

Tactile Display of Vibratory Information in Teleoperation and Virtual Environments

Dimitrios A. Kontarinis and Robert D. Howe

Division of Applied Sciences

Pierce Hall

Harvard University

Cambridge, MA 02138

Tel. (617) 496-8359, Fax (617) 495-9837

Email: howe@das.harvard.edu

ABSTRACT

This paper investigates the use of tactile displays for conveying task-related vibrations in teleoperation and virtual environments. Vibration displays can be implemented with inexpensive, open loop devices that can be added to many existing systems to improve performance. We describe the design of our prototype vibration sensing and display system, and experimentally demonstrate the utility of this type of tactile feedback. We also delineate the kinds of tasks where high frequency vibratory feedback is important. In inspection and exploration tasks the detection of vibrations can be the fundamental goal of the task, while in some manipulation tasks vibrations can enhance performance by reducing reaction times or permitting minimization of forces. Design guidelines for implementation of vibration displays, based on simple mechanical models, are also presented.

1 INTRODUCTION

People make use of high frequency vibrations in many common manipulation tasks. We feel with our fingers for a rattle indicating a loose screw, or for a gritty sensation indicating a damaged ball bearing. Vibrations are often designed into the function of manufactured items, such as snap closures on clothing or calculator keys with detents. In many tasks, contact between hard surfaces is accompanied by copious vibrations; a typical example is the clatter of aligning a steel wrench with a bolt. Texture perception also relies on high frequency vibrations, and professional machinists determine the roughness of machined surfaces by comparing the "feel" of a newly-cut surface to calibration surfaces of known roughness.

Despite their importance in manipulation, vibrations have received little attention in haptic interface research. Most work on teleoperation has focused on force reflection, with bandwidths typically limited to a few Hz (Sheridan, 1992). Work has recently appeared on tactile display of shape in teleoperation and virtual environments (Cohn, Lam, & Fearing, 1992; Hasser & Weisenberger, 1993; Kontarinis & Howe, 1993), as well as sensory substitution aids for the blind such as the Optacon (Bliss, Katcher, Rogers, & Shepard, 1970). These devices often use vibrating pin elements to stimulate the skin. In each of these cases the vibratory stimulus is designed to provide information about another physical parameter such as shape or optical intensity, and no attempt is made to relay information about vibrations that occur as part of a task. Similarly, Massimino and Sheridan (1993) used a vibrotactile display to relay force information in teleoperation, rather than to portray the vibrations at the remote manipulator.

A few previous studies have examined the display of vibrations in a task-related context. Hawkes (1987) used an acceleration sensor to detect vibrations in the finger tips of a remote manipulator, but displayed the resulting signal to the teleoperator in audio form through a loudspeaker. Minsky et al. (1990) developed a system that can provide vibratory information in virtual environments. This joy-stick based device simulates the mechanical interactions, including vibrations, produced by stroking a stylus over various surface textures and features.

In this paper we explore the use of tactile displays for conveying task-related vibratory information in teleoperation and virtual environments. This is a low-cost, open loop display modality that can be easily added to many existing systems to improve performance. Our goal here is to demonstrate the utility of this type of device, and to delineate the kinds of tasks where high frequency vibratory feedback is important. In some tasks, the detection of vibrations is the fundamental goal of the task, while in others, vibrations can enhance task performance by reducing reaction times or permitting minimization of forces. We also show that in simple positioning tasks, or tasks that are limited by precise control of forces, vibration feedback may not improve performance, even though vibrations are generated by the task.

Because the device is designed to stimulate the human tactile sensing system, we begin this paper with a review of pertinent aspects of human tactation. Of particular note here is the fact that high frequency vibrations are poorly localized on the skin, suggesting that a single vibratory display for each finger is adequate. We then describe

the design of our prototype display, which consists of a voice coil actuator working against a freely supported mass to generate inertial forces. For experimental testing of the display we use the planar force-reflecting teleoperated hand developed in our laboratory, with skin acceleration sensors in the finger tips of the remote manipulator to detect task-generated vibrations. Subjects use this system to perform several manipulation tasks, and the results demonstrate the range of tasks for which this type of display is beneficial. Finally, we discuss the design of vibratory display systems to improve the performance of haptic interfaces for teleoperation and virtual environments.

1.1 Human Tactile Sensing

The design of effective high frequency tactile displays requires an understanding of the mechanisms through which humans sense vibrations. The primary receptors for vibrations are located in the skin. Kinesthetic receptors, which include muscle spindles, tendon organs, and joint receptors, probably do not play a large role in high frequency vibration sensing (Clark & Horch, 1986). In particular, the soft subcutaneous tissue under the contact areas of the fingers and palm acts as a mechanical low pass filter which attenuates small amplitude vibrations before they reach these receptors. There are four types of specialized mechanoreceptor nerve endings in the glabrous skin of the human hand which play important roles in manipulation tasks (Johansson & Vallbo, 1983). They may be categorized by two criteria: the size of their active areas and their response to static stimuli. Nerve endings with small receptive fields are called Type I units, while those with large fields Type II. Units that respond to static stimuli are denoted SA (for slowly adapting), while those with no static response are denoted FA or RA (for fast or rapidly adapting).

It is difficult to draw definitive conclusions about mechanoreceptor function since their response is nonlinear and time-varying, and measured sensitivity varies greatly with stimulus size, shape, and duration. However, a few generalizations are helpful. Particularly relevant are the frequency response data of Johansson, Landström, and Lundström (1982). Although all types of units respond to high frequency vibrations to some extent, both SAI and SAII units respond primarily to frequencies below ca. 30 Hz, and are thus relatively unimportant for high frequency vibration sensing. In contrast, FAI units are most easily excited by frequencies in the 10-60 Hz range, and FAII units by frequencies above ca. 60 Hz, with significant response up to at least 1 kHz. The variation of threshold with frequency for these receptors suggest that the FAI units respond to the rate of skin deformation, while the FAII units respond to acceleration of the skin (Gescheider, Verillo, & Checkosky, 1988).

Johansson and Vallbo (1983) summarize spatial response properties for these mechanoreceptors. FAI units are located near the skin surface, and have 3-4 mm diameter receptive fields. This small receptive field size and their high density on the finger tips suggest that they provide spatial information about skin deformation. FAII units are deeper in the subcutaneous tissue and have large receptive fields, at least 20 mm in diameter. Often a single unit will respond to vibrations applied anywhere on

a finger or half of the palm. This suggests that FAII unit responses do not localize vibratory stimulus on the skin surface.

These properties imply that the FAII units are the most important receptors for vibratory information above about 60 Hz. Since these receptors do not exhibit a localized response, we can provide high frequency vibration information with a single vibration display for each finger tip. For lower frequencies, an array-type display may be appropriate to provide spatially resolved stimuli for the FAI and SA receptors (Cohn, et al., 1992; Hasser & Weisenberger, 1993; Kontarinis & Howe, 1993) .

