Proceedings of the IEEE 1994 International Conference on Robotics & Automation, pp. 3205-10, San Diego, CA, May 1994.

Factors Affecting the Z-Width of a Haptic Display

J. Edward Colgate
J. Michael Brown

Department of Mechanical Engineering
Northwestern University
2145 Sheridan Rd, Evanston, Illinois 60208

Abstract

This paper addresses the performance of force-reflecting interfaces ("haptic displays"). We suggest that an important measure of performance is the dynamic range of achievable impedances — "Z-Width" — and that an impedance is achievable if it satisfies a robustness property such as passivity. Several factors affecting Z-Width — sample-and-hold, inherent interface dynamics, displacement sensor quantization, and velocity filtering — are discussed. A set of experiments designed to evaluate these factors is described, and experimental results are presented. A striking result is that inherent interface damping exerts an overwhelming influence on Z-Width.

1. Introduction

In recent years haptic interfaces (also know as manipulanda and hand controllers) have been developed for an impressive array of applications. For instance, Mussa-Ivaldi et al. [10] describe a two degree-of-freedom manipulandum for studies of multijoint human limb movement, and Adelstein and Rosen [1] a two degree-of-freedom manipulandum for studies of involuntary tremor. Bejczy and Salisbury [2] introduce a six degree-of-freedom hand controller for use in space telerobotics, while Jacobsen, et al. [8] have developed a 22 dof force-reflecting exoskeleton for use in underwater telerobotics. Virtual reality has also provided significant impetus for haptic interface development, as the molecular docking work of Brooks, et al. [3] and the virtual sandpaper system developed by Minsky [9] attest.

A haptic interface may be thought of as a device which generates mechanical *impedances*. "Impedance," here, should be understood to represent a dynamic (history-dependent) relationship between velocity and force. For instance, if the haptic interface is intended to represent manipulation of a point mass, it must exert on the user's hand a force proportional to acceleration; whereas if it is to represent squeezing of a spring, it must generate a force proportional to displacement.

In the physical world, impedances vary widely. For instance, while holding a pencil, the perceived impedance is that of a low mass rigid body, but when pressing a pencil against a writing surface, the perceived impedance is that of a stiff viscoelastic body. In one case, the pencil provides almost no resistance to motion, in the other case almost complete resistance to motion (at least in the direction normal to the surface). The challenge of designing a haptic interface is to build a single programmable device which can exhibit a comparably broad *dynamic range* of impedances (or at least a "Z-Width" which is perceived to be comparably broad).

Our group has for some time been studying the problem of virtual wall implementation as a representative task featuring both very high impedance (when in contact with the wall) and very low impedance (when out of contact). A wall, moreover, is an example of a unilateral constraint — a ubiquitous form of kinematic constraint in the physical world. Because of this, we feel that understanding how to implement a virtual wall which "feels good" and is robust is a problem of fundamental importance in the area of haptic display.

This paper will not address the psychophysics of what makes a virtual wall "feel good" except to say that one important factor seems to be dynamic range. An excellent article on this topic has recently been written by Rosenberg and Adelstein [11]. We will present instead some of our findings, both theoretical and experimental, concerning achievable dynamic range. In short, we will address the question of how to build a haptic interface capable of exhibiting a wide range of mechanical impedances while preserving a robust stability property. We begin by discussing the issue of robustness.

2. Robustness

Both physical and virtual systems of significant complexity are characterized by *interaction*. Indeed, the excitement surrounding virtual reality is due in large part to the promise of interactive capabilities approaching those of the physical world. But here, an important distinction must be drawn. The interaction of physical systems obeys flawlessly a set of underlying laws, while the interaction of virtual systems obeys similar laws only approximately. The consequences of approximate obedience can be profound. For instance, in the physical world we could scarcely even contemplate the possibility that, upon bolting together two steel beams, the entire assembly would exhibit sustained or growing oscillations. But this is precisely what might occur in a virtual world if appropriate laws are not enforced to govern interaction.

To ensure robust interactive behavior, as in the example of the two beams, the physical world relies heavily upon the property of *passivity*. The steel beams are obviously examples of passive systems, neither being able to provide energy to the other. It is well-known that the coupling of passive systems is guaranteed to be stable. It is by no means established that virtual worlds must rely on passivity, but surely some comparable underlying property is essential for stability/robustness.

