# M-Width: Stability and Accuracy of Haptic Rendering of Virtual Mass

Nick Colonnese
Mechanical Engineering
Stanford University
Stanford, California 94305-2232
Email: ncolonnese@stanford.edu

Allison Okamura
Mechanical Engineering
Stanford University
Stanford, California 94305-2232
Email: aokamura@stanford.edu

*Abstract*—In many physical human-robot interaction scenarios, such as haptic virtual environments for training and rehabilitation, it is desirable to carefully control the apparent inertia of a robot. Inertia compensation can be used to mitigate forces felt by the user during free-space motion, and rendering of additional inertia is desired for particular rehabilitation and training procedures. Many factors influence the stability and accuracy of rendering for haptic display of a pure mass, including device mechanical properties, sample rate, control structure, and human behavior. Inspired by the "Z-Width" approach to haptic device stability and performance analysis, we introduce "M-width", which we define as the dynamic range of virtual masses renderable in a stable manner. We identify the important parameters for system stability, find stability boundaries, and describe the expected accuracy of the haptic rendering for a canonical haptic system. These results serve as a design tool for creating stable and accurate haptic virtual environments, establish limits of performance, and lay the groundwork for new controllers to improve mass rendering.

## I. INTRODUCTION

In many applications of haptic virtual environments, we desire the user to feel a very specific set of mechanical properties, which arise from inherent device dynamics, programmed virtual environment dynamics, and (potentially unexpected) interactions between the two that depend on control structure, sampling rate, and other system properties. Consider the example application shown in Figure 1. Here, researchers wish to test the hypothesis that damage to the cerebellum (a region of the brain that plays an important role in motor control) holds a dynamic model of the body used in feed-forward planning of movements. Patients with cerebellar damage who exhibit ataxia (uncoordinated movements) are asked to make fast reaching movements in a backdrivable, planar exoskeleton robot. If a rendered inertia could be found that eliminates the ataxia, this would provide significant evidence toward the theory that cerebellar damage results in a specific bias in internal modeling of body dynamics, and it would open up a host of potential new rehabilitation therapies for patients with ataxia. However, modifying the effective inertia of the user/robot combination to be a specific value evokes fundamental challenges in haptic device control.

Haptic interfaces should be simultaneously stable and accurate. Unstable behavior can cause damage to a haptic device, injure the human user, and generate unrealistic and unexpected
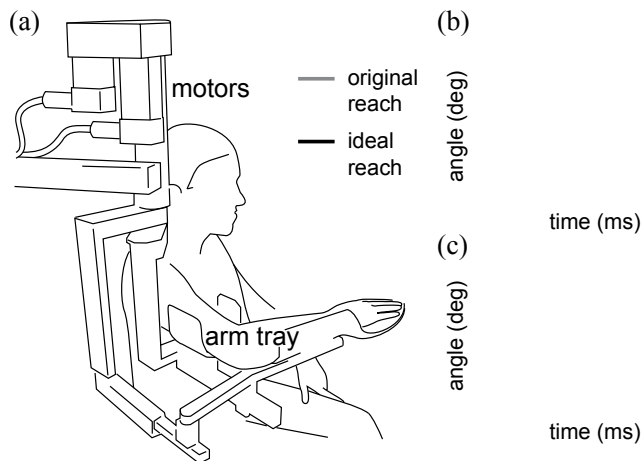


Fig. 1. Example application of inertia rendering in rehabilitation. (a) A user makes single-joint reaches to a 30° target in an exoskeleton robot, during which rendered inertia alters his movement. (b) A patient with hypermetria (tendency to overshoot) is predicted to benefit from decreased effective arm inertia. (c) A patient with hypometria (tendency to undershoot) is predicted to benefit from increased effective arm inertia. While neuromechanical models predict these effects, implementing inertia rendering requires improved understanding of system stability and accuracy.

haptic sensations. Accurate haptic virtual environments should feel exactly as desired, with no unwanted effects of haptic device dynamics (e.g., the inherent inertia of the mechanism). In the rehabilitation example given above, stability and accuracy are both necessary for the scientific validity of the study. Yet, there exists a classic tradeoff between ensuring the stability of a haptic display and the range of impedances that it can render. Mass rendering (presenting a force proportional to acceleration) presents a particular challenge because measuring acceleration on a typical impedance-type haptic device is noisy; it involves double discrete differentiation of a noisy position signal. While system improvements (e.g., the addition of new sensors) can improve mass rendering, an improved understanding of the performance of a canonical haptic device will yield insight as to what modifications in hardware, sensing, and control will improve mass rendering. The long-term goal of this work is to develop stable and accurate mass rendering for rehabilitation applications, and predict the limits of performance. Here, we begin by developing a framework

for analyzing the stability and accuracy of mass rendering with theory applicable to a one-degree-of-freedom linear system.

Colgate and Brown [5] define "Z-Width" as the dynamic range of impedances that can be rendered passively. Passive systems are incapable of generating a net amount of energy, and the coupling of passive systems is guaranteed to be stable. However, conditions for passivity can be conservative compared to conditions for stability. Although Z-width is a general term spanning all impedances, it has generally been discussed as relating to haptic rendering of virtual springs and dampers. In addition, Colgate and Schenkel [6] derived a general theorem for the passivity of haptic interfaces.