1.2 Tasks and vibrations

We now consider the kinds of tasks where display of vibration information will be useful. For this purpose tasks may be divided into three categories: perceptual tasks where detection of vibrations is the fundamental goal of the task; manipulation tasks where vibrations indicate the state of the task; and tasks where vibrations are not directly important. In the first case, many inspection tasks require the detection of vibrations for their successful completion. Touch inspection differs from visual inspection in that it requires mechanical interaction with the object under inspection. Examples here include tasks where vibrations are produced as parts are manipulated, such as detection of looseness in an assembly of parts or damage to a ball bearing. Also in this category is texture perception, where the objective is determining the smoothness of a surface, or detecting the presence of contaminants such as dirt or liquids – the sorts of tasks where we stroke a finger over a surface to generate vibrations.

In the second category, vibrations can enhance performance by indicating the mechanical state of the hand-object system. This can reduce reaction times or permit minimization of forces. Vibrations can indicate that contact has occurred between a grasped tool and a surface in the environment; examples include bringing a wrench into contact with a bolt, or lowering a grasped object into contact with a table top. Rapid detection of contact permits the human to stop the motion, which reduces contact forces and prevents dislodging the object. This can be particularly important in delicate probing tasks, where detection of the vibrations that indicate the earliest instant of contact can be essential.

Finally, in some tasks vibration information may be unimportant. In simple positioning tasks, or tasks that are limited by precise control of forces, vibration feedback may not improve performance, even though vibrations are generated by the task. However, it may still be useful to provide vibratory information, as it adds to the subjective "feel" of the system for the user. This is particularly true in virtual environments, where vibrations may significantly contribute to the sense of remote presence, even though they may not improve a particular performance measure for a given task.

In the following sections we present three sets of teleoperation experiments aimed at confirming the function of tactile display of high frequency vibrations in each of the cases discussed above. In each experiment we compare subjects' performance with and without vibration feedback. The first experiment tests the utility of the vibration display in an inspection task: subjects attempt to distinguish a worn ball bearing set

using the vibrations generated when it is rotated. The second experiment demonstrates the importance of vibrations in manipulation tasks. Subjects pierce a thin membrane while attempting to minimize contact forces. The results are evaluated in terms of reaction times and applied force magnitudes. In the third experiment subjects perform a close-fit peg-in-hole assembly task as quickly as possible. Here performance is limited by successful control of contact forces, rather than detection of the mechanical state of the system.

2 EXPERIMENTAL APPARATUS

Subjects used a force-reflecting teleoperated hand system to perform the experimental tasks. Sensors in the finger tips of the remote manipulator measured the vibrations generated during task execution, and prototype high frequency vibrotactile displays relayed these vibrations to the subjects operating the system.

2.1 High frequency vibration display and sensor

From the human factors discussed above, a high frequency vibration tactile display device should produce mechanical vibrations in the range of 60-1000 Hz with variable amplitude and frequency. Since this is in the audio frequency range, miniature loudspeakers were easily modified for this application. These prototype high frequency vibration displays consist of 0.2 watt loudspeakers mounted "upside down," with the outer cones and metal frames removed. The remaining structure containing the magnet, coil, and central diaphragm is then attached to the master manipulator near the operator's fingers (Figure 1). The base containing the permanent magnets was thus free to move in space. Passing current through the coil generates a force against the magnet, which accelerates the 35 gram base and produces an inertial reaction force against the manipulator. This inverted mounting results in a higher moving mass compared to the usual audio configuration, providing larger inertial forces. The range of motion is 3 mm, and the displays can, for example, produce up to 0.25 N peak force at 250 Hz. Vibrations are transmitted to the fingers of the human operator through aluminum bracket "handles" mounted at the ends of the master manipulator finger links, as shown in Figures 1 and 2.

A number of different tactile sensors can detect task-related vibrations at the remote robot finger tips (Howe & Cutkosky, 1992) . We use skin acceleration sensors, which consist of miniature instrumentation accelerometers mounted on the inner surface of the rubber finger tip skin of the slave fingers, shown in Figure 1 (Howe & Cutkosky, 1989) . A layer of foam rubber beneath the skin provides passive compliance to improve grasp stability and isolate the sensor from vibrations in the robot mechanism. In the mounting configuration used here, the skin acceleration sensors have their greatest sensitivity to vibrations in the vertical direction; however, the compliant skin and foam readily couple vibrations in other directions to create an omnidirectional sensor.

One advantage of the skin acceleration sensor is its excellent sensitivity to vibrations at the frequencies we are concerned with here. Acceleration sensing is also appropriate because the human FAII mechanoreceptors that are particularly important

in this context appear to respond to acceleration of the finger tip skin. Furthermore, the vibration display produces inertial forces which will in turn accelerate the operator's finger tip skin, reproducing the same physical quantity that the sensor is measuring at the remote manipulator.

To characterize the response of the sensor–display system we applied vibrations to the remote manipulator finger tips with a high bandwidth vibration test shaker. The skin acceleration sensor detected the induced vibrations, and the resulting signal drove the display through a small audio amplifier. The output vibrations produced by the display were recorded by an accelerometer on the master finger tip handle near the location where the operator's fingers rest. The gain of the amplifier can be adjusted to control the relative sensitivity of the system; here we set it to produce approximately the same average amplitude vibrations at the input and output of the system. The measured response amplitude was flat to within 7 dB across the frequency range of interest. There appear to be several weak mechanical resonances in the structure of the remote manipulator finger tip and master handles, but the relative amplitudes of these features are small, probably due to the passive damping of the rubber in the remote finger tip and subcutaneous tissue in human finger tips. Overall, the response of the system was adequate for the purpose of testing the effectiveness of vibration display.

2.2 Teleoperated Hand System

These experiments use a force-reflecting teleoperated hand system developed in our laboratory for the study of tactile sensing and display (Howe, 1992). This system trades a limitation on the number of degrees of freedom for a clean and simple mechanical design, which results in good control of fine forces and motions. The system is designed to execute tasks that humans usually accomplish with a precision pinch grasp between the thumb and index finger. For most tasks the operator's wrist rests on the table top and the operator makes contact with the master only at the tips of the thumb and index finger (Figure 2).

Both master and remote slave manipulators are identical two-fingered hands with two degrees of freedom in each finger, so finger tip position or force can be controlled within the vertical plane. The mechanism uses a direct-drive, parallel linkage design, which minimizes friction, backlash, and moving mass. Two axis strain gauge force sensors measure finger tip forces on both master and slave hands. The controller uses a conventional bilateral force reflection control scheme. The joint angles of the master manipulator are the command inputs for position control of the joints of the slave manipulator. Conversely, the forces measured at the slave finger tips are the command inputs for force control of the master. The measured slave position bandwidth is 18 Hz and the master force reflection bandwidth is greater than 80 Hz. Further details of the manipulator system design and performance are presented in (Howe, 1992; Howe & Kontarinis, 1992). To ensure that high frequency information was provided primarily by the tactile display, we used a 30 Hz second-order low pass digital filter to limit the frequency content of the force sensor signals from the slave manipulator. This relatively high frequency helped to minimize rise rate and delay differences between the force and vibration displays. The resulting rise time delay for the force feedback system was

15 ms (measured from force application at the slave finger tip to force production against at the operator's finger tip), as expected for this filter. These filtered force signals were used in all of the force feedback cases unless otherwise stated below. To preclude passive transmission of vibrations, the master and remote manipulators were situated on separate tables.

