Haptography: Capturing and Recreating the Rich Feel of Real Surfaces

Katherine J. Kuchenbecker, Joseph Romano, and William McMahan

Abstract. Haptic interfaces, which allow a user to touch virtual and remote environments through a hand-held tool, have opened up exciting new possibilities for applications such as computer-aided design and robot-assisted surgery. Unfortunately, the haptic renderings produced by these systems seldom feel like authentic re-creations of the richly varied surfaces one encounters in the real world. We have thus envisioned the new approach of haptography, or haptic photography, in which an individual quickly records a physical interaction with a real surface and then recreates that experience for a user at a different time and/or place. This paper presents an overview of the goals and methods of haptography, emphasizing the importance of accurately capturing and recreating the high frequency accelerations that occur during tool-mediated interactions. In the capturing domain, we introduce a new texture modeling and synthesis method based on linear prediction applied to acceleration signals recorded from real tool interactions. For recreating, we show a new haptography handle prototype that enables the user of a Phantom Omni to feel fine surface features and textures.

1 Introduction

When you touch objects in your surroundings, you feel a rich array of haptic cues that reveal each object's geometry, material, and surface properties. For example, the vibrations and forces experienced by your hand as you stroke a piece of fabric or write on a sheet of corrugated cardboard are easily identifiable and distinct from those generated by gripping a foam ball or tapping on a hollow bronze sculpture. Humans excel at eliciting and interpreting haptic feedback during such interactions, naturally leveraging this wealth of information to guide their actions in the physical world (Klatzky and Lederman, 2003).

Motivated by the richness and usefulness of natural haptic feedback, we have envisioned the new approach of **haptography**. Like photography in the visual

---

Katherine J. Kuchenbecker · Joseph Romano · William McMahan
Haptics Group, GRASP Lab, University of Pennsylvania, Philadelphia, PA, USA
e-mail: `{kuchenbe,jrom,wmcmahan}@seas.upenn.edu`
website: `http://haptics.grasp.upenn.edu`

C. Pradalier, R. Siegwart, and G. Hirzinger (Eds.): Robotics Research, STAR 70, pp. 245–260.
springerlink.com

domain, haptography enables an individual to quickly record the feel of an interesting object and reproduce it at another time and/or place for someone to interact with as though it was real. The idea for haptography was first articulated by Kuchenbecker in 2008, and this paper provides an overview of its goals and methods. Haptographic technology involves highly sensorized handheld tools, haptic signal processing for model synthesis, and uniquely actuated haptic interfaces, all focused on capturing and recreating the rich feel of real surfaces.

Once these capabilities are available, a wide variety of practical applications will benefit from haptography. For example, it will provide a fast, simple way to store the current feel of a physical object (such as a unique marble statue or a dental patient's tooth), compare it with a database of other recordings, and analyze surface changes over time. Haptographs will also allow a wide range of people to touch realistic virtual copies of objects that are not directly accessible, such as archaeological artifacts and merchandise being sold online. Furthermore, haptography has the potential to significantly increase the realism of medical simulators and video games by incorporating object models built from quantitative contact data captured during real interactions. Beyond virtual environments, haptography can have a beneficial impact on teleoperation, where the operator uses a haptic interface to control the movement of a remote robot and wants to feel the objects being manipulated as though they were locally present. Finally, the haptographic focus on recording, analyzing, and recreating everything felt by the human hand will probably yield new insights on the sense of touch, which may help robotic hands achieve human-like dexterity and sensitivity in interactions with real physical objects.

Enabling the art and science of haptography requires us to answer two main questions: *How can we characterize and mathematically model the feel of real surfaces?* and *How can we best duplicate the feel of a real surface with a haptic interface?* Building on knowledge of the human haptic sensory system, haptography research uses measurement-based mathematical modeling to derive perceptually relevant haptic surface models and dynamically robust haptic display methods. The following sections of this paper explain the envisioned system paradigm, our initial work on capturing the feel of surfaces, and our continuing work on recreating such surfaces realistically.

2 Overview of Haptography

Despite its ubiquitous importance in human life, we currently lack a formal method for analyzing and understanding the feel of touch-based interaction with physical objects. Furthermore, fundamental surface modeling and device design choices prevent the vast majority of existing haptic interfaces from compellingly duplicating the feel of real objects.

Target Interactions. Direct-touch haptography would enable an individual to capture natural interactions between their fingertip and an interesting surface and then recreate that exact feel with a programmable tactile interface that can be freely explored. While fascinating and useful, there are currently many technological

Table 1 The user experience of haptography can be understood via an analogy to photography.

Photography	Haptography
digital SLR camera	highly sensorized handheld tool
interchangeable lenses	interchangeable tool tips
framing a shot and focusing the camera	exploring an object's surface
planar distribution of light intensities	stream of positions, forces, and accelerations
optics and the human eye	haptics and the human hand
LCD monitor	uniquely actuated handheld tool
viewing the digital image	freely exploring the digital model
spatial resolution and focus	high frequency accelerations

challenges that preclude the realization of such an ambitious objective; it will take many years for today's most promising noninvasive tactile sensors, e.g., (Sun et al, 2007), and high resolution tactile displays, e.g., (Koo et al, 2008), to mature to the level needed for such an approach. Thus, we focus our research on **tool-mediated contact**, where the user touches the target surface through an intermediate tool such as a metal probe, a ball-point pen, or a surgical scalpel.

Restricting haptography to tool-mediated interactions is not as limiting as it might initially seem. First, many everyday activities are conducted with a tool in hand, rather than with bare fingertips; tools extend the hand's capabilities for a specific task and protect it from damage. Second, humans are surprisingly good at discerning haptic surface properties such as stiffness and texture through an intermediate tool (Klatzky and Lederman, 2008; Yoshioka and Zhou, 2009). This acuity stems partly from the human capability for distal attribution, in which a simple hand-held tool comes to feel like an extension of one's own body (Loomis, 1992).

