our previous work has addressed exactly this problem in the context of mobile robot localization with considerable success. (Roumeliotis 1999a, 1999b). Both model consistency and latency issues are addressed by this approach. Further, the latency issues are also addressed by intelligent control of force (cf. previous section).

Mutual exclusion is an interesting problem not directly addressed by the above technique. The basic underlying problem (Dijkstra, 1965) has been studied extensively as a resource conflict problem in the literature. Processes that share access to the same memory successfully (due to locks or atomic read/writes) are said to have a "critical section" which is the code that allows access to shared variables. In our work these might be the physical properties of the objects in the virtual environment. We plan to use a well-known technique based on semaphores (Lamport, 1977) to implement basic locks to avoid read-write problems in shared data. The approach is classical but its application to this domain would be novel.

Haptics for the Visually Impaired: A New Application Area

We are undertaking a new application, "Soundscapes for the Blind," in collaboration with Los Angeles area artist Marcos Lutyens. As Lutyens notes: “For the visually impaired who cannot use their eyes to situate themselves in their surroundings, touch and hearing become the main means of connection to the outside world. Sounds become the lynch-pin around which a scene is deciphered, and descriptive words often substitute solid objects that are out of reach. So, in a sense, words and sounds begin to have a volume and a mass, as their interpretation and meaning substitute visual experience. "Soundscapes" is a haptic environment that allows blind people to touch and feel their own words and other familiar sounds, as three-dimensional objects in space. This will be both an entertaining and educational tool for visually impaired children, as words, noises and sounds can be explored and played with by means of haptic devices connected to microphones and voice processing programs that together create a fully interactive environment.”

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, Hitachi America, Inc., and the Annenberg School for Communication.

References

Adams, R. J., and B. Hannaford, **“**Stable Haptic Interaction with Virtual Environments,” *IEEE Trans. Robot. Automat.*, 15(3), 465-474 (1999).

Aloliwi, B., and H. K. Khalil, “Robust Control of Nonlinear Systems with Unmodeled Dynamics,” *Proc. of the 37th Conf. on Decision and Contr.*, pp. 2872-2873 (1998).

Anderson, B. D. O., T. S. Brinsmead, F. de Bruyne, J. P. Hespanha, D. Liberzon and A. S. Morse, **“**Fundamental Issues in Multiple Model Adaptive Control,” to appear in the George Zames Special Issue of the *Int. J. on Robust and Nonlinear Control* (forthcoming).

Arcak, M., and P. V. Kokotovic, “Further Results on Robust Control of Nonlinear Systems with Input Unmodeled Dynamics,” *Proc. 1999 American Contr. Conf.*, pp. 4061-4065 (1999).

Carignan, C. R., and K. R. Cleary, “Closed-loop Force Control for Haptic Simulation of Virtual Environments,” *Haptics-e, The Electronic Journal of Haptics Research,* 1 (2) http://www.haptics-e.org/Vol_01/he-v1n2.pdf (2000).

Chen, E., “Six Degree-of-freedom Haptic System for Desktop Virtual Prototyping Applications,” *Proc. First International Workshop on Virtual Reality and Prototyping*, pp. 97-106 (1999).

Colgate, J. E., P. E. Grafing, M. C. Stanley and G. Schenkel, “Implementation of Stiff Virtual Walls in Force-reflecting Interfaces,” *Proc. IEEE Virtual Reality Annual International Symposium*, pp. 202-208 (1993).

Colgate, J. E., and G. Schenkel, “Passivity of a Class of Sampled-data Systems: Application to Haptic Interfaces,” *J. of Robotic Systems*, 14(1), 37-47 (1997).

Dijkstra, E. W. “Solution of a Problem in Concurrent Programming Control”. *Communications of the ACM* 8, 9 (Sept), pp. 569 (1965).

Freeman, R. A., and P. V. Kokotovic, “Robust Nonlinear Control Design: State-Space and Lyapunov Techniques,” *Systems & Control: Foundations & Applications*. Birkhäuser: Boston (1996).

Fujii, S., J. P. Hespanha, and A. S. Morse, **“**Supervisory Control of Families of Noise Suppressing Controllers,” *Proc. 37th Conf. on Decision and Contr.*, pp. 1641-1646 (1998).

Gillespie, B., and M. Cutkosky, “Stable User-specific Rendering of the Virtual Wall,” *Proc. ASME International Mechanical Engineering Conference and Exposition*, Vol. 58, pp. 397-406 (1996).

Hajian, A. Z., and R. D. Howe, “Identification of the Mechanical Impedance at the Human Finger Tip,” *ASME Journal of Biomechanical Engineering*, 119(1), 109-114 (1995).

Hespanha, J. P., “Single-camera Visual Servoing.” Submitted to the 39th Conf. on Decision and Contr., Jan. 2000 (2000).