3 EXPERIMENTAL PROCEDURES

Six male and one female subjects, 21-32 years old, participated in the experiments. All subjects had sufficient experience in operation of the teleoperated hand system so that learning effects were not significant. In each task subjects were seated in front of the master manipulator and directly viewed the slave manipulator at a distance of about 1.5 m. Subjects also wore headphones playing white noise to mask any audible feedback produced by the tasks.

3.1 Ball bearing Inspection Task

In this experiment subjects used the system in an illustrative inspection task, the identification of a damaged ball bearing set. In machine maintenance procedures workers often test ball bearing sets by rotating them with their finger tips. The tactile detection of vibrations indicates wear or contamination of the bearing. We performed two series of experiments with slightly different protocols. The first protocol was aimed at demonstrating the ability to use the information provided by the vibration display to discriminate between worn and new bearings, while the second required identification of a worn bearing without comparison in a limited time.

In both cases, subjects were presented with a pair of ball bearing sets (Figure 3), one of which was new while the other was worn and produced vibrations when rotated (although rotational torque was the same for the two bearings). Subjects used one finger of the teleoperated hand to rotate each bearing to detect its condition. Between trials the spatial order of the bearings could be altered without the knowledge of the subject. Subjects performed the task under four feedback conditions: no haptic feedback; force feedback only; vibratory feedback only; and both vibratory and force feedback. In the first protocol the force feedback signal from the remote manipulator was not filtered and thus was equal to the full bandwidth of our manipulator. In this experiment five subjects rotated both bearings and made a forced choice as to which one was worn. In the second protocol, three subjects were asked to gently inspect the top bearing only and after 5 seconds make a forced choice as to whether that bearing was worn (i.e., no comparison). Subjects completed a total of 80 trials in the first protocol and 120 trials in the second. In both experimental series subjects were informed as to the type of feedback provided.

3.2 Puncturing Task

This experiment is designed to test the proposition that when vibrations indicate the progress of a manipulation task, vibrotactile display can improve performance by minimizing reaction times or force magnitudes. The task we selected is designed to emulate the key aspects of medical procedures such as biopsies and catheterizations,

where a needle must penetrate a thin and relatively stiff layer of tissue but avoid damaging soft tissue underneath. The experimental set up is shown in Figure 4. A 0.05 mm thick plastic membrane (cellophane tape) is mounted on a rectangular frame 0.8 mm in height and approximately 10 mm x 30 mm across. Under the frame is a layer of soft latex rubber. This test structure is mounted on a force sensor that measures the applied force normal to the surface. Subjects were asked to pierce through the tape using a sharp needle held between the fingers of the slave manipulator without exerting excessive force that could damage the underlying rubber layer. They were also instructed that their performance would be penalized for applying excessive force.

Three subjects participated in these experiments. They were presented with the four different combinations of visual, filtered force, and vibratory feedback described for the previous experiment. The order of feedback combinations was randomized to minimize sequence effects. In total, for each feedback case there were 38 sequences comprised of 152 trials. Subjects were permitted to practice the task until they became proficient, and each 10 trial session lasted for about 30 minutes.

The force on the apparatus was sampled at 10 kHz, then digitally filtered (first forward and then backward to eliminate phase delays) with a 250 Hz, ten-pole low pass filter, and finally decimated by a factor of ten. This force information was used to evaluate the subjects' reaction time to penetration of the membrane. The piercing of the membrane was easily recorded as a transient in the force trace recorded from the force sensor under the apparatus (Figure 5). Defining the time at which the subjects started retracting required an objective measure of when the force began to significantly and monotonically decrease, independent of small-scale noise on the force signal. We defined the reaction time as the first point of the trace (after piercing) that a line of slope of 2 N/s was tangent to the force and greater than all points of the force signal for the remainder of the trial. For comparison, we also determined the reaction time in the first 8 trials by selecting the point at which the force trace showed a significant decrease, indicating that subjects started retracting. This measurement resulted in slightly faster response times, but there is no indication that either method favored the vibratory display cases.

The force records were also used to evaluate how much more force the subject applied to the apparatus than required to successfully execute the task. For this purpose excessive force is defined as the portion of the force signal after the piercing transient that exceeds the level required to pierce the membrane. For example, the force record in Figure 5 shows excessive force above the piercing force threshold line. We measured both the fraction of trials that resulted in excessive force and the sum of excessive readings, calculated as

$$S = \frac{1}{n} \sum_{i > i_{\text{pierce}}} (f_i - f_{\text{pierce}}) \Delta t$$

where the f_i are force samples that exceed the piercing force f_{pierce} , Δt is the sampling period, n is the number of trials for the feedback case, and the summation begins after the piercing transient.

3.3 Peg-in-Slot Task

The final experiment involves an assembly task where precise control of contact forces is essential. This task is a peg-in-hole insertion; for the planar case considered here, the task becomes insertion of a cylindrical peg into a rectangular slot perpendicular to the plane of motion (Figure 6). This ensures that only the two force components and one torque component in the plane may be generated by contact between the peg and hole. The peg is a ground steel cylinder 12.70 mm in diameter, attached to an aluminum block which provides a flat grasping surface for the slave manipulator finger tips. The two adjacent holes or slots are constructed from precision ground steel machinist's parallels, with an opening width of 12.71 mm and a depth of 11.91 mm. This tight clearance emphasizes the role of accurate force alignment in task execution (Whitney, 1982).

In each trial, the subject begins by grasping the peg and holding it between the slave manipulator finger tips above and in between the two holes. The subject then proceeds to insert the peg. When the peg reaches the bottom of the hole an electrical contact closes, lighting an indicator visible to the operator. The operator then extracts the peg and proceeds to insert the peg in the second hole. The number of successful insertions within the 30 second data acquisition time is then recorded. The objective is to determine minimum completion time for the task, and operators are instructed to perform the task "as quickly as possible." Before each trial subjects are informed of the type of feedback that will be available for the particular trial. For this experimental series three subjects performed a total of 160 trials.

4 EXPERIMENTAL RESULTS

4.1 Ball Bearing Task

Results for the first experimental protocol, where subjects performed the inspection task using comparative touch, are presented in Figure 7(a). With no haptic feedback the correct response rate was no better than chance (50%), as expected since it was impossible to detect the bad bearing visually. The use of force feedback raised the correct response rate to 80% ($p \sim 0.1$)¹, while with vibratory feedback subjects were able to select the bad bearing every time, either with or without force feedback ($p < 0.025$).