This paper introduces "M-Width", the dynamic range of virtual mass renderable in a stable manner. We allow for positive virtual mass, corresponding to mass rendering, as well as negative virtual mass, corresponding to mass compensation. The definition of M-width differs from Z-width in three key ways: (1) we consider BIBO (bounded input, bounded output) stability, not passivity – to avoid an overly conservative implementation, (2) it models the human operator as a specific impedance, not a general passive element, and (3) the desired virtual environment is a pure mass with motion data filtering.

This work builds upon significant prior work related to the stability of haptic displays, e.g. [1, 6, 7, 10]. Virtual mass rendering has been previously explored for specific medical robotics scenarios, including dynamic compensation for a surgical teleoperator [14], and inertia compensation for a lower-leg exoskeleton [3]. Gil et al. [11] explore inertia compensation by force feedforward for an impedance haptic device using additional sensors for acceleration estimation. Adams and Hannaford [2] present a general virtual coupling approach that could be used to render mass that is guaranteed to be passive, but do not consider mass explicitly. Hannaford and Ryu [12] extend the passivity domain, at the necessary expense of reducing the quality of haptic rendering. Brown and Colgate [4] establish conditions for passive positive mass simulations.

While recognizing that many scenarios require treatment of nonlinear models, in this paper we consider linear models – which are relevant to the problem of dysmetria (under- or overshooting targets) in single-jointed movements, and enable the use of classic linear control theory.

## II. System Models

In this section we introduce system models for a human interacting with a haptic device, where the goal is to render a pure mass. There are many different haptic system architectures that can be used to render inertia, but in this paper we are interested in an impedance-type device with position sensing only. Three system models are introduced. One is a hybrid model, containing both continuous and discrete elements. This model is the closest to reality, in that it captures the control of a physical haptic device through a computer containing A/D (analog to digital) and D/A (digital to analog) components. Entirely continuous and discrete models are also introduced. The continuous model is used for the identification
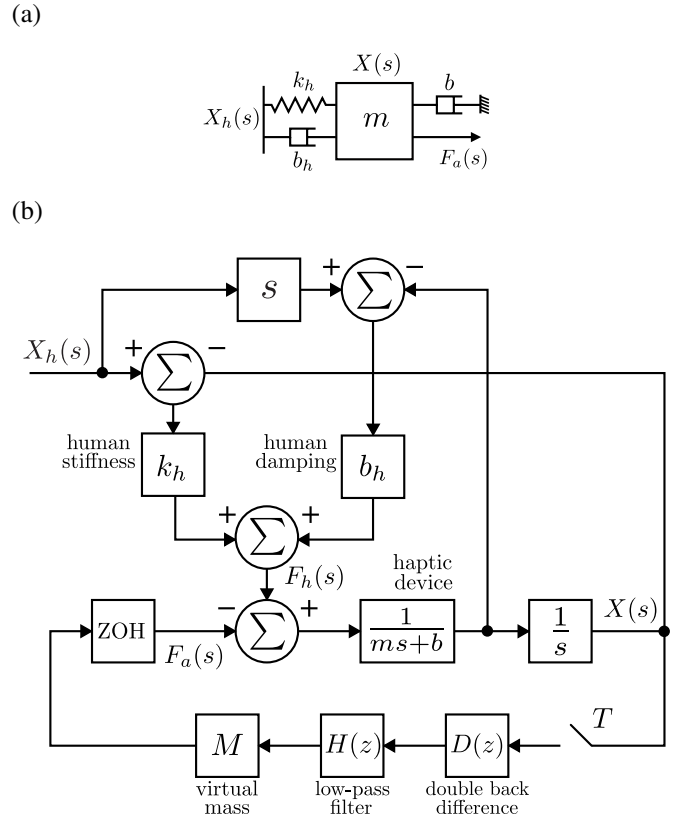
(a)

(b)

Fig. 2. (a) Schematic of a human interacting with a haptic device implementing virtual mass. (b) Block diagram for this system, considering human parameters, device dynamics, sampling, and ZOH.

of important parameters with respect to stability, and for use in a quality of rendering analysis. We show that the discrete model results in stability bounds that are identical to the hybrid model stability bounds.

We consider a haptic device described by a mass, $m$, and viscous damper, $b$, acted upon by two external forces: the force applied by the operator, $F_h(s)$, and the force applied by the actuator implementing the virtual environment, $F_a(s)$ (Figure 2). The system is equipped with only one sensor measuring the position of the mass, $X(s)$. An estimate of the acceleration of the device is formed by performing a double back difference on the sampled position signal and then low-pass filtering. The actuator force is the product of the acceleration estimate and the desired virtual mass.

### A. Hybrid Model

The system equation of motion is:

$$F_h(s) - F_a(s) = X(s)(ms^2 + bs). \tag{1}$$

The force of the operator is determined by the motion of the operator $X_h(s)$ via the hand impedance,

$$F_h(s) = b_h(sX_h(s) - sX(s)) + k_h(X_h(s) - X(s)), \tag{2}$$

where $b_h$ and $k_h$ are positive constants corresponding to the hand damping and stiffness, respectively. The continuous

position signal $X(s)$ is sampled with a constant sampling period of $T$ to get a sampled position $X(k)$. The sampled position measurements go through a discrete double back differencing operator,

$$D(z) = \frac{(z-1)^2}{(Tz)^2}, \tag{3}$$

and then a discrete low-pass filter,

$$H(z) = \frac{(1 - e^{-\omega_0 T})z}{z - e^{-\omega_0 T}}, \tag{4}$$

to form an estimate of the acceleration signal. The discrete low-pass filter has cut-off frequency $\omega_0$ (rad/s). The low-pass filter exists for practical considerations because the double differentiation of the sampled position signal is expected to be noisy. Other filters could be used; we chose one of the simplest possible filters to provide information about baseline performance. In a Z-width analysis, larger virtual damping enables the stable display of larger virtual environment stiffness. Here, a greater degree of filtering enables the stable display of a larger span of virtual masses. We will see, however, that filtering can result in inaccurate mass rendering.

