

M-Width: Stability, Noise Characterization, and Accuracy of Rendering Virtual Mass

Nick Colonnese and Allison M. Okamura

Abstract

In certain robot control and physical human-robot interaction scenarios, it is desirable to carefully control the apparent mass of a robot. Manipulating the apparent mass can be accomplished through virtual mass rendering, where the actuators of the robot produce forces proportional to measured acceleration. Many factors influence mass rendering, including device mechanical properties, sample rate, control structure, filtering, environment dynamic coupling parameters, and delay. Inspired by the “Z-Width” approach to sampled-data robotic system passivity, which represents the infinite-dimensional passive impedances of a robotic haptic display, we establish “M-Width”: the passive range of pure virtual mass. In this paper, we identify important parameters for system passivity and stability, present passivity and stability boundaries, predict noise limit cycles and establish conditions for their existence, and describe the expected accuracy of rendered virtual mass. We construct explicit parameter regions to describe the three-way trade-off that occurs between passivity/stability, noise, and accuracy. Experimental data gathered with a Phantom Premium 1.5 robot validates the theoretical analysis. These results serve as a general design tool for manipulating the effective mass of a robot, which is particularly relevant for rehabilitation robotics, robotic exoskeletons, and haptic display applications.

Index Terms

passivity, stability, sampled-data systems, physical human-robot interaction, mass control, exoskeletons, rehabilitation robotics, kinesthetic haptic interfaces, quantization, transparency

I. INTRODUCTION

FOR many applications of back-drivable robotic systems, we want the robot to behave as if it has a very specific apparent mass. Manipulating the apparent mass can be accomplished through *virtual mass rendering*, where the actuators of the robot produce forces proportional to the acceleration of the robot. Increasing the apparent mass of a robot is useful for particular training and rehabilitation scenarios, and decreasing the apparent mass is useful for mitigating forces felt during free space motion in exoskeleton

and haptic display applications. Controlling the apparent mass to be a specific value evokes fundamental challenges in robotic control.

Robotic systems used for physical interaction with humans should be simultaneously stable, free of noisy force signals that can generate unrealistic and unexpected audio or haptic sensations, and be accurately controlled to have the desired apparent mechanical properties. Direct mass rendering (presenting a force proportional to acceleration) is particularly challenging because acceleration measurements on a typical robot (equipped with only optical encoders to sense position) involve double discrete differentiation of a quantized position signal. We contrast the *direct* mass rendering explored in this paper from the *indirect* method using a virtual coupling connected to a proxy virtual mass (Brown and Colgate, 1998; Adams and Hannaford, 1999). While system modifications (e.g., the addition of new sensors such as tachometers or accelerometers) can improve mass rendering, understanding the performance of a canonical robot yields insight as to what modifications in hardware, sensing, and control will improve mass rendering.

A. *Prior Work*

This work builds upon significant prior research related to the stability of sampled-data robotic and haptic systems, as well as research to control apparent mass. In this paper, we expand on our preliminary results for mass rendering (Colonnese and Okamura, 2012) by including an analysis on passivity, presenting stability results with respect to a *general* mass, viscous damping, and spring coupling environment instead of a *specific* viscous damping and spring environment, including a characterization of noise from position quantization, and experimentally validating all theoretical analyses on a Phantom Premium 1.5 robot.

Robust stability conditions for back-drivable robotic and haptic systems were established by Colgate and Brown (1994) and Colgate and Schenkel (1997) in terms of closed-loop robot passivity. These works placed special emphasis on virtual spring and damper rendering (corresponding to robot actuators commanding forces proportional to position and velocity, respectively) for implementing “virtual walls.” Abbott and Okamura (2005) and Diolaiti et al. (2006) performed noise characterization for virtual spring rendering. Gil et al. (2009a) proposed uncoupled stability (the robot with no external coupling environment) results for virtual spring and damper rendering, and Hulin et al. (2008) analyzed the sensitivity of the uncoupled stability of virtual spring and damper rendering to environment coupling. Gil et al. (2009b) explored mass compensation by force feedforward for an impedance haptic device using additional sensors for acceleration estimation to manipulate the multi-dimensional mass properties of the robot to be approximately spherical in its workspace (Khatib, 1995). Brown and Colgate (1998) establish conditions for passive positive mass simulations using the indirect mass rendering architecture. These previous research investigations established design guidelines for virtual spring and damper rendering, as well as mass rendering using system modifications, but they did not present results for direct mass rendering.

Virtual mass rendering has been previously explored for many specific robotics scenarios. Mahvash and Okamura (2007) include mass compensation for a haptic-feedback surgical teleoperator using model-based feed-forward terms to cancel the unwanted dynamic properties of the manipulators. Their analysis showed that 70% of the mass of the slave manipulator could be canceled during low-frequency motion, but high gains to achieve complete cancellation could make the teleoperator unstable. Grow et al. (2010) and Bhanpuri et al. (2014) tested a hypothesis that cerebellar damage causes a mismatch between the brain's modeled dynamics and the actual body dynamics by conducting reaching experiments in patients with ataxia (uncoordinated movement) in a planar back-drivable robot where the robotic system included mass rendering to change the dynamic properties of the arm-robot combination by increasing or decreasing the apparent mass of the robot. Aguirre-Ollinger et al. (2010, 2012) used mass compensation in exoskeletons to enable faster movements of the lower extremities. They found that subjects with higher limb masses could accommodate comparatively larger levels of mass reduction, and that robust stability of the coupled robot and human system was sensitive to damping. These prior mass rendering studies were driven by specific applications, and general results concerning fundamental issues for mass rendering were not presented.

B. Contributions

In this paper we address direct mass rendering by analyzing a canonical one-degree-of-freedom robotic or haptic device. To our knowledge, this work is the first to establish general results for (1) passivity, (2) stability, (3) noise characterization, and (4) closed-loop accuracy for mass rendering. By constructing explicit parameter regions for passive, stable, quantization error noise-free, and accurate robot operation, we describe the the three-way trade-off which occurs between passivity/stability, noise, and accuracy. In addition to presenting these parameter regions, we also present the following main results:

- Passivity is extremely conservative compared to coupled stability for typical coupling environments (Section III)
- Positive mass stability is generally not sensitive to environment coupling (Section IV)
- Negative mass stability is sensitive to environment coupling (Section IV)
- Stability depends more on the *ratio* of virtual to device and environment mass, $\frac{M}{(m+m_e)}$, compared to the absolute value of the virtual mass M (Section IV)
- Physical damping in the robot, b , and environment damping, b_e , expand the range of stable virtual mass (Section IV)
- Position quantization error can cause limit cycles, but cannot cause instability. In other words, if a system is bounded-input, bounded-output (BIBO) stable without quantization error, the system will still be BIBO stable with the addition of quantization error (Section V)

- Increasing the robot Coulomb friction c , decreasing the position sensing quantization interval, Δ , and more aggressive low-pass filtering allow a larger quantization noise-free range of virtual mass (Section V)
- The closed-loop robot renders the desired dynamics closely only at frequencies smaller than the cut-off frequency of the low-pass filter (Section VI)

We consider “M-Width” to be a natural and important extension of the well studied “Z-Width” of robotic and haptic displays (Colgate and Brown, 1994). 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. We imagine Z-Width representing the infinite dimensional passive impedances of a robotic or haptic display, where M-Width is the passive range of pure virtual mass, B-Width is the passive range of pure virtual damping, K-Width is the passive range of pure virtual stiffness, etc.

C. Paper Organization

The organization of the paper is as follows. In Section II we propose several models of a robotic system rendering mass. In Section III, we generate passivity bounds and introduce “M-Width”, the dynamic range of virtual mass that can be rendered passively. Because passivity can be conservative compared to stability, we also analyze closed-loop robot and coupled environment stability in Section IV, where we use a second-order model of the environment to identify important parameters for stability and form stability boundaries. In Section V, we characterize the noise introduced to the system by quantization of the position measurement. Using a describing function analysis, we predict limit cycles due to sensor resolution, then perform an energy analysis to establish conditions for which limit cycles cannot occur. In Section VI, we analyze how the closed-loop robot impedance compares to the ideal apparent mass. Finally, in Section VII, the theoretical stability, noise characterization, and accuracy results are compared to experimental data gathered with a Phantom Premium 1.5 robot.

II. SYSTEM MODELS

We begin by introducing system models for a robot interacting with an environment, where the goal is to manipulate the apparent mass of the robot by rendering virtual mass. The mass can be positive (additional mass rendering) or negative (mass compensation). Many different system architectures can be used to render mass; in this paper we are interested in back-drivable “impedance-type” robots with position sensing. Robots of this type are frequently used in rehabilitation, exoskeleton, and haptic applications because of their ease of movement, simplicity, and safety compared to other robot types that require the actuators to accomplish movement. Three system models are introduced. One is a *nonlinear hybrid* model,

which contains elements that are both linear and nonlinear, and both continuous and discrete. This model is the closest to reality, in that the device friction is nonlinear, the position sensor measurement (usually corresponding an optical encoder) is quantized, and the control of a physical robot is performed through a computer containing A/D (analog to digital) and D/A (digital to analog) components. We also introduce *linear continuous* and *linear discrete* models. The linear continuous model is used for the identification of important parameters with respect to stability, finding limit cycles, and for use in analysis of quality of rendering. The linear discrete model is used to generate stability bounds and characterize error from quantization.

A. Nonlinear Hybrid Model

We consider a one-degree-of-freedom robot described by a mass, m , viscous damper, b , and Coulomb friction, c , acted upon by two external forces: the force applied by the environment, $F_e(s)$, and the force applied by the actuator implementing the control law, $F_a(s)$ (Figure 1 (b)). The environment dynamics are modeled by a mass-spring-damper,

$$Z_e(s) = \frac{F_e(s)}{V(s)} = m_e s + b_e + \frac{k_e}{s} \quad (1)$$

where m_e , b_e and k_e are non-negative values corresponding to the environment mass, damping, and stiffness, respectively. For some relevant applications, the environment dynamic coupling is a human operator. In this case, the second-order model represents the non-volitional biomechanics that are important for stability. Because environment dynamics can change significantly, from no interaction to soft or rigid dynamics, we will consider the set of likely environment impedances.