Results for the second protocol (no comparison) are shown in Figure 7(b). The correct response rate was 53% with visual feedback only, again indicating that subjects were guessing as to the condition of the bearing. The use of filtered force feedback raised the correct response rate to 73% ($p \sim 0.1$). When subjects were presented with vibratory feedback without force feedback they responded correctly 66% of the time ($p < 0.05$). Due to the absence of force feedback, subjects had difficulties in manipulating the ball bearing in the 5 seconds allowed, which resulted in the production of minimal or faulty tactile cues. Finally, with vibratory and filtered force feedback, subjects were able to select the bad bearing 90% of the time ($p < 0.025$).

¹ p values were obtained from a matched pair t-test. Appendix 1 contains the complete results for these tests.

4.2 Puncturing Task

Figure 8 presents reaction time results for the puncturing task experiment. In the force feedback only case the 15 ms delay introduced by the low pass filter has been subtracted from the reaction times. The width of the histogram in the no haptic feedback case in Figure 8 suggests that subjects had difficulty reliably detecting the puncture event without haptic feedback. The presence of either force or vibration feedback significantly decreased mean reaction times by approximately a half ($0.005 < p < 0.025$). This accords with previous reports that visual reaction times are considerably slower than tactile reaction times in manual tasks (Boff & Lincoln, 1988).²

The combination of both force and vibration feedback was similarly significant in further reducing reaction time by roughly 50 ms. In both the force only and vibration only histograms most of the reaction times were less than about 220 ms, with a few trials showing considerably longer response times. In the combined force and vibration case, all trials resulted in reaction times of less than 220 ms. This is reflected in the reduction of the standard deviation by a factor of two. These distributions suggest that in a few trials the subjects failed to detect the transient event when it was conveyed by only one feedback modality. The provision of both vibration and force feedback may provide redundant cues that eliminate these errors.

HAPTIC FEEDBACK	NONE	FORCE ONLY	VIBRATION ONLY	FORCE AND VIBRATION
FRACTION OF TRIALS WITH EXCESSIVE FORCE	63 %	18 %	23 %	5 %
SUM OF EXCESSIVE FORCE	210 mN-s	7 mN-s	18 mN-s	3 mN-s

Table 1. Summary of force results for the puncturing task experiment.

Figure 5 shows a typical force record for this experiment and Table 1 lists excessive force results. As with reaction times, the provision of either type of haptic feedback alone significantly decreases excessive force, with roughly the same improvement from either type of feedback. The combination of both vibration and force feedback results in another significant improvement in performance.

4.3 Peg-in-Slot Task

Figure 9 shows the completion time results for the peg-in-slot task. The clear conclusion is that force feedback substantially decreases completion times, while vibration feedback does not have a significant effect. Observation revealed that most of the time in each trial was spent either in the transport phase (peg freely supported by

² The gain from slave sensor to master display was set to unity for both vibration and force feedback systems; however, in these initial experiments we do not address the effects of magnitude variation, which may affect reaction times.

the fingers moving between slots) or in the insertion-extraction phase (peg in contact with both sides of a slot). Subjects often made the transition between the transport and insertion-extraction phases by lowering the peg into contact with the flat surface adjacent to the target slot and dragging the peg along the surface until it dipped into the slot. This search phase generates abundant vibrations which were reproduced by the vibration display, but this phase represented only a relatively small fraction of the total completion time.

Subjective reports of the subjects nevertheless indicate that the "feel" of the system was quite different with vibration feedback. Subjects asserted that the vibration display helped them when force feedback was also available since the system felt more "complete." When the vibration display was absent subjects felt that the system was "dead" and that "something was missing," suggesting that vibrations had become an expected part of the task. Subjects preferred to use the vibration display and indicated that it improved their perception of the situation at the remote manipulator.

5 DISCUSSION

5.1 Force versus vibration feedback

These experiments provide insight into the comparative advantages of force and vibration feedback, particularly in the second protocol of the ball bearing task (Figure 7b) and the puncturing task (Figure 8 and Table 1). When provided with force feedback alone subjects had difficulty detecting vibrations, as required for successful completion of the task. Although the force feedback bandwidth was sufficient to convey some vibratory information, the time constraints in these tasks apparently impeded its use. Conversely, with vibration feedback alone subjects had difficulty manipulating objects, so that they were often unsuccessful in rotating the ball bearing with the slave finger tip, or applied excessive force after piercing the membrane.

This suggests that force and vibration feedback are best considered complimentary modalities which improve performance for different reasons. However, vibrations are generated by tasks that involve contact, which are the sorts of tasks where force feedback is most helpful. This implies that while vibration display alone can be useful, the combination with force feedback is most beneficial, as illustrated by the large improvement in performance when both force and vibration feedback was available.

As used in these experiments the teleoperated hand system had a force reflection bandwidth of at least 30 Hz, so the beneficial effects of the vibration feedback may be more pronounced with a conventional manipulator with lower bandwidth. This possible improvement may be mitigated by the fact that force feedback is not as sensitive as vibration feedback at high frequencies due to the difference in sensor location and construction. The skin acceleration sensor is located directly at the point of contact, and the compliant foam layer ensures a good connection with the environment while isolating the sensor from vibrations in the manipulator structure. The low moving mass of the sensing element permits the sensor to readily follow high frequency vibrations. The force sensor is situated between the finger tip and manipulator links, and the mass and compliance of the finger tip acts as a mechanical low pass filter that

attenuates the vibrations produced at its outer surface. Similar considerations apply to force reflection at the master manipulator; see Section 6.2 below. Thus even with high bandwidth force reflection small, fast vibrations are not as well sensed and displayed with force feedback, as confirmed by observation of the force and vibratory feedback signals in these experiments.

5.2 Task properties and vibrations

The importance of high frequency information varies greatly among tasks. For these initial studies we obviously selected applications where vibrations are important in order to establish the functionality of high frequency sensing and display. The ball bearing experiment confirms the utility of vibration display in inspection tasks such as detecting looseness in mechanical assemblies or assessing surface texture. In fact, we can provide far greater vibration sensitivity with this system than humans have when manipulating objects directly with their fingers, since the amplifier driving the display can have arbitrary gain. The limit to this "superhuman" sensitivity is set by the vibration sensor noise level and the operator's ability to generate appropriate tactile stimuli.

One of the most important roles of vibrations is in indicating transitions in complex tasks. Such tasks consist of a sequence of phases or subtasks separated by discrete events. These events indicate the status of the task as it proceeds and permit the sequential coordination of activities that lead to completion of the task. In the puncturing task, one such transition is from applying increasing force against the membrane without motion before piercing, to decreasing the applied force and pulling back after piercing. Transitions are important to detect because they usually trigger a change of trajectory (e.g. from pushing down to pulling up) or a change of control mode (e.g. from position control to force control). Other examples of transition-marking events include making or breaking contact and starting or stopping sliding. These transient events are often best detected through vibratory signals. Westling and Johansson (1987) have shown that human FAII mechanoreceptors are particularly responsive to transient events in manipulation, and Howe et al. (1990) have demonstrated the use of this kind of event detection in robotic manipulation.

