

# **PART I**

## **INTRODUCTION**

# CHAPTER 1

## HAPTIC DISPLAY

This thesis addresses the “haptic display” of virtual environments. A haptic display is a device which allows a human operator to manipulate and feel virtual environments. Virtual environments, as considered in this thesis, may be thought of as computer simulations with which a human operator can interact in real-time, and which create for the human operator the illusion of interaction with a physical environment. “Interaction”, in this thesis, generally refers to haptic interaction; i.e., involving the hand and arm.

Haptic display is an emerging technology [31] with potentially important applications in human operator training and in design prototyping [1,9,12,16,26,27]. At the time of this writing, however, very few of these applications have been realized. One of the reasons for this is simply the state of the display technology. Most available haptic displays cannot generate convincing illusions of even simple environments (e.g., virtual walls), and no available display has demonstrated a convincing illusion of a complex environment, such as a hand tool. Displays are hampered by limitations in mechanical design, electromechanical interface, computer hardware, software development tools, and simulation algorithms.

This thesis focuses principally on simulation algorithms. Inappropriately designed algorithms contribute significantly to one of the major shortcomings of haptic display: instability. “Instability”, here, refers to mechanical instability of the operator/haptic display/virtual environment system at times when the associated operator/physical environment system would be stable. The long term objective to which this thesis will contribute is that of *guaranteed stability* in the haptic display of virtual environments.

## 1.1 Instability and Passivity

Instability is a property of the operator/display/simulation *system*. Unfortunately, at least two pieces of this system, the operator and the simulation, are very complicated. The human operator, in addition to being a complex dynamic system, is highly uncertain and difficult to model. Therefore, to make analyses tractable, a standard approach is to require that the display/simulation subsystem appear *passive* to the user. In other words, it should be impossible for the operator to extract energy from the haptic display on a continuing basis. There are, however, many reasons why a haptic display might not be passive.

The basic difficulty with implementing passive virtual environments may be understood with the example of a virtual spring. An ideal physical spring is a lossless system; therefore, if energy is stored in the spring by squeezing, then removed by releasing, precisely as much energy will be removed as was stored. Now consider the virtual spring. Because it is implemented in discrete time, the force provided by the spring will not increase smoothly with deflection. Instead, the force will be repeatedly “held” at a

constant value until updated<sup>1</sup>. Because of this, the average force during squeezing will be slightly less than for a physical spring of identical stiffness, and the average force during release will be slightly greater. This is illustrated in Figure 1. The result is not only that the exquisite balance of stored energy and released energy is lost, but that the spring always acts to store or *generate* energy, never to dissipate energy.

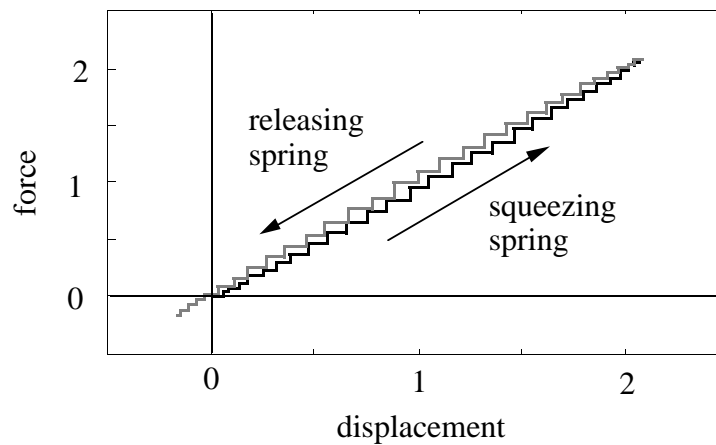


Fig. 1. Illustration of energy generation by a virtual spring. Energy is given by the area under the curve.

Of course, any real implementation of a virtual environment includes some dissipation. It can be shown, however, that any implementable discrete-time simulation of a virtual dissipator will be capable of producing energy via a haptic display. This stems from the fact that there is always delay and information loss in the sample-compute-hold process.

Another example of a passive physical behavior which can cause energy generation when experienced via a haptic display is *unilateral constraint*. Unilateral constraints, which

<sup>1</sup>This assumes a zero-order hold. Similar difficulties exist for other types of holds.

occur whenever two rigid bodies collide, cannot act as energy sources. Thus, two physical systems which are stable in isolation and also stable when rigidly coupled together, are stable when allowed to collide. The fact that collisions do not generate energy has been used to advantage, for instance, in modeling rigid body contact [54]. For discrete time (or sampled-data) systems, however, it is possible that unilateral constraints will be directly responsible for instability. This problem will be further addressed in Chapter 3.