Haptographic Process. Haptography intentionally parallels modern photography, but the interactive nature of touch sets the two apart in several ways. To help clarify the differences, Table 1 lists analogous elements for the two domains, and Fig.1 depicts an overview of the haptic capturing and recreating processes.

A haptographer begins by identifying an object with a unique or important feel; a museum curator might select an interesting historical relic, and a doctor might target

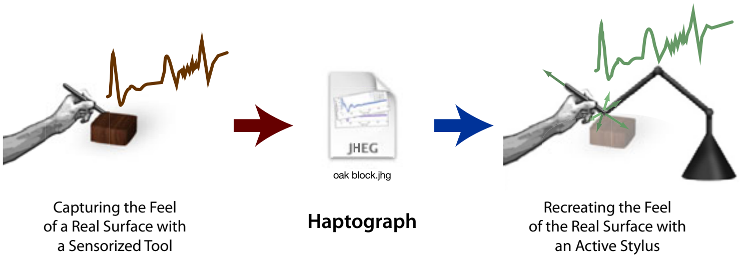

Fig. 1 The envisioned approach of haptography will enable individuals to quickly capture, analyze, and recreate the exquisite feel of any surface they encounter in the real world.

an in vivo sample of tissue or bone. Working in his or her standard surroundings, the haptographer attaches a chosen tool tip to a highly sensorized hand-held instrument. For real-time haptography in teleoperation, the slave robot's tool is instrumented in the same way as a hand-held haptography tool. The operator then uses the tool to explore the object's surface via natural motions. The system collects multiple data streams throughout this interaction—all of the touch-based sensations that can be felt by a human hand holding a tool—including quantities such as the translation and rotation of the stylus and the object, the forces and torques applied to the object's surface, and the three-dimensional high frequency accelerations of the tool tip.

In teleoperation, the haptic interface held by the user seeks to recreate the measured sensations as they occur, and the challenge lies in perfecting this connection. In non-real-time applications, the haptographic processor needs to distill the recorded data into a general surface model so that future users can explore the virtual object in a natural way. Because the global shape of an object can be captured efficiently with optical sensors or reconstructed via computer-aided design (CAD) tools, haptography focuses on capturing surface attributes that are not readily apparent by sight, such as friction, texture, stiffness, and stickiness. Section 3 describes this identification problem in more detail, along with our preliminary work on texture modeling. We plan to store haptographs of different surface-tool interactions in a public online database so that virtual environment designers can apply haptographic surface swatches to chosen areas of synthetic objects.

An acquired haptographic model can be explored via any kinesthetic haptic interface, but the flexibility of the interaction and the quality of the haptic response will greatly depend on the mechanical, electrical, and computational design of the chosen system. Commercially available haptic devices generally struggle to duplicate the full feel of real surfaces. Thus, the second major aim of haptography research is to discover and refine high fidelity methods for rendering haptographs. As described in Section 4, tool-mediated haptography centers on the use of a dedicated high frequency vibration actuator, and we have tested this approach through creation of a prototype system. We want any individual to be able to use this "haptography handle" to explore 3D virtual surfaces and feel rich, natural sensations that are indistinguishable from those one would experience when touching the original item.

The Key to Haptographic Realism. Researchers studying virtual and remote environments have long sought to replicate the feel of real objects with a haptic interface. Arguably, the most important advance toward this goal came in 1994 when Massie and Salisbury presented the Phantom haptic interface. The design of this device evolved from three essential criteria, namely that "free space must feel free," "solid virtual objects must feel stiff," and "virtual constraints must not be easily saturated" (Massie and Salisbury, 1994, p. 296). Prioritization of these goals and clever mechanical design yielded a lightweight, easily backdrivable, three-degree-of-freedom robot arm actuated via base-mounted brushed DC motors equipped with high resolution optical encoders and smooth capstan cable drives. This invention inspired a wave of similar impedance-type haptic interfaces, many of which are now widely available as commercial products. Such systems are typically programmed to

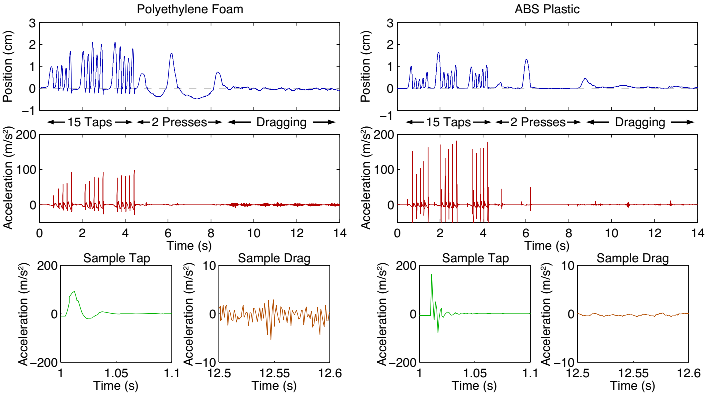

Fig. 2 Sample data from interactions with two real materials through a stylus instrumented with an accelerometer. One can quickly observe that the plastic is stiffer and smoother than the foam.

generate interaction forces via one-sided linear springs: when the measured device tip position intersects a region occupied by a virtual or remote object, the device outputs a restoring force that is proportional to the penetration vector.