The peg-in-slot task included a number of transition events that generated copious vibrations, including initial contact of the peg with the slot, sliding of the peg along the edge of the slot, and contact of the peg with both slot faces. Nevertheless, the presence of vibratory feedback did not significantly improve performance in this experiment (Figure 9). In this experiment performance was evaluated exclusively on the basis of the completion time. Since the task typically took several seconds to complete, the small advantage in reaction time provided by the vibration display seen in the puncturing task was not significant here. The tight clearance between peg and slot necessitated accurate coordination of insertion forces to prevent wedging or jamming of the peg. The vibration display provides minimal information about this aspect of the task, and so failed to improve performance. In contrast, further informal experiments suggest that when subjects are asked to minimize contact forces, vibration feedback can improve performance.

These assembly experiments also indicate that humans can modulate the generation of vibrations by controlling grasp stiffness. When working quickly, subjects generally used a firm grip which produced large contact transients, with concomitant large amplitude vibrations. When minimizing forces, subjects often used a compliant grip that reduced the magnitude of the impact and thus reduced vibrations. The compliant finger tips on our slave manipulator tend to decrease grasp stiffness and thus vibrations; if a remote tool is rigidly connected to a slave robot then the importance of vibrations will be further accentuated.

Fundamental questions remain about the nature of vibrations generated in particular manipulation tasks and the methods that can be used to evaluate their importance. Further study of human tactile sensation may help to determine parameter ranges such as optimum amplitudes and useful frequencies to sense and display, and methods of combining this modality with other tactile displays. Improved understanding of the relationship between vibrations and task type can help ascertain which applications will benefit from vibration sensing and display. This will also permit us to learn the appropriate circumstances for vibration display in virtual environments.

5.3 Virtual environments

Since humans are accustomed to feeling vibrations in many manipulation tasks, vibration display may increase the sense of remote presence in virtual environments (Loomis, 1992; Rheingold, 1991). Subjective reports in these experiments indicate that vibrations contributed to subjects' conscious perception of the remote environment, and the absence of vibrations was noted as detrimental. The preceding discussion of tasks provides some guidelines on the sorts of events where vibratory information can enhance the realism of virtual environments. In training applications the provision of realistic vibrations may help users learn to distinguish correct and erroneous task progress.

In practical terms, the addition of vibratory displays to a haptic interface for virtual environments can be simple and inexpensive. Our initial observations of vibratory signal for many transition events indicate that the appropriate waveforms may be readily synthesized from exponentially decaying sums of sine waves. Alternatively, the vibratory waveforms for each type of event could be recorded from hand-held vibration sensors in trial executions of typical tasks. Because the frequencies of interest are within the audio frequency range, a sound generation board may be used to drive the display; the sound board may then be triggered by the controller to reproduce the appropriate event without imposing an additional computational burden.

6 DISPLAY DESIGN CONSIDERATIONS

6.1 Display design

One attractive feature of this display modality is its low implementation cost. The display device is simple to construct from a small loudspeaker costing only a few dollars, and a low power commercial audio amplifier is more than adequate to drive the

display. We use a piezoelectric accelerometer in compliant finger tips as our vibration sensor, but there are a number of less expensive vibration sensors that can be robustly mounted on conventional robot grippers for many applications (Howe & Cutkosky, 1992). A vibration sensing and display system can be added to many existing teleoperated systems, with or without force reflection. Because the display is an open-loop system that acts at frequencies above the manipulator bandwidth, it does not induce stability problems in the manipulator controller.

The vibration sensing and display system described here has not been optimized, and a number of central issues remain to be addressed. Among these are the optimal choice of amplitude and frequency range. For many applications a single display on each finger may be adequate, but with further development it may be possible to use a single device to convey both object shape and high frequency vibration information. We have developed a tactile shape display which consists of an array of pin elements that are raised against the operator's finger tip to approximate the object shape sensed at the slave manipulator's finger tip (Kontarinis & Howe, 1993). The vibratory signal may be imposed on the low frequency shape information, and could provide both spatially unresolved high frequency vibrations and localized vibrations at intermediate frequencies. Spatially resolved intermediate frequencies are particularly appropriate for stimulating FAI mechanoreceptors, and could provide an optimum means of conveying slip information (Edin, Howe, Westling, & Cutkosky, 1993; Howe, 1992; Johansson & Westling, 1987).

6.2 Combining force reflection and vibration display

In conventional force reflection, the low frequency components of the force applied by the slave manipulator are measured by a force sensor, and the master manipulator displays this force to the human operator. These experiments suggest that high frequency information may be effectively conveyed by specialized sensors and displays that have good response at high frequencies and minimal response at low frequencies. This approach divides haptic feedback into separate displays of low and high bandwidth. In this section we discuss one important consideration in the design of a vibration display system, the compliance of the master manipulator structure near the finger tips. This compliance effectively decouples the finger tips from the rest of the manipulator at high frequencies, which permits a relatively small vibration display device to produce adequate vibrations on the operator's fingers.

To examine this aspect of display design, we use a simple linear lumped-parameter one-degree-of-freedom model of the master manipulator (Figure 10). Here M_{LINK} is the effective moving mass of the manipulator inwards of the compliance (including linkages, actuators, etc.), K is the stiffness of the compliant element (in our manipulator this is due to the flexibility of the master finger tip handles), and M_{TIP} is the effective mass outwards of the compliance (including the moving portions of the master handles and human operator's finger tip). The actuators of the manipulator apply the force $F_{actuator}$ to the linkages, and the vibration display applies force $F_{display}$ to the finger tip. Position of the linkage is denoted $x(t)$, and of the finger tip $y(t)$. The equations of motion for this system in the Laplace domain are

$$(M_{link} s^2 + K) X(s) - K Y(s) = F_{actuator}(s)$$

$$(M_{tip} s^2 + K) Y(s) - K X(s) = F_{display}(s).$$

We wish to compare the ability of the actuator and the vibration display to generate vibrations at the finger tip. We can solve the above equations to find the transfer functions $s^2 Y(s)/F_{actuator}$ and $s^2 Y(s)/F_{display}$ from each force to tip acceleration, and then equate the accelerations. The resulting ratio of force amplitudes required to produce a given vibration amplitude as a function of frequency ω is

$$\left| \frac{F_{display}}{F_{actuator}} \right| = \left| \frac{\omega_n^2}{\omega_n^2 - \omega^2} \right|$$

This is an undamped second-order system with a resonance at $\omega_n = \sqrt{K/M_{link}}$. At low frequencies ($\omega \ll \omega_n$) the ratio is near unity, which indicates that the actuator and display are equally efficient, while above resonance the vibratory display is increasingly effective.

Physically, this is because the compliance decouples the small tip mass from the large link mass at high frequencies. This permits the small forces generated by the display to produce significant accelerations. The decoupling also means that at high frequencies the actuator must produce large displacements at the links to induce vibrations at the tip (i.e., $|x| \gg |y|$). At low frequencies the actuator is well-coupled to the tip, which provides the means for conventional force reflection, while the display can generate only small-amplitude vibrations due to its limited authority and the strong coupling to the large link mass.