In principle, this conclusion applies whether or not a virtual environment is connected to a haptic display. Haptic display, however, makes the need for robustness more acute. There are two reasons. The first is that the human tactile sensory apparatus is extremely receptive to small amplitude mechanical vibrations in the 100 Hz - 1 KHz range (while vision is not). The second is that the human is also a dynamic system. Thus, even though a non-passive virtual environment may be stable, interaction with a human via a haptic interface may *cause* instability. In our studies of virtual walls, we have had many experiences with human operators adjusting their own behavior until oscillations resulted.

The robustness property obeyed by virtual systems could, in principle, be passivity. As discussed in Section 3, a haptic interface connected to a virtual environment may indeed be passive. In our studies of virtual walls, we have found that passivity provides an extremely useful intellectual framework for understanding the robustness problem. Therefore, we will appeal to passivity frequently in this paper. We have also, however, found that passivity is too conservative a property to demand of a haptic display. In other words, even if a haptic display (connected to a virtual environment) is not passive, interaction with a human operator may not destabilize it. This point will also be discussed.

3. Sampled data and inherent dynamics

In this section, we will consider the one degree-of-freedom haptic interface shown in Figure 1. It will be represented by the model shown in Figure 2. The objective of the section is to study how both sampling and the inherent dynamics of the display affect the achievable dynamic range. As a point of departure, an impedance will be considered achievable if it can be implemented passively. The following theorem, proven in [5], is useful:

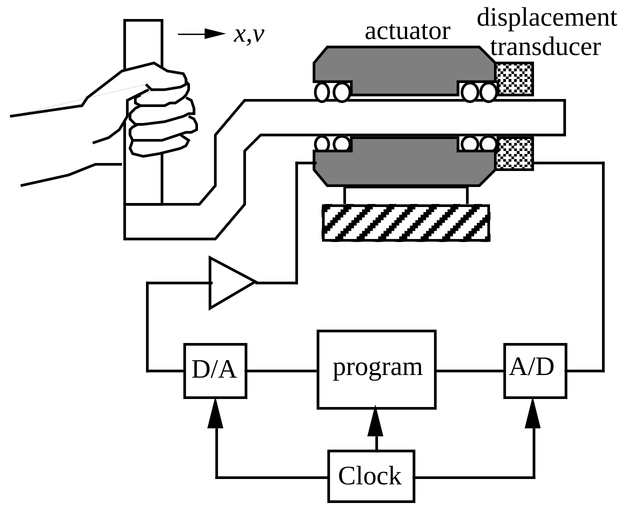

Figure 1. Schematic of a one degree-of-freedom haptic interface.

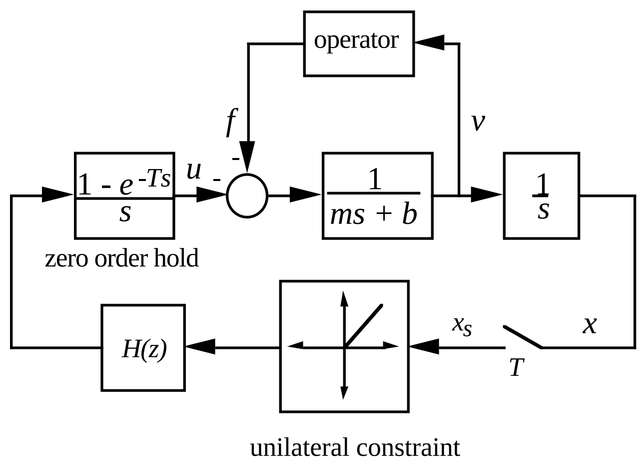

Figure 2. Model of a one degree-of-freedom haptic interface. m is the inherent mass of the display, b is inherent damping, v is velocity, x is position, x_s is the sampled position, T is the sampling rate, u is the control effort, and f is the force applied by the operator.

Theorem — A necessary and sufficient condition for passivity of the haptic interface model in Figure 2 is:

$$b > \frac{T}{2} \frac{1}{1 - \cos \omega T} \mathrm{Re}\left\{ (1 - e^{-j\omega T})\, H(e^{j\omega T}) \right\} \quad \text{for} \quad 0 \le \omega \le \omega_N \tag{1}$$

Here, b is the inherent damping of the display, T is the sampling rate, $H(z)$ a pulse transfer function representing the virtual environment, and $\omega_N = \pi/T$.