One conclusion that can be drawn from the above is that passivity is a hopeless objective unless the haptic display incorporates some amount of inherent, physical damping. This point has been addressed at length in [13], and will be further reviewed in Chapter 2. Even more importantly, however, it must be understood that, in the conversion from continuous time to discrete time, governing laws (e.g., energy conservation, Tellegen's Theorem) may be lost. It is critical that these laws, and not just the accuracy of the equations, be considered in implementation.

## 1.2 Specific Objectives

This thesis addresses two problems that arise in attempting to ensure the passivity of a haptic display.

- Unilateral constraints are ubiquitous in physical interaction. In physical environments, unilateral constraints do not generate energy; however, in virtual environments experienced via haptic display, unilateral constraints may generate energy. The first objective of this work is to find useful, non-conservative

stability criteria for sampled data feedback systems (i.e., virtual environment / haptic display systems) involving unilateral constraints.

- The real-time simulation of multibody dynamic systems is a challenging task. One way to achieve sufficiently high update rates is to employ multiple, parallel processors. A difficulty with this approach, however, is that interprocessor communication delays make passivity difficult to guarantee. The second objective of this work is to develop an approach to the design of computational networks which is isomorphic to physical network modeling, and which ensures passivity despite communication delays.

## CHAPTER 2

### BACKGROUND

This chapter provides a brief review of the coupled problems of *stability* and *performance* of haptic displays. The following two design objectives underlie this discussion:

- 1) A haptic display should be free of instability and high frequency oscillations (limit cycles) [37]. Instability destroys the illusion of interaction with a physical environment, while limit cycles mask it.
- 2) A haptic display should be able to implement a high dynamic range of impedances (high “Z-width” [11]). Consider the example of a pencil. When moved in free space, the impedance of the pencil is very small, but when pressed against a tabletop, the impedance is very large in the normal direction. A successful haptic display should therefore be able to span a large range of impedances.

An effective way to understand the issues of stability and performance is to consider the benchmark problem of a “virtual wall”. Figure 2.1 is a block diagram of a

single-axis system for the haptic display of a virtual wall. The system, as considered here, includes the human operator, and has been discussed in depth elsewhere [11,13].

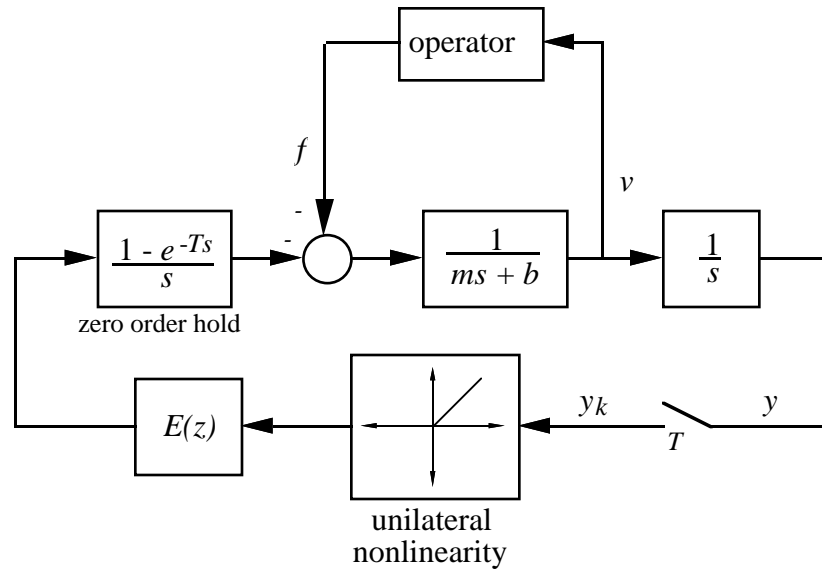


Fig. 2.1. Simplified model of a haptic display system. Here,  $m$  is the inertia and  $b$  is the damping, while  $y$  is the position and  $v$  is the velocity of the haptic display hardware.  $f$  is the total force generated by the operator.  $T$  is the sampling period of the virtual environment simulation, and  $E(z)$  is a pulse transfer function representing the behavior within the virtual wall.

This particular model has been designed to be only as complicated as necessary to display the behavior of interest. It should not be considered an accurate or complete physical model. Indeed, it is a simple enough model that analytical conditions can be found for the passivity of the display/simulation subsystem (i.e., not including the human operator). These conditions are [13]:

$$b > \frac{T}{2} \frac{1}{1 - \cos \omega T} \operatorname{Re} \left( 1 - e^{-j\omega T} \right) E(e^{j\omega T}) \quad \text{for } 0 \leq \omega \leq \omega_N \quad (2.1)$$

where,  $b$  is the inherent damping of the display,  $T$  is the sampling time,  $E(z)$  a pulse transfer function representing the virtual environment, and  $\omega_N = \pi/T$  is a Nyquist frequency.