The addition of compliance at the manipulator tip for a vibration display requires careful consideration. The linkages will inevitably show some compliance, which is responsible for the "lowest resonant mode" that is often specified as a structural design performance criterion. This built-in compliance alone may be adequate, provided that the effective moving tip mass is sufficiently small, and that the stiffness is sufficiently low that the resonance is at an appropriate frequency. Alternatively, it may be necessary to add a compliant element near the finger tip. The resulting resonant frequency should be above the bandwidth of the force reflection controller to ensure stability, but low enough that the vibration display works against a low mechanical impedance.

7 CONCLUSIONS

These experiments represent a first attempt to define the role of high frequency sensing and display in teleoperation and virtual environments. The results demonstrate that vibratory information can play an important role in manipulation. Humans make use of this information, and sensing and display of vibrations can improve performance in teleoperation. A simple and inexpensive sensing and display system

can be readily added to existing teleoperated manipulation systems and haptic interfaces for virtual environments. Further work will be directed at understanding the relationship between vibrations and manipulation tasks and optimizing sensor and display device characteristics.

ACKNOWLEDGMENTS

Funding for this work was provided by the Office of Naval Research under ONR Grant No. N00014-92-J-1887.

REFERENCES

- Bliss, J. C., Katcher, M. H., Rogers, C. H., & Shepard, R. P. (1970). Optical-to-tactile image conversion for the blind. *IEEE Trans. on Man-Machine Systems*, MMS - 11(1), 58-65.
- Boff, K. R., & Lincoln, J. E. (Eds.). (1988). *Engineering Data Compendium: Human Perception and Performance*. Ohio: H.G. Anderson Aerospace Medical Research Laboratory.
- Clark, F. J., & Horch, K. W. (1986). Kinesthesia. In K. R. Boff, L. Kaufman, & J. P. Thomas (Eds.), *Handbook of Perception and Human Performance*, John Wiley and Sons.
- Cohn, M. B., Lam, M., & Fearing, R. S. (1992). Tactile feedback for teleoperation. In H. Das (Ed.), *Telem manipulator Technology*, Proc. SPIE 1833, 240-254.
- Edin, B., Howe, R., Westling, G., & Cutkosky, M. (1993). A physiological method for relaying frictional information to a human teleoperator. *IEEE Trans. on Systems, Man, and Cybernetics*, 23(2), 427-432.
- Gescheider, G. A., Verillo, R. T., & Checkosky, C. M. (1988). Four channels mediate the mechanical aspects of touch. *J. Acoust. Soc. Am.*, 84(5), 1680-1694.
- Hasser, C., & Weisenberger, J. M. (1993). Preliminary evaluation of a shape memory alloy tactile feedback display. H. Kazerooni, B. D. Adelstein, & J. E. Colgate (Eds.), *Proceedings of Symposium on Haptic Interfaces for Virtual Environments and Teleoperator Systems in the ASME Winter Annual Meeting*, New Orleans, LA, 49, 73-80.
- Hawkes, G. S. (1987). *Apparatus providing tactile feedback to operators of remotely controlled manipulators* (U.S. Patent No. 4,655,673). Filed on May 10, 1983.
- Howe, R. D. (1992). A force-reflecting teleoperated hand system for the study of tactile sensing in precision manipulation. *Proceedings of 1992 IEEE International Conference on Robotics and Automation*, Nice, France, 1321-1326.
- Howe, R. D., & Cutkosky, M. R. (1989). Sensing skin acceleration for texture and slip perception. *Proceedings of 1989 IEEE International Conference on Robotics and Automation*, Scottsdale, AZ, 145-150.
- Howe, R. D., & Cutkosky, M. R. (1992). Touch sensing for robotic manipulation and recognition. In O. Khatib, J. J. Craig, & T. Lozano-Perez (Eds.), *Robotics Review 2*, Cambridge, Mass: MIT Press.

Howe, R. D., & Kontarinis, D. A. (1992). Task performance with a dextrous teleoperated hand system. H. Das (Ed.), In *Telemanipulator Technology*, Proc. SPIE 1833, 199-207.

Howe, R. D., Popp, N., Akella, P., Kao, I., & Cutkosky, M. (1990). Grasping, manipulation, and control with tactile sensing. *Proceedings of 1990 IEEE International Conference on Robotics and Automation*, Cincinnati, Ohio, 1258-1263.

Johansson, R. S., Landström, U., & Lundström, R. (1982). Responses of mechanoreceptive afferent units in the glabrous skin of the human hand to sinusoidal skin displacements. *Brain Research*, 244, 17-25.

Johansson, R. S., & Vallbo, Å. B. (1983). Tactile sensory coding in the glabrous skin of the human hand. *Trends in Neuroscience*, 6(1), 27-32.

Johansson, R. S., & Westling, G. (1987). Signals in tactile afferents from the fingers eliciting adaptive motor responses during precision grip. *Experimental Brain Research*, 66, 141-154.

Kontarinis, D. A., & Howe, R. D. (1993). Tactile Display of Contact Shape in Dextrous Manipulation. H. Kazerooni, B. D. Adelstein, & J. E. Colgate (Eds.), *Proceedings of Symposium on Haptic Interfaces for Virtual Environments and Teleoperator Systems, ASME Winter Annual Meeting*, New Orleans, LA, 49, 81-88.

Loomis, J. M. (1992). Distal attribution and presence. *Presence*, 1(1), 113-119.

Massimino, M. J., & Sheridan, T. B. (1993). Sensory substitution for force feedback in teleoperation. *Presence*, 2(4), 344-352.

Minsky, M., Ouh-young, M., Steele, O., Brooks, F. P. J., & Behensky, M. (1990). Feeling and Seeing: Issues in force display. *Proceedings of Symp. on Interactive 3D Graphics*, 235-243.

Rheingold, H. (1991). *Virtual Reality*. New York: Summit.

Sheridan, T. B. (1992). *Telerobotics, Automation, and Human Supervisory Control*. Cambridge, MA: MIT Press.

Westling, G., & Johansson, R. S. (1987). Responses in glabrous skin mechanoreceptors during precision grip in humans. *Experimental Brain Research*, 66, 128-140.

Whitney, D. E. (1982). Quasi-static assembly of compliantly supported rigid parts. *Journal of Dynamic Systems, Measurement, and Control*, 104, 65-77.

APPENDIX 1: STATISTICAL ANALYSES

I) BALL BEARING TASK.