While haptic interfaces designed and programmed in this way do succeed at conveying the global shape of virtual and remote items, the surfaces of these objects typically have only a weak haptic resemblance to real objects. Instead, haptically rendered shapes tend to feel soft and undefined, strangely slippery or peculiarly active and vibratory, and largely devoid of the coherent dynamic cues that one associates with natural surface properties. In one study targeted at this problem, human subjects used a Phantom to blindly tap on the stiffest possible spring-based virtual surface, among other real and virtual samples (Kuchenbecker et al, 2006). The spring-based surface received a realism rating of two on a scale from one to seven, where a seven denotes the feel of a real wooden block. Clearly something important is missing from these traditionally rendered haptic surfaces: we believe this deficiency stems from a reliance on haptic object models and interface hardware that prioritize low-frequency behavior over the **naturalistic high frequency accelerations** that give real objects their distinctive feel.

Human haptic capabilities are inherently asymmetric, allowing motion at just 8 to 10 Hz (Loomis and Lederman, 1986) and vibration perception up to 1000 Hz (Bell et al, 1994). As illustrated in Fig. 2, tool-mediated interactions with hard and textured objects create vibrations that strongly excite the Pacinian corpuscle mechanoreceptors in the glabrous skin of the human hand (Bell et al, 1994). It is clear that high frequency accelerations are a rich source of feedback during tool use, encoding information about surface material, surface texture, tool design, downward force, and tool velocity; a user naturally expects a haptic virtual environment to provide these same cues, but they are generally absent. When appropriate acceleration transients were added to the spring-based virtual surfaces in (Kuchenbecker et al, 2006), subjects responded with realism ratings of five out of seven, a significant improvement. Haptography is thus founded on the belief that only a haptic interface that authentically recreates these salient accelerations will be able to fool a human into believing that a virtual object is real, or a remote object is present.

3 Capturing the Feel of Real Surfaces

The first aim of haptography is to enable an individual to quickly and easily capture the feel of a real surface through a hand-held or teleoperated tool. This process yields a stream of interaction data that is distilled into a mathematical model of the surface's salient haptic properties for further analysis and subsequent re-creation.

Prior Work. Haptography takes a nontraditional approach to modeling object surfaces. The most carefully controlled attribute of a typical haptic virtual object is its global shape (Salisbury et al, 1995, 2004). As in computer graphics, these geometric models are usually composed of a polygonal mesh that defines the surface of the object. Contact with this mesh is then rendered as a spring force that seeks to push the user's virtual tool perpendicularly out of the surface. The high computational load incurred by fine meshes is avoided by blending the orientation of larger adjacent facets so that the normal force varies smoothly across the surface. The behavior of the coupling impedance can be modulated somewhat to change the feel of the surface, although nonidealities (e.g., position sensor quantization) cause instability at high stiffness and damping values. The additional surface properties of texture and friction are included in such models as a parametric relationship between the virtual tool's motion (position and velocity) and an additional force that is added to the spring response. For example, (Salisbury et al, 1995) use a sum of cosines for synthetic texture, (Minsky and Lederman, 1996) simulate roughness with a variety of lateral-force look-up tables based on surface location, and (Basdogan et al, 1997) create height-field textures inspired by the "bump map" approach from computer graphics. Numerous other hand-tuned surface representations have been developed, but most struggle to capture the rich variety of sensations caused by contact with real objects because they are not based on physical principles.

Rather than relying on hand-tuned parametric relationships, haptography derives virtual object models from real-world data. The idea of using a stream of physical measurements to create a haptic virtual environment is not new, but the central involvement of the human haptographer and the focus on high frequency accelerations are significant departures from previous work. The first discussion of a measurement-based modeling approach occurs in MacLean's 1996 paper on the "Haptic Camera," a fully automated one-degree-of-freedom probe that interacts with a mechanical system while recording position and force data to enable automatic fitting of a piecewise linear dynamic model. Autonomous interaction and identification techniques have since been applied to several other simple mechanical systems, such as switches and buttons (Colton and Hollerbach, 2005), and also to whole object contact through ACME, the robot-based Active Measurement Facility (Pai et al, 2000). In contrast to a robot, a human haptographer holding an instrumented tool can quickly and safely explore the surfaces of almost any physical object with natural motions that are fine-tuned in real time. However, there have been only a few previous efforts to generate haptic surface models from data recorded with a hand-held tool, typically involving either simple parametric relationships or direct playback (Okamura et al, 2008). For example, (Okamura et al, 2001) fit a decaying sinusoid model to acceleration transients recorded from a series of taps, while (Kuchenbecker et al, 2006) explicitly stored such recordings. Others have created hand-held tools fitted with sensors, e.g., (Pai and Rizun, 2003), but little has been done to distill the resulting data into surface models.

Our Approach. Given the limitations of traditional models and the success of several previous data-driven studies, we believe a sensorized hand-held tool and a sophisticated signal processing algorithm can be used to create very accurate models of the tool-mediated feel of any real surface. Haptographic probes are designed to record high bandwidth haptic data (tool position, velocity, acceleration, force, etc.) during natural human exploration of a surface. The recorded signals must then be segmented by interaction state (e.g., no contact, stationary contact, and sliding contact) and analyzed to yield mathematical models for the characteristic feel of each state, as well as for the transitions between states. The high sensory bandwidth of the human hand makes us believe that the realism of a haptographic surface model will strongly depend on its ability to encapsulate the high frequency accelerations of tool–surface contact. A full suite of haptographic capturing tools and algorithms will require a significant body of research, such as the physics-based modeling of tapping in (Fiene and Kuchenbecker, 2007); here, we present a new, general method for modeling the response of real textured surfaces felt through a tool.

Texture Modeling. We have developed a new method for using recorded data to obtain a predictive model of the tool accelerations produced during real texture exploration. As one can determine through quick experimentation, dragging a certain hand-held tool across a given surface creates vibrations that vary with both normal force and scanning velocity, as well as contact angle, hand configuration, and grip force. We are beginning our characterization of this complex dynamic system by analyzing the vertical acceleration experienced by a hand-held stylus as it is dragged across a variety of surfaces under different conditions. Our data set was recorded using a custom designed data collection apparatus from (Yoshioka, 2009) in well-controlled human subject trials where mean contact force, scanning velocity, and the other relevant variables were all held constant. The data collection system allows for precision recording of probe–texture interaction data including all contact forces and torques, three-dimensional tool acceleration, tool velocity, and the subject's grip force at a rate of 5000 Hz. For each recorded trial, we start by seeking a model that can generate an optimal prediction of the next real value in the acceleration time series given the previous n data points. We have found that this problem is best addressed with forward linear prediction, a common technique from system identification.