The robot is equipped with only one sensor measuring the position of the mass, $X(s)$, which is quantized with resolution Δ , and then sampled with a constant sampling period of T . To obtain an estimate of the acceleration of the device, the sampled position measurements go through a discrete double back differencing operator,

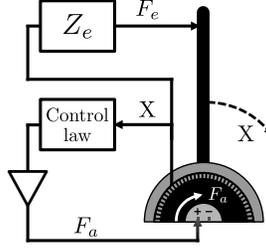
$$D(z) = \frac{(z-1)^2}{(Tz)^2}, \quad (2)$$

and then n first-order discrete low-pass filters,

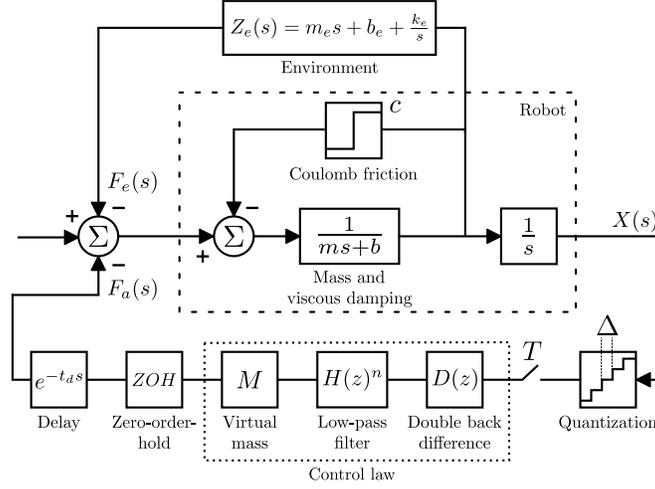
$$H(z) = \frac{(1 - e^{-\omega_0 T})z}{z - e^{-\omega_0 T}}, \quad (3)$$

where ω_0 (rad/s) represents the cut-off frequency of the low-pass filter. Other filters could be used; we chose one of the simplest possible filters to provide information about baseline performance. The force of the actuator is the product of the filtered acceleration estimate and a desired virtual mass, M . It

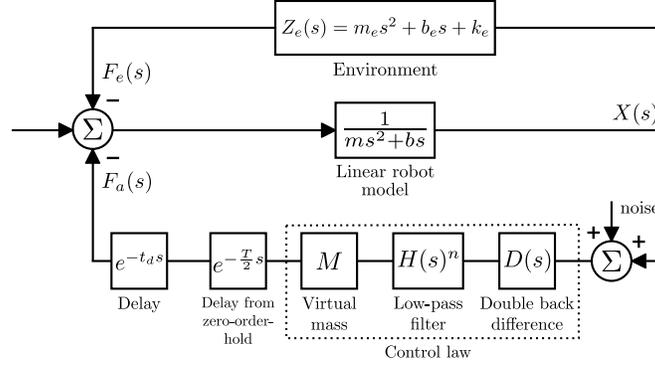
(a)



(b)



(c)



(d)

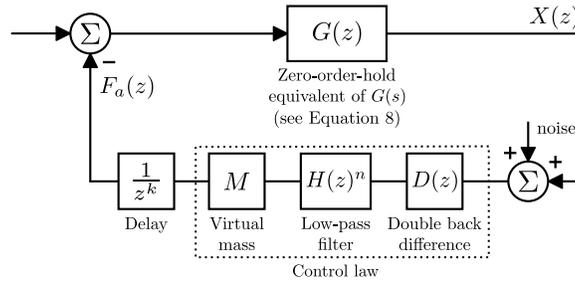


Fig. 1. (a) Schematic of a robot implementing virtual mass interacting with an environment. For many applications, the environment is a human operator. (b) The *nonlinear hybrid model* for this system (Section II-A) considers environment and device dynamics, quantization, sampling, ZOH, and delay. (c) The *linear continuous model* (Section II-B) is created from the nonlinear hybrid model by omitting the nonlinear elements and converting the discrete elements to continuous ones. (d) The *linear discrete model* (Section II-C), has the same input-output behavior as the linear hybrid model at sample times.

is held constant for the duration of the sampling period with a zero-order hold (ZOH), resulting in a continuous-time staircase signal. At sample time k , the actuator force is

$$F_a(k) = M H(z)^n D(z) X(k). \quad (4)$$

To model time-delayed actuator forces due to amplifier dynamics or transport time, the actuator force is delayed by t_d .

B. Linear Continuous Model

An entirely linear continuous model shown in Figure 1(c) approximates the nonlinear hybrid system shown in Figure 1(b). The linear continuous model is desirable because of its tractability for accuracy analysis and finding important parameters with respect to stability using conventional linear control systems theory. The stability of the two systems are similar, but not identical.

To represent the system linearly, the Coulomb friction of the device is removed, and the quantization element is replaced with additive noise. To represent the system continuously, the discrete elements are converted to continuous ones. The continuous representation of $D(z)$, $D(s)$, is found using the bilinear (Tustin) mapping:

$$D(z) \rightarrow D(s) = \frac{4s^2}{(sT + 2)^2}. \quad (5)$$

The continuous representation of $H(z)$, $H(s)$, is a first-order low-pass filter with cut-off frequency ω_0 (rad/s) with unity gain at DC.

$$H(z) \rightarrow H(s) = \frac{\omega_0}{s + \omega_0}. \quad (6)$$

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

$$\text{ZOH} \rightarrow e^{-\frac{sT}{2}}. \quad (7)$$

C. Linear Discrete Model

We also represent the system with an entirely linear discrete model shown in Figure 1(d). The linear discrete model is useful for stability and noise propagation analyses in which the discrete elements can be considered explicitly. The linear discrete model is made linear from the hybrid model by the same steps taken to make the continuous system linear. To make the system discrete, the continuous elements in the linear hybrid model are replaced with discrete ones. 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,

$$\frac{F_e(s)}{X(s)} = G(s) = \frac{1}{(m + m_e)s^2 + s(b + b_e) + k_e}. \quad (8)$$

Then, the ZOH, $G(s)$, and the sample elements are converted into a discrete element, $G(z)$, using a zero-order hold equivalent (Franklin et al., 2001),

$$G(z) = \frac{z-1}{z} \mathcal{Z} \left\{ \frac{G(s)}{s} \right\}. \quad (9)$$

The continuous delay element is replaced by a discrete one,

$$e^{-t_d s} \rightarrow \frac{1}{z^k}, \quad (10)$$

where k is a nonnegative integer representing the length of the delay in number of sample times. The input-output characteristics of the linear continuous and discrete elements (at sample times) are the same. For this reason, the stability of the linear hybrid and discrete models are essentially identical.

III. SYSTEM PASSIVITY

Passivity is commonly used in robotic control design to conservatively ensure stability when the dynamics of the environment are unknown. If the closed-loop robot is shown to be passive, and the environment is assumed strictly passive, the feedback interconnection of the two is passive, and therefore necessarily stable (Colgate, 1992). For the case when the environment is a human user, the user violates passivity due to volitional control; however, the human is only active at relatively low frequencies compared to frequencies relevant to coupled stability, and in practice it is sufficient that the passive bio-mechanics are stable when coupled to the robotic system (Griffiths et al., 2011).

In this section we use the passivity theorem of Colgate and Schenkel (1997) to generate passivity bounds for a robotic system rendering direct mass. The theorem applies to the uncoupled linear hybrid model (Figure 1(b) without Coulomb friction, quantization, and the environment). It states that a necessary and sufficient condition for passivity of the closed-loop robot (without quantization) is

$$b > \frac{T}{2} \frac{1}{1 - \cos(\omega T)} \Re\{(1 - e^{j\omega T})Q(e^{j\omega T})\} \quad \text{for } 0 \leq \omega \leq \frac{\pi}{T} \quad (11)$$

where $Q(z)$ represents the discrete-time feedback control law.

When no low-pass filter is included in the system, and $Q(z) = D(z)M$, we can derive a symbolic condition for passivity,

$$M < \left(\frac{1}{2}\right) bT \quad (M > 0), \quad (12)$$

$$|M| < \left(\frac{8}{9}\right) bT \quad (M < 0). \quad (13)$$

A derivation of this result is included in Appendix B. This unintuitive result establishes that larger virtual mass is passive for *larger* sample times. This condition is also extremely restrictive because the passive

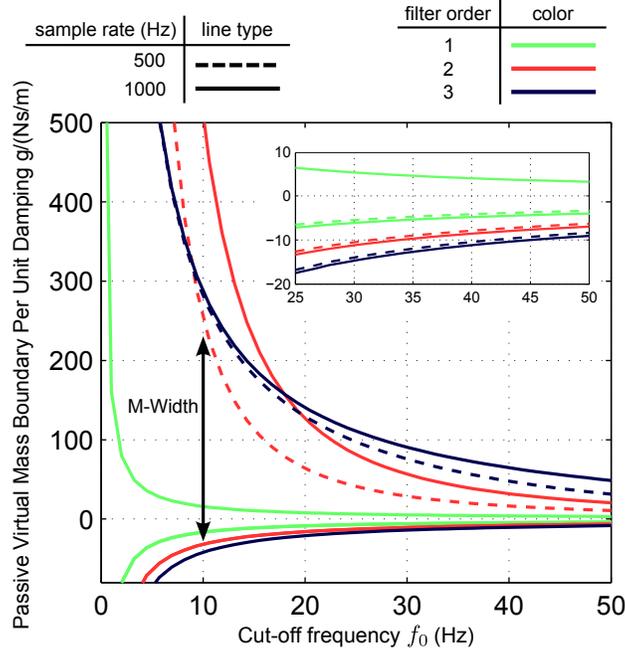


Fig. 2. M-Width: passivity bounds for rendering positive and negative virtual mass. Low-pass filtering expands the passive virtual positive mass region significantly, but not the negative mass region. For cut-off frequencies satisfying common noise rejection or accuracy objectives, the passive range of virtual mass is on the order of tens of grams.

virtual mass range is restricted to grams for most currently existing back-drivable robots, where $b < 1$ Ns/m, and $T \approx 0.001$ s.

When the low-pass filter is included in the feedback loop, and $Q(z) = D(z)H(z)^n M$, the passivity region expands greatly. We did not find a tractable symbolic condition for passivity, but the passivity region can be generated numerically. Figure 2 shows the passive range of virtual mass per unit damping for a given sample time, cut-off frequency, and filter order. Given system parameters, the “M-Width” of the robotic system is defined as the maximum positive virtual mass minus the minimum negative virtual mass. Note that this analysis ignores quantization, which can affect passivity. In Section V we consider quantization and establish conditions for which it can cause a non-passive system.

IV. SYSTEM STABILITY

Because passivity can be conservative with respect to stability, we now analyze the system considering both the linear continuous and discrete systems (Figures 1(c) and 1(d)) from which bounded-input, bounded-output (BIBO) stability can be determined. First, we identify important parameters for system stability using the continuous model. This is useful as a design tool because it identifies which key parameters determine mass rendering stability. We then form quantitative stability boundaries using the discrete model; these are equivalent to the stability boundaries of the linear hybrid model. This establishes

the stability limits of the system rendering mass using this architecture.

A. Effects of Parameters on Stability

To find important parameters with respect to stability, we analyze the characteristic polynomial of the continuous system shown in Figure 1(c). System stability is determined using the Nyquist criterion, and we observe the effect of the parameters on stability by examining how the form of the Bode plot changes. The characteristic polynomial of the system is

$$1 + L(s), \quad (14)$$

where

$$L(s) = \frac{M}{(m + m_e)} P(s) D(s) H(s)^n d, \quad (15)$$

$$P(s) = \left[\frac{1}{s^2 + 2\zeta\omega_n s + \omega_n^2} \right], \quad (16)$$

$$\omega_n = \sqrt{\frac{k_e}{(m + m_e)}}, \quad (17)$$

$$\zeta = \frac{(b + b_e)}{2\sqrt{k_e(m + m_e)}}, \quad (18)$$

$$d = e^{-s(\frac{1}{2}T + t_d)}, \quad (19)$$

and $D(s)$ and $H(s)$ are defined by Equations (5) and (6), respectively.

It is convenient to define ω^* 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

$$GM = \frac{1}{|L(\omega^*)|} \quad (20)$$

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. Each parameter affects the gain margin differently.