Table A1. Comparison task: Matched pair t-test (5 Subjects).

df=4	Vibration Only to No Haptic	Force Only to No Haptic	Vibration Only to Force Only	Vibr. & Force to Force Only	Vibr. & Force to Vibration Only	Vibr. & Force to No Haptic
<i>t</i>	3.1623	1.5	2.1381	2.1381	∞	3.1623
<i>p</i> ~	<0.025	0.1	0.05	0.05	0	<0.025

Table A2. No Comparison task: Matched pair t-test (3 Subjects).

df=2	Vibration Only to No Haptic	Force Only to No Haptic	Vibration Only to Force Only	Vibr. & Force to Force Only	Vibr. & Force to Vibration Only	Vibr. & Force to No Haptic
<i>t</i>	4	2.0	0.76	1.8898	7	5.5
<i>p</i> ~	<0.05	0.1	>0.25	0.1	0.01	<0.025

II) PIERCING TASK

Table A3. Piercing task: Matched pair t-test (3 Subjects).

df=2	Vibration Only to No Haptic	Force Only to No Haptic	Vibration Only to Force Only	Vibr. & Force to Force Only	Vibr. & Force to Vibration Only	Vibr. & Force to No Haptic
<i>t</i>	7.2667	7.3270	0.7195	3.0849	3.4501	10.9042
<i>p</i> ~	<0.025	<0.025	>0.25	0.05	0.05	0.005




Figure 1.




Figure 2.




Figure 3.

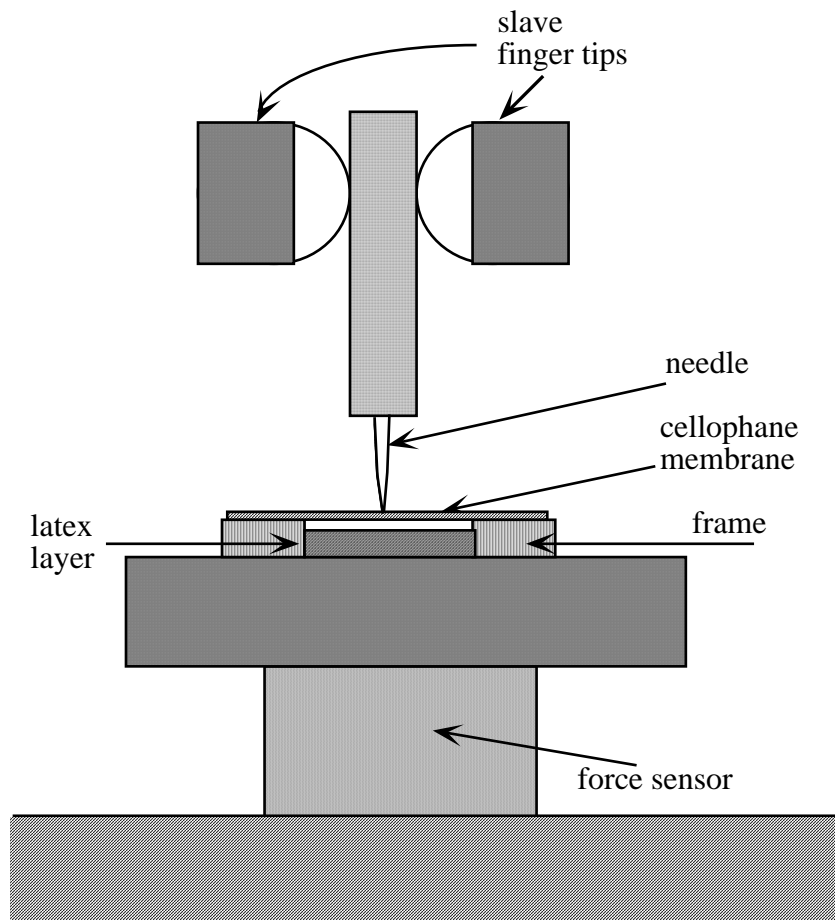


Figure 4.

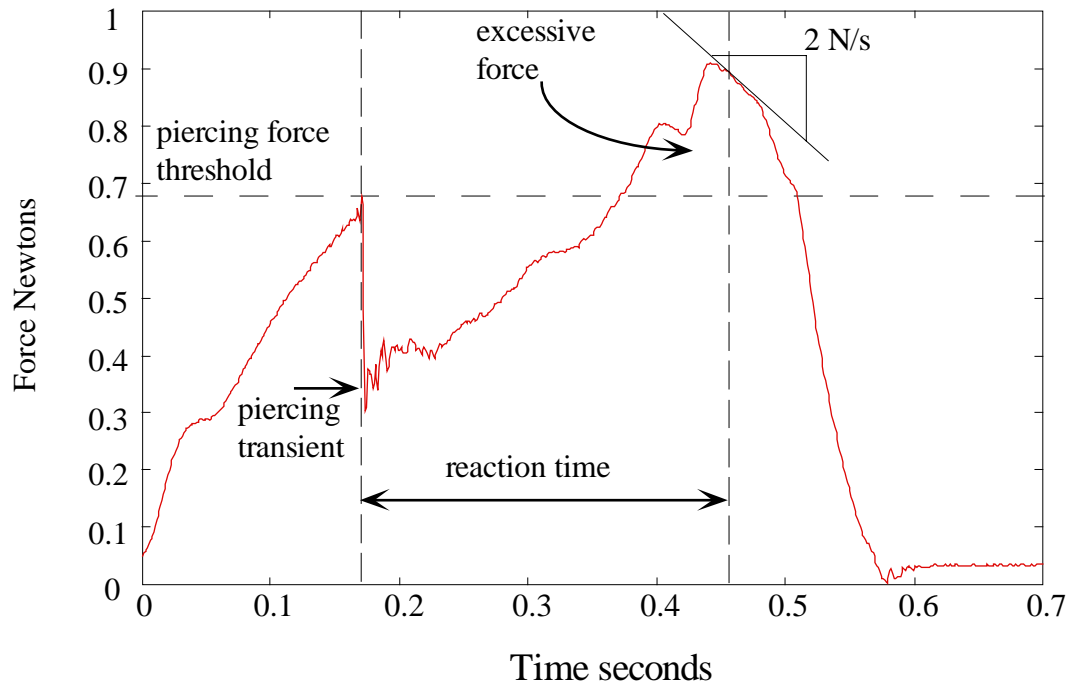


Figure 5.




Figure 6.