Forward Linear Prediction. The speech synthesis community has known for over thirty years that the complex dynamic vibrations created by the human vocal tract can be modeled by a form of the Wiener filter, the forward linear predictor (Atal and Hanauer, 1971). The standard procedure in speech synthesis is to treat the vocal tract response as an unknown filter that shapes a white noise excitation signal, which comes from air passed through the system by the glottis. The output is the spoken sound wave, which can be recorded with a microphone. Similarly, we record contact vibrations with an accelerometer and treat the dynamic response of the tool–texture interaction as a filter we wish to identify.

Fig. 3(a) shows the block diagram used for this system identification process, and the following mathematical analysis follows the conventions of (Benesty et al, 2008) as closely as possible. Our input signal $\mathbf{a}(k)$ is the original recorded time series of accelerations. The filter's output vector is defined as $\hat{\mathbf{a}}(k)$, which represents the forward linear prediction. $H(z)$ is assumed to be an IIR filter of length n of the form $H(z) = [-h_1 z^{-1} - h_2 z^{-2} \ldots -h_n z^{-n}]$. The residual of these two signals is the error vector $\mathbf{e}(k)$, and the transfer function $P(z)$ is:

$$\frac{E(z)}{A(z)} = 1 - H(z) = P(z) \tag{1}$$

We define the vector of filter coefficients as $\mathbf{h} = [h_1\ h_2\ h_3\ \ldots\ h_n]^T$, and the n-length time history of our input signal as $\mathbf{a}(k-1) = [a(k-1)\ a(k-2)\ \ldots\ a(k-n)]$. We then write the residual at each step in time with the following difference equation:

$$e(k) = a(k) - \hat{a}(k) = a(k) - \mathbf{h}^T \mathbf{a}(k-1) \tag{2}$$

Optimal filter values of $\mathbf{h}$ can be found by defining a suitable cost function. We use the standard choice of mean-square error, $J(\mathbf{h}) = \mathrm{E}\{\mathbf{e}^2(k)\}$, where $\mathrm{E}\{\cdot\}$ denotes mathematical expectation, as defined by (Benesty et al, 2008). When the gradient of $J(\mathbf{h})$ is flat, $\mathbf{h}$ is at an optimal value, $\mathbf{h}_o$. By algebraic manipulation we can derive the following result for the gradient:

$$\frac{\partial J(\mathbf{h})}{\partial \mathbf{h}} = -2\mathrm{E}\{(e(k)\mathbf{a}(k-1))\} \tag{3}$$

When the gradient is flat at $\mathbf{h}_o$, the error is at a minimum $\mathbf{e}_o(k)$, and we can simplify the problem to:

$$\mathrm{E}\{e_o(k)\mathbf{a}(k-1)\} = \mathbf{0}_{nx1} \tag{4}$$

By substituting values for the cross-correlation matrix $(\mathbf{R} = \mathbf{a}(k-1)\mathbf{a}^T(k-1))$ and the cross-correlation vector $(\mathbf{p} = \mathbf{a}(k-1)a(k))$ into (4), we arrive at the Wiener-Hopf equation:

$$\mathbf{R}\,\mathbf{h}_o = \mathbf{p} \tag{5}$$

Assuming non-singular $\mathbf{R}$, the optimal forward predictor coefficients can be found by simply inverting the cross-correlation matrix, such that $\mathbf{h}_o = \mathbf{R}^{-1}\mathbf{p}$. Alternatively, we can use a more efficient recursive method, such as the Levinson-Durbin algorithm (Durbin, 1960), to solve for $\mathbf{h}_o$ from (5). For demonstration, Fig. 4 shows a sample plot of $\mathbf{a}(k)$, $\hat{\mathbf{a}}(k)$, and $\mathbf{e}(k)$ for the optimal filter $H(z)$ of order $n = 120$.

Signal Generation. The previous section details a process for finding the linear transfer function $H(z)$ that is best able to predict the acceleration response of a

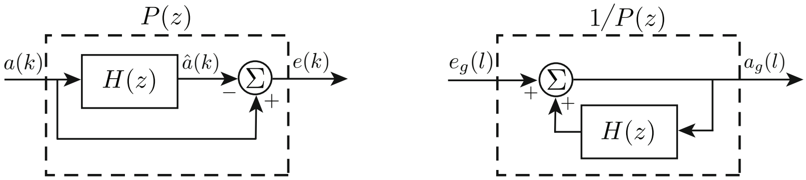

(a) LPC forward prediction model.

(b) Acceleration synthesis model.

Fig. 3 Block diagrams for prediction of the next contact acceleration $\hat{a}(k)$ given the recorded series $a(k)$ and synthesis of an acceleration signal $a_g(l)$ from the white noise input $e_g(l)$.

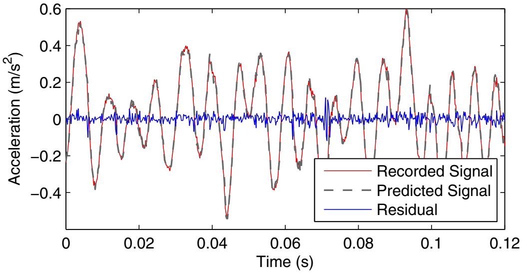

Fig. 4 Forward prediction signal generation and residual error. The data shown are from a real sample of organza fabric mounted on a stiff substrate and touched with a plastic probe at a velocity of 4.0 cm/s and a downward force of 1.5 N. The linear prediction filter $H(z)$ includes 120 coefficients ($n = 120$).

texture based on its previous n acceleration values. Subtracting the predicted response from the recorded signal removes almost all its spectral components, leaving only the noise signal $\mathbf{e}(k)$, which is ideally white and Gaussian. This section describes how to reverse this process and obtain a completely new (but spectrally similar) acceleration signal based on a white noise input, given an identified filter $H(z)$.