The force of the actuator is the product of the acceleration estimate and a desired virtual mass, $M$, which can be positive (mass rendering), or negative (mass compensation). The actuator force is held constant for the duration of the sampling period with a zero-order hold (ZOH), resulting in a continuous-time staircase actuator force. At sample time $k$ the actuator force is

$$F_a(k) = MH(z)D(z)X(k). \tag{5}$$

### B. Continuous Model

It is possible to convert the general hybrid system shown in Figure 2 into an entirely continuous one. The continuous model is desirable because of its tractability for finding important parameters with respect to stability using conventional linear control systems analysis. The two systems are similar, but not identical. To represent the system continuously, the discrete elements are converted to continuous ones. The continuous representation of $D(z)$, $D(s)$, is found using bilinear (Tustin) mapping:

$$D(z) \rightarrow D(s) = \frac{4s^2}{(sT+2)^2}. \tag{6}$$

The continuous representation of $H(z)$, $H(s)$, is a first-order low-pass filter with cut-off frequency $\omega_0$ (rad/s) with unity gain at DC.

$$H(z) \rightarrow H(s) = \frac{\omega_0}{s + \omega_0}. \tag{7}$$

The zero-order hold is modeled as a time delay equal to half the sample period.

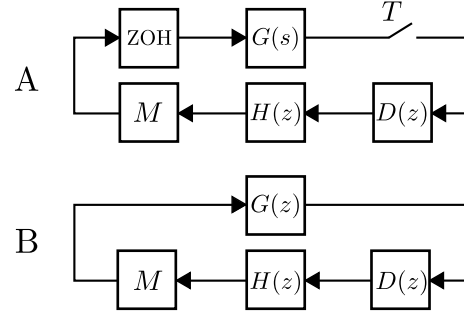$$\text{ZOH} \rightarrow e^{-\frac{sT}{2}}. \tag{8}$$



Fig. 3. Steps for converting the hybrid system into an entirely discrete system. A: From the hybrid model shown in Figure 2, we find a continuous transfer function $G(s)$. B: Then, the continuous elements are converted into an equivalent element $G(z)$, resulting in an entirely discrete system.

### C. Discrete Model

It is also possible to represent the system with an entirely discrete model. The discrete model is useful for stability analysis in which the discrete elements can be considered explicitly. Figure 3 illustrates the process.

First, a transfer function from $F_a(s)$ to $X(s)$, $G(s)$, is formed combining only the continuous elements into a single transfer function:

$$\frac{F_a(s)}{X(s)} = G(s) = \frac{1}{ms^2 + s(b + b_h) + k_h}. \tag{9}$$

Then, the ZOH, $G(s)$, and the sample elements are converted into a discrete element, G(z) using a zero order hold equivalent [9]

$$G(z) = \frac{z-1}{z} \mathcal{Z}\left\{ \frac{G(s)}{s} \right\}. \tag{10}$$

The input-output characteristics of the continuous and discrete elements are the same. For this reason, the stability of the hybrid and discrete models are identical.

### III. SYSTEM STABILITY

We now analyze the stability of the system shown in Figure 2, considering the all-continuous and discrete representations of the hybrid system from which BIBO (bounded input, bounded output) stability can be determined. First, we identify important parameters with respect to stability using the continuous model. This is useful as a haptic system design tool because it identifies what key parameters determine mass rendering stability. Then, we form quantitative stability boundaries using the discrete model equivalent to those of the hybrid model. This establishes the limits of a haptic system rendering mass using this architecture. Finally, we discuss how the continuous and discrete models differ.

### A. Effects of Parameters on Stability

To find important parameters with respect to stability we analyze the characteristic polynomial of the continuous system. We can determine system stability using a Bode plot. By examining how the form of the Bode plot changes with

system parameters we observe the effect of the parameters on stability. The characteristic polynomial of the system is

$$1 + L(s), \tag{11}$$

where

$$L(s) = \frac{M}{m} \left[ \frac{s^2}{s^2 + 2\zeta\omega_n s + w_n^2} \right] \left[ \frac{4s^2}{(sT+2)^2} \right] \left[ \frac{\omega_0}{s + \omega_0} \right] e^{-\frac{sT}{2}} \tag{12}$$

$$\omega_n = \sqrt{\frac{k_h}{m}} \tag{13}$$

$$\zeta = \frac{(b + b_h)}{2\sqrt{k_h m}}. \tag{14}$$

We see that $L(s)$ can be formed by the product of the ratio of $M$ to $m$, two poles from a second order system described by $\omega_n$ and $\zeta$, four zeros at the origin, two poles on the real axis at $-2/T$, a first-order low-pass filter with cut-off frequency $\omega_0$ (rad/s), and a linear phase delay depending on $T$. It is convenient to define $\omega^*$ as the frequency at which the phase of $L(s)$ is $-180°$. The stability of the system is determined by the value of the gain margin of $L(s)$. The gain margin, $GM$, is defined to be

$$GM = \frac{1}{|L(\omega^*)|} \tag{15}$$

A gain margin greater than one corresponds to a stable system, and a gain margin of less than one corresponds to an unstable system. The Bode plot for positive and negative $M$ have the same magnitude, and the phase for negative $M$ is $180°$ lower than for positive $M$, at every frequency. The effects of these system parameters on the gain margin of $L(s)$ is summarized in Table I.