We separate our analysis into two cases. The first case is general and considers system parameters that can take any values. The second case considers the common parameter set for typical back-drivable robot systems. All current impedance-type haptic displays (Diolaiti et al., 2006), lower-extremity and upper-extremity rehabilitation robots (Bhanpuri et al., 2014; Aguirre-Ollinger et al., 2012), and master-side teleoperator suites system parameters satisfy the conditions to be considered “common.”

TABLE I
EFFECT OF PARAMETERS ON GAIN MARGIN
OF $L(s)$ FOR POSITIVE AND NEGATIVE M

Parameter	Gain Margin
$ M/(m + m_e) $	↓
b, b_e	↑
k_e	↓
low-pass filtering (compared to no filter)	↑
T, t_d	↓

Table entries represent the change in gain margin as the parameter increases. In some cases, the change in gain margin is marginal.

1) *General Case*: From examining the loop polynomial $L(s)$, the general effects of the parameters on the stability of the system are:

- $|M/(m + m_e)|$ is a gain of $L(s)$. There exists a maximum stable value of $|M/(m + m_e)|$ for both positive and negative M .
- b and b_e expand the stable range of virtual mass, i.e., robot damping and environment damping affect stability in the same way to stabilize the system.
- k_e reduces the stable range of virtual mass.
- low-pass filtering expands the stable range of virtual mass compared to no low-pass filtering.
- T and t_d reduce the stable range of virtual mass, i.e., slower sampling and time-delay are de-stabilizing.

The general effects of the parameters on stability are summarized in Table I. Note that for systems that are not well described by our model, e.g., robots with many coupled degrees-of-freedom, or flexible body dynamics, the parameters may affect stability differently.

2) *Common Case*: All system parameters can theoretically take any non-negative value, but for most relevant back-drivable robotics they have certain bounds. The following qualitative analysis considers the “typical” range of parameters where two conditions are assumed to hold. First, the sampling rate is sufficiently fast compared to the dynamics of the open-loop robot. Specifically, the phase lag induced from the sample-and-hold at the inverse of the mechanical time constant of the robot, $\frac{m}{b}$, is less than a threshold, $T \ll \frac{m}{b}$, which implies $\frac{bT}{m}$ is sufficiently small. This is a common assumption used for robots characterized by low damping and fast sampling (Gil et al., 2009a). Second, the low-pass filtering is not extremely aggressive compared to the dynamics of the coupled robot-environment system, meaning ω_0 is not far below ω_n . We will see in Section VI that the cut-off frequency strongly affects the accuracy

bandwidth of the closed-loop robot, and that the case of very strong filtering is not particularly relevant. The reason for making these assumptions is that they are true for most currently existing back-drivable robots and coupling environments, and they help to simplify the analysis, i.e., the assumptions allow statements about which parameters are particularly important for positive and negative M rendering.

- ω_n is the frequency at which the magnitude and phase of the second order system transitions. For positive M , ω_n has little effect on the gain margin. For negative M , increasing ω_n can increase or decrease the gain margin depending on ζ .
- ζ affects the magnitude and phase of the second order system around ω_n . For positive M , ζ does not significantly affect the gain margin. For negative M , increasing ζ increases the gain margin.
- ω_0 affects the margin differently depending on the sign of the virtual mass M and the filter order n .
 - 1) positive M , $n = 1, 2$:
 ω^* is mainly determined by T and t_d . As ω_0 is increased, corresponding to less aggressive filtering, the gain margin is decreased.
 - 2) positive M , $n \geq 3$:
 ω^* is mainly determined by ω_0 , and ω_0 does not significantly affect the gain margin.
 - 3) negative M :
 ω^* is mainly determined by ω_n and ζ , and ω_0 does not significantly affect the gain margin.

The effects of these system parameters on the gain margin of $L(s)$ are summarized in Table II. Sample robot and environment parameters are displayed in Table III. In some cases the environment interaction is a human user. Typical human-arm dynamics are listed in Gil et al. (2004).

Figure 3 shows how $L(s)$ varies with the filter order and sample time for positive and negative virtual mass. The robot parameters are that of a medium device and medium environment (Table III), $|M/m| = 1$, and $\omega_0 = 180$ (rad/s). For positive virtual mass, ω^* depends strongly on the filter order and sample time, however, for negative virtual mass, ω^* is nearly the same for all filter orders and sample times. This qualitative result applies to all robots and environments in Table III.

Figure 4 shows how $L(s)$ varies with different environments for positive and negative virtual mass. The parameters are the same as those for Figure 3 with filter order $n = 3$. For positive virtual mass, ω^* and the gain margin for each environment are similar (corresponding to different m_e , ω_n , and ζ values), but for negative virtual mass, ω^* and the gain margins for the environments vary widely.

These results show that (for conditions in which the assumptions of low robot damping, fast sampling, and no extreme low-pass filtering hold), control feedback structure, in terms of the cut-off frequency, filter order, sample rate, and delay, are important for positive M stability, but not as important for negative M stability. Conversely, environment dynamic interaction, parameterized by m_e , ω_n , and ζ , do not strongly

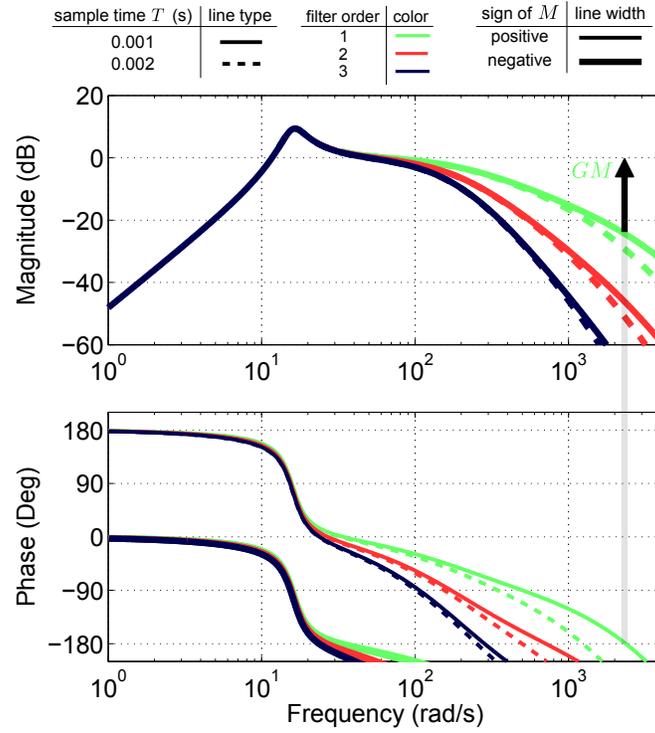


Fig. 3. Bode plots of $L(s)$ for different sampling times and filter orders for positive and negative virtual mass. The gain margin for the positive M system with $T = 0.001$ (s) and filter order $n = 1$ is shown as an example. For positive M , the gain margins of the systems are significantly different for the different sample times and filter orders, but for negative M , the gain margins are very similar.

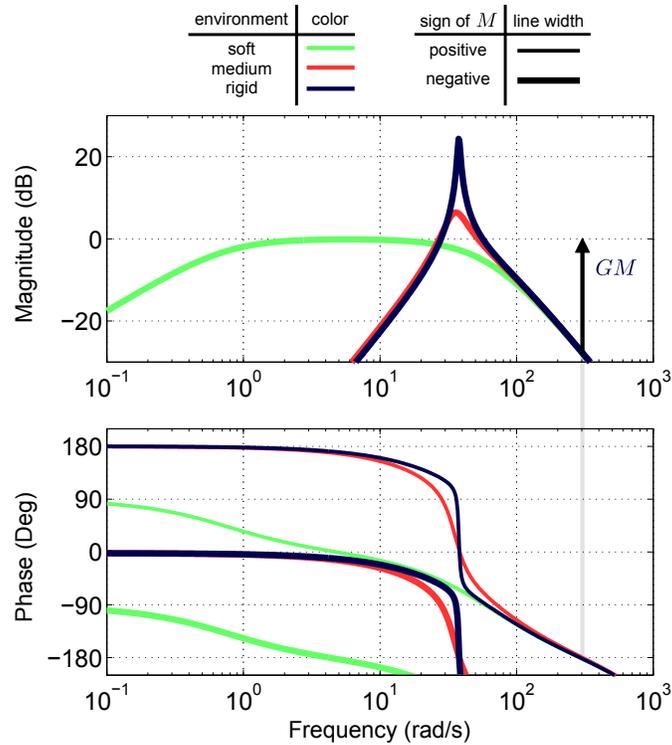


Fig. 4. Bode plots of $L(s)$ for different environments (Table III) showing the effect of environmental coupling on stability for positive and negative virtual mass. The gain margin for the positive M rigid environment is shown as an example. The gain margins for positive virtual mass are essentially identical, but for negative virtual mass, the gain margins vary strongly for the different environments.

TABLE II
EFFECT OF PARAMETERS ON GAIN MARGIN OF $L(s)$
FOR TYPICAL BACK-DRIVABLE ROBOT SYSTEMS

Parameter	Gain Margin for positive M	Gain Margin for negative M
ω_n ($\zeta \lesssim 0.5$)	-	\Downarrow
ω_n ($\zeta \gtrsim 0.5$)	-	\Uparrow
ζ	-	\Uparrow
ω_0 ($n = 1, 2$)	\Downarrow	-
ω_0 ($n \geq 3$)	-	-

Table entries represent the change in gain margin as each parameter *increases* for typical back-drivable robot parameters.

TABLE III
REPRESENTATIVE ROBOT AND ENVIRONMENT PARAMETERS

Device size	m	b	Environment	m_e	b_e	k_e
small	50	0.05	soft	100	1.0	10
medium	200	0.15	medium	200	5.0	500
large	500	1.00	rigid	50	1.0	1000

Device size	Environment					
	soft		medium		rigid	
	ω_n	ζ	ω_n	ζ	ω_n	ζ
small	8.2	0.43	44.7	0.23	100.0	0.05
medium	5.8	0.33	35.4	0.18	63.2	0.04
large	4.1	0.41	26.7	0.16	42.6	0.04

The device parameters were inspired by Diolaiti et al. (2006), and the environment parameters by Kuchenbecker et al. (2003) and Willaert et al. (2011). Masses m, m_e are in units of g, damping coefficients b, b_e in Ns/m, stiffness k_e in N/m, system natural frequency ω_n in rad/s, and system damping ratio ζ is dimensionless.

affect positive virtual mass stability, but can directly determine negative virtual mass stability. If the assumptions do not hold, then all the parameters can significantly effect the stability of the system. In other words, it is possible to choose parameter values in which environment coupling strongly affects positive mass rendering, or filtering significantly affects negative mass stability.

The virtual mass to robot plus environment mass ratio, $M/(m + m_e)$, is a gain of the system and directly affects the gain margin of $L(s)$. The range of virtual mass ratios for which the system is stable (in the BIBO sense) can be determined by evaluating the gain margin for $L(s)$ and $-L(s)$. For example, the Bode plot of Figure 3 at a sample rate of 1000 Hz and a filter order of 3, is about 20 dB (10) at ω^* .