As seen in Fig. 3(b), the input signal $\mathbf{e}_g(l)$ is a white noise vector that we generate in real time. The output vector is $\mathbf{a}_g(l)$, a synthesized acceleration signal with spectral properties that are very close to those of the real data signal $\mathbf{a}(k)$ for which the filter $H(z)$ is tuned; higher order filters generally result in a better spectral match. By rewriting (1), we can formulate this new transfer function as follows:

$$\frac{A_g(z)}{E_g(z)} = \frac{1}{1 - H(z)} = \frac{1}{P(z)} \tag{6}$$

We now observe that the difference equation for the synthesized acceleration is:

$$\mathbf{a}_g(l) = \mathbf{e}_g(l) + \mathbf{h}^T \mathbf{a}_g(l-1) \tag{7}$$

During texture synthesis, we generate white noise with a Gaussian distribution of amplitudes and apply it to (7). One should note that the signal power of the white noise input is important for creating an acceleration signal $\mathbf{a}_g(l)$ with the proper magnitude. The power of the generated noise signal $P\{\mathbf{e}_g(l)\}$ must be equivalent to

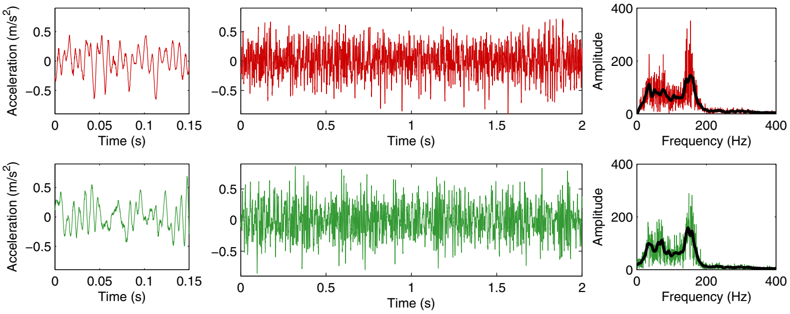

Fig. 5 Time- and frequency-domain views of a recorded acceleration and a signal synthesized to emulate that interaction using our novel texture-modeling techniques. The real setup and the synthesis filter are the same as those used in Fig. 4.

that of the power remaining in the residual signal, $P\{\mathbf{e}(k)\}$, after filter optimization. We have achieved good results when applying this acceleration modeling technique to data from many surfaces and at many levels of downward force and translational velocity. Fig. 5 shows one such sample in both the time and frequency domains.

Future Work. We are encouraged by the expressiveness and versatility of linear prediction for synthesizing realistic texture acceleration waveforms, and we are in the process of investigating many additional aspects of this approach. For example, how many filter coefficients are required to capture the haptically salient properties of an individual texture trial? And how should one synthesize accelerations for values of downward force and scanning velocity that were not explicitly tested? Currently, we fit data sparsely sampled from this parameter space and then use two-dimensional linear interpolation to choose coefficients of $H(z)$ for unique combinations of these parameters. In the future we intend to look into interpolating between filter poles or cepstral coefficients, both of which are directly related to the filter coefficients $\mathbf{h}$. More generally, we need to develop methods for processing data from interactions that are less controlled and for making models that go beyond texture to include other salient surface properties. In addition to increasing our knowledge of tool–surface contact, we hope that these haptographic modeling methods can be used to provide sensations that closely mirror those of real interactions, and also to evaluate the fidelity of virtually rendered haptic surfaces.

4 Recreating the Feel of Real Surfaces

The second aim of haptography is to enable an individual to freely explore a virtual or remote surface via a haptic interface without being able to distinguish its feel from that of the real object being portrayed. Realistically recreating haptographic models requires haptic device hardware and control algorithms that excel at delivering high frequency tool accelerations without impeding free-space hand motion.

Prior Work. During contact with a virtual or remote object, traditional haptic systems employ the device's actuators (usually base-mounted DC motors) to apply a force to the user's hand based on tool tip penetration, which is inherently a slowly changing signal. The mechanical elements that connect the motor to the hand would ideally be massless, frictionless, and infinitely stiff, but no real device can meet these requirements; instead, the dynamics of the intervening linkages, joints, and cables distort the output of the motors, which especially interferes with the display of any high frequency vibrations (Campion and Hayward, 2005; Kuchenbecker et al, 2006). Still, some previous work has shown that vibrations displayed with such actuators can improve the perceived realism of spring-based virtual surfaces, but these approaches require either extensive human-subject tuning (Okamura et al, 2001) or exhaustive device characterization (Kuchenbecker et al, 2006). Furthermore, high frequency base motor actuation is susceptible to configuration-based and user-based changes in the system's high-frequency dynamics, so it cannot achieve the consistent, high fidelity feel needed for haptography.