- $|M/m|$ is a gain of the system; changing it directly affects the gain margin. There exists a maximum stable value of $|M/m|$.
- $\omega_n$ is the frequency at which the magnitude and phase of the second order system transitions. For positive $M$, if $\omega_n \ll \omega^*$ then $\omega_n$ has no effect on the gain margin. For negative $M$, $\omega^*$ is nearly $\omega_n$, so increasing $\omega_n$ increases the gain margin.
- $\zeta$ affects the magnitude and phase of the second order system around $\omega_n$. The lower the value of $\zeta$, the higher the "spike" in magnitude around $\omega_n$, and the smaller the span for the phase to transition. For positive $M$, if $\omega_0 \ll \omega^*$ the value of $\zeta$ does not affect the gain margin. For negative $M$, increasing $\zeta$ increases the gain margin.
- $\omega_0$ affects the gain margin by introducing a pole at frequency $\omega_0$. Lower cut-off frequencies, corresponding to more aggressive filtering, increase the gain margin.
- $T$ affects how quickly phase lag is added to the system. Larger sample periods result in more phase lag added at a given frequency. For positive $M$, $T$ has a large effect on the gain margin; a smaller $T$ corresponds to a higher gain margin. For negative $M$, $T$ has little effect on the gain margin.

TABLE I
EFFECT OF PARAMETERS ON GAIN MARGIN OF $L(s)$

| Parameter | Gain Margin (positive M) | Gain Margin (negative M) |
|:---:|:---:|:---:|
| $|M/m|$ | $\Downarrow$ | $\Downarrow$ |
| $\omega_n$ | - | $\Uparrow$ |
| $\zeta$ | - | $\Uparrow$ |
| $T$ | $\Downarrow$ | - |
| $\omega_0$ | $\Downarrow$ | $\Downarrow$ |

for $\omega_n \ll \omega^*$

TABLE II
HUMAN AND DEVICE PARAMETERS

| Human Parameters [13] | | Device Parameters [7] | |
|:---:|:---:|:---:|:---:|
| $k_h$ 700 | (N/m) | $m$ 70 | (g) |
| $b_h$ 5 | (Ns/m) | $b$ .005 | (Ns/m) |

Figure 4 shows how $L(s)$ varies with the sample time, $T$, for positive desired virtual mass. (The Bode plots for negative desired virtual mass would have $-180°$ lower phase at every frequency.) The parameter values are $|M/m| = 1$, $\omega_n = 100$ (rad/s), $\zeta = 0.5$, and $\omega_0 = 30$ (rad/s). These values were chosen to represent a human interacting with a Phantom Premium haptic device (SensAble, Inc.) based on a human model from Kuchenbecker et al. [13] and device model from Diolaiti et al. [7] (Table II). Figure 4 demonstrates that the sample rate significantly affects the gain margin for positive virtual mass because the value of $\omega^*$ is mainly determined by the sample rate. However, for negative virtual mass, $\omega^*$ is nearly $\omega_n$, so the sample rate has little effect on the gain margin.

Figure 5 shows how $L(s)$ varies with $\omega_n$ for positive virtual mass. For positive virtual mass, $\omega_n$ and $\zeta$ do not affect the gain margin of $L(s)$, as long as the phase transition of the second order system reaches its asymptote before $\omega^*$. However, for negative virtual mass they do affect the gain margin. Thus, the haptic device and operator parameters have more effect on stability for inertia compensation than inertia rendering.

The virtual inertia to device inertia ratio, $M/m$, is a gain of the system and directly affects the gain margin of $L(s)$. If other system parameter values are set, we can determine the range of $M/m$ for which the system is stable. This analysis shows how much inertia rendering is possible. For example, the Bode plot of Figure 4 at a sample rate 1000 Hz is about $-45$ dB at $\omega^*$, so for positive $M$, values of $M/m$ up to 45 dB are stable. For negative $M$, the magnitude of the system is about $-15$ dB at $\omega^*$, so values of $|M/m|$ up to 15 dB are stable. This establishes a range of stable *commanded* virtual mass values, but we will see in Section IV that the rendered mass of the system may not actually match the commanded value. This analysis is based on the continuous model whose stability boundaries are similar, but not identical, to the hybrid and discrete model stability boundaries.
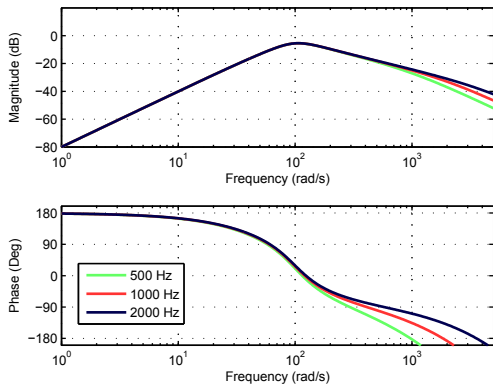
Fig. 4. Bode plots of $L(s)$ with positive virtual mass $M$ for various sample rates. Here $M/m = 1$, $\omega_n = 100$ (rad/s), $\zeta = 0.5$, and $\omega_0 = 30$ (rad/s). The Bode plots for negative virtual mass have $180°$ lower phase at every frequency. The sample rate affects the gain margin greatly for positive virtual mass, but not for negative virtual mass.
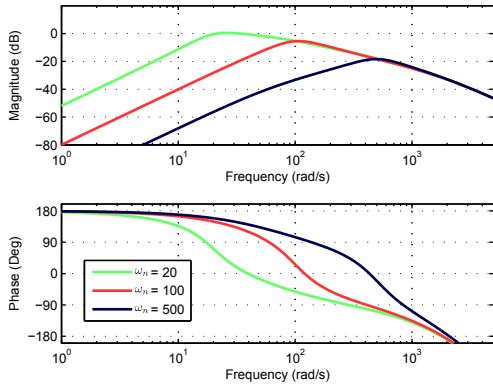