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

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

Hespanha, J. P., D. Liberzon, A. S. Morse, B. D. O. Anderson, T. S. Brinsmead and F. de Bruyne, **“**Multiple Model Adaptive Control, Part 2: Switching,” Submitted for publication in the *Int. J. on Robust and Nonlinear Control* (2000).

Hespanha, J. P., D. Liberzon and A. S. Morse, “Logic-based Switching Control of a Nonholonomic System with Parametric Modeling Uncertainty,” *Syst. & Contr. Lett. Special Issue on Hybrid Systems*, 38 (3), 167-177 (1999).

Hespanha, J. P., and A. S. Morse, “Certainty Equivalence Implies Detectability,” *Syst. & Contr. Lett*, 36 (1), 1-13 (1999a).

Hespanha, J. P., and A. S. Morse, “Supervisory Control of Integral-input-to-state Stabilizing Controllers,” *Proc. 1999 European Contr. Conf.* (1999b).

Hogan, N., “Controlling Impedance at the Man/machine Interface,” *Proc. IEEE International Conference on Robotics and Automation*, 3, pp. 1626-1631 (1989).

Jiang, Z.-P. and D. J. Hill, “A Robust Adaptive Backstepping Scheme for Nonlinear Systems with Unmodeled Dynamics,” *IEEE Trans. Automat. Contr.*, 44 (9), 1705-1711 (1999).

Jiang, Z.-P., I. Mareels and J.-B. Pomet, “Controlling Nonlinear Systems with Input Unmodeled Dynamics,” *Proc. 35th Conf. on Decision and Contr.*, pp. 805-806 (1996).

Krstic, M., J. Sun and P. V. Kokotovic, “Robust Control of Nonlinear Systems with Input Unmodeled Dynamics,” *IEEE Trans. Automat. Contr.*, 41 (6), 913-920 (1996).

Lamport, L., “Concurrent Reading and Writing,” *Communications of the ACM* 20 (11), 806—811 (1977).

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).

Mårtenson, B., “The Order of any Stabilizing Regulator is Sufficient a priori Information for Adaptive Stabilization,” *Syst. & Contr. Lett.*, 6 (2), 87-91 (1985).

McLaughlin, M. L., S. G. Goldberg, N. B. Ellison, and J. Lucas, “Measuring Internet Audiences: Patrons of an Online Art Museum,” in S. Jones (Ed.), *Doing Internet Research* (pp. 163-178). Sage Publications, Inc. (1999).

McLaughlin, M. L., G. Sukhatme, J. Hespanha, C. Shahabi, A. Ortega, and G. Medioni, “The Haptic Museum,” *Proc. EVA 2000 Conference on Electronic Imaging and the Visual Arts*, Florence, Italy, March 2000 (2000a).

Miller, B. E., J. E. Colgate and R. A. Freeman, “Passive Implementation for a Class of Static Nonlinear Environments in Haptic Display,” *Proc. IEEE International Conference on Robotics and Automation*, Vol. 4. pp. 2937-42 (1999).

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).

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).

Murray, R. M., Z. Li, and S. S. Sastry, *A Mathematical Introduction to Robotic Manipulation*, CDC Press. Boca Raton, Florida (1994).

Roumeliotis, S. I., G. S. Sukhatme, and G. A. Bekey, "Fault Detection and Identification in a Mobile Robot using Multiple-Model Estimation,” *Proc. 1998 IEEE International Conference on Robotics and Automation*, Leuven, Belgium, pp. 2223-2228 (1998).

Roumeliotis, S.I., G. S. Sukhatme, and G. A. Bekey, "Circumventing Dynamic Modeling: Evaluation of the Error-State Kalman Filter applied to Mobile Robot Localization,” *Proc. 1999 IEEE International Conference on Robotics and Automation*, Detroit (1999a).

Roumeliotis, S. I., G. S. Sukhatme, and G. A. Bekey, "Smoother based 3-D Attitude Estimation for Mobile Robot Localization,” *Proc. 1999 IEEE International Conference on Robotics and Automation*, Detroit (1999b).

Sukhatme, G. S., J. S. Brizius, and G. A. Bekey, “Mobility Evaluation of a Wheeled Microrover using a Dynamic Model,” *Advanced Robotics*, 12 (5), 579-592 (1998).

Sukhatme, G. S., and M. J. Mataric, “Embedding Robots into the Internet,” *Communications of the ACM* May (2000).

Tsai, J. C., and J. E. Colgate, “Stability of Discrete Time Systems with Unilateral Nonlinearities,” *Proc. of the ASME International Mechanical Engineering Conference and Exposition* (1995).

Wilson, J. P, R. J. Kline-Schoder, M. A. Kenton, and N. Hogan, “Algorithms for Network-Based Force Feedback,” Salisbury, J. K. and Srinivasan, M. A. (Eds.), *Proc. PHANTOM Users Group Workshop*, MIT. http://www.sensable.com/community/PUG99/02_wilson.pdf (1999).