5-12-2009

The AirWand: Design and Characterization of a Large-Workspace Haptic Device

Joseph M. Romano

University of Pennsylvania, jrom@seas.upenn.edu

Katherine J. Kuchenbecker

University of Pennsylvania, kuchenbe@seas.upenn.edu

Follow this and additional works at: https://repository.upenn.edu/meam_papers

Part of the [Mechanical Engineering Commons](#)

Recommended Citation

Romano, Joseph M. and Kuchenbecker, Katherine J., "The AirWand: Design and Characterization of a Large-Workspace Haptic Device" (2009). *Departmental Papers (MEAM)*. 223.
https://repository.upenn.edu/meam_papers/223

Suggested Citation:

Romano, Joseph M. and Katherine J. Kuchenbecker. (2009). *The AirWand: Design and Characterization of a Large-Workspace Haptic Device*. 2009 IEEE International Conference on Robotics and Automation. Kobe, Japan. May 12-17, 2009.

©2009 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/meam_papers/223
For more information, please contact repository@pobox.upenn.edu.

The AirWand: Design and Characterization of a Large-Workspace Haptic Device

Abstract

Almost all commercially available haptic interfaces share a common pitfall, a small shoebox-sized workspace; these devices typically rely on rigid-link manipulator design concepts. In this paper we outline our design for a new kinesthetic haptic system that drastically increases the usable haptic workspace. We present a proof-of-concept prototype, along with our analysis of its capabilities. Our design uses optical tracking to sense the position of the device, and air jet actuation to generate forces. By combining these two technologies, we are able to detach our device from the ground, thus sidestepping many problems that have plagued traditional haptic devices including workspace size, friction, and inertia. We show that optical tracking and air jet actuation successfully enable kinesthetic haptic interaction with virtual environments. Given an appropriately large volume high-pressure air source, and a reasonably high speed tracking system, this design paradigm has many desirable qualities when compared to traditional haptic design schemes.

Disciplines

Engineering | Mechanical Engineering

Comments

Suggested Citation:

Romano, Joseph M. and Katherine J. Kuchenbecker. (2009). *The AirWand: Design and Characterization of a Large-Workspace Haptic Device*. 2009 IEEE International Conference on Robotics and Automation. Kobe, Japan. May 12-17, 2009.

©2009 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

The AirWand: Design and Characterization of a Large-Workspace Haptic Device

Joseph M. Romano and Katherine J. Kuchenbecker

Haptics Group, GRASP Laboratory

University of Pennsylvania, USA

{jrom, kuchenbe}@seas.upenn.edu

Abstract—Almost all commercially available haptic interfaces share a common pitfall, a small shoebox-sized workspace; these devices typically rely on rigid-link manipulator design concepts. In this paper we outline our design for a new kinesthetic haptic system that drastically increases the usable haptic workspace. We present a proof-of-concept prototype, along with our analysis of its capabilities. Our design uses optical tracking to sense the position of the device, and air jet actuation to generate forces. By combining these two technologies, we are able to detach our device from the ground, thus sidestepping many problems that have plagued traditional haptic devices including workspace size, friction, and inertia.

We show that optical tracking and air jet actuation successfully enable kinesthetic haptic interaction with virtual environments. Given an appropriately large volume high-pressure air source, and a reasonably high speed tracking system, this design paradigm has many desirable qualities when compared to traditional haptic design schemes.

I. INTRODUCTION

The AirWand (Fig. 1) is a new portable haptic device that strives to vastly increase the size of haptic interaction workspaces. The user of a common commercial haptic system is relegated to sitting at a desk, holding a handle, and probing within a small volume of about 0.006 m^3 . This workspace limitation stems from a common design concept: current kinesthetic haptic devices are almost all some type of grounded rigid-link robot. While its dominance among

haptic device models is a testament to its cost effectiveness and usefulness, the realm to which we want to take haptic interactions has outgrown the workspace limitations of rigid-link robots.