A viable alternative can be found in (Kontarinis and Howe, 1995)'s approach to teleoperation, where high frequency slave tool accelerations were overlaid on standard low frequency force feedback via an inverted speaker mounted near the user's fingertips. The slave acceleration was amplified by an empirically determined gain to drive the actuator, and the system's acceleration output was reported to vary by a factor of 2.24 across the frequency range of interest. Human subject tests indicated that this simple dual-actuator feedback strategy increased user performance in several tasks and also improved the "feel" of the interface, one of the main goals of haptography. Since this encouraging early work, several groups have created interesting active styli meant to be used without a force-feedback device, e.g., (Yao et al, 2005). The only project closer to our interests is that of (Wall and Harwin, 2001), who made a vibrotactile display stylus to study the effect of device output bandwidth on virtual grating perception. Their system uses a voice-coil actuator between the stylus and the end-effector of a desktop haptic device, with a controller that seeks to regulate actuator displacement using high-resolution measurements from a parallel LVDT sensor. The associated human-subject study found that the additional actuator between the hand and the haptic device significantly improved the rendering of virtual gratings but also reduced the system's ability to render stiff springs.

Our Approach. Considering the limitations of base-mounted motors and the results others have achieved with auxiliary actuators, we believe that haptographic models can be excellently recreated by attaching a high bandwidth bidirectional linear actuator to the handle of a typical haptic interface. This "haptography handle" should be designed and controlled to enable the system to significantly accelerate the user's hand at high vibrotactile frequencies (20–1000 Hz) in real time, while it is being held and moved around by the user. Imposing a grounded force at the handle is very challenging, so instead we attach an additional mass to the handle through a spring and a sliding joint. The auxiliary actuator applies equal and opposite forces on the handle and this mass, thereby pushing and pulling them relative to one another. Such a system can be carefully controlled only by understanding its mechanical dynamics and their impact on the user's experience. One final benefit to this approach is that we believe it will require only one linear actuator (rather

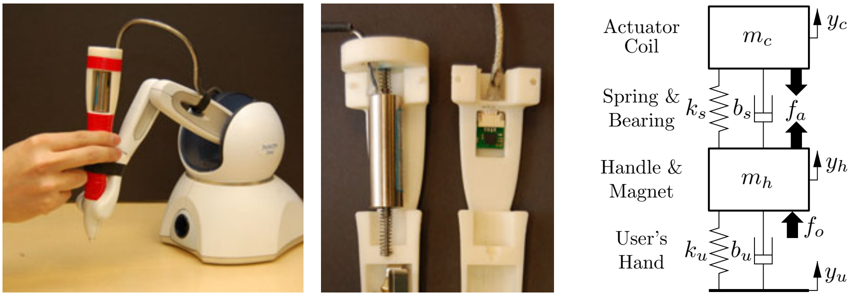

(a) Haptography handle in use (b) Handle interior (c) Dynamic model

Fig. 6 A prototype haptography handle for use with the SensAble Phantom Omni. The voice coil actuator applies a high frequency force f_a between the coil and the magnet to accelerate the handle.

than three) because the human hand is not particularly sensitive to the direction of high-frequency vibrations (Bell et al, 1994).

Sample Implementation: Haptography Handle for the Phantom Omni. To evaluate the merits of our approach, we developed the prototype shown in Fig. 6 to act as an interchangeable handle for the Phantom Omni, a widely available impedance-type haptic device from SensAble Technologies, Inc.

Prototype. At the heart of our design is an NCM02-05-005-4JB linear voice coil actuator from H2W Technologies. We have installed this actuator in a moving coil configuration, where the permanent magnet core is rigidly attached to a handle and the coil is free to slide relative to this core along internal jeweled bearings. Additionally, we place compression springs at both ends of the coil to center it within the actuator's travel limits. The actuator is driven with a high bandwidth linear current amplifier, and it can output a peak force of 6.6 N. For more details on this actuator and our experimental procedures, please consult (McMahan and Kuchenbecker, 2009b), which describes an earlier prototype. Mounting this haptography handle to an Omni allows for measurement of the position and velocity of the handle, as well as the exertion of low-frequency forces, via the Omni's base-mounted encoders and DC motors. The addition of a dedicated voice coil actuator gives this low cost haptic device the capability of providing the high frequency contact accelerations that are essential to haptography.

System Dynamics. In order to accurately control the handle accelerations felt by the user, we must characterize the dynamics of our system. We use the lumped-parameter model shown in Fig 6 to represent our system: m_c is the mass of the actuator coil, k_s is the combined stiffness of the centering springs, b_s represents viscous friction in the linear bearings, f_a is the electromagnetic force exerted by the actuator, m_h is the effective mass of the handle and magnet, and f_o represents the forces provided by the Omni. The user is modeled as a parallel spring and damper

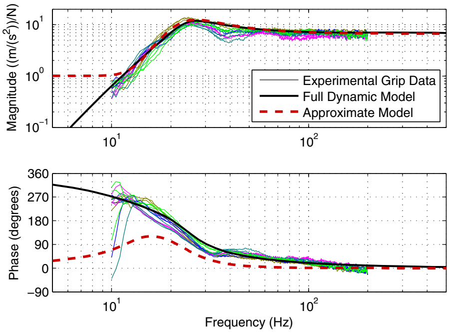

Full Dynamic Model

Parameter	Value
m_c	0.018 kg
k_s	399 N/m
b_s	1.2 N/(m/s)
m_h	0.145 kg
k_u	1000 N/m
b_u	20 N/(m/s)

Fig. 7 Frequency-domain system identification validates the structure of our dynamic model and enables the selection of appropriate values for parameters that cannot be directly measured.

(k_u and b_u) that connect the handle mass to the hand's set-point position, y_u. We can then derive the transfer function from actuator force to handle acceleration:

$$\frac{A_h(s)}{F_a(s)} = \frac{m_c s^4}{(m_c s^2 + b_s s + k_s)(m_h s^2 + (b_s + b_u)s + (k_s + k_u)) - (b_s s + k_s)^2} \tag{8}$$

Note that the Omni force f_o and the hand set-point y_u are both low frequency and thus will not affect the high frequency accelerations felt by the user.

