The 2010 IEEE/RSJ International Conference on
Intelligent Robots and Systems
October 18-22, 2010, Taipei, Taiwan

Multi-DOF Equalization of Haptic Devices for Accurate Rendering at High Frequencies

Robert P. Wilson, Sonny Chan, J. Kenneth Salisbury, and Günter Niemeyer

Abstract— **Previous work has shown that high frequency content is important for realistic haptic feedback, while stability considerations limit the ability of closed-loop control to effectively generate high frequencies. Open-loop playback of high frequencies offers a promising way to generate rich contact transients and textures, but complex high frequency dynamics cause distortion. This paper explores the equalization and dynamic decoupling of multi-DOF haptic devices for accurate open-loop playback. Toward this end, a user study is performed to determine the frequency limit of human force direction sensitivity at 35Hz. This information together with experimental system identification techniques is used to develop a strategy for equalization in different frequency bands. Finally, MIMO equalization is accomplished through online simulation of the system model under the control of an LQR tracking controller.**

I. INTRODUCTION

During interaction with their environment, humans make significant use of high frequency haptic feedback to gauge material and surface properties. Investigation of these high frequencies for realistic haptic simulation has been performed, in particular, for conveying rigid contact [1], [2] and textures [3]. In both cases, rendering high frequency forces in a closed-loop architecture is shown to be limited by the nonidealities of system digitization and, especially for textures, excitation of a device's nonrigid body dynamics.

Several efforts propose avoiding these difficulties through open-loop playback of a desired signal. By rendering these high frequency signals open-loop, stability is no longer a concern and rich contact transients or textures can be superimposed over closed-loop forces to substantially improve realism. In [4], desired impact transients are generated using a model-based approach while [5] proposes a sample-based approach that plays back recorded contact accelerations. Both approaches supply command torques to a haptic device's motors to render the desired acceleration at the end effector. Here, using the motors preserves the ability to render this content with directionality, however the device dynamics must be considered. This makes accurate reproduction of desired endpoint accelerations a challenge and, due to the existence of high frequency dynamic coupling, multiple degree of freedom(DOF) signals can become particularly distorted. As an alternative approach, [6] uses a voice coil attached to the stylus to play high frequency waveforms, avoiding device dynamics. This is limited to one-DOF signals, however. The extent to which directionality is perceptually important at higher frequencies remains largely unexplored.

In this paper we look at how to equalize and dynamically decouple a haptic device for accurate playback of arbitrary high frequency signals in multiple degrees of freedom. We pay particular attention to two aspects of haptic simulation. First, over what range of frequencies can humans discriminate direction? And second, what are the capabilities of the haptic device with respect to replicating multi-DOF signals versus frequency? It is discovered that a frequency band exists where humans remain sensitive to direction, but where non-rigid body dynamic coupling distorts the output accelerations. In order to flatten and dynamically decouple the frequency responses of the device in this band, we implement a multiple input multiple output (MIMO) plant deconvolution based on a MIMO system model obtained through a black box system identification(sysID).

In order to arrive at a well informed MIMO equalizer, we first review existing approaches for dynamic decoupling and plant inversion used in haptics. That happens in Section II. In Section III, we investigate directionally of haptic forces as a function of frequency in terms of both human perception and device capability. Namely, we identify the frequency threshold of human force discrimination via a user study and perform a MIMO sysID of an omega.3 3-DOF haptic device from Force Dimension. Finally, we discuss equalizing the device in different frequency bands and demonstrate MIMO equalization implemented on the omega in Section IV.

II. BACKGROUND

In haptics, a few options for dynamic decoupling and equalization of devices currently exist. Here, the inversion of rigid body dynamic models and a technique called acceleration matching are briefly reviewed.

A. Rigid Body Dynamic Inversion

Very commonly used throughout robotics, analytic models of rigid body dynamics can be used to compensate for the actual dynamics of a device. While not covered in detail here, the operational space framework [7] achieves dynamic decoupling through inversion of the operational space dynamics. The success of this approach relies on the accuracy of the model and will break down at frequencies beyond the onset of nonrigid dynamic effects, making it suitable for low frequency decoupling.