Haptic technology has vast potential to improve a variety of human-computer interaction applications. A short list of these applications includes patient rehabilitation, task training and evaluation for skills such as machine maintenance or surgery, interactive gaming, and remote teleoperation. However, most of these tasks require a workspace larger than the shoebox-sized volume provided by most commercial devices. In fact, a majority of our daily motions take place outside this restrictive space. In order to realize the full potential of haptic technology, it is important to increase the usable device workspace such that it matches the workspace of the real task we wish to perform. While software fixes such as “clutching” of small-workspace haptic devices exist to allow us to explore larger environments [9], they do not provide any additional freedom of the user’s motion, and thus do not adequately recreate large arm motions.

A. Past Research

Researchers have developed a wide variety of solutions in attempting to increase the workspace size of haptic devices. One possible solution to this problem is to take the current small-workspace designs that are successful, and scale up the mechanism size. Zinn et al. [13] show in their DM^2 design how impedance principles, which are highly successful in small desktop haptic devices, can be modified to work in larger rigid-link robots by distributing the device weight and force load among several different motors. Van der Linde et al. [11] are able to use an admittance configuration to precisely control the motion of a large rigid-link robot by placing a force sensor on the end of their *HapticMaster* design. These designs must contend with problems caused by high inertia and friction common to large rigid-link robot arms. By mounting small-workspace haptic devices to mobile platforms, several different groups have managed to increase the workspace size of existing commercial devices, [7], [1]. However, this method also creates several challenging problems by coupling the haptic system’s dynamic response to the response capabilities of the mobile platform.

Buoguil et al. [2] designed a cable-based system for large workspace haptics, *SPIDAR*, where cables from the

Fig. 1. The one-degree-of-freedom AirWand handheld tool. Notice the infrared markers placed on the front face of the tool for optical tracking, and air exit ports aligned along the tool axis used to generate forces. Air inlet hoses have been removed for image clarity.

corners of a room are attached to a central object that the user is holding. By adjusting the tension in each cable the system displays various forces and torques to the user. The authors had a good deal of success with this device, but had difficulty accurately resolving the orientation of the user-held object. Various cables strung about the workspace also act to limit the mobility of the user in the space.

More recently, several groups have explored the use of air jets as a means of conveying kinesthetic information. Air jets are capable of generating ungrounded forces, which makes them a useful tool when trying to create large-workspace haptics. Xu et al. developed a system where a single air jet was mounted on the wrist of a user, and applied short force perturbations to the user's arm by quickly opening and closing the air jet valve in order to study human arm properties [12]. Suzuki and Kobayashi developed a force feedback interface where air nozzles are mounted inside of a table, and the user holds an "air receiver", a cup-shaped object. By optically tracking the object, the system controls the air output from the table such that different jets create various forces as the ejected air interacts with the cup [10].

The research most related to the work presented in this paper is the *AirGlove* system developed by Gurocak et al. [5]. The *AirGlove* consists of six jets attached about the wrist, oriented along three orthogonal axes. By controlling the air flow through each valve, Gurocak et al. command three degrees of force output at the user's wrist. A magnetic tracking device resolves the orientation of the user's wrist in space. Gurocak et al. performed several psychophysical experiments on the human perception of weight using this system. The air actuator design and fluid flow analysis discussed in the following sections closely aligns with that done in [5], but we have taken a significantly different approach in the modeling and analysis of our system.

II. AIRWAND DESIGN

Our goal for the AirWand project was to develop a proof-of-concept prototype to analyze the feasibility of optical tracking and air jets for use in large-workspace haptics. Our current design is a one-degree-of-freedom force feedback handheld tool that can be tracked in three Cartesian dimensions (Fig. 1). An optical tracking system triangulates the device location based on infrared markers placed on the tool, while one-dimensional force output is accomplished by exhausting pressurized air jets through two opposing nozzles along the tool's axis. This sensing and actuation scheme is expressed in Fig. 2.

A. Optical Tracking

We use an Optotrak 3020 optical tracking system to measure the handheld tool position [8]. Three infrared markers are equally spaced on the tool face (Fig. 1), and are used to track the centerpoint along the tool axis where the front air jet exits the device. The centerpoint is found by calculating the mean of the three position vectors obtained from the infrared markers.

Fig. 2. AirWand sensing and actuation overview. The optical tracking system images the handheld tool's infrared markers, sending their position data back to the host PC. Software running on the PC then decides on the appropriate command signal to send to the air actuation system, which is connected to the handheld tool via flexible tubing. This command signal determines the amount of air flow out of the valve, and thus the force experienced by the user.