Thus, for positive M , values of $M/(m + m_e)$ up to 10 are stable. For negative M , the magnitude of the system is about -5 dB (0.56) at ω^* , so values of $|M/(m + m_e)|$ up to 0.56 are stable. This establishes a range of stable *commanded* virtual mass values, but we will see in Section VI that the *closed-loop apparent* mass of the system may not actually match the commanded value. Also, this analysis is based on the continuous model whose stability boundaries are similar, but not identical, to the linear hybrid and discrete model stability boundaries.

B. Stability Boundaries

In this section, analytical stability boundaries are formed using the discrete model. First, we present a general condition to determine stability given system parameters. Then, we analyze the relationship between passivity and stability for the system. Finally, we form stability boundaries for important parameters in the cases of positive and negative virtual mass.

1) *General Stability Condition*: The discrete model shown in Figure 1(d) will be stable (ensuring the stability of the non-linear hybrid model shown in Figure 1(b)), if and only if

$$GM[L(z)] > 1, \quad (21)$$

where

$$L(z) = G(z)D(z)H(z)^n z^{-k} M. \quad (22)$$

The evaluation of the gain margin of $L(z)$ is straightforward for given system parameters. However, because system parameters m_e , ω_n , and ζ change with environmental interaction, stability of the system requires stability of the *set* of possible parameter values. Software to generate the stability boundaries can be found in Appendix A (Extension 1).

2) *Interaction With the Environment*: Environmental interaction enters the system from m_e , b_e , and k_e , and the effect on stability is parametrized by m_e , ζ , and ω_n . The passivity results of Section III establish necessary and sufficient conditions for which the coupled closed-loop robot and *any* passive impedance are stable, but because real environments are capable of rendering only a subset of passive impedances, these results can be overly conservative. Figure 5 shows the stable ranges of $|M|/m$ for four possible environments: (1) any passive environment, (2) uncoupled (no environment), (3) spring impedance environment, and (4) spring-damper impedance environment. The device and feedback loop parameters are that of Figure 3. Informally, a high-stiffness spring can be interpreted as a “worst-case,” coupling environment (Colgate and Schenkel, 1997) (e.g., a rigid physical wall), and the spring-damper as a human operator impedance because human bio-mechanics necessarily add damping when adding stiffness (Diolaiti et al., 2006; Kuchenbecker et al., 2003). We see that for positive mass rendering, spring

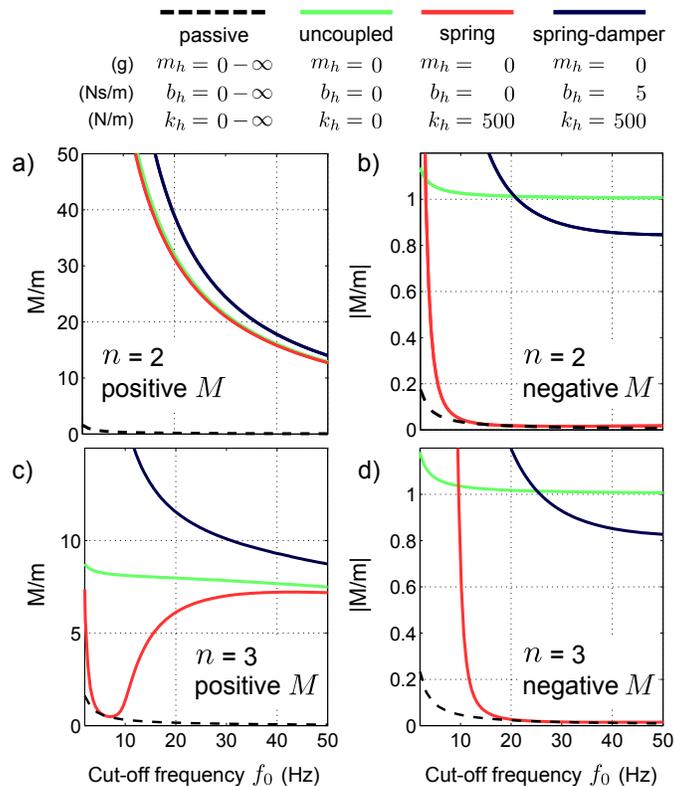


Fig. 5. Stable ranges of $|M|/m$ for positive and negative virtual mass for two filter orders ($n = 2$ and 3) and different environment impedances. The stable regions are *under* the curves. (a) Positive M , $n = 2$, (b) negative M , $n = 2$, (c) positive M , $n = 3$, (d) negative M , $n = 3$ (the top row is for $n = 2$ and the bottom row is for $n = 3$, and left column is for positive M and the right column is for negative M). For positive M , the spring-damper stability region are larger than the uncoupled regions, but not for negative M .

stability is similar to uncoupled stability except if very aggressive filtering is applied. Spring stability for negative virtual mass, in contrast, is similar to passivity. For positive M , the uncoupled stability region is a subset of the spring-damper stability region, but not for negative M . This qualitative result does not depend on the order of the filter. This analysis suggests that for positive mass rendering, except for extremely rigid coupling dynamics, the environment *aids* stability. I.e., the uncoupled closed-loop robot has a “baseline” stability that can be improved by the environment. Therefore, to analyze the stability of positive virtual mass rendering, we can consider the stability of the uncoupled ($m_e = b_e = k_e = 0$) system. This is similar to the technique used to analyze the stability of virtual spring and damper rendering, in which for reasonable environment coupling conditions, the environment is assumed to aid stability (Hulin et al., 2008; Gil et al., 2009a). For negative mass rendering, however, environmental interaction is important for coupled stability and must be analyzed directly.

3) *Positive and Negative Virtual Mass Stability Boundaries*: Figure 6 shows stability boundaries for a family of systems with *positive virtual mass* and no delay. Large values of M/m are stable, for example, a M/m ratio of 25 for a filter order of $n = 2$, and a ratio of 8 for $n = 3$, at a 1 kHz sample rate and

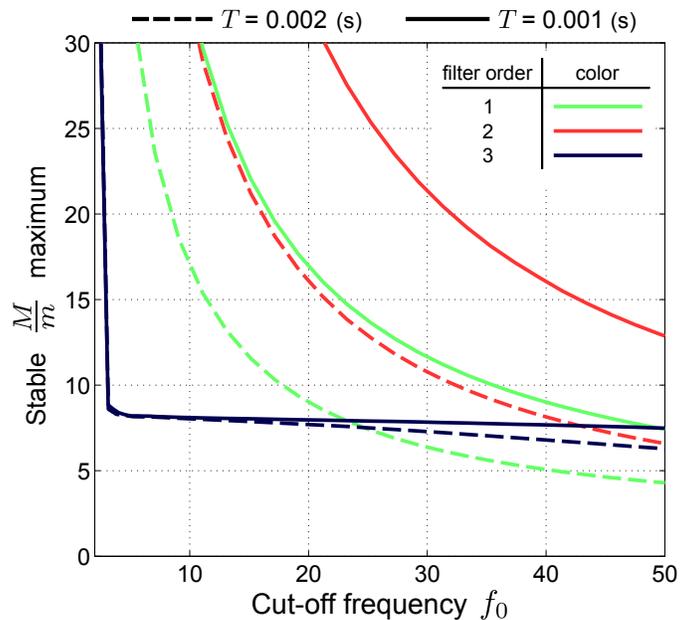


Fig. 6. Positive M uncoupled stability boundaries with no time delay. The stable regions are *under* the curves. The stability boundaries are largely determined by the sample time and cut-off frequency for filter orders 1 and 2, but not for higher filter orders.

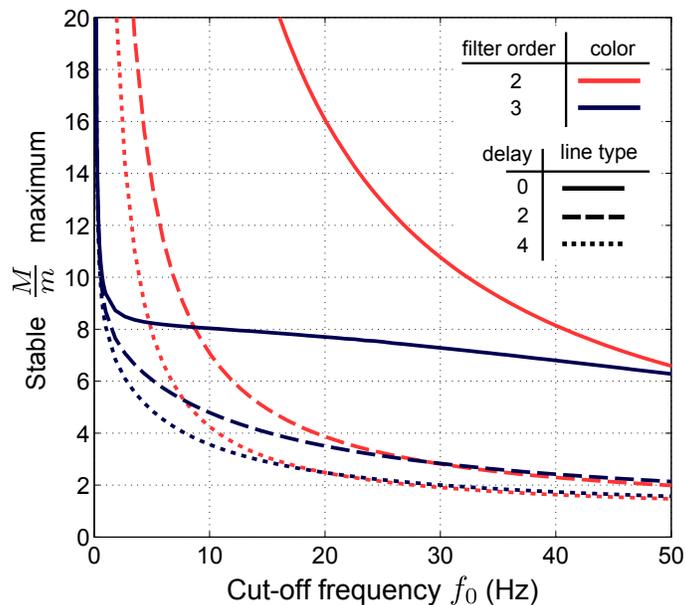


Fig. 7. Positive M uncoupled stability boundaries with time delay. The stable regions are *under* the curves. Small delays significantly reduce the stable range of virtual mass to device mass.

a cut-off frequency of 25 Hz. Stability for positive mass rendering is sensitive to delay in the feedback loop. Figure 7 displays the effect of delay on the system with $n = 2$ and 3 at a sample rate of 1 kHz. When delay is introduced, the stability boundaries reduce significantly. We will see in Section V that for most current robot parameters, noise, not stability, is the limiting factor for positive mass rendering.

Unlike positive virtual mass, interaction from the environment can have a strong effect on negative

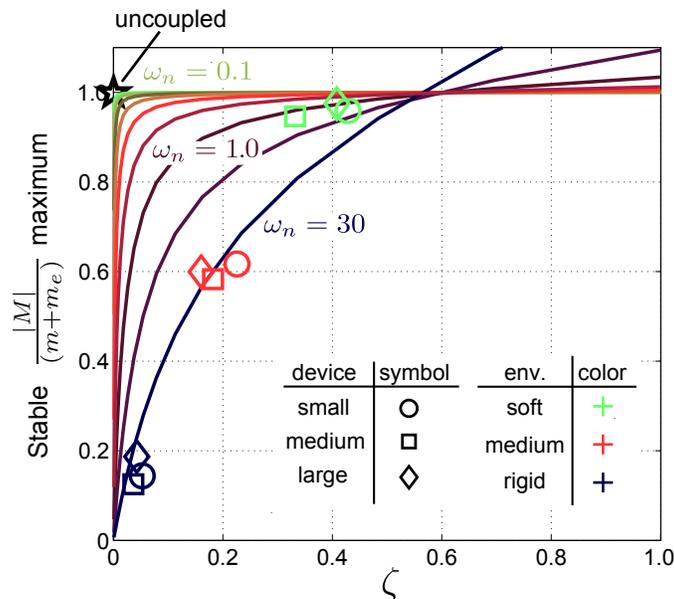


Fig. 8. Negative M stability boundaries versus ω_n and ζ showing the effect of environment coupling on stability. The stable regions are *under* the curves. Larger mass compensation is stable for larger ζ . High-stiffness, low-damping coupling environments can be unstable with small amounts of mass compensation (20%), but low-stiffness, high-damping environments are stable with large amounts (90%). *Completely* compensating the (commanded) mass of the uncoupled robot is stable.

virtual mass stability. Figure 8 shows the stability boundaries for *negative virtual mass*. Interestingly, with no environment interaction, it is stable to *completely* compensate the (commanded) mass of the robot, but coupled environment dynamics can drive the system unstable for low values of ζ . High-stiffness, low-damping, “rigid” coupling environments can be unstable with small amounts of mass compensation (20%), but low-stiffness, high-damping, “soft” environments are stable with large amounts of mass compensation (90%). However, we will see in Section VI and VII that the apparent mass may not match the commanded value, and that the closed-loop robot can be stable, but be non-passive at frequencies relevant to operator input.