We empirically validate and tune this transfer function by sending a repeating 10–200 Hz swept sinusoid force command to the linear voice coil actuator and recording the resulting accelerations at the handle with an accelerometer. We performed three trials of this test with five different users, each lightly holding the stylus with their right hand in a three-fingered pinch grip. Frequency-domain analysis of these tests guides the selection of model parameters. Fig. 7 shows the empirical transfer function estimates from the grip experiments, as well as the parameters chosen for the full dynamic model and its frequency-domain response.

This model enables us to design a *dynamically compensated controller* targeted at good acceleration tracking; our present controller consists of a feedforward term that inverts our estimate of the transfer function $A_h(s)/F_a(s)$ in order to determine the proper actuator force needed to achieve a desired handle acceleration. A careful look at (8) shows that naively inverting this transfer function will result the placement of four poles at the origin, which corresponds with a quadruple integrator in the controller. A controller with a quadruple integrator has infinite gain at steady-state and very high gain at low frequencies. These large gains pose a problem because they will immediately saturate the maximum force and deflection capabilities of our linear actuator. As a result, we approximate this transfer function with one that has finite DC gain, but still manages to capture the magnitude response of the full dynamic model in the important frequency range of 20-1000 Hz. The frequency-domain response of this approximate model is also shown in Fig. 7.

Teleoperation Testing. We tested our handle's performance at recreating realistic contact accelerations by conducting master-slave teleoperation experiments; the

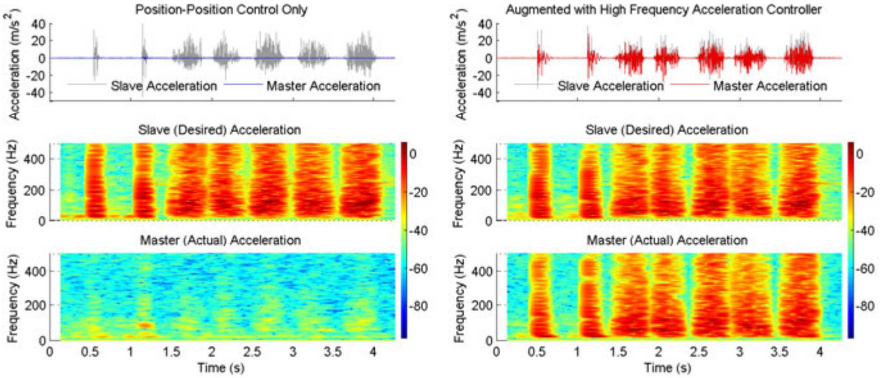

Fig. 8 Time- and frequency-domain results for the teleoperation experiments.

operator (grasping the haptography handle) uses a master Omni to command a slave Omni to perform exploratory tapping and dragging motions on a remote piece of unfinished plywood through a position-position controller. This configuration allows us to obtain real contact accelerations from an accelerometer mounted to the slave Omni's end effector and to render these accelerations to the user in real-time via the haptography handle. This experiment also serves as a proof-of-concept demonstration for haptography's potential use in teleoperation applications.

Attempting to recreate only the high frequency accelerations measured along the longitudinal axis of the slave's tool tip, we ran the experiment twice: once without using the voice coil actuator and once driving it with our dynamically compensated controller. In both cases, the operator tapped twice on the surface and then laterally dragged the tool tip five times. Fig. 8 shows time domain plots of the slave (desired) and master (actual) accelerations recorded during these experiments, as well as spectrograms of these signals. Visual inspection of these plots shows that the Omni's native motors and the implemented position-position controller do a poor job of transmitting high frequency accelerations to the user. However, augmenting the system with our dedicated vibration actuator and dynamically compensated controller provides a substantial improvement. Without this actuation, the normalized RMS error between actual and desired acceleration spectrograms is 100%, while auxiliary actuation brings this strict error metric down to 48%. Still, there is room for further refinement of the controller, as one can observe a general trend of underactuation and also some phase lag at lower frequencies.

Hands-On Demonstration. To obtain qualitative feedback about the feel of this system, we demonstrated the haptography handle in bilateral teleoperation at the 2009 IEEE World Haptics Conference (McMahan and Kuchenbecker, 2009a). Conference attendees were invited to use the master–slave Omni system to remotely explore textured samples both with and without acceleration feedback from the dedicated actuator. The demonstration was well received and participants provided a great deal of positive feedback, especially that the accelerations allowed them to feel small details and surface textures that were not detectable with only the position-position controller. Several participants thus commented that their hand felt "numb" when they explored the samples without haptographic feedback. The contact accelerations were also noted to make the surfaces feel "harder" even though the normal force provided by the Omni remained constant. This demonstration was honored to be selected by a panel of experts for the conference's Best Demonstration award.

Future Work. As we continue this research, we hope to improve the fidelity of our haptographic rendering by investigating more sophisticated acceleration controllers. We are also working to determine the perceptually correct mapping of three-dimensional accelerations to a one-dimensional actuator. Lastly, we are preparing to run human subject experiments to study the perceptual requirements for discrimination of realistic contact accelerations, as well as the potential benefits the approach of haptography may have on common applications for haptic interfaces.

Acknowledgements. We thank Rebecca Pierce for her initial work on haptographic data capture. We thank Takashi Yoshioka for providing texture contact data and for many useful conversations on haptography. This work was supported by the University of Pennsylvania and the National Science Foundation (grant IIS-0845670).