The Optotrak system has a usable pyramid-shaped workspace that begins at ~ 1.5 m from the unit and increases in cross-sectional area with distance until the marker signal strength falls off at ~ 6 m. The volume of this workspace can be conservatively estimated at 15 m³. The maximum RMS positional error per marker is 0.1 to 0.9 mm, depending on the location within the workspace.

The host PC (Fig. 2) retrieves the current position of the markers from the Optotrak system at a rate of 650 Hz using the vendor-supplied C++ programming interface. Our software uses this information to compute the centroid of the handheld device. Our design is capable of handling tool rotations of up to $\sim 170^\circ$ about the X and Y tool axes (Fig. 1), and can handle any rotation about the Z axis. The system is not currently robust in the event of marker occlusions.

B. Air Jet Actuation

We use high-pressure air jets to generate forces to display to the user. Fig. 1 shows our one-degree-of-freedom tool that has two air jets aligned along the longitudinal axis of the tool, indicated as air exits. These two jets are used to create forces along the longitudinal axis, in both the positive and negative Z direction.

Several components are necessary to create a system capable of generating and controlling air jet output. Below we discuss the important design selections we made for our system, which are connected as shown in Fig. 3.

1) *Fluid Analysis*: Because the AirWand is designed to be an impedance-type haptic device, it is necessary to develop equations that govern the force output based on the input control signal. We formulate these equations by using a control volume analysis (Fig. 4), similar to that defined in Gurocak et al. [5], although an important distinction is that our control volume does not cut through any physical wall structures as seen in [5]. Rather, our control volume boundaries are defined as the surfaces infinitely close to, but not intersecting, the physical system boundaries.

Conservation of momentum enables us to obtain an equation for the net thrust force the system produces, which acts opposing the reaction force experienced by the control volume :

$$\vec{F} = -\vec{R} = -(\dot{m}_1 \vec{v}_1 + \dot{m}_3 \vec{v}_3) \quad (1)$$

Where $\dot{m}_i$ is the mass flow rate at the i th point in the flow, and v_i is the velocity. By continuity we know that:

$$|\dot{m}_1| = |\dot{m}_3| = \rho_1 \vec{A}_1 \cdot \vec{v}_1 = \rho_3 \vec{A}_3 \cdot \vec{v}_3 = \dot{m} \quad (2)$$

Where ρ_i and A_i are the density and cross-sectional area, respectively. When $|\vec{A}_1| = |\vec{A}_3|$ and $\rho_1 \gg \rho_3$, then $|\vec{v}_1| \ll |\vec{v}_3|$ and the force equation (1) is reduced, as noted in Gurocak et al. [5], to:

$$\vec{F} \approx -\dot{m} \vec{v}_3 \quad (3)$$

Thus the force magnitude we can generate with an air jet is the product of the mass flow rate through the valve and the air speed. In order to derive an expression for the mass flow rate through the system, we use compressible air flow analysis as presented in [6], [3], and outlined below. Starting with the definition of mass flow rate, $\dot{m} = \rho A v$, and applying the ideal gas law and definition of Mach number M , we arrive at:

$$\dot{m} = AM \sqrt{\gamma RT} \frac{P}{RT} = A \sqrt{\frac{\gamma}{R}} M \frac{P}{\sqrt{T}} \quad (4)$$

Where γ is the specific heat ratio for air, R is the specific gas constant for air, and T and P are the respective absolute temperature and pressure of the air. Now applying isentropic flow relationships from [6] for both temperature and pressure, we can substitute and obtain:

$$\dot{m} = \frac{AP_1}{\sqrt{T_1}} \sqrt{\frac{\gamma}{R}} M \left(1 + \frac{1}{2} (\gamma - 1) M^2 \right)^{-\frac{\gamma+1}{2(\gamma-1)}} \quad (5)$$

We can obtain constraints on the velocity of air within our system using choked flow theory. Choked flow exists when a gas has reached its maximum possible velocity, the speed of sound v_{sonic} , and any increase in the upstream pressure can only lead to an increase in the gas's density ρ , not to its velocity v_2 . Air flow is choked under the following conditions:

$$\frac{P_1}{P_2} < \left(\frac{\gamma + 1}{2} \right)^{\frac{\gamma}{\gamma-1}} \quad (6)$$

Fig. 3. Air actuation system diagram.