V. NOISE CHARACTERIZATION

Direct mass rendering requires an acceleration estimate output from double differentiation on position. Characterization of the resulting noise and its impact on mass rendering is the focus of this section. Error is introduced into the system by quantization of the position, usually corresponding to an optical encoder. This error can be amplified, manifesting as noise that causes rumble, i.e., inconsistent vibrations that affect the quality of the rendering. Because quantization introduces finite error, it cannot cause a stable system to go unstable in the BIBO sense (Franklin et al., 1997); however, noise can be so large that practically speaking the system is “unstable,” and the limiting factor in the mass rendering is noise.

In this section we first predict the existence and qualities of sustained oscillations (limit cycles) with

a describing function analysis of quantization using the continuous model. Then, we establish conditions for which the Coulomb friction of the device will prevent limit cycles using transfer function norms and the sub-multiplicative property.

A. Describing Function Analysis

Describing function analysis (Slotine and Li, 1991) provides a tool to analyze “average” behavior in which the frequency, amplitude, and stability of self-sustaining oscillations from quantization can be predicted.

Figure 9 shows the continuous system of Figure 1(c) with all the linear components lumped together,

$$L(s) = D(s)H(s)^n MG(s)e^{-j(T/2+t_d)}, \quad (23)$$

and the quantizer in the feedback loop. For convenience we define $\bar{L}(s)$ as $L(s)$ with the virtual mass M factored out

$$L(s) = \bar{L}(s)M. \quad (24)$$

If we assume a sinusoidal input to the quantizer of

$$A \sin(\omega t), \text{ for } A > 0, \omega > 0, \quad (25)$$

then the input-output properties of the non-linear quantizer can be approximated by its linear describing function. The describing function of quantization is given by

$$N(A) = \frac{2}{\pi A} + \frac{4}{\pi A^2} \sum_{l=1}^{\lfloor A \rfloor} \sqrt{A^2 - l^2}, \quad (26)$$

where A is the amplitude in quantization intervals (Slotine and Li, 1991). Equation (26) is illustrated in Figure 9 (d). Note that quantization is static and odd with respect to position, so $N(A)$ is real and not a function of frequency ω (Khalil, 2001).

From the Nyquist criterion, sustained oscillations are predicted to occur (without Coulomb friction) if

$$\bar{L}(j\omega)MN(A) = -1 \text{ for } A > 0, \omega > 0. \quad (27)$$

The magnitude, frequency, and stability of the limit cycles can be determined by finding conditions for which Equation (27) is satisfied. Figure 9(c) displays the Nyquist contour of $L(j\omega)$ for positive ω , and $1/N(A)$ and $-1/N(A)$ for positive A . The parameters for $L(j\omega)$ are that of Figure 3, but its general form is not sensitive to these parameters. Limit cycles are predicted to occur if $\bar{L}(j\omega)M = -1/N(A)$ for positive virtual mass, or if $\bar{L}(j\omega)|M| = 1/N(A)$ for negative virtual mass. Figure 9 shows that Equation (27) is satisfied for both signs of the virtual mass, and that limit cycles are predicted to occur for both

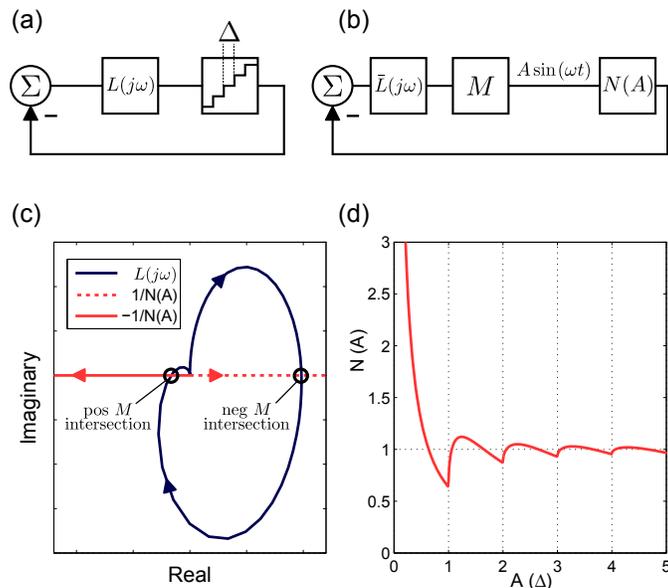


Fig. 9. (a) Block diagram of the lumped linear system and quantizer. (b) Block diagram of the system and quantizer describing function. (c) Nyquist contour of $L(j\omega)$ without Coulomb friction showing the existence of limit cycles for positive and negative M . The arrows of $L(j\omega)$ represent the direction of increasing ω , and the arrows of $1/N(A)$ and $-1/N(A)$ represent the direction of increasing A . (d) Quantizer describing function.

positive and negative M . This confirms that “noise” limit cycles occur because of quantization and not Coulomb friction. In fact, we will find that Coulomb friction can actually prevent these limit cycles.

The solution sets can be solved for two separate cases of the input amplitude to quantization. The stability of the limit cycles is determined using an extension of the Nyquist criterion (Khalil, 2001). The first case is large amplitude ($A > \Delta$), where $N(A)$ tends to unity. In this case, the limit cycles are unstable, and the analysis reduces to the standard stability analysis which is covered in Section IV. The second case is small input to the quantizer ($A < \Delta$), in which the limit cycles are stable. For both positive and negative virtual mass, there is a unique frequency, ω^{LC} , in which Equation (27) is satisfied

$$\omega^{LC} = \omega \quad \text{such that} \quad \bar{L}(j\omega)M = -1/N(A). \quad (28)$$

Real robotic systems, however, have Coulomb friction which may prevent these limit cycles, as well as change their frequency. It is possible to form a describing function for $L(j\omega)$ including Coulomb friction in which its effect can be analyzed directly (Diolaiti et al., 2006). This approach requires solving coupled nonlinear equations for which finding solutions, or guaranteeing that none exist, is difficult. In our analysis, this technique applied to mass rendering is prohibitively cumbersome, so instead we take another approach. Limit cycles can only occur if the energy created from quantization error is greater than the energy dissipated. If the energy dissipation of the robot is greater than the quantization error “noise energy” for all possible trajectories of the robot, then limit cycles cannot occur, and noise will not be a

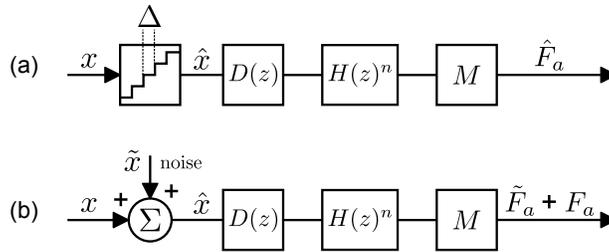


Fig. 10. Quantization noise modeling. (a) Quantization operation: true position x is quantized with quantization interval Δ , to output quantized position \hat{x} , which propagates to produce actuator force derived from quantized position \hat{F}_a . (b) Quantization modeling: noise $\tilde{x} \in \mathcal{R}[-\Delta/2, \Delta/2]$ is added to the true position, which propagates to produce “noise” actuator force \tilde{F}_a .

limiting factor.

B. Worst-case Energy From Quantization and Noise Boundaries

In this section we bound the maximum energy generated from quantization error, and establish a sufficient condition for which the quantization error “noise energy” will be dissipated by the friction of the robot.

The quantizer introduces error at the position measurement, which can be modeled as additive noise. Figure 10 shows a model of quantization as a noise source. True position x is quantized with quantization interval Δ , to output quantized position \hat{x} , which propagates through part of the feedback loop to produce the actuator force derived from the quantized position \hat{F}_a . The error on position from quantization, $\tilde{x} = \hat{x} - x$, also propagates to produce the actuator error force \tilde{F}_a .

The net generated energy of the closed-loop robot considering the energy created from quantization error and dissipation of the device over a sample period is

$$E_{\text{total}}(x, \tilde{x}) = \tilde{E}_{\text{gen}}(x, \tilde{x}) - E_{\text{diss}}(x). \quad (29)$$

The generated energy from quantization error over a sample period is

$$\tilde{E}_{\text{gen}}(x, \tilde{x}) = \int_0^T \tilde{f}_a(\tilde{x}) \dot{x} dt, \quad (30)$$

where

$$\tilde{f}_a(\tilde{x}) = \sum_{k=0}^p q(k) \tilde{x}(p-k) \quad (31)$$

is the convolution sum of the position error \tilde{x} with the impulse response of the haptic control law $Q(z)$, $q(k)$, with p non-zero terms. The actuator “noise force,” $\tilde{f}_a(\tilde{x})$, is held constant over the sample period by the zero-order hold. The dissipated energy of the robot from Coulomb friction and viscous damping over a time period is

$$E_{\text{diss}}(x) = \int_0^T (c \operatorname{sgn}(\dot{x}) + b \dot{x}) \dot{x} dt. \quad (32)$$

With respect to the trajectory of true position x over a sample period, the energy dissipation of Coulomb friction is minimized for monotonic trajectories with no stops of finite time (Abbott and Okamura, 2005). Because of this, the nonlinear switching nature of Coulomb friction is not present. An expression for the minimum energy dissipation of viscous damping is obtained from the Cauchy-Schwarz theorem. Therefore the energy generation of the closed-loop robot must be less than

$$\tilde{E}_{\text{gen}}(\tilde{x}, x) - E_{\text{diss}}(x) = \sum_{k=0}^p q(k)\tilde{x}(p-k)|\delta x| - c|\delta x| - \frac{b}{T}|\delta x|^2 \quad (33)$$

where δx represents the distance travelled over a sample period. Note that the maximum energy generation with respect to the trajectory of x is a function of the true position at sample times only.

Because the magnitude of the position error at each time step is, at maximum, half the quantization interval, $|\tilde{x}| \leq \frac{\Delta}{2}$ (Bertram, 1958), the actuator noise force can be bounded

$$\sum_{k=0}^p q(k)\tilde{x}(p-k) \leq \sum_{k=0}^p |q(k)| \left| \frac{\Delta}{2} \right| = \|Q(z)\|_{\ell_1} \frac{\Delta}{2}, \quad (34)$$

where $\|Q(z)\|_{\ell_1}$ is the ℓ_1 norm on the discrete-time control law $Q(z)$ (Zhou et al., 1996).