Fig. 5. Bode plots of $L(s)$ with positive virtual mass $M$ for various values of $\omega_n$. Here $M/m = 1$, $\zeta = 0.5$, $\omega_0 = 30$ (rad/s), and the sample rate is 1000 Hz. The Bode plots for negative virtual mass $M$ have $180°$ lower phase at every frequency. For positive virtual mass $\omega_n$ and $\zeta$ do not affect the gain margin of $L(s)$ as long as the phase transition of the second order system reaches its asymptote before $\omega^*$, but for negative virtual mass they do affect the gain margin.

## B. Stability Boundaries (M-width)

In this section, analytical stability boundaries are formed using the continuous and discrete models. Stability boundaries are also formed using numerical simulation. Figure 6 displays stability boundaries for a family of systems. The maximum minus the minimum stable value of the (desired) virtual to device mass ratio, $M/m$, scaled by the device mass, give the M-width of the system: the dynamic range of virtual masses renderable in a stable manner. The maximum stable value for rendering positive mass is predominately determined by the sample time. For rendering negative inertia, the maximum stable value is mainly determined by $\omega_n$ and $\zeta$. Reducing the cut-off frequency of the filter increases the M-width.

## C. Continuous and Discrete Model Comparison

How do the continuous and discrete models compare? We now examine whether the stability intuition that we get from the continuous model is applicable to the discrete model,
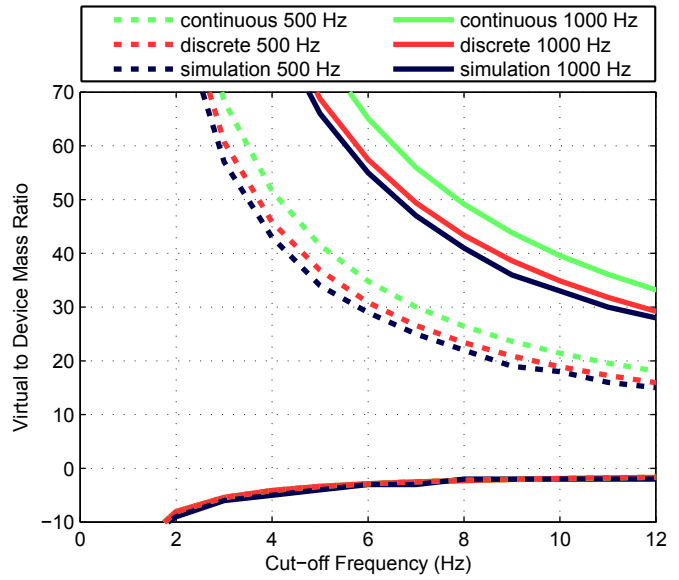


Fig. 6. Examining M-width: Stability boundaries from analysis and numerical simulation. Lines mark boundaries between stable and unstable behavior; the origin side is stable. Given the parameters in Table II, a cut-off frequency of the low-pass filter, $f_0$ (Hz), and sample rate, the plots display the range of desired virtual mass to device mass ratios, $M/m$, that are stable. The maximum minus the minimum stable value of virtual to device ratio scaled by the device mass is the M-width.
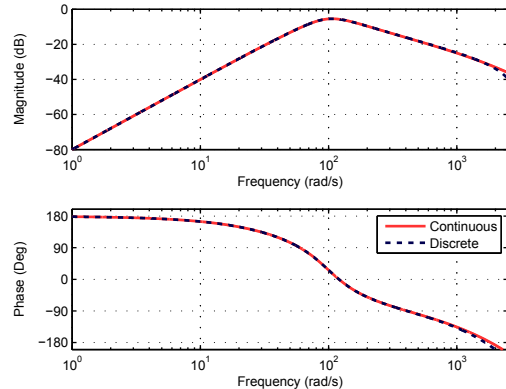


Fig. 7. Bode plots of the loop polynomials for the continuous and discrete systems. The plots are similar, but the discrete system has more phase lag at higher frequencies.

which shares the same stability boundaries as the hybrid one. Figure 7 displays the Bode plots of $L(s)$ and $L(z)$ for the continuous and discrete models, respectively. The system parameter values are the same as in Figure 4 with a sample time of 1000 Hz. The Bode plots are very similar, but the discrete model has more slightly more phase lag at higher frequencies. Thus, the continuous model will appear stable for parameters at which the discrete model is unstable. However, the general effects of parameters for the models are similar, so the intuitive parameter guidelines apparent in the continuous model are still useful.

## IV. ACCURACY OF HAPTIC RENDERING

We are interested not only in stability, but also in the accuracy of the haptic rendering. We investigated this using two different approaches. The first approach is to analyze the input-output relationship from the position of the human $X_h(s)$ to the force of the human $F_h(s)$. The second approach is to perform a system identification on the simulated system assuming a model of pure mass. This results in an estimate for the mass of the system, and this estimate can be compared to the ideal case.