The specific case of interest here is the "virtual wall." We will consider a common implementation composed of a virtual spring and damper in mechanical parallel, together with a unilateral constraint operator. A velocity estimate is obtained via backward difference differentiation of position, giving the following transfer function within the wall:

$$H(z) = K + B\frac{z - 1}{Tz} \qquad (2)$$

where $K > 0$ is a virtual stiffness, and B is a virtual damping coefficient (we will allow B to be positive or negative). A condition for passivity is found by inserting (2) into (1). After some manipulation:

$$b > \frac{KT}{2} - B \cos \omega T \quad \text{for } 0 \le \omega \le \omega_N \qquad (3)$$

(3) can be further reduced to:

$$b > \frac{KT}{2} + |B| \qquad (4)$$

The following conclusions may be drawn from this analysis:

- To achieve passivity, some physical dissipation is essential.
- With other variables (b and B) fixed, the maximum achievable virtual stiffness is proportional to the sampling rate.
- The maximum achievable virtual damping for zero stiffness is independent of the sampling rate.
- With other variables (b and T) fixed, higher virtual stiffnesses can be achieved at *lower* values of virtual damping.

While these are potentially useful guidelines for design, it is important to recognize that passivity is a conservative design requirement. This is because, even if an interface is active, a human operator may not be able to destabilize it. As an example, we rarely find instances in which the instability caused by coupling to a human operator exceeds 150-200 Hz, but examination of (3) shows that, for positive B, passivity is most readily violated at the Nyquist frequency, which may be 500 Hz or greater. At the Nyquist frequency, the factor $\cos \omega T$ is negative, but at less than half the Nyquist frequency, it is positive. If it is positive at the frequency of instability, then the implications of the last bullet point above are completely reversed: higher virtual stiffnesses can be achieved at *higher* values of virtual damping.

With the above caveat in mind, the implications for haptic interface design can be discussed. For instance, to implement stiff, dissipative walls (high K, B), it is apparently helpful to maximize b and minimize T. Fast sampling is a standard objective, but maximizing damping goes against conventional wisdom [7]. It is generally argued that the dynamics of a haptic interface should be dominated by the virtual environment (which is, after all, the programmed behavior we wish to display) rather than any inherent dynamics (which is considered parasitic). In other words, the interface hardware should be "transparent." Unfortunately, the notion of transparency places focus on mimicking the governing equations (e.g., state equations) of physical systems, but not on obeying underlying physical laws (such as conservation of energy). Adding physical damping helps the sampled-data system to behave as physical law would dictate.

But is there a cost to additional damping? Is the behavior *inside* the wall improved at the cost of the behavior *outside* the wall? The answer is no. The reason for this answer is that, as seen in equation 4, one may introduce negative virtual damping outside the wall. In fact, since $K = 0$, one may select $B = -b$, resulting in zero net damping (although this is borderline passive, and perfect cancellation is difficult to achieve in practice). This importance of physical damping will be further elaborated in the discussion of experimental results.

4. Sensor quantization and velocity filtering

One of the more commonly used position sensors in haptic displays is the optical encoder. Encoders are reasonably rugged and easy to interface, and are extremely linear and free of dynamics. Unfortunately, the output of an encoder is quantized, and it is well-known that quantization can lead to limit cycles in digital control systems [4]. Of course, the angle quanta are typically quite small and therefore cause little practical problem, *unless* the quantized signal is differentiated (e.g., to obtain a velocity signal for virtual damping).

Differentiation is notorious for amplification of high frequency noise. In the context of an encoder-based, sampled-data control system, the consequences of differentiation are easily understood. Suppose that a quantum is Δ radians and the sampling period is T seconds. Then the resolution of a finite difference differentiator is Δ/T rad/sec. If an 8000 count/rev encoder is used and T = 0.001 sec, then the velocity resolution is 45°/sec! If we further assume a 0.1 m lever arm, then the smallest measurable translational velocity at the tip of this lever arm is 7.8 cm/sec. Clearly, this is an extremely high velocity compared to that which would be desirable when contacting a wall.