Plugging in Equation (34) into Equation (33), we obtain an upper bound on the worst-case energy generation of the closed-loop robot from quantization error

$$\left[\tilde{E}_{\text{gen}}(\tilde{x}, x) - E_{\text{diss}}(x) \right]_{\max} = \|Q(z)\|_{\ell_1} \frac{\Delta}{2} |\delta x| - c|\delta x| - \frac{b}{T}|\delta x|^2. \quad (35)$$

As δx approaches zero, the energy dissipation from viscous damping is negligible compared to the energy dissipation from Coulomb friction. A sufficient condition for no energy generation of the closed-loop robot due to quantization error over any sample period is

$$c \geq \|Q(z)\|_{\ell_1} \frac{\Delta}{2}. \quad (36)$$

For mass rendering, $Q(z) = D(z)H(z)^n M$, and this condition becomes

$$c \geq \|D(z)H(z)^n\|_{\ell_1} |M| \frac{\Delta}{2}. \quad (37)$$

Inequality (37) provides an analytical expression for the quantization “noise-free” region. Software to generate the quantization noise-free boundaries can be found in Appendix A (Extension 1). If the Coulomb friction in the robot, c , is sufficiently large, then quantization noise cannot generate energy, and will not be perceptible. It is a *sufficient* condition for no energy generation based on a worst-case analysis, and can be conservative. In other words, there are conditions where Inequality (37) is not satisfied, but trajectories of x exist such that quantization error will not generate energy, e.g., a case in which the velocity of the robot is large enough so that the energy dissipation due to viscous damping is significant. Figure 11 shows noise boundaries for filter orders $n = 1, 2, 3$ versus f_0 . These results are similar to those of Abbott and

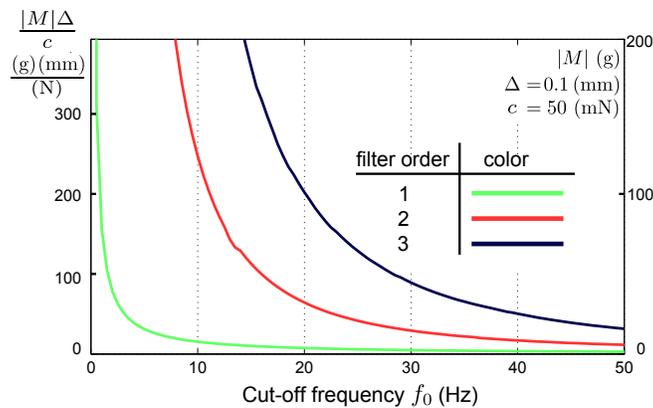


Fig. 11. Noise thresholds for mass rendering generated by Inequality (37). For values of $M\Delta/c$ under the boundary, we predict noise limit cycles will not occur.

Okamura (2005) and Diolaiti et al. (2006), which established that for a spring rendering with stiffness K , the effect of quantization error noise will be perceptible only if the energy leaks from quantization will not be entirely dissipated by Coulomb friction, $c > K \frac{\Delta}{2}$.

Increasing the robot Coulomb friction c , decreasing the position sensing quantization interval, Δ , and more aggressive low-pass filtering allow a larger quantization noise-free range of virtual mass. Unlike stability, where the *ratio* of virtual mass to device plus environment mass is important, quantization error noise scales *directly* with the virtual mass.

VI. ACCURACY OF CLOSED-LOOP RENDERING

We are interested not only in stability and noise characterization, but also in the accuracy of the mass rendering. We investigate this using two different approaches. The first approach is presented in this section and analyzes the theoretical closed-loop impedance of the robot to present how the system feels at different frequencies. The second approach is to perform a time-based system identification on experimental data, which is presented in Section VII.

For the analysis we use the continuous model (Figure 1(c)) and examine accuracy of the closed-loop robot by comparing the impedance of three systems. The first system is the ideal case in which the apparent dynamics of the robot are only the desired mass, which is the sum of the device and virtual mass ($m + M$). The virtual mass can increase (positive M) or decrease (negative M) the apparent mass. The closed-loop impedance of this “ideal” system is

$$Z_{\text{ideal}}(s) = \frac{F_e(s)}{V(s)} = (m + M)s. \quad (38)$$

The second system is the open-loop robot, that is, the robot with no feedback loop. The open-loop robot impedance is given by

$$Z_{\text{open-loop robot}}(s) = \frac{F_e(s)}{V(s)} = ms + b. \quad (39)$$

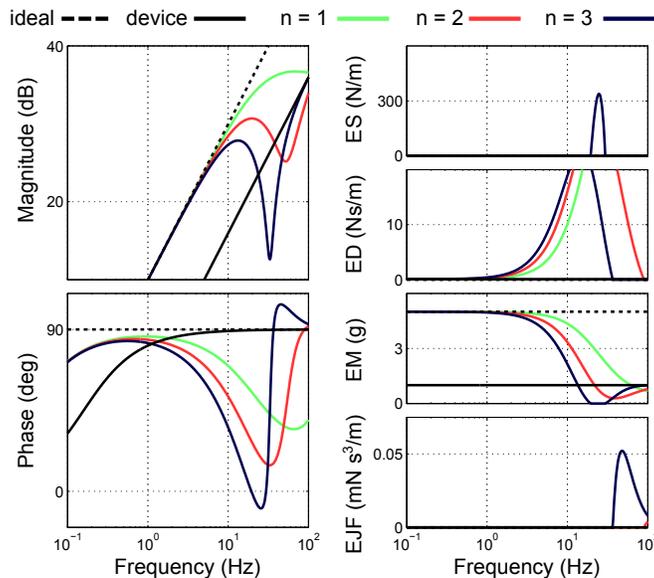


Fig. 12. Bode (left column) and effective impedances (right column) of the systems for *positive* virtual mass $M = 4m$. ES is effective stiffness, ED, effective damping, EM, effective mass, and EJF, effective jerk feedback. The full-feedback closed-loop robot system has impedances that are close to ideal at low frequencies with respect to the cut-off frequency, but deviate from ideal for higher frequencies. At high frequencies, the full-feedback system has lower effective mass than desired, and significant effective damping. ED is energy *dissipating*, EJF is energy *generating*, while ES and EM are energy *conserving*.

The third system is the closed-loop robot with the feedback loop implementing a virtual mass,

$$Z_{\text{closed-loop robot}}(s) = \frac{F_e(s)}{V(s)} = \frac{1 + G(s)D(s)H(s)^n M}{sG(s)}. \quad (40)$$

Accuracy can be analyzed by examining how the impedance of the full system with feedback (the third case) compares to the other two. For high-quality, the system impedance should be similar to ideal (the first case). We present the system impedances with their Bode plots, as well as their “effective impedances” (Colonnese et al., 2014). Effective impedances present the same information as other frequency response tools, such as Bode or Nyquist plots, but they express the system response by the use of physical analogs.

For the analyses we use the device model from the Phantom Premium 1.5 (Table IV), a cut-off frequency of 25 Hz, a sample rate of 1 kHz, and investigate filter orders of 1, 2, and 3.

Figure 12 shows the Bode and effective impedance plots for positive virtual mass, $M = 4m$. For low frequencies (less than 3 Hz) the system with full-feedback is similar to ideal for all filter orders. From the Bode plots, the magnitude and phase of the full-feedback system is close to the magnitude and phase of the ideal system. From the effective impedances, the effective mass ratio is close to 5 ($(m + M)/m = (m + 4m)/m = 5$), and the effective damping is just that of the open-loop robot, b . For higher frequencies, however, the full-feedback system deviates from ideal behavior. From the Bode plots, the magnitude and phase of the full-feedback system differs from the magnitude and phase of the ideal system. From the effective impedances, the full-feedback system has less effective mass than ideal,

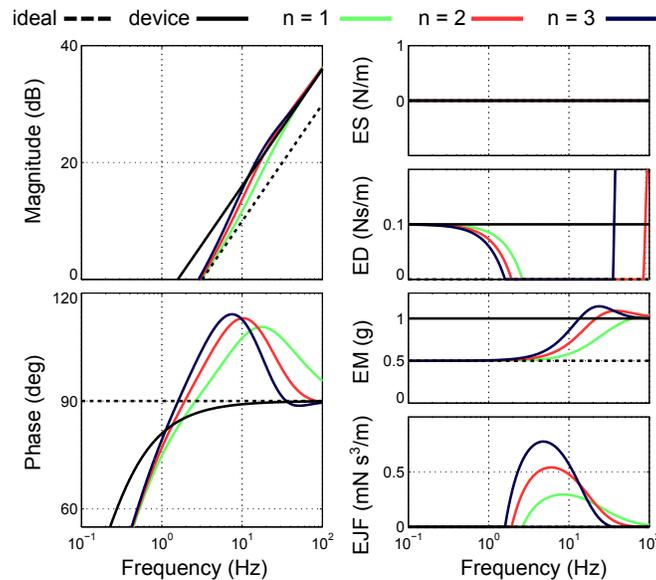


Fig. 13. Bode (left column) and effective impedances (right column) of the systems for *negative* virtual mass $M = -0.5m$. ES is effective stiffness, ED, effective damping, EM, effective mass, and EJF, effective jerk feedback. The full-feedback closed-loop robot system has impedances that are close to ideal at low frequencies with respect to the cut-off frequency, but deviates from ideal for higher frequencies. At high frequencies, the full-feedback system has higher effective mass than desired, *reduced* effective damping, and is non-passive (corresponding to frequencies that the phase of the impedance is greater than 90° , or that have non-zero effective jerk feedback.)

and has significant effective damping. The higher the filter order, the larger the difference between the full system and ideal impedance. For very high frequencies, the full-system feedback loop contributes relatively little force compared to the open-loop dynamics, and the Bode plot and effective impedances are similar to that of the open-loop device.

Figure 13 shows the Bode and effective impedance plots for negative virtual mass, $M = -0.5m$. Like the positive virtual mass case, the system with full-feedback is similar to ideal for all filter orders at low frequencies. From the Bode plots, the magnitude and phase of the full-feedback system is close to the magnitude and phase of the ideal system. From the effective impedance plots, the effective mass ratio is close to 0.5, and the effective damping is just that of the open-loop device, b . Also like the positive mass case, the full-feedback system deviates from ideal behavior at high frequencies (more than 3 Hz). From the Bode plots, the magnitude and phase of the full-feedback system differs from the magnitude and phase of the ideal system. The system phase is greater than 90° representing frequencies where the systems are non-passive. From the effective impedances, the full feedback system has larger effective mass, *reduced* effective damping, and significant jerk (the derivative to acceleration) feedback. Jerk feedback, which can be interpreted as negative damping, is non-zero when the system impedance is between 90° and 270° and represents frequencies at which the closed-loop robot is capable of generating energy. Like positive M , the negative M full-feedback system is similar to the open-loop device at very high frequencies.

This analysis shows that the cut-off frequency and filter order play a powerful role in determining the

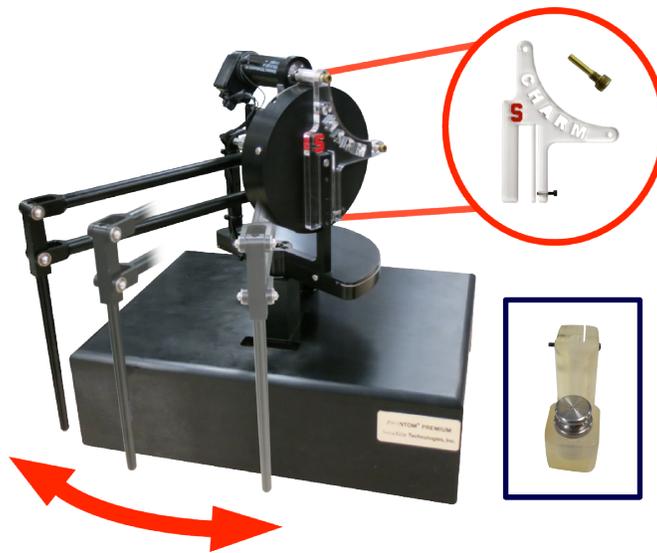


Fig. 14. Experiments were performed on a single joint of the Phantom Premium 1.5. All other joints were mechanically constrained by a custom-made “motor stop” displayed in the circle. Environment mass (displayed in the rectangle) was added by affixing a weight to the Phantom arm stylus.

accuracy of the rendering. The higher the cut-off frequency, or lower the filter order, (corresponding to less aggressive filtering), the larger the frequency span in which the system behaves like the ideal. If the cut-off frequency is low enough (to satisfy stability or noise rejection objectives) so that non-ideal system impedance behavior is present at frequencies of interaction, positive mass renderings will feel lighter and more damped than ideal, and negative mass renderings will feel heavier and “jerkier” than ideal, potentially with active behavior.