### A. Accuracy Analysis by Transfer Function Comparisons

We now analyze the accuracy of the haptic rendering by comparing the admittance of three systems using the continuous model. The first system is the ideal case in which the operator feels only a mass. The value of this mass is the sum of the device mass and virtual mass:

$$G_{ideal}(s) = \frac{F_h(s)}{X_h(s)} = (m + M)s^2. \qquad (16)$$

The second system is the transfer function of the device alone:

$$G_{device}(s) = \frac{F_h(s)}{X_h(s)} = \frac{b_h s(s + b/m)(s + k_h/b_h)}{s^2 + 2\zeta\omega_n s + \omega_n^2}. \qquad (17)$$

The third is the transfer function of the device with the haptic loop implementing a virtual mass:

$$G_{sys}(s) = \frac{F_h(s)}{X_h(s)} = \frac{(b_h s + k_h)(ms^2 + bs + D(s)H(s)M)}{ms^2 + (b + b_h)s + k_h + D(s)H(s)M}. \qquad (18)$$

A reasonable metric for the accuracy of the rendering is how the Bode plot of the full system (the third case) compares to the other two. For high-quality haptic rendering, the system plot should be more similar to the first case (ideal). To aid this analysis, we introduce effective mass ratio and effective damping ratio.

$$\text{effective mass ratio} = \frac{1}{m}\frac{-\mathfrak{Re}\{G(s)\}}{w^2} \qquad (19)$$

$$\text{effective damping ratio} = \frac{1}{(b + b_h)}\frac{\mathfrak{Im}\{G(s)\}}{w} \qquad (20)$$

These quantities can be plotted against frequency to identify components of the system response corresponding to mass and damping. The plots contain the same information as the Bode plots, but allow for a more easy comparison.

Figure 8 shows the Bode and effective mass and damping ratio plots of a family of systems where the human and device parameters used are those listed in Table II, the cut-off frequency of the low-pass filter, $f_0$, is 10 (Hz), and the sample rate is 1000 Hz. The system with full feedback behaves similarly to the ideal system for frequencies less than the cut-off frequency of the filter, and similar to the device system for larger frequencies. This analysis shows that the cut-off frequency of the filter plays a powerful role in determining the accuracy of the rendering. The higher the cut-off frequency, corresponding to less aggressive filtering, the larger the frequency span in which the system behaves like the ideal.
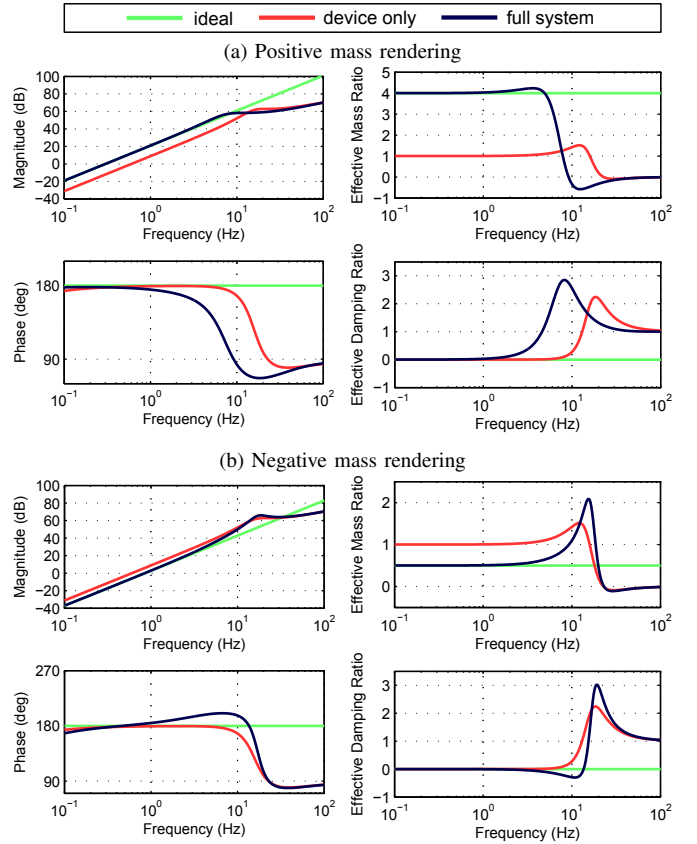


Fig. 8. Measuring accuracy of haptic rendering using transfer functions. (a) Bode and effective mass and damping ratio plots of three systems with positive virtual mass M = 3m. (b) Bode and effective mass and damping ratio plots of three systems with negative virtual mass M = -.5m. The system with full feedback behaves similar to the ideal system for frequencies less than the cut-off frequency of the filter, and similar to the device system for frequencies past.