The “virtual wall” is usually implemented as a composition of a virtual spring and damper in mechanical parallel, together with a unilateral constraint operator [10]. A velocity estimate is obtained via backward difference differentiation of position, giving the following transfer function within the wall:

$$E(z) = K + B \frac{z - 1}{Tz} \quad (2.2)$$

where  $K > 0$  is a virtual stiffness, and  $B > 0$  is a virtual damping coefficient. It is interesting to note that this is a discrete-time passive transfer function; yet, the sampled-data system including the haptic display and sample-and-hold will be passive for certain parameter values only. The sampled-data passivity condition can be found by inserting equation (2.2) into equation (2.1) and doing some algebraic manipulation [13]:

$$b > \frac{KT}{2} - B \cos \omega T \quad 0 \leq \omega \leq \omega_N \quad (2.3)$$

This can be maximized at  $\omega = \omega_N$  to arrive at:

$$b > \frac{KT}{2} + |B| \quad (2.4)$$

From this equation, we can draw certain conclusions:

- 1) To achieve passivity, some physical dissipation ( $b$ ) is essential.
- 2) Given fixed physical and virtual damping, the maximum achievable virtual stiffness is proportional to the sampling rate.
- 3) The achievable virtual damping is independent of the sampling rate.

This example shows how stability and performance (Z-width) are interrelated. Instability limits the maximum achievable impedance. Effective measures for increasing the dynamic range are increasing the sampling rate, and increasing the inherent physical damping. In [11], Brown and Colgate show the efficacy of both of these measures, as well as the importance of sensor resolution and velocity signal filtering.

Salcudean and Vlaar [47] have also studied the stability of virtual walls. They have worked principally with magnetically levitated haptic displays, modeled as shown in Figure 2.2. Rather than study passivity (their implementation cannot be passive, due to the absence of damping), they have studied the stability of the display/simulation system alone. Their virtual wall is implemented as a digital PD controller:

$$f_k = -Ky_{k-1} - \frac{B}{T}(y_{k-1} - y_{k-2}) \quad (2.5)$$

where  $f_k$  is force exerted by the wall,  $y$  is the displacement into the wall,  $K$  and  $B$  are the stiffness and damping of the wall, and  $T$  is the sampling period. By ignoring the unilateral nonlinearity, a stability region is found in terms of normalized stiffness ( $T^2K/2m$ ) and damping ( $TB/2m$ ). Perhaps more interestingly, they suggest that the perceived stiffness of the virtual wall can be increased by adding a “braking pulse” to the above expression upon initial contact. The height of this pulse is velocity dependent; it can, therefore, be viewed as a form of “skin damping” that augments the wall damping during initial contact.

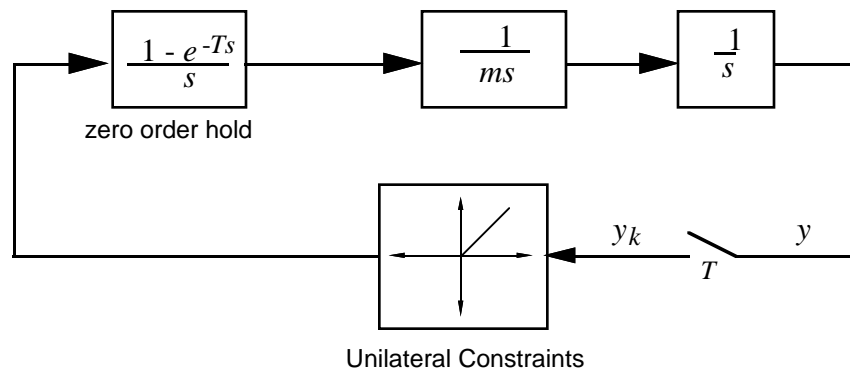


Fig. 2.2. Model of a magnetically levitated haptic display

In conclusion, the stable, high performance haptic display of virtual environments is a mix of electromechanical design, software design, parameter tuning, and psychophysics. Very little is known about how to extend the above results to complex environments, or even how to account for simple nonlinearities like the unilateral nonlinearity that is so important in modeling collisions. One encouraging step was recently reported by Colgate, et al. [14], who showed that their result for the passivity of a virtual wall could be extended to the passivity of a broad class of linear virtual environments, if

properly implemented. Their approach was to constrain the virtual environment to be discrete-time passive, then to connect this environment to the haptic display via a “virtual coupling” consisting of a virtual spring ( $K$ ) and damper ( $B$ ). They showed that the conditions for passivity of the display/simulation system would then be the same as for the display of a virtual wall having stiffness  $K$  and damping  $B$ . This suggests that, rather tune many parameters in a complex simulation, it would be profitable to ensure simply that the simulation is discrete-time passive. This is a much more tractable problem.

The remainder of this thesis addresses two issues addressed above. First, the problem of unilateral nonlinearities is taken up; then, the problem of discrete-time passivity is taken up.