How can a better velocity estimate be obtained? One way to improve resolution is to sample more slowly! Unfortunately, this runs contrary to the goal of high stiffness as discussed above. Another approach is to filter the velocity estimate digitally. This will be discussed briefly below. A third approach is to use higher resolution encoders. Simulation experience suggests that encoder resolution has little effect on the *existence* of limit cycles, but considerable effect on the amplitude of limit cycles [6]. A fourth approach is to use analog sensors (for position, velocity, or both), although these sensors suffer from noise as well.

As an aid in understanding the effects of filtering, consider a virtual wall implementation in which the differentiator is cascaded with a first order low pass filter of time constant τ. If a backwards difference mapping is used, the transfer function of the wall is:

$$H(z) = K + B\frac{z-1}{(\tau + T)z - \tau} \quad (5)$$

It is easily shown that the resolution of the filtered differentiator is $\Delta/(T+\tau)$. Thus, the slower the filter, the better the velocity resolution, as might be expected. In practice, we generally find that better than an order of magnitude improvement may be obtained in resolution with no obvious performance cost (the experiments described below provide a quantitative assessment).

One might also expect that the cost of filtering would be that the haptic display becomes less passive. In general, this is true because filters introduce delay. For the first order filter considered here, however, the condition for passivity is considerably less restrictive than without a filter (compare to equation 4):

$$b > \frac{KT}{2} + \frac{BT}{2\tau + T} \qquad (B \geq 0) \quad (6a)$$

$$b > \frac{KT}{2} - B \qquad (B < 0) \quad (6b)$$

5. Experiments

Experiments were performed using a one degree of freedom manipulandum. The device is powered by a DC brushless motor mounted to a sturdy table so that the motor shaft points upward. Attached to the shaft is a crank handle ($r \approx 0.15$m) which the user may grab with his/her hand . The motor shaft is also coupled via steel tape to a rotary viscous damper ($b \approx 0.22$ Nm-sec/rad). The motor shaft is equipped with two encoders (8000 cpr and 900,000 cpr) for position sensing. Motor currents are supplied by a PWM amplifier, and voltage inputs to the amplifier are provided by a 12-bit A/D converter.

To describe the range of achievable impedances (Z-Width), objective criteria for success had to be developed. An impedance was considered achievable if a person could not elicit visually apparent sustained oscillations from the manipulandum handle (see Figure 3 for comparison).

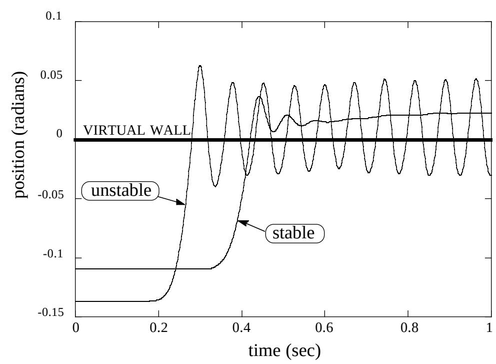

Figure 3. Samples of stable and unstable interactions with a virtual wall.

Preliminary experiments showed that results varied greatly with the type of grip used by the subject. Generally, it is much easier to destabilize a virtual wall when pushing against the handle with one finger than when enveloping the handle in a full-fingered grip. To ensure consistent results, subjects were required to place four fingers on the handle at all times.

Subject learning was embraced rather than avoided. Since normal interaction with a haptic interface would include learning, we felt it important to keep that aspect intact. Subjects were allowed to familiarize themselves with the device before the experiment began. Once they felt comfortable interacting with and destabilizing walls, we started collecting data. With each new configuration, the subjects were allowed practice time to adjust to the new settings. Subjects were given as much time and as many attempts as they desired to generate instability, so that a given trial was ended only when the subject labeled the interaction as stable or unstable.

To eliminate fatigue, subjects were not allowed to work with the device for more than one hour at a time, with several hours rest before starting again. They were allowed to progress through the trials at what they considered a comfortable pace. To ensure consistency of results, some parameter values were repeated at the end of each configuration. Certain configurations were also repeated for the same reason. Three subjects were used to gauge how parameters varied from person to person.