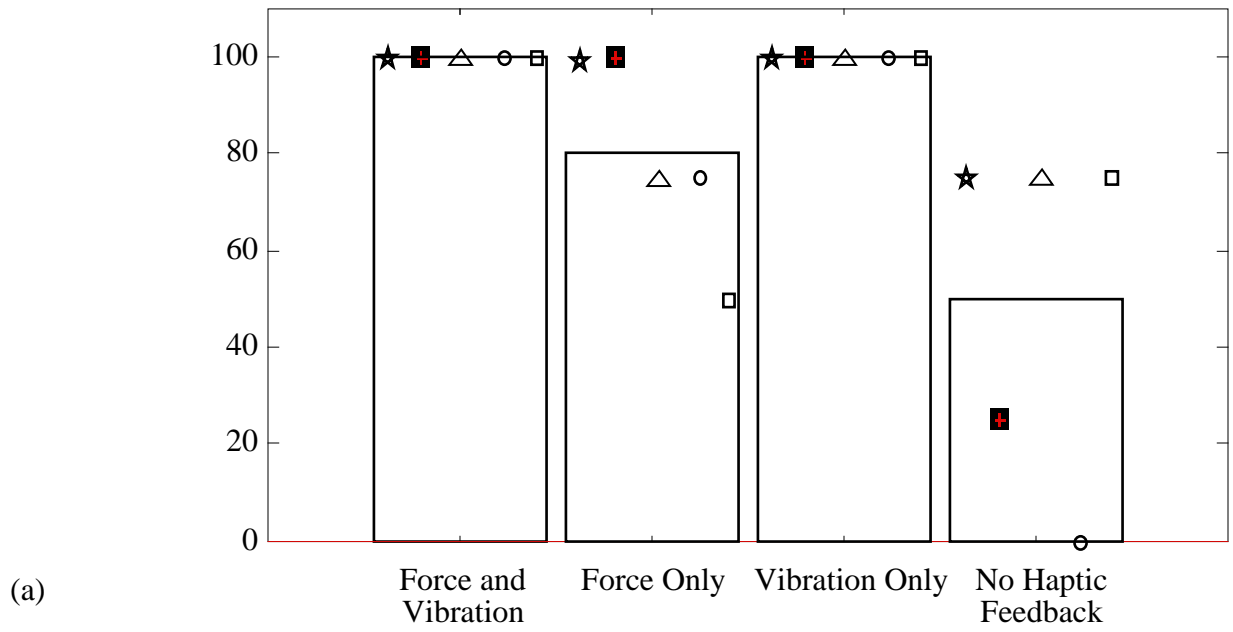


Figure 7.

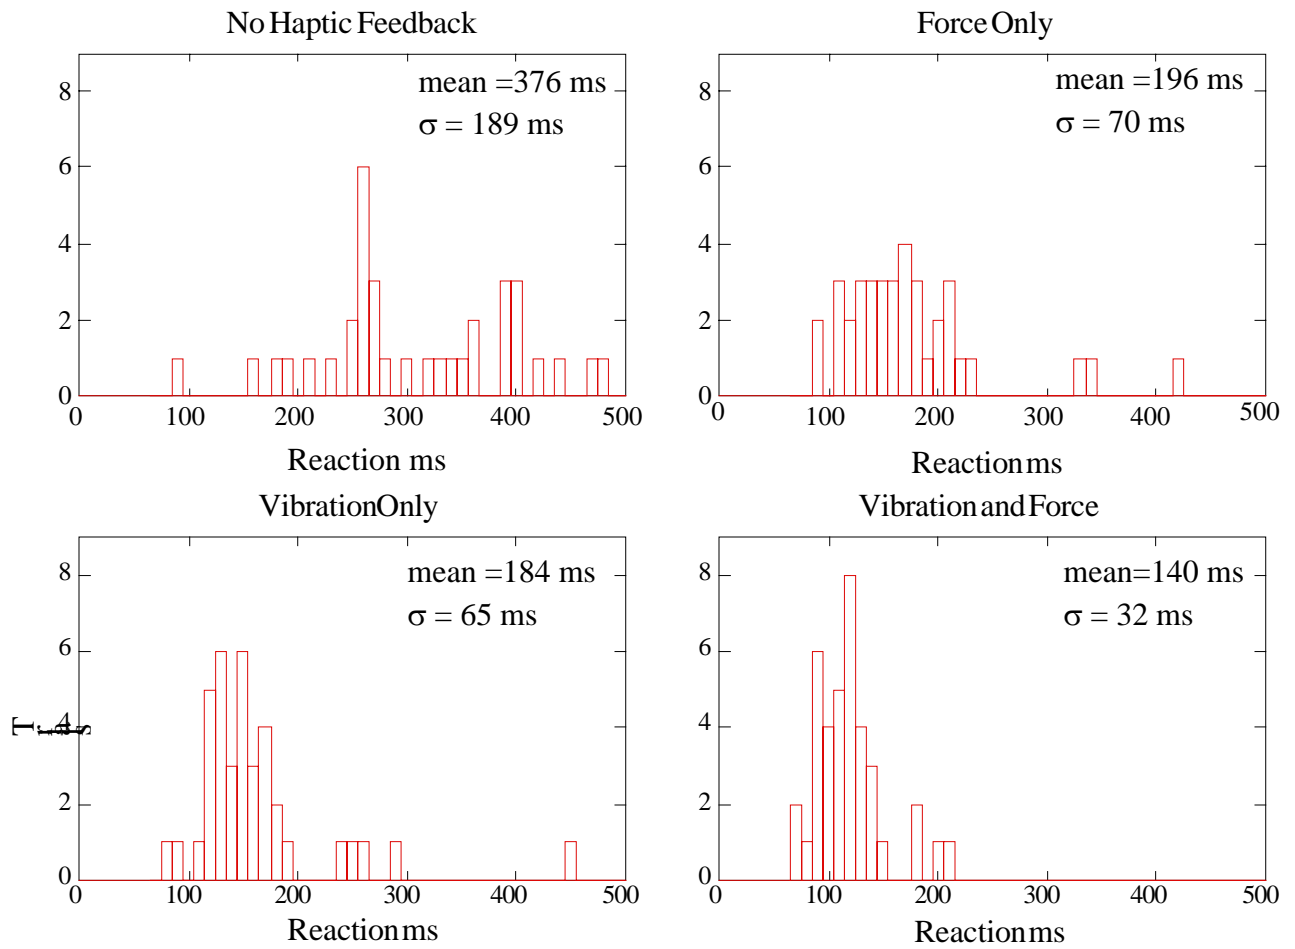


Figure 8.

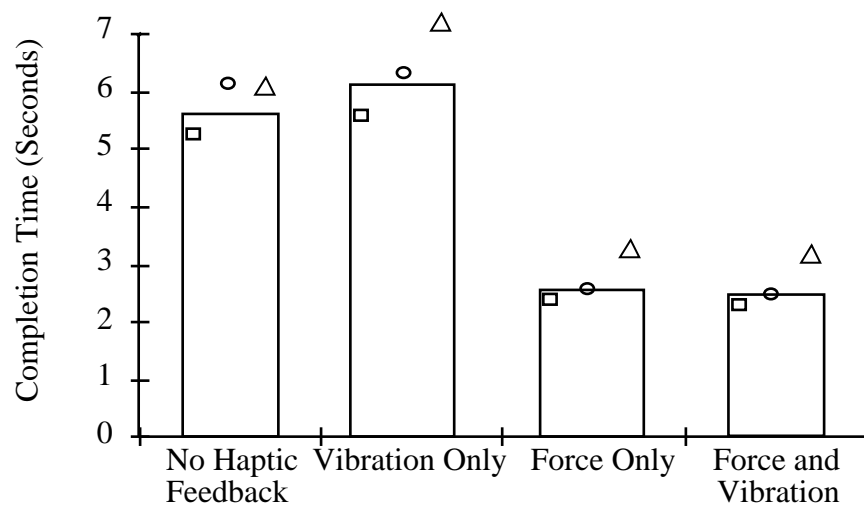


Figure 9.




Figure 10.