Fig. 4. Control volume for the AirWand system. Air flows from the upstream supply at (1) toward the handheld tool. (2) is a point at the mouth of the valve exit. (3) is a point significantly downstream from the valve exit where forces from the volume expansion of the valve exit are negligible.

Using a value of $\gamma = 1.4$ we calculate that the flow is choked when:

$$P_1 * 0.528 > P_2 \quad (7)$$

We use a compressor that can supply a nominal upstream gauge pressure P_1 of 130 *psi* (896.3 *kPa*). We know that for an infinitely small cross sectional valve opening, the flow will be choked. Using empirical results from Gurocak et al. [5], who have used the same valve in their system, the flow is choked for a considerable range of valve operation. Knowing that the velocity through the valve will be sonic, we can substitute $M = 1$ in equation (5) and rearrange to yield:

$$\dot{m} = C_p A P_1 \sqrt{\frac{\gamma}{RT_1} \left(\frac{2}{\gamma + 1} \right)^{\frac{\gamma+1}{\gamma-1}}} \quad (8)$$

Note that we have introduced an additional variable at this stage, the discharge loss coefficient due to rapid expansion, C_p . C_p represents a value between 0 and 1 which quantifies the loss when ejecting the air to atmospheric conditions. C_p is commonly a function of density (ρ), velocity (v), and pressure (P), but since these variables are all coupled in our system we can reduce the definition of C_p as follows:

$$C_p = f(P, \rho, v) = f(P) \quad (9)$$

2) *Tubing*: We sought to build a system that could deliver an air jet force of up to 10 N. The amount of force a tube permits is solely dependent on its inner diameter. Frictional loss coefficients are negligibly small and are thus ignored. By combining equations (3) and (8), we can solve for the required value of A :

$$A = \frac{F}{C_p P_1 v_3} \left(\frac{\gamma}{RT_1} \left(\frac{2}{\gamma + 1} \right)^{\frac{\gamma+1}{\gamma-1}} \right)^{-\frac{1}{2}} \quad (10)$$

By substituting in values of $F = 10$ N, $v_3 = v_{sonic} = 343$ m/s, $\gamma = 1.4$, $T_1 = 298$ K, $R = 287 \frac{m^2}{s^2 K}$, $C_p = 1.0$ [3], and $P_1 = 896.3$ *kPa*, we can calculate A to be 0.0000139 m^2 , or an inner diameter of 4.2 *mm*. We selected a tube with an inner diameter of 6 *mm* in order to overshoot this value. Additional tubing selection criteria are flexibility, durability, and burst pressure. We chose a polyethylene tubing from MSC Industrial Co. with a burst pressure of 160 *psi* and a minimum bending radius of 20 *cm*.

3) *Air Source*: By rearranging equation (3) to solve for mass flow rate, $\dot{m} = \frac{F}{v_{sonic}}$, we calculate the mass flow rate to achieve 10 N to be $0.0292 \frac{kg}{s}$. Upon evaluating current compressor technology, we concluded that a compressor capable of $0.0292 \frac{kg}{s}$ was too physically large, too audibly dissatisfying, and too expensive for our feasibility study. Our system uses a Hitachi EC79 compressor as the air source. It is used to compress air, and also to maintain a reservoir of pressurized air so that a significant volume of air is ready at any one time. The EC79 is capable of continuously supplying 2.7 cubic feet per minute (cfm), or $0.00127 \frac{m^3}{s}$ of volumetric air flow Q , at a pressure of 90 *psi* (620.5 *kPa*). Mass flow rate can be calculated simply as a product of density and volumetric flow rate:

$$\dot{m} = \frac{P}{RT} Q \quad (11)$$

Substituting values of $P = 620.5 \text{ kPa}$, $R = 287 \frac{m^2}{s^2 K}$, $T = 298 \text{ K}$, and $Q = 0.00127 \frac{m^3}{s}$ yields a value of $\dot{m} = 0.0092 \text{ kg/s}$. Placing this value into equation (3) yields a maximum sustainable force of $F = 3.16 \text{ N}$. Larger instantaneous forces are achievable, since we are able to empty the compressor reservoir at a more rapid rate than it can refill, although doing this is undesirable since it affects the available force at later times. The EC79 is manufactured with an inline pressure regulating valve with an undesirable slow-regulating response. We removed this regulator in order to allow rapid access to the unregulated compressed air supply tank.