For each of the four factors discussed above, two conditions were examined:

Damper	engaged	disengaged
Sampling rate	high (1 KHz)	low (100 Hz)
Encoder resolution	high (900K cpr)	low (8K cpr)
Velocity filter	first order, 30 Hz cutoff	none

Sixteen "configurations" (combinations of the above conditions) were possible. All of these configurations were studied for each subject. Within each configuration, the maximum achievable stiffness was found for the entire range of achievable damping. These data lead to plots of maximum virtual stiffness (K) vs. maximum virtual damping (B), showing the Z-width for different configurations of the device.

Results — Figures 4a-d show sixteen Z-width plots for one of the subjects. These figures make it clear that physical damping can play a pivotal role in increasing the Z-width, regardless of the configuration. In all cases, the addition of physical damping increased both the maximum stiffness and the maximum damping. The figures also show that in order to achieve high stiffness, high update rate is needed. However, high update rates exacerbate noise due to differentiation of the position signal, making large damping coefficients difficult to achieve (Figures 4a,b). To achieve higher damping, the update rate can be slowed down (at the expense of stiffness) or a digital filter can be used to attenuate high frequency noise. With the proper digital filter, the velocity signal can be smoothed out to allow large damping in addition to high stiffness (Figure 4c).

Limitations — We believe that, if a virtual environment is intended to emulate a physical counterpart, successful implementation involves three steps :

1) Eliminate gross instabilities (i.e. determine the Z-width where gross instabilities are absent). Subjects can

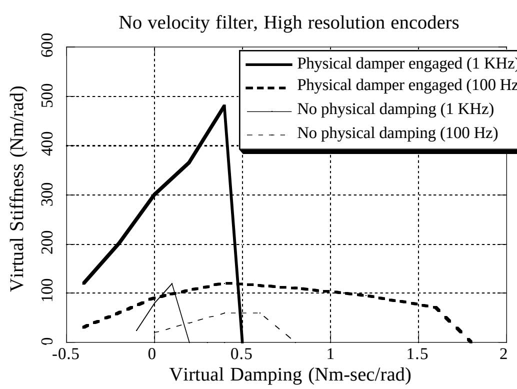

4a.

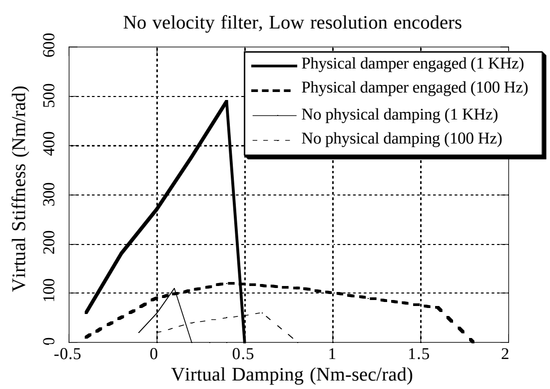

4b.

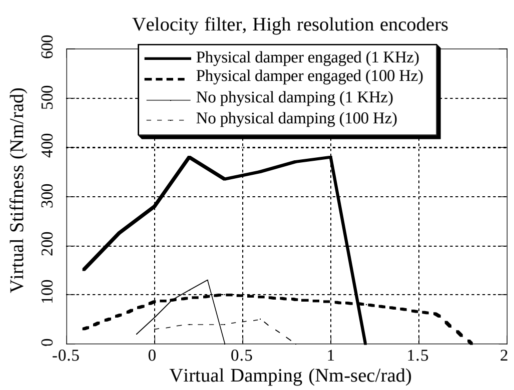

4c.

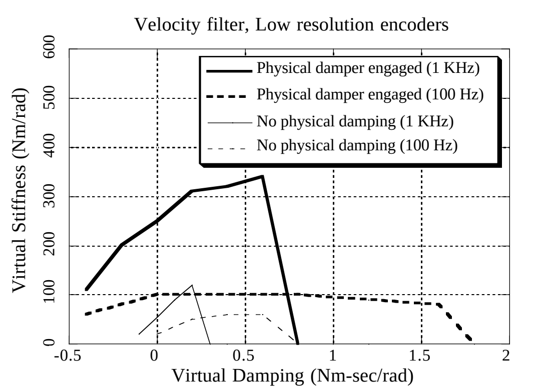

4d.

Figure 4a-d. Z-width for one subject (sixteen haptic interface configurations). Note how physical damping improves performance for both virtual stiffness and damping in all cases. Also note in parts c and d how digital filtering helps increase virtual damping at high update rates.

successfully interact with a virtual environment that passes through this step. However, the subjects may be able to distinguish between it and a physical system.