### B. Accuracy Analysis by System Identification

Another way to measure the accuracy of haptic rendering is to compare the perceived mass of the system to the ideal case. Ideally, at every instant of time the system should follow

$$f_h(t) = (m + M)\ddot{x}, \qquad (21)$$

where $\ddot{x}$ is the acceleration of the device. We can form an estimate for the perceived mass of the system, $\hat{m}$, by performing a system identification on simulated data assuming a model of pure mass. We can compare this estimate to the ideal case, $(m + M)$. Numerical simulation of the system can yield numerous "measurements" for model fitting.

$$[\ddot{x}]\,\hat{m} = [F_h], \qquad (22)$$

where $\hat{m}$ is a scalar and $[\ddot{x}]$ and $[F_h]$ and are vectors of the samples. The best estimate of $\hat{m}$, in the least squares sense, is

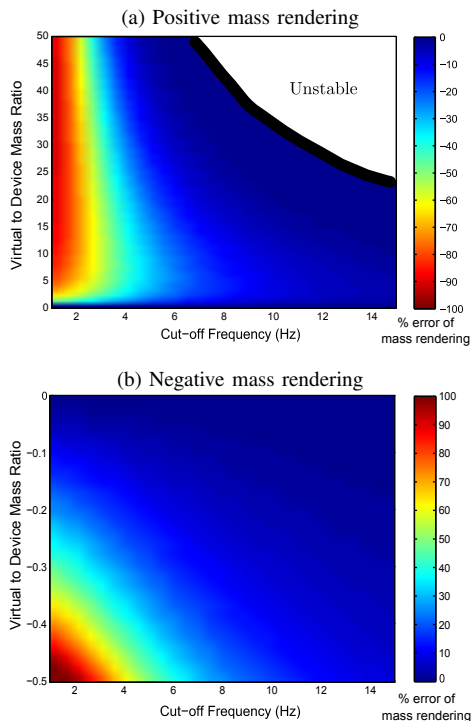$$\hat{m} = \frac{[\ddot{x}]^T\,[F_h]}{[\ddot{x}]^T\,[\ddot{x}]}. \qquad (23)$$

Fig. 9. The error of the estimate of the mass of the system, $\hat{m}$, with respect to the ideal case, $(m+M)$, versus the cut-off frequency of the low-pass filter, $f_0$ (Hz), for various mass ratios. (a) The case of *positive* virtual mass. (b) The case of *negative* virtual mass.

We define the error as the difference between the estimate and the ideal value scaled by the ideal value:

$$\text{error} = \frac{\hat{m} - (m + M)}{(m + M)}. \qquad (24)$$

If the error is negative, the system would feel less massive than desired, and if the error is positive, the system would feel more massive than desired. This method outputs a scalar for the error by combing errors in magnitude and phase, and as such has no physical meaning. Despite this, it is a useful metric to establish conditions in which the commanded and rendered system mass are similar or different.

Figure 9 shows the error in mass rendering as a function of cut-off frequency for various mass ratios. The motion input to the simulated system was a chirp signal with frequencies of 0.1 to 5 Hz. For both positive and negative mass rendering, as the cut-off frequency of the filter is increased, corresponding to less aggressive filtering, the perceived mass approaches the ideal. Also, generally, the larger the mass ratio, the larger the error. For this system, the human parameters and sample period did not significantly affect the simulation.

## V. DISCUSSION

### A. Design Trade-offs

Our results illustrate the classic trade-off between stability and accuracy of haptic rendering. For the case of mass rendering, we demonstrated that accurate rendering can be achieved for a small ratio of virtual mass to device mass and high cut-off frequency of the low-pass filter. However, these requirements generate restrictive stability bounds, as shown in Figure 6. In practice, an upper bound for the cut-off frequency of the low-pass filter exists for noise reduction. This results in a rendering accuracy limit based on the noise present in acceleration measurement.

For systems in which $\omega_n \ll \omega^*$, we established that sampling rate is an important parameter for implementing positive virtual mass, but not negative virtual mass. For positive virtual mass, the stability of the system is largely determined by the delay in the haptic feedback loop. For negative virtual mass, the stability of the system is largely determined by the user behavior, so the time delay introduced by sampling has a small relative effect.

The M-width of a haptic device scales with the device mass. This means that to display large virtual mass, large device mass is desired. This is at odds with a primary goal of haptic device design: to create mechanisms that result in minimal free-space dynamics. It is possible that the inertia of a more massive device could be canceled, but mass compensation stability is highly dependent on user behavior. For haptic environments that aim to render large virtual inertia, such as the rehabilitation scenario presented in Figure 1, this presents a design challenge, which we would like to circumvent through the use of clever controllers and improved sensing.

With the results of this paper, we can now consider various changes to the haptic system design, including mechanical parameters of the device, position and/or acceleration sensors, filters, and novel controllers, and predict their effects on stability and accuracy of mass rendering. For applications in which perception of inertia (as opposed to display of specific inertia values) is the goal, we can also consider the use of multimodal feedback to alter inertia perception. Dominjon et al. [8] found that the ratio of the amplitude of movements of a user's hand to the amplitude of a virtual cursor, as well as gravity, play a large role in inertia perception.

### B. Influence of the Human on Stability

The parameters of the system affected by the human mechanical properties are $\omega_n$ and $\zeta$. These parameters, which also depend on the physical properties of the haptic device, increase with increased human impedance or co-contraction, as occurs when an operator strongly grasps a haptic device. For positive virtual mass, if $\omega_n \ll \omega^*$, the human does not contribute to stability. As $\omega_n$ gets closer to $\omega^*$, however, $\omega_n$ and $\zeta$ affect stability more strongly. This can be interpreted in the following way: if the user is grasping weakly, then the stability of the system does not depend on the user. However, as the user grasps the haptic device more strongly, the values of $\omega_n$ and $\zeta$ increase, and the system becomes more stable because of the contribution of the human. For positive virtual mass the system has a "baseline" stability without any interaction from the human, that can only be improved by the user. This is not the case for negative virtual mass, where $\omega_n$ and $\zeta$ affect stability directly. This confirms the results of Gil et al. [11]:
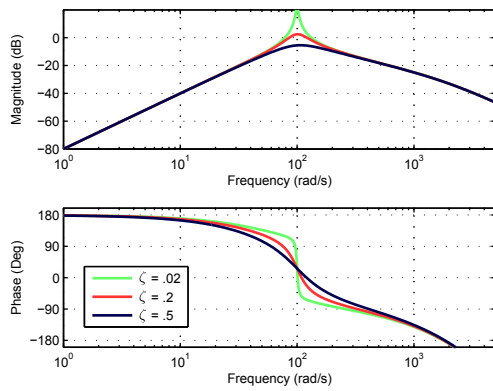
Fig. 10. Bode plots of $L(s)$ with positive virtual mass $M$ for various values of $\zeta$. Here $M/m = 1$, $\omega_n = 100$ (rad/s), $\omega_0 = 30$ (rad/s), and the sample rate is 1000 Hz. The value of $\zeta$ does not affect the stability of the system for positive virtual mass, but has an extreme affect on stability for negative virtual mass.