4) *Flow Valve and PC Communication*: Flow rate for our system is controlled by a Festo MPYE-5-1/4-010-B bidirectional proportional flow control valve. This valve was selected for three major reasons. First, the valve can handle a high volume of flow. The manufacturer quotes the volumetric flow rate of the valve at 1400 standard liters per minute (slpm), or $0.0233 \frac{m^3}{s}$. Second, was the fast response time of the valve. The critical frequency (the frequency at which the valves response drops below the $-3dB$ point) is quoted at 90 Hz. Third, the input voltage maps directly to the cross-sectional area of the valve opening: +5 V and +10 V indicate fully closed and fully open valve area respectively. As we will see in our system analysis below, this mapping is of critical importance when solving equation (3). Command signals are sent to the valve from the host PC using a Sensoray 626 analog I/O card in the form of an analog voltage level.

5) *Pressure Sensor*: As can be seen in equation (8) above, estimating the flow rate of the system is dependent on knowing the upstream pressure. Since we decided to use a small compressor, it is impossible to maintain a constant value of upstream pressure. Thus, we added a pressure sensor to our system to detect this value. The pressure sensor selected was a model SPT4V0500PG5W02 from Invensys. The sensor provides a voltage output signal that is proportional to static pressure, which is read into the host PC though an analog input channel on the Sensoray 626 card.

Fig. 5. The air jet characterization setup.

III. ACTUATOR CHARACTERIZATION

A test apparatus was created to study our air jet system performance and construct a mathematical model of its behavior. The test apparatus consists of the actuation system listed above, with the addition of two components. The first is a specially designed tube-holding device (Figure 5), which was created to orient the exiting air jet in a reliable manner. The tube-holding device sits atop our second additional component, a force sensor. The force sensor selected was an iLoad Analog sensor from Loadstar. This sensor outputs a voltage signal that is proportional to the load applied normal to its top face, and it was interfaced to the computer in the same manner as the pressure sensor described above. The force sensor is used to record the normal force the tube holder imparts when the valve is open. This is the force a user would feel when holding a haptic device powered by air jets.

Our goal in this characterization is to create an open-loop predictive control scheme to map our controlled voltage to the force output of the device. Our model does use sensory feedback in the form of upstream pressure information, but does not close the control loop using any sensed force information, due to the difficulty of obtaining accurate force output information during operator usage.

A. Parameter Fitting

It is necessary to obtain a model that expresses the relationship between valve command voltage and air jet force in order to accurately control the force output of the device. We note that if upstream temperature is constant, the entire square root term of equation (8) is also constant, and we can rewrite equation (3) as:

$$F = C_p A P_1 v_{sonic} \kappa \quad (12)$$

Where $\kappa = 0.0023 \text{ s/m}$, a result of replacing $R = 287 \frac{m^2}{s^2 K}$, $T = 298 \text{ K}$, and $\gamma = 1.4$ in the square root term of (8). The valve manufacturer states that the value of A seen in equation (8) is a linear function of voltage. We develop a mapping for A by observing the pressure drop in the system

Fig. 6. A square wave voltage command signal along with air-supply pressure, and model/actual force as a function of time. The air pressure (middle) can be seen to drop as a result of repeated opening and closing of the valve. The effect of this pressure drop is that the same command signal results in lower forces being generated as the pressure decreases. Our model has proven capable of tracking these pressure-dependent changes in force. The unfiltered force represents the model response without the effects of the low-pass filter and moving average filter discussed in our modeling section. The final force represents the model's behavior with all the effects discussed in the modeling section included.

as a function of the valve command signal. The value of A increases linearly with the command voltage to its max value of 50 mm^2 at a rate of $12.63 \frac{\text{mm}^2}{\text{V}}$, with an initial deadband of $\pm 1.02 \text{ V}$.

It is also necessary to determine the form of the equation that governs our loss coefficient, C_p . Fig. 6 shows the pressure and real force output of the system when given a square wave voltage input. This plot represents one of several data sets that were collected in our work at understanding the relationship of these variables. Other data sets collected include sine wave inputs and triangle wave inputs, all at a variety of frequencies and amplitudes. C_p was experimentally fit to our data sets. As noted earlier, C_p was known to be a function of upstream pressure. We found that a nonlinear mapping of $C_p = (8.8 \times 10^{-6})P_1^{0.83}$ worked best for our system.

While the gross shape of the force output was well matched, several discrepancies between our model and the recorded data were noticed: incorrect force peaks were being predicted at the start of each step function input, and a time lag existed between the real data and our predicted results, as seen in Fig. 6. A low-pass filter was placed on the calculated

Fig. 7. Block diagram representing the mathematics used in implementing equation (12). The valve area A is estimated using a deadband with a linear gain and a low-pass filter. C_p is calculated via a non-linear function of pressure P_1 . A Bartlett-Hanning window is used as a weighted moving average filter to add in a time-delay effect to the final force calculation.

Fig. 8. A sample desired force trajectory that closely simulates the kind of force response that a user would generate when using our system. The command force is the actual force trajectory that was sent to the valve by our software program. The predicted model force is the response we anticipated seeing based on the command force and system pressure. The real and predicted data closely coincide until forces above 6 N are attempted. At this point we leave the regime of choked flow by opening the valve area too much, and our model diverges from the recorded forces.

valve area value to correct the force peak prediction problem. This filter is representative of the fact that the valve has a finite 90 Hz bandwidth and cannot move instantaneously. This filter corrects for the model behavior during fast, high amplitude signals. A weighted moving average filter was placed on the predicted force to model the time delay effects seen in the real data. It is believed that this effect is a result of the delay involved with pressurization changes along the length of the tubing that leads to the exit nozzle. We believe that a variable amount of time, depending on the length of tubing between the air supply and force output device, is necessary in order for these pressure signals to propagate through the system and appear in the force output. The final force prediction seen in Fig. 6 shows that these steps generate a model that closely agrees with the physical data. Many additional data sets were tested in order to check this model's validity, including sine, triangle, and step wave functions of varying frequencies and amplitudes. Fig. 7 shows a block diagram of the various model steps discussed above.

B. Force Response Verification

Returning to our original motivation for valve characterization, we can rearrange equation (8) to solve for cross sectional area A (and thus voltage, since they have a piecewise linear relationship) as a function of force and pressure:

$$A = \frac{F_{desired}}{C_p P_1 v_{sonic} \kappa} \quad (13)$$

Therefore, when a force output is desired for haptic feedback, we can use the current pressure value to decide the appropriate voltage to command to the valve. Fig. 8 shows a predicted force signal, along with the actual output created using our inverted equation above (13). Several different trials with desired force amplitudes ranging from 0-10 N, and frequencies from 0.5-4 Hz were tested on our system. These tests were used to determine an RMS error of $\pm$

0.137 N when the force was constrained to the choked flow range where our model is accurate (~ 0 -6 N). The 0.137 N discrepancy is most likely due to various sources of error in our data collection process including zeroing of the load cell used in experiments, imperfect parameters for our value of C_p , and unmodeled nonlinear valve response behavior. The system was found to have an open-loop bandwidth of 8.7 Hz, which was determined by running high frequency sinusoidal voltage input commands through the model while assuming constant pressure conditions.

IV. INTEGRATED VIRTUAL ENVIRONMENT

A demonstration application was created using C++ and OpenGL in order to create a final example of the fully integrated system. The virtual environment, pictured in Fig. 9, consists of a ball that represents the tip of the haptic tool in the virtual environment, and surfaces with which the user can make contact. The surfaces are programmed to behave like linear springs, outputting air jet forces that are proportional to surface penetration.

V. CONCLUSION

Our final prototype system exhibited many traits that would be beneficial to large-workspace haptics, but it also highlighted many of the challenges involved with this design. The final device workspace is roughly 15 m^3 