2) Eliminate awareness of high frequency oscillations (limit cycles). While not dangerously unstable, limit cycles detract from the illusion of a physical environment. As an example, our results show little effect due to encoder resolution, but our subjects reported a significantly better “feel” with the higher resolution encoders. At the end of this step, the virtual environment should feel like a physical system, though not necessarily the desired one. Virtual environments that pass through steps one and two are "implementable."

3) Perform psychophysical experiments to match the virtual environment to the desired physical one. Virtual environments that pass all three steps should provide a high degree of realism.

The experiments reported here address the first step in the process. The second step is the subject of current research. The main difficulty with this step is that high frequency oscillations are always present, but do not always detract from the illusion of a physical environment. Thus, careful psychophysical experimentation is needed. The final step is the subject of future research.

6. Conclusions

Factors affecting the dynamic range of haptic displays have been discussed. Toward achieving very high impedances, the following suggestions are made:

- Maximize inherent damping. In our experience, this is the least expensive and highest payoff measure available. Negative virtual damping may be used to extend the lower limit of impedance.
- Maximize sensor resolution. This is particularly important if position measures will be differentiated to provide a velocity estimate.
- Maximize sampling rate (with the caveat that faster sampling exacerbates the velocity estimation problem unless appropriate filtering is used).
- Filter the velocity signal. A first order low pass filter improves subject impressions of wall quality. We have not yet attempted to design an “optimal” filter.

7. Acknowledgments

The authors gratefully acknowledge the support of the National Science Foundation, grants MSS-9022513 and IRI-9213234, as well as valuable discussions with Paul Millman, Gerd Schenkel, and Mike Stanley.

8. References

1. Adelstein, B.D. and M.J. Rosen. *Design and Implementation of a Force Reflecting Manipulandum for Manual Control Research.* ASME Winter Annual Meeting. Anaheim, California. pp. 1-12, 1992.
2. Bejczy, A. K. and J. K. Salisbury. *Controlling Remote Manipulators Through Kinesthetic Coupling.* Computers in Mechanical Engineering, vol. 2, no. 1, pp. 48-60, 1983.
3. Brooks, F. et al. *Haptic Displays for Scientific Visualization.* Computer Graphics, vol. 24, no. 4, pp. 177-185, 1990.
4. Claasen, T., W. F. G. Mecklenbrauker and J. B. H. Peek. *Frequency Domain Criteria for the Absence of Zero-Input Limit Cycles in Nonlinear Discrete-Time Systems, with Applications to Digital Filters.* IEEE Transactions on Circuits and Systems, vol. CAS-22, no. 3, pp. 232-239, 1975.
5. Colgate, J. E. and G. G. Schenkel. *Passivity of a Class of Sampled-Data Systems: Application to Haptic Interfaces.* To appear in Proceeding of the American Control Conference, 1994.
6. Colgate, J. E., M. C. Stanley and G. G. Schenkel. *Dynamic Range of Achievable Impedances in Force Reflecting Interfaces.* Telemanipulator Technology and Space Robotics. Boston. SPIE, 1993.
7. Fischer, P., R. Daniel and K. V. Siva. *Specification and Design of Input Devices for Teleoperation.* IEEE International Conference on Robotics and Automation. Cincinnatti, Ohio. pp. 540-545, 1990.
8. Jacobsen, S. C., et al. *Design of a Multiple Degree of Freedom, Force Reflective Hand Master/Slave with a High Mobility Wrist.* Third Topical Meeting on Robotics and Remote Systems. Charleston, SC. 1989.
9. Minsky, M., M., et al. *Feeling and Seeing: Issues in Force Display.* Computer Graphics, vol. 24, no. 2, pp. 235-243, 1990.
10. Mussa-Ivaldi, F. A., N. Hogan and E. Bizzi. *Neural, Mechanical, and Geometric Factors Subserving Arm Posture in Humans.* Journal of Neuroscience, vol. 5, no. 10, pp. 2732-2743, 1985.
11. Rosenberg, L. B. and B. D. Adelstein. *Perceptual Decomposition of Virtual Haptic Surfaces.* IEEE Symposium on Research Frontiers in Virtual Reality. San Jose, CA. 1993.