VII. STABILITY, NOISE CHARACTERIZATION, AND ACCURACY EXPERIMENTS

In this section we present experimental results and compare the data to theoretical predictions. The experiments were conducted on a single joint of the Phantom Premium 1.5 robot (the first revolute joint), while the other degrees of freedom were mechanically constrained by a custom-made fixture (Figure 14). A system identification procedure was performed to find the mass, viscous damping, and Coulomb friction of the open-loop device, similar to the procedure described in Section VII-C. The estimates for the robot averaged over ten trials are shown in Table IV.

A. Stability Experiments

Experimental stability boundaries were formed to compare to theoretical ones. Environment coupling mass was added by affixing a weight to the Phantom Premium arm stylus (Figure 14). Because physical spring and damper environment coupling is difficult to implement and control precisely, virtual spring and

TABLE IV
ESTIMATED ROBOT MECHANICAL PARAMETERS

	Estimated mass	Estimating damping	Estimated Coulomb friction
	\hat{m}	\hat{b}	\hat{c}
Objective Function	(g)	(Ns/m)	(N)
ℓ_2 , sum of squares	94	0.13	0.11
ℓ_1 , sum of absolute values	88	0.14	0.08

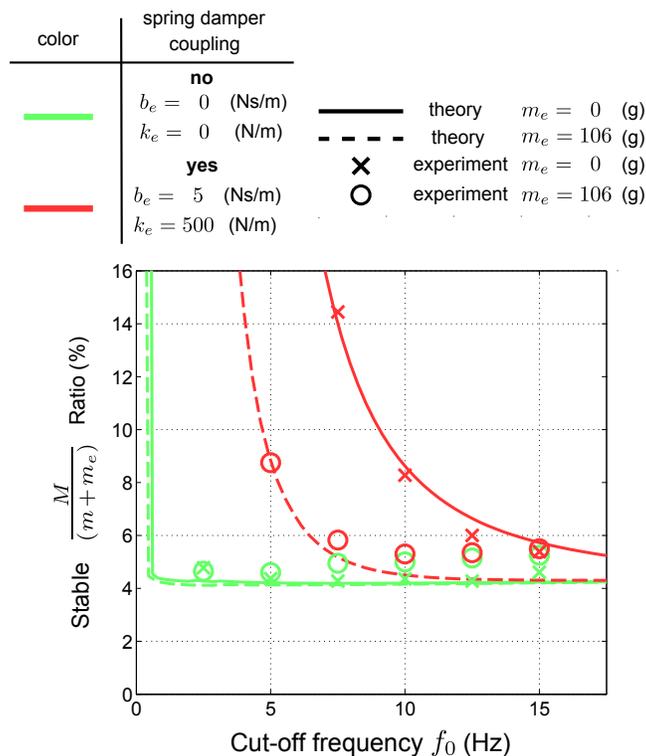


Fig. 15. Positive M experimental data for a fourth-order filter and various cut-off frequencies.

damper coupling was used. For an experiment, the magnitude of the virtual mass was increased in small intervals starting from zero, and the system response to a single strong impulse was measured. Instability was detected at the virtual mass value where growing oscillations were detected. The theoretical bounds were found using the general stability condition (21) and the device model of the ℓ_2 objective function of Table IV.

Positive M experimental stability bounds were found for a fourth-order filter with various cut-off frequencies, a sample rate of 1000 Hz, and no external delay. Figure 15 shows the theoretical and experimental data. The theoretical predictions are close to the experimental data.

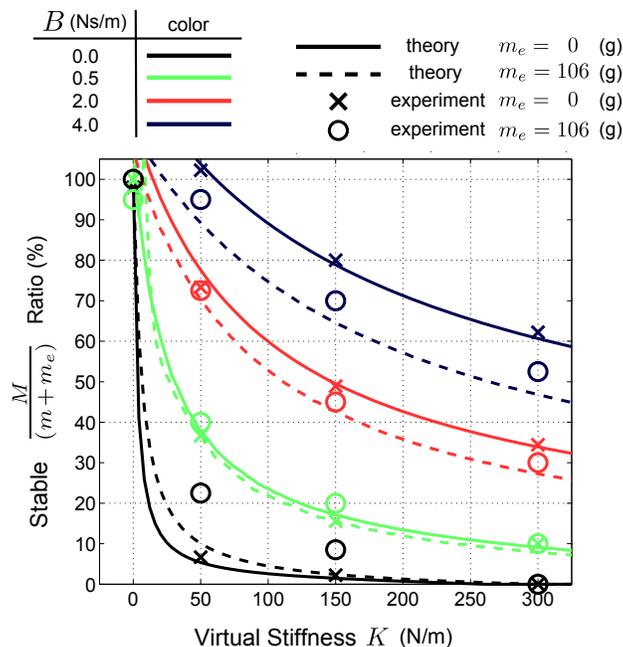


Fig. 16. Negative virtual mass experimental data. The theoretical predictions are close to the experimental data. Smaller K and larger B allow a larger range of stable mass compensation.

Negative M experimental stability bounds were found for systems with different values of the virtual stiffness, K , and virtual damping, B , for a sample rate of 1000 Hz, a cut-off frequency of 30 Hz, a filter order of 3, and no external delay. Experimental results are shown in Figure 16. The experimental data show that it is indeed stable to compensate nearly all of the robot plus environment mass for low-stiffness environment coupling. Also predicted by theory, larger ranges of mass compensation are stable for lower stiffness and higher damping.

B. Noise Characterization Experiments

Noise experiments were conducted for different filter orders and cut-off frequencies. Similar to the stability experiments, the magnitude of the virtual mass was increased in small intervals starting from zero, and the system response to a single strong impulse was measured. A “noise boundary” was detected when sustained (bounded) oscillations were detected. For negative M , instability occurs at smaller virtual mass values than noise limit cycles for the parameters of the Phantom Premium 1.5, so the experiments were all conducted with positive virtual mass. Figure 17 displays the theoretical and experimental noise boundaries. The theoretical noise boundaries were formed with the natural quantization of the Phantom Premium, $\Delta = 52.5\mu\text{m}$, and the Coulomb friction from the ℓ_2 objective in Table IV, $c = 0.11$ N.

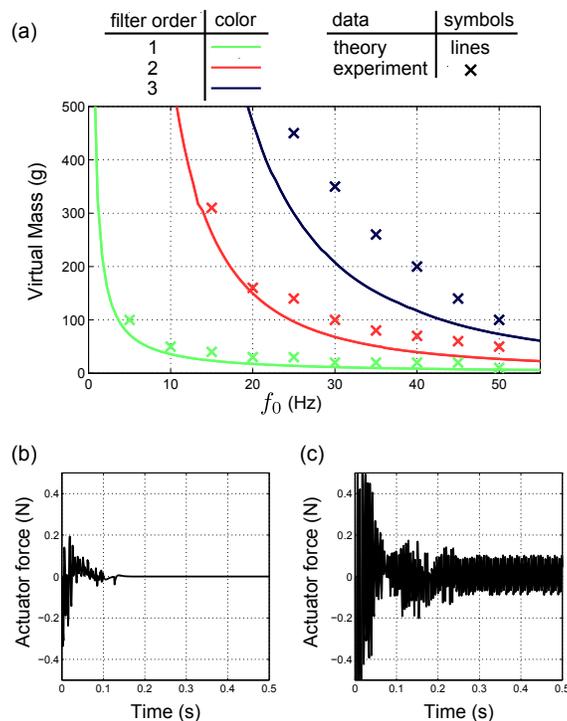


Fig. 17. (a) Experimental noise boundaries. (b), (c) Actuator force time responses to impulse input for $n = 2$, and $f_0 = 30$ Hz, for (b) no limit cycles at $M = 50$ g and (c) limit cycles at $M = 150$ g.

C. Accuracy Experiments

Here we describe the accuracy analysis done by performing a time-based system identification on experimental data resulting in estimates for the mass and viscous damping of the closed-loop system. The experiments varied the commanded virtual mass (0 to $5m$ for positive M , 0 to $-0.5m$ for negative M) for two cut-off frequencies: 5 and 15 Hz. The Phantom Premium 1.5 robot was excited with an exogenous chirp force from 0.1 to 5 Hz over 20 seconds, and the position was measured.

For the experiments we modeled the closed-loop robot as a mass and viscous damper. With this model, ideally, at every instant of time the system should satisfy

$$f(t) = \hat{m}\ddot{x}(t) + \hat{b}\dot{x}(t), \quad (41)$$

where $f(t)$ is the external force, $\ddot{x}(t)$ and $\dot{x}(t)$ are the acceleration and velocity of the device, respectively, and \hat{m} and \hat{b} represent the mass and damping of the system, respectively. The velocity and acceleration signals used in the system identification were generated by numerical differentiation and smoothed with non-causal low-pass filters with no phase lag.

For a particular experiment \hat{m} and \hat{b} were computed via optimization,

$$\hat{m}_{\ell_2}, \hat{b}_{\ell_2} = \operatorname{argmin} \sum (f_i - \hat{m}_{\ell_2}\ddot{x}_i - \hat{b}_{\ell_2}\dot{x}_i)^2, \quad (42)$$

$$\hat{m}_{\ell_1}, \hat{b}_{\ell_1} = \operatorname{argmin} \sum |f_i - \hat{m}_{\ell_1}\ddot{x}_i - \hat{b}_{\ell_1}\dot{x}_i|, \quad (43)$$

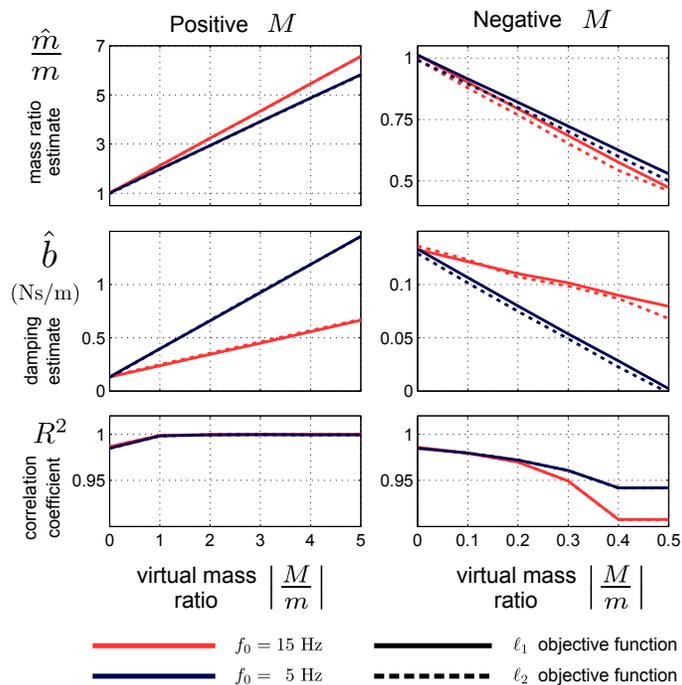


Fig. 18. System identification experiments for varying virtual mass and a cut-off frequency. Both signs of M have estimated mass ratios close to ideal, however, the estimated damping is affected: larger damping for positive M and reduced damping for negative M .

where f_i , \ddot{x}_i , and \dot{x}_i represent the signal values at a specific instant in time, and the subscripts ℓ_2 and ℓ_1 refer to the sum of squares and sum of absolute values objective functions respectively. The classic pseudo-inverse was used to solve the ℓ_2 optimization problem (Strang, 1988). For the ℓ_1 problem, we used CVX, a package for solving convex programs (Boyd and Vandenberghe, 2004), Grant and Boyd (2012).