References

Atal, B.S., Hanauer, S.L.: Speech analysis and synthesis by linear prediction of the speech wave. Journal of the Acoustic Society of America 50(2), 637–655 (1971)

Basdogan, C., Ho, C.H., Srinivasan, M.A.: A ray-based haptic rendering technique for displaying shape and texture of 3D objects in virtual environments. Proc. ASME Dynamic Systems and Control Division, DSC 61, 77–84 (1997)

Bell, J., Bolanowski, S., Holmes, M.H.: The structure and function of Pacinian corpuscles: A review. Progress in Neurobiology 42(1), 79–128 (1994)

Benesty, J., Sondhi, M.M., Huang, Y. (eds.): Springer Handbook of Speech Processing. Springer, Berlin (2008)

Campion, G., Hayward, V.: Fundamental limits in the rendering of virtual haptic textures. In: Proc. IEEE World Haptics Conference, pp. 263–270 (2005)

Colton, M.B., Hollerbach, J.M.: Identification of nonlinear passive devices for haptic simulations. In: Proc. IEEE World Haptics Conference, pp. 363–368 (2005)

Durbin, J.: The fitting of time series models. Revue de l'Institut International de Statistique / Review of the International Statistical Institute 28(3), 233–244 (1960)

Fiene, J.P., Kuchenbecker, K.J.: Shaping event-based haptic transients via an improved understanding of real contact dynamics. In: Proc. IEEE World Haptics Conference, pp. 170–175 (2007)

Klatzky, R.L., Lederman, S.J.: Touch. In: Healy, A.F., Proctor, R.W. (eds.) Handbook of Psychology. Experimental Psychology, ch. 6, vol. 4, pp. 147–176. John Wiley and Sons, Chichester (2003)

Klatzky, R.L., Lederman, S.J.: Perceiving object properties through a rigid link. In: Lin, M., Otaduy, M.(eds.) Haptic Rendering: Algorithms and Applications, vol. 1, pp. 7–19 (2008)

Kontarinis, D.A., Howe, R.D.: Tactile display of vibratory information in teleoperation and virtual environments. Presence 4(4), 387–402 (1995)

Koo, I.M., Jung, K., Koo, J.C., Nam, J.D., Lee, Y.K., Choi, H.R.: Development of soft-actuator-based wearable tactile display. IEEE Transactions on Robotics 24(3), 549–558 (2008)

Kuchenbecker, K.J.: Haptography: Capturing the feel of real objects to enable authentic haptic rendering (invited paper). In: Proc. Haptic in Ambient Systems (HAS) Workshop, in conjunction with the First International Conference on Ambient Media and Systems (2008)

Kuchenbecker, K.J., Fiene, J.P., Niemeyer, G.: Improving contact realism through event-based haptic feedback. IEEE Transactions on Visualization and Computer Graphics 12(2), 219–230 (2006)

Loomis, J.M.: Distal attribution and presence. Presence: Teleoperators and Virtual Environments 1(1), 113–119 (1992)

Loomis, J.M., Lederman, S.J.: Tactual perception. In: Boff, K.R., Kaufman, L., Thomas, J.P. (eds.) Handbook of Perception and Human Performance. Cognitive Processes and Performance, ch. 31, vol. II, pp. 1–41. John Wiley and Sons, Chichester (1986)

MacLean, K.: The 'Haptic Camera': A technique for characterizing and playing back haptic properties of real environments. Proc. ASME Dynamic Systems and Control Division, DSC 58, 459–467 (1996)

Massie, T.H., Salisbury, J.K.: The PHANToM haptic interface: A device for probing virtual objects. Proc. ASME Dynamic Systems and Control Division, DSC 55(1), 295–301 (1994)

McMahan, W., Kuchenbecker, K.J.: Displaying realistic contact accelerations via a dedicated vibration actuator. In: Proc. IEEE World Haptics Conference, pp. 613–614 (2009a)

McMahan, W., Kuchenbecker, K.J.: Haptic display of realistic tool contact via dynamically compensated control of a dedicated actuator. In: Proc. IEEE/RSJ International Conference on Intelligent RObots and Systems (2009b)

Minsky, M., Lederman, S.J.: Simulated haptic textures: Roughness. Proc. ASME Dynamics Systems and Control Division, DSC 58, 421–426 (1996)

Okamura, A.M., Cutkosky, M.R., Dennerlein, J.T.: Reality-based models for vibration feedback in virtual environments. IEEE/ASME Transactions on Mechatronics 6(3), 245–252 (2001)

Okamura, A.M., Kuchenbecker, K.J., Mahvash, M.: Measurement-based modeling for haptic rendering. In: Lin, M., Otaduy, M.(eds.) Haptic Rendering: Algorithms and Applications, ch. 21, pp. 443–467 (2008)

Pai, D.K., Rizun, P.: The WHaT: a wireless haptic texture sensor. In: Proc. IEEE Haptics Symposium, pp. 3–9 (2003)

Pai, D.K., Lang, J., Lloyd, J., Woodham, R.J.: ACME, a telerobotic active measurement facility. In:Experimental Robotics VI, Proc. Int. Symposium on Experimental Robotics. LNCS, vol. 250, pp. 391–400. Springer, Heidelberg (1999)

Salisbury, K., Conti, F., Barbagli, F.: Haptic rendering: Introductory concepts. IEEE Computer Graphics and Applications 24(2), 24–32 (2004)

Salisbury, Z., Zilles, C.B., Salisbury, J.K.: A constraint-based god-object method for haptic display. Proc. IEEE/RSJ International Conference on Intelligent RObots and Systems 3, 146–151 (1995)

Sun, Y., Hollerbach, J.M., Mascaro, S.A.: EigenNail for finger force direction recognition. In: Proc. IEEE International Conference on Robotics and Automation, pp. 497–502 (2007)

Wall, S.A., Harwin, W.: A high bandwidth interface for haptic human computer interaction. Mechatronics 11(4), 371–387 (2001)

Yao, H.Y., Hayward, V., Ellis, R.E.: A tactile enhancement instrument for minimally invasive surgery. Computer-Aided Surgery 10(4), 233–239 (2005)

Yoshioka, T.: Probe–texture interaction dataset. Personal communication (2009)

Yoshioka, T., Zhou, J.: Factors involved in tactile texture perception through probes. Advanced Robotics 23, 747–766 (2009)