R.P. Wilson, S. Chan, and J.K. Salisbury are with the Department of Mechanical Engineering, Stanford University, Stanford, CA 94035, USA `rpwilson@stanford.edu, sonny@cs.stanford.edu, jks@robotics.Stanford.edu`

G. Niemeyer is a senior research scientist at Willow Garage and with the Faculty of Mechanical Engineering, Stanford University, Stanford, CA 94035, USA `gunter.niemeyer@stanford.edu`

978-1-4244-6676-4/10/$25.00 ©2010 IEEE

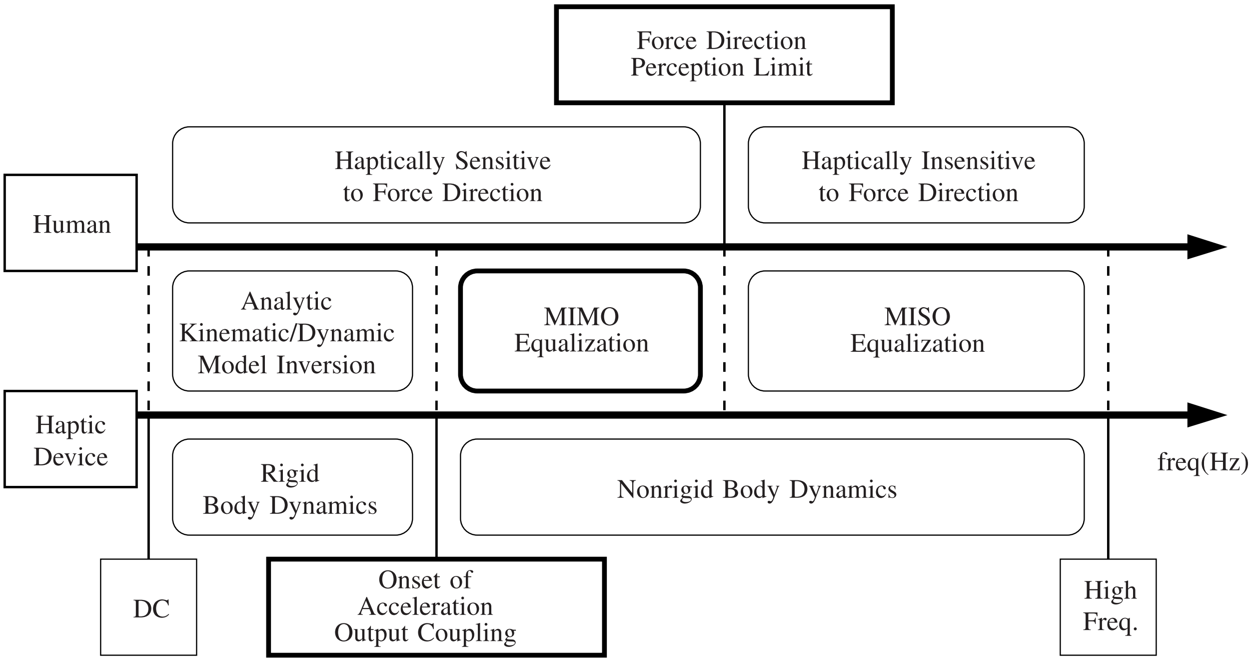

Fig. 1. MIMO EQUALIZATION SHOULD BE USED ON THE FREQUENCY BAND BETWEEN RIGID BODY MODEL BREAKDOWN AND THE LIMIT OF DIRECTION PERCEPTION

B. Acceleration Matching

A method for high frequency equalization called acceleration matching is described in [5] to improve the fidelity of open-loop sample playback. Here, standard SISO system identification techniques are used to generate an experimental transfer function estimate (ETFE) mapping input force to end effector acceleration for a one-DOF haptic simulation. A transfer function polynomial is then fit to the spectral ETFE data over much of the device's bandwidth. Transforming the desired acceleration trajectory with this transfer function's inverse yields force inputs that, after passing through the dynamics of the device, will faithfully reproduce the desired transient. In effect, the transfer function inverse deconvolves the haptic device's dynamics, flattening its frequency response. Since this is a one-DOF approach, dynamic coupling is not an issue.