that the mechanical properties of the operator strongly affect stability. For example, if the user grasps the haptic device weakly, the system can have a strong resonance peak at $\omega_n$ which could make the system go unstable for very small amounts of mass compensation. This effect is shown in Figure 10, which displays the effect of $\zeta$ on the Bode plot of $L(s)$.

*C. Model Choice*

Many models could be used to describe a haptic system for inertia rendering; our choice of the model shown in Figure 2 was driven by relevance as a practical design tool, tractability for finding analytical solutions, and our ideas about the most significant factors in system stability and accuracy. In future work, we may wish to capture additional system properties, such as quantization of signals occurring due to A/D and D/A conversions and encoder position measurements. Quantization is a non-linear phenomenon, and in position measurement it can lead to sustained oscillations and instability in the system [1, 7]. Nonlinear Coulomb friction has also been shown to have a substantial effect on the passivity of haptic devices. Comparing our theoretical and simulation results to the stability and accuracy of haptic rendering on a variety of haptic devices will determine the broad applicability of our approach and whether improved models are needed.

## VI. CONCLUSION

The main results of our analysis are the identification of important parameters for system stability, stability boundaries (M-width), and analysis of the accuracy of mass rendering using impedance-type haptic devices. We verify our analytical results in numerical simulations, while recognizing that experiments with a variety of haptic devices are needed to validate the broad applicability of our approach. Our results could serve as a design tool for creating virtual environments, controllers, and device/sensor designs to improve rendering of virtual mass. This could apply to inertia compensation for more massive haptic interfaces, as well as rendering increased or decreased inertia in rehabilitation and motor training scenarios.

## REFERENCES

[1] J. J. Abbott and A. M. Okamura. Effects of position quantization and sampling rate on virtual-wall passivity. *IEEE Transactions on Robotics*, 21(5):952–964, 2005.

[2] R. J. Adams and B. Hannaford. Stable haptic interaction with virtual environments. *IEEE Robotics and Automation*, 15(3):465–474, 1999.

[3] G. Aguirre-Ollinger, J. E. Colgate, M. A. Peshkin, and A. Goswami. Design of an active one-degree-of-freedom lower-limb exoskeleton with inertia compensation. *The International Journal of Robotics Research*, 30(4):486–499, 2010.

[4] J. M. Brown and J. E. Colgate. Minimum Mass for Haptic Display Simulations. In *Proc. ASME Dynamic Systems and Control Division*, pages 249–256, 1998.

[5] J. E. Colgate and J. M. Brown. Factors affecting the Z-Width of a haptic display. In *Proc. IEEE International Conference on Robotics and Automation*, pages 3205–3210, 1994.

[6] J. E. Colgate and G. G. Schenkel. Passivity of a class of sampled-data systems: Application to haptic interfaces. *Journal of Robotic Systems*, 14(1):37–47, 1997.

[7] N. Diolaiti, G. Niemeyer, F. Barbagli, and J. K. Salisbury. Stability of haptic rendering: Discretization, quantization, time delay, and coulomb effects. *IEEE Transactions on Robotics*, 22(2):256–268, 2006.

[8] L. Dominjon, A. Lécuyer, J. M. Burkhardt, P. Richard, and S. Richir. Influence of control/display ratio on the perception of mass of manipulated objects in virtual environments. In *Proceedings of IEEE Virtual Reality*, pages 19–25, 2005.

[9] G. F. Franklin, D. J. Powell, and A. Emami-Naeini. *Feedback Control of Dynamic Systems*. Pearson Prentice Hall, 5th edition, 2006.

[10] J. J. Gil, A. Avello, A. Rubio, and J. Florez. Stability analysis of a 1 dof haptic interface using the routh-hurwitz criterion. *IEEE Transactions on Control Systems Technology*, 12(4):583–588, 2004.

[11] J. J. Gil, A. Rubio, and J. Savall. Decreasing the Apparent Inertia of an Impedance Haptic Device by Using Force Feedforward. *IEEE Transactions on Control Systems Technology*, 17(4):833–838, 2009.

[12] B. Hannaford and J.-H. Ryu. Time-Domain Passivity Control of Haptic Interfaces. *IEEE Transactions on Robotics*, 18(1):1–10, 2002.

[13] K. J. Kuchenbecker, J. G. Park, and G. Niemeyer. Characterizing the human wrist for improved haptic interaction. In *Proc. ASME Int. Mechanical Engineering Congress and Exposition*, volume 2, 2003. paper number 42017.

[14] M. Mahvash and A. M. Okamura. Enhancing transparency of a position-exchange teleoperator. In *World Haptics*, pages 470–475, 2007.