Figure 18 shows the results of the system identification experiments. For positive virtual mass, the estimated mass ratio (\hat{m}/m) increased linearly with increasing virtual mass. For increasing virtual mass, the estimated damping of the system increases linearly. The lower cut-off frequency results in higher estimated damping. For negative virtual mass, the estimated mass ratio decreases with with larger $|M|$. The estimated damping also reduces with increasing $|M|$, and the reduction for $f_0 = 5$ is much larger than for $f_0 = 15$. These experimental results match the the theoretical predictions of Section VI. For all experiments, there was no large difference between the ℓ_1 and ℓ_2 estimates, showing that the results are not sensitive to the optimization objective function.

VIII. CONCLUSIONS AND FUTURE WORK

The main results of our research are the identification of important parameters for system passivity and stability, the generation of passivity and stability boundaries, noise characterization, and an analysis of

the mass rendering accuracy of the closed-loop robot. We verify our analytical results with experimental data using a Phantom Premium 1.5. Our results serve as a design tool for robotic devices and controllers to improve rendering of virtual mass. This applies to mass compensation for exoskeletons and haptic interfaces, as well as rendering increased or decreased mass in rehabilitation and motor training experiment scenarios.

Using the results provided, system parameters to fulfill stability, noise rejection, and accuracy objectives can be chosen to create high-quality mass rendering for robotic systems. Because trade-offs exist between these objectives (e.g., more aggressive filtering results in less noise, but also less accuracy), there is no single optimal system design.

For positive virtual mass rendering, large M/m ratios are stable and insensitive to environment dynamic coupling instability. Therefore, the main trade-off occurs between noise rejection and accuracy. Figures 6 and 12 provide design aids in terms of system parameter regions for achieving noise and accuracy objectives, respectively; the amount of filtering and virtual mass can be varied for the desired properties. Because the indirect positive mass rendering method (Brown and Colgate, 1998; Adams and Hannaford, 1999) does not require performing a double differentiation on position, it is less sensitive to noise, and may outperform the direct method. How the stability and accuracy of the direct and indirect methods compare is subject for future research.

For negative M , the stability of the system relies relatively little on the filtering in the feedback loop and more on the nature of environment interaction. Because of this, noise rejection and accuracy objectives are easy to accomplish compared to stability. Mass compensation design should select the percentage of the open-loop mass to cancel depending on the interaction the closed-loop robot will have with the environment. For many relevant applications, such as rehabilitation and haptic simulations where the environment coupling is a human operator, in which additional mass is introduced, and where human damping increases with human stiffness, (Diolaiti et al., 2006; Kuchenbecker et al., 2003), high percentages of mass compensation can be stable (60 - 80%). However, because large mass compensation can create non-passive robot behavior at low frequencies (Figure 13), the rendering can feel unnatural. For high-stiffness, low-damping and mass, “rigid” coupling environments, such as manipulation of hard physical objects with an exoskeleton, even a small amount of mass compensation (10%) can be unstable.

We chose a very basic low-pass filter type to simplify the analysis. The qualitative results for stability, noise characterization, and accuracy extend to other filter types, where the cut-off frequency and filter order correspond to the bandwidth and roll-off of the frequency response of other low-pass filter types, respectively, presented here. The general stability condition, Equation (21), and general quantization error energy condition, Inequality (37), are applicable to any low-pass filter type.

Section V establishes that Coulomb friction is helpful for dissipating noise from encoder error, but most back-drivable robotic devices are designed to have as little friction as is practical to reduce free space motion forces. This presents a trade-off that can likely be avoided only with the addition of additional hardware or sensors.

Many models could be used to describe a robotic system for mass rendering; our choice of the model shown in Figure 1 was driven by relevance as a practical design tool, tractability for finding 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 robot-environment coupling vibrational modes and multiple degrees of freedom.

APPENDIX A

INDEX TO MULTIMEDIA EXTENSIONS

The multimedia extensions to this article can be found online by following the hyperlinks from www.ijrr.org

Extension	Type	Description
1	Code	<code>ComputeMassRenderingStabilityAndNoiseLimits.m</code> Matlab function that takes in robotic system parameters and outputs stability limits and quantization error noise-free limits for mass rendering.
2	Code	<code>MassRenderingExamples.m</code> Matlab example code that shows how Extension 1 can be used for typical applications.

APPENDIX B

PASSIVITY CONDITION WITH NO LOW-PASS FILTER

Here we derive the passivity condition for pure mass rendering with no low-pass filter (Inequalities 12 and 13). As stated in Section III, the uncoupled linear hybrid model (Figure 1(b) without Coulomb friction, quantization, and the environment) will be passive if and only if

$$b > \frac{T}{2} \frac{1}{1 - \cos(\omega T)} \Re\{(1 - e^{j\omega T})Q(e^{j\omega T})\} \text{ for } 0 \leq \omega \leq \frac{\pi}{T}, \quad (44)$$

where $Q(z)$ represents the discrete-time feedback control law. When no low-pass filter is included in the system

$$Q(z) = D(z)M = \frac{(z-1)^2}{(Tz)^2}M. \quad (45)$$

Substituting Equation (45) into Inequality (44) and using simple algebra, the passivity condition can be simplified to

$$b > \frac{M}{T} \sin\left(\frac{\omega T}{2}\right)^2 \left(8 \sin\left(\frac{\omega T}{2}\right)^2 - 6\right) \text{ for } 0 \leq \omega \leq \frac{\pi}{T}. \quad (46)$$

For positive M , the condition is satisfied over $0 \leq \omega \leq \frac{\pi}{T}$ if

$$M < \frac{bT}{2}, \quad (47)$$

with the right side of Inequality (46) maximized over ω for $\omega = \frac{\pi}{T}$. For negative M , the inequality is satisfied over $0 \leq \omega \leq \frac{\pi}{T}$ if

$$M < \frac{8bT}{9}, \quad (48)$$

with the right side of Inequality (46) maximized over ω for $\omega = \frac{4}{3T}$.

ACKNOWLEDGEMENTS

This work was supported in part by National Science Foundation grant 1217635.

REFERENCES

- 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.
- R. J. Adams and B. Hannaford. Stable haptic interaction with virtual environments. *IEEE Transactions on Robotics and Automation*, 15(3):465–474, 1999. ISSN 1042-296X. doi: 10.1109/70.768179.
- 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.
- G. Aguirre-Ollinger, J. E. Colgate, M. A. Peshkin, and A. Goswami. Inertia compensation control of a one-degree-of-freedom exoskeleton for lower-limb assistance: Initial experiments. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 20(1):68–77, 2012.
- J. E. Bertram. The effect of quantization in sampled-feedback systems. *Transactions of the American Institute of Electrical Engineers, Part II: Applications and Industry*, 77(4):177–182, 1958.
- Nasir Bhanpuri, Allison M. Okamura, and Amy J. Bastian. Predicting and correcting ataxia using a model of cerebellar function. *Brain*, 2014.
- S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, New York, NY, USA, 2004.
- 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.

- J. E. Colgate. Strictly positive real admittances for coupled stability. *Journal of the Franklin Institute*, 329(3):429–444, 1992.
- 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.
- 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.
- N. Colonnese and A. Okamura. M-width: Stability and accuracy of haptic rendering of virtual mass. In *Proceedings of Robotics: Science and Systems*, Sydney, Australia, July 2012.
- N. Colonnese, S. M. Sketch, and A. M. Okamura. Closed-loop stiffness and damping accuracy of impedance-type haptic displays. In *IEEE Haptics Symposium*, pages 97–102, Feb 2014.
- 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.
- G. F. Franklin, M. L. Workman, and D. Powell. *Digital Control of Dynamic Systems*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 3rd edition, 1997.
- G. F. Franklin, D. J. Powell, and A. Emami-Naeini. *Feedback Control of Dynamic Systems*. Prentice Hall PTR, Upper Saddle River, NJ, USA, 4th edition, 2001.
- J. Gil, T. Hulin, C. Preusche, E. Sánchez, and G. Hirzinger. Stability boundary for haptic rendering: Influence of damping and delay. *Journal of Computing and Information Science in Engineering*, 9(1): 1–8, 2009a.
- 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.
- 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, 2009b.
- M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming. <http://cvxr.com/cvx>, 2012.
- P. G. Griffiths, R. B. Gillespie, and J. S. Freudenberg. A fundamental linear systems conflict between performance and passivity in haptic rendering. *IEEE Transactions on Robotics*, 27(1):75–88, 2011.
- D. I. Grow, N. H. Bhanpuri, S. K. Charles, A. J. Bastian, and A. M. Okamura. A proposed method for correcting coordination deficits: Models and simulation. In *IEEE International Conference on Robotics and Automation, Workshop on Softening Rehabilitation Robotics*, Anchorage, Alaska, May 2010.
- T. Hulin, C. Preusche, and G. Hirzinger. Stability boundary for haptic rendering: Influence of human operator. In *IEEE International Conference on Intelligent Robots and Systems*, pages 3483–3488, 2008.

- H. K. Khalil. *Nonlinear Systems*. Prentice Hall, 3 edition, 2001.
- Oussama Khatib. Inertial properties in robotic manipulation: An object-level framework. *International Journal of Robotics Research*, 14:19–36, 1995.
- 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.
- M. Mahvash and A. M. Okamura. Enhancing transparency of a position-exchange teleoperator. In *World Haptics*, pages 470–475, 2007.
- J. E. Slotine and W. Li. *Applied Nonlinear Control*. Pearson, Upper Saddle River, NJ, 1991.
- G. Strang. *Linear Algebra and Its Applications*. Brooks Cole, 1988.
- B. Willaert, M. Franken, H. Van Brussel, and E.B. Vander Poorten. On the use of shunt impedances versus bounded environment passivity for teleoperation systems. In *IEEE International Conference on Robotics and Automation*, pages 2111–2117, May 2011.
- K. Zhou, J. C. Doyle, and K. Glover. *Robust and Optimal Control*. Prentice-Hall, Upper Saddle River, NJ, 1996.