III. HIGH FREQUENCY MULTI-DOF HAPTICS

In this paper we consider extending the acceleration matching technique described above to multi-DOF haptic simulations. Now, high frequency multi-DOF desired acceleration trajectories are to be played open-loop through a MIMO haptic device. In addition to flattening the system's frequency response, dynamic coupling between the separate degrees of freedom must now be considered if acceleration directions are to be accurately conveyed. To this end, the SISO models used above are insufficient, and MIMO state space models capable of capturing coupling effects must be used for deconvolving the plant instead. Unfortunately, these models quickly grow to prohibitively high order when trying to fit large frequency bands and highly complex dynamics. Therefore, aside from actually fitting a model and implementing the equalization, much of the problem lies in finding the minimum frequency range over which a MIMO approach must be implemented.

Figure 1 outlines several frequency bands of interest from both the human and device perspective and illustrates the region where we expect MIMO equalization will be used to the best effect. Specifically, we expect that MIMO equalization will be required above the frequency where dynamic coupling starts to become an issue, and that it will not be required above the user's frequency limit of direction sensitivity. Of course, a MIMO equalizer could be used outside this region at a greater computational cost, but we feel this frequency band best leverages a knowledge of the user's perceptual needs and the device's performance to minimize such costs. Indeed, if the perceptual limit were to occur before the onset of coupling, MIMO equalization may not be needed at all, allowing a SISO approach such as acceleration matching or the one-DOF, stylus-mounted, vibration source approach of [6] to be used instead.

With this road map in place, the actual perception limit and coupling onset frequencies are determined. For the limit of direction discrimination a user study is performed. The onset of cross coupling is determined via an experimental sysID.

A. User Study: Force Direction Sensitivity vs Frequency

In order to select the frequency range over which directional sample playback should be preserved, it is necessary to have an estimate of the upper frequency bound on human force direction sensitivity. In [8], it is noted that despite extensive research on discrimination thresholds for force magnitude, displacement, and stiffness, force direction discrimination has largely not been explored. The ability of humans to discriminate between static forces is then investigated. Discrimination as a function of frequency has not yet been investigated.

To get a conservative estimate of the upper bound we ran a user study involving a worst case force discrimination task carried out at five different frequencies (f=5,15,25,35,45Hz). The study was carried out on the omega.3, a high performance, three-DOF, parallel-linkage haptic device. Specifically, we were interested in identifying the frequency on the omega, in particular, above which direction was no longer detectable, given a maximum allowable amplitude.

The procedure utilized a forced-choice, one-up three-down adaptive method [9]. In each trial the user was randomly presented with one of two orthogonal sinusoidal forces and asked to identify it by selecting 1 or 2. At the start of each new trial, the force amplitude was faded in to prevent startup effects from providing discrimination cues. Subjects were allowed to evaluate each trial as long as they wanted before providing an answer. An image was supplied illustrating the two potential force directions with corresponding oriented line segments and their respective response (1 or 2). Orthogonal directions were selected as they represent a worst case direction disparity appropriate for obtaining a conservative discrimination estimate. Additionally, these orthogonal directions were selected so that they lied across a line of symmetry in the omega's kinematics in order to minimize discrimination cues that might be caused by unbalanced kinematics. For each frequency the commanded force amplitude was initialized to 1N. The amplitude of the sinusoidal forces was then increased after each incorrect response and decreased after three consecutive correct responses. The max amplitude was saturated at 3N, because in pre-study trials, many users complained of discomfort at higher amplitudes, especially at higher frequencies. On the first three reversals (correct followed by incorrect and vise versa), the amplitude was modified by 4dB and by 1dB for the remaining ten reversals. The detection threshold was estimated by the mean of the last ten peaks and valleys. After thirteen reversals the test was repeated for the next frequency. The order that frequencies were presented was randomized.

Ten subjects, eight male and two female, participated in the study. All were graduate students and right handed. Before beginning the study, each subject performed a training run at 5Hz to learn the task and find a comfortable grip. Subjects were instructed to try to maintain a constant grip throughout the study. Figure 2 shows the individual subject force thresholds for each frequency, as well as the group means and one-way ANOVA 95% confidence intervals for multiple comparisons. In order to make sure the results weren't colored by dramatic variations in cross coupling from one frequency to the next, normalized acceleration ellipsoids were computed along one of the directions used in the user study. These ellipsoids were constructed by commanding a constant amplitude sinusoid along the test direction at several frequencies and looking at the mean of the squared measured accelerations rotated from the accelerometer frame to the test direction at each frequency. Figure 3 shows ellipses corresponding to each ellipsoid's two largest axes, providing a graphical comparison of relative coupling at different frequencies.

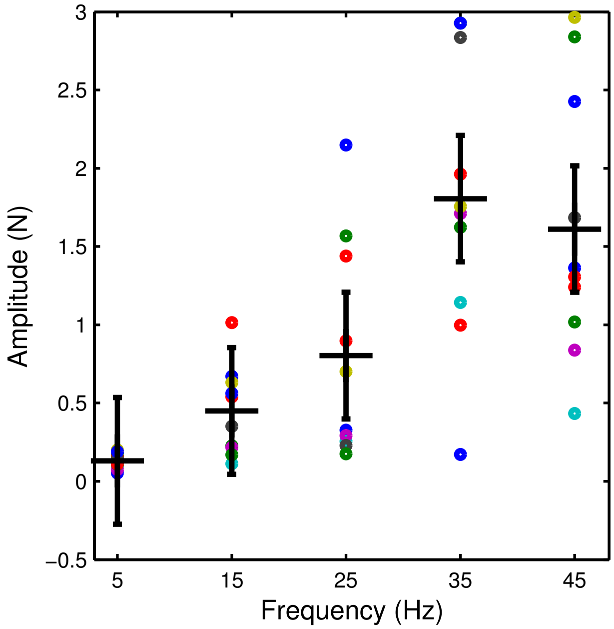

Fig. 2. FORCE THRESHOLDS FOR DIRECTION DETECTION AT SEVERAL FREQUENCIES. ERROR BARS INDICATE THE GROUP MEAN AND ONE-WAY ANOVA 95% CONFIDENCE INTERVALS FOR MULTIPLE COMPARISONS.

Although more subjects are needed to establish statistical significance, the results suggest a drop in performance between 5Hz and 15Hz, as well as a significant increase standard deviation. Many subjects noted after performing the study that they thought all of the frequencies except for 5Hz 'felt like vibration'. The increased spread in performance combined with this anecdotal evidence may suggests that there is a roll-off in kinesthetic detection between 5Hz and 15Hz, and a transition to less consistent tactile detection. Coupling is nearly identical for 5, 10, and 15Hz and is likely not a factor. Further apparent performance loss from 15Hz to 25Hz may be caused by the roll-off of directional skin stretch tactile sensors [10] or an increase in cross coupling.

At 35Hz there is a statistically significant ($p < 0.05$) increase in the group mean and a very large spread in perfomance emerges. While some subjects remain able to converge, others are unable to detect direction and diverge to the amplitude saturation at 3N. Coupling at 35Hz is actually less than at 25Hz. One possibility is a roll-off in compressive stress tactile sensors around 30Hz [10]. At 45Hz the results are fairly similar. Informally, tests performed up to 200Hz revealed some subjects' ability to detect direction to a similar threshold level. However, the ability of some to discriminate at high frequency may not be a strict detection of a direction, but a combination of strategic gripping by these users and vibrational differences born from imperfect

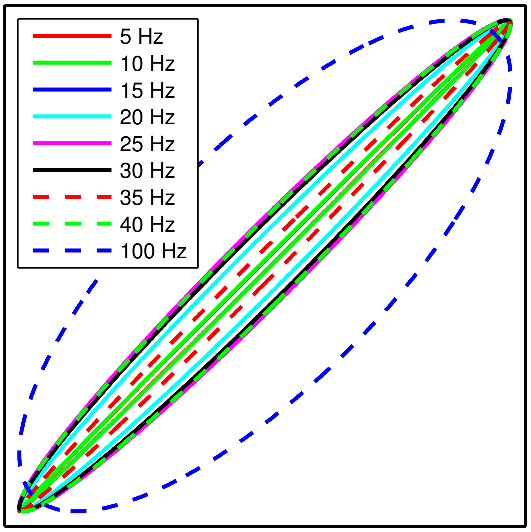

Fig. 3. VISUALIZATION OF DIRECTIONAL COUPLING ON THE OMEGA AT SEVERAL FREQUENCIES.

testing hardware. Anecdotally, subjects who were successful at high frequency all employed the same cognitive strategy: maintain a very light grip and search for vibrational differences across different fingers. When asked if they thought they would notice a particular direction if not asked to try and distinguish between two signals, even successful users unanimously agreed that 35Hz and 45Hz feel like undirected vibration.

This study represents a pragmatic first pass at identifying force direction sensitivity vs frequency. It is potentially specific to the omega.3 device, colored by its coupling effects and mismatch between the output mirrored kinematics used for the orthogonal directions. From the perspective of selecting a frequency band for MIMO equalization on an omega.3, however, we feel that 35Hz is a conservative direction discrimination frequency threshold.

B. *System Frequency Response and Coupling*

To characterize the location of the major resonance modes and the onset of cross coupling, we did a system identification of the omega.3 from Cartesian force commands to Cartesian end effector accelerations. An ADXL335 3-axis accelerometer was used with a range of $\pm 3g$ and bandwidth of 500Hz. This sensor was mounted on the omega.3's handle. The system ID was performed from 1Hz to 500Hz using a separate thirty second long swept sinusoid (repeated to eliminate transients) for each input channel.

Instead of identifying the system under a human grip, which cannot be easily held constant over the duration of a single trial let alone over multiple different trials, we digitally implemented a weak spherical stiffness about the handle at the center of the workspace. The natural frequency of the resulting mass-spring was kept low at about 3.5Hz. This is very repeatable and prevents the device from moving around its workspace or striking the edge of the workspace during sweeps.

After completing the swept sinusoids, the ETFE from each input to each output was computed and can be seen in Fig. 4. The results indicate that the system experiences large resonances and cross coupling starting around 80Hz, and that there is a lower frequency resonance near 11Hz that introduces noticeable coupling. The slope below 10Hz shows that the system is dominated by frictional and damping terms. Apparent coupling at very low frequency is likely inaccurate, as there insufficient energy below 3.5Hz for reliable identification. From these results we identify the frequency of coupling onset at around 10Hz.

IV. EQUALIZED HAPTIC RENDERING STRATEGY

Having identified the frequency ranges shown in Figure 1 for the omega.3, we propose leveraging different equalization approaches at different frequencies.

At low frequencies the conventional closed-loop rendering approach remains the best choice. Since the user actively interacts with the environment at these speeds, closed-loop feedback is important. Also, at such low frequencies, digital non-idealities are not problematic and dynamic effects are minimal. A simple kinematic model and perhaps a rigid body dynamic model will produce accurate quasi-static forces.

At high frequencies the user is not sensitive to force direction and a one-DOF approach similar to acceleration matching may be used. In fact, a separate SISO equalizer may be implemented for each DOF. Since the DOF(s) used to generate desired frequency content largely doesn't matter, the DOF(s) responsible for playback may be shifted away from motors close to saturation to under-utilized motors. Alternatively, using a voice coil or vibrating motor mounted directly at the stylus makes a lot of sense in this region.

At middle frequencies where replication of multi-DOF accelerations is desired, the MIMO equalization described below is employed. Since open-loop playback will occur through both this equalizer and the SISO high frequency equalizer, the desired acceleration trajectories must be split for each band. The specifics of how this is done will depend on the method used to generate trajectories. As an example, sample-based approaches might bandpass filter recorded multi-DOF accelerations and then perform an appropriate reduction to one-DOF for the upper band.

A. *MIMO Equalization*

As mentioned in Section III, we implement MIMO equalization by fitting a state space model to MIMO input-output data and then use this model to deconvolve the plant. For the omega.3, the input data is the commanded Cartesian forces supplied to the kinematic model, and the output data is Cartesian accelerations measured at the handle. Here the kinematic model is included in the sysID. Based on the frequency bounds determined earlier, we fit the model over the slightly conservative range from 5Hz to 35Hz, using swept sinusoids to generate input-output data on this range. The parameterized state space model fit is performed

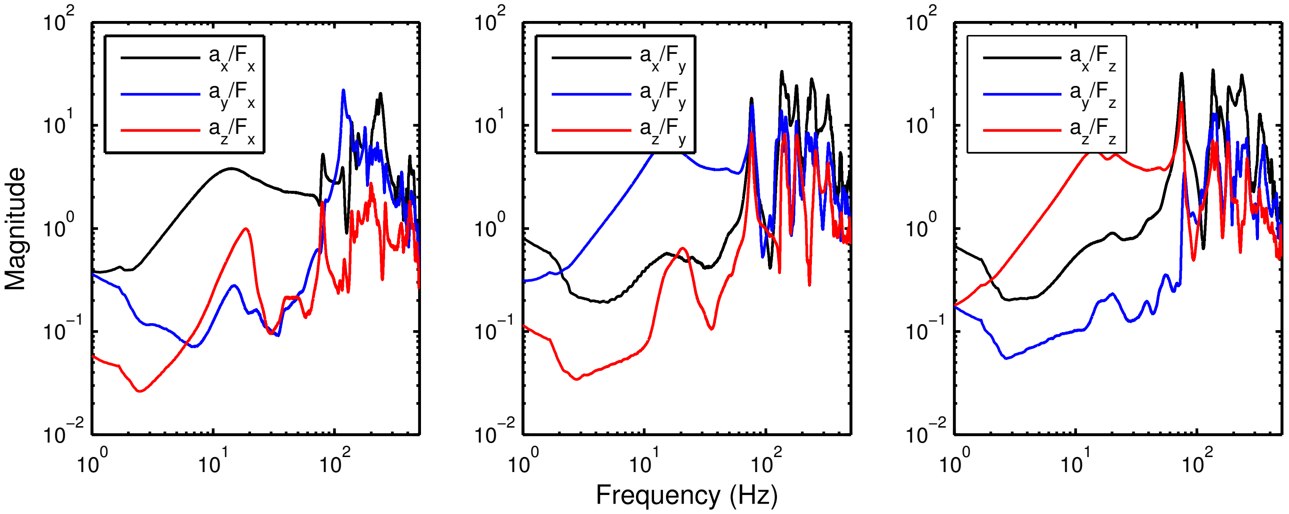

Fig. 4. EXPERIMENTAL TRANSFER FUNCTION ESTIMATES FOR ALL INPUT/OUTPUT CHANNELS (1 HZ TO 500 HZ)

using the n4sid algorithm implemented in Matlab's System Identification Toolbox. A 26 order state space model is fit, trading off quality of the model and order.

Given this model, we ideally want to find the set of feedforward force commands that will ultimately produce the desired output. Assuming that the desired output acceleration trajectory is known and that there is no acceleration feedback, this time series of feedforward forces may be found via the least squares solution to the block Toeplitz system of equations relating the input and output time series'. This approach is computationally expensive, however, and not well suited for real time application.

Equalization is instead implemented by dynamically simulating the experimental model under the control of a standard LQR tracking controller. Using both the estimated state space model and the known output trajectory, the time varying LQR controller gains are computed by the Riccati recursion. In the tracking LQR formulation, the resulting controller contains both feedforward and feedback terms. Setting the feedback terms to zero, the experimental state space model and feedforward LQR control equations are simulated forward in time to generate the feedforward control inputs that should produce the desired output in the absence of disturbances, deconvolving the system. This LQR approach yields the same result as least squares, but can be computed much faster and is more suitable for real time usage.

The need to know the desired output acceleration time series is satisfied by many existing approaches for generating high frequency haptic content, such as the playback of recorded acceleration samples or model-based signal synthesis. Additionally, the desired trajectory may transformed at any time step during playback, provided the Riccati recursion is recomputed from the final time back to the current time. This is particularly useful in the case of recorded sample playback, where a directional desired acceleration is recorded in only one configuration. For example, if a real world impact acceleration transient against a common surface is recorded only in one direction, for an impact against an arbitrarily oriented haptic simulation of this surface the entire recorded time series is first rotated to align with the contact normal. The LQR controller is then solved online by the Riccati recursion once at impact using the rotated output vector. We have implemented this approach in realtime at 1000Hz.

Figure 5 displays the ETFEs of the un-equalized system as originally identified on the first row. The second row shows ETFEs of the equalized system. The resonance that exists near 11Hz is essentially removed and cross coupling peaks are attenuated by up to an order of magnitude. The third row shows ETFEs of the equalized system near the edge of the workspace instead of at the center. Since the accuracy of the model inevitably decreases away from the sysID configuration, the degradation in performance is expected, however the response remains fairly flat and coupling still shows minor improvement. For better performance throughout the workspace, interpolation between models at several points could be considered.

V. CONCLUSION

Toward equalized open-loop playback on multi-DOF haptic devices, we have proposed a framework that handles feedback differently in different frequency bands. At low frequency, closed-loop control allows for active exploration at human interaction speeds. Between 5Hz and 35Hz we found that directional forces are detectable on the omega.3 haptic device and device equalization should be performed. A MIMO equalizer based on simulated LQR control of a state space model runs at haptic rates and provides a substantially flatter and more decoupled frequency response than without equalization. At frequencies above 35Hz, we suggest that, at least on the omega.3, there is little perceptible direction content, and MISO equalization largely identical to that in [5] should be used.

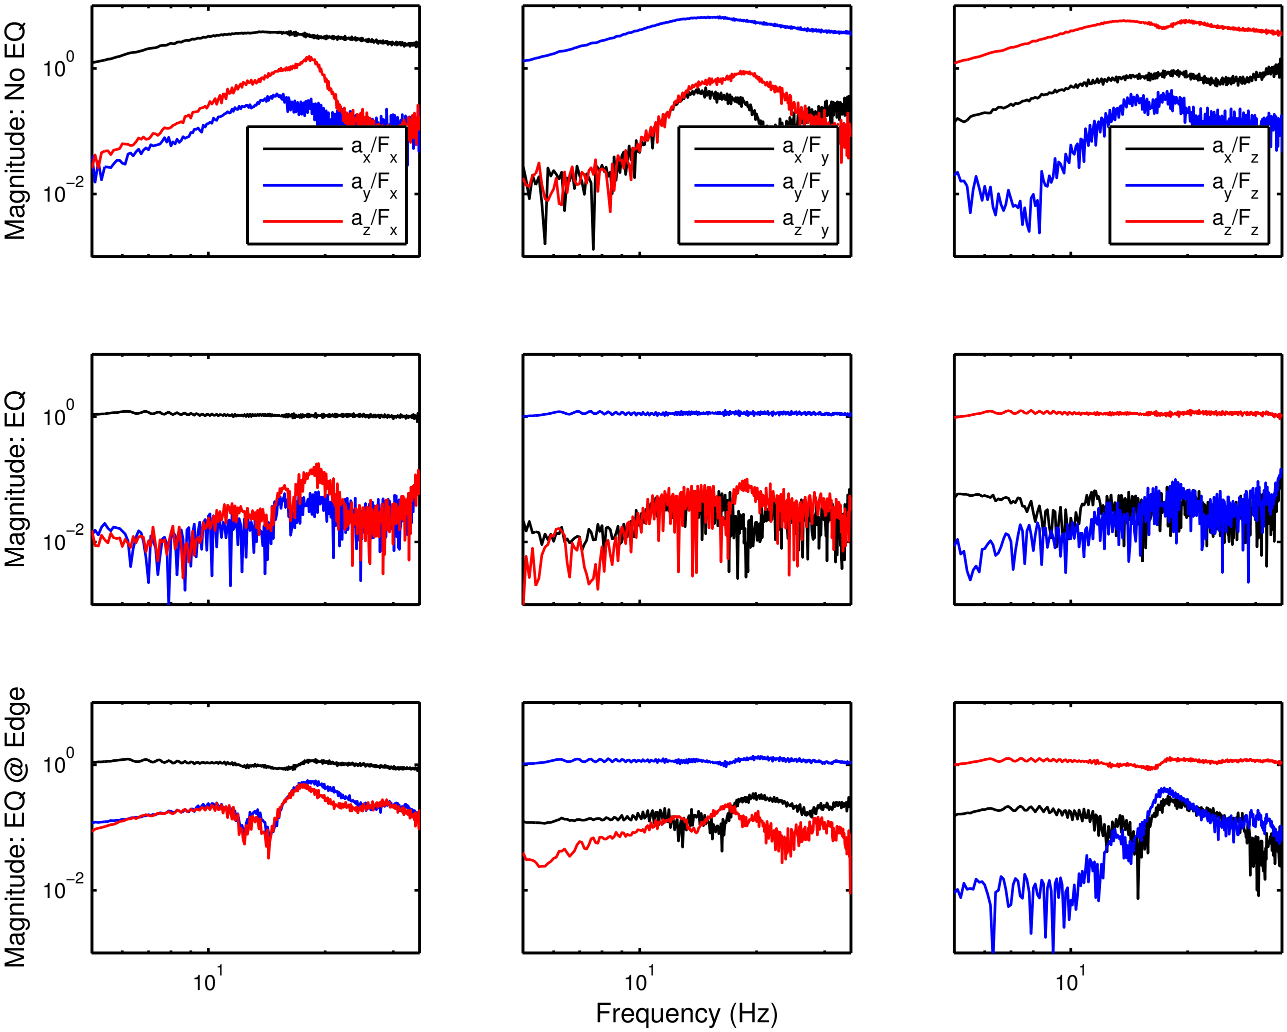

Fig. 5. EXPERIMENTAL TRANSFER FUNCTION ESTIMATES FOR ALL INPUT/OUTPUT CHANNELS WITH AND WITHOUT MIMO EQUALIZATION (5 HZ TO 35 HZ). EQUALIZATION PROVIDES A FLAT FREQUENCY RESPONSE AS WELL AS REDUCED COUPLING

Future work includes a more psychoacoustically-motivated investigation of human force direction sensitivity vs frequency. In particular, we plan to use a physically reconfigurable single-voice coil based device to present different direction forces while holding all other perceptible variables constant. Also, we plan to investigate equalization under a time varying human grip.

VI. ACKNOWLEDGMENTS

We would like to thank Francois Conti for providing access to an omega.3 haptic device.

REFERENCES

[1] D.A. Lawrence, L.Y. Pao, A.M. Dougherty, M.A. Salada, and Y. Pavlou. Rate-hardness: a new performance metric for haptic interfaces. *IEEE Transactions on Robotics and Automation*, 16(4):357–371, Aug 2000.

[2] J. Edward Colgate and J. Michael Brown. Factors affecting the z-width of a haptic display. In *IEEE Conference on Robotics and Automation*, pages 3205–3210, 1994.

[3] Gianni Campion and Vincent Hayward. Fundamental limits in the rendering of virtual haptic textures. *World Haptics Conference*, 0:263–270, 2005.

[4] Allison M. Okamura, Associate Member, Mark R. Cutkosky, and Jack Tigh Dennerlein. Reality-based models for vibration feedback in virtual environments. In *Proceedings of the ASME Dynamic Systems and Control Division*, pages 1117–1124, 2001.

[5] Katherine J. Kuchenbecker, Jonathan Fiene, and Gunter Niemeyer. Improving contact realism through event-based haptic feedback. *IEEE Transactions on Visualization and Computer Graphics*, 12:219–230, 2006.

[6] P. Wellman and R.D. Howe. Toward realistic vibrotactile display in virtual environments. In *Proc. ASME Dynamic Systems and Control Division*, volume 57, pages 713–718, 1995.

[7] O. Khatib. The operational space framework. *JSME International Journal. Ser. C: Dynamics, Control, Robotics, Design and Manufacturing*, 36(3):277–287, 1993.

[8] Federico Barbagli, Ken Salisbury, Cristy Ho, Charles Spence, and Hong Z. Tan. Haptic discrimination of force direction and the influence of visual information. *ACM Trans. Appl. Percept.*, 3(2):125–135, 1 06.

[9] H Levitt. Transformed up-down methods in psychoacoustics. *J. Acoustical Soc. Am.*, 49:467–477.

[10] D.A. Kontarinis and R.D. Howe. Display of high-frequency tactile information to teleoperators. In *Proceedings of SPIE*, volume 2057, page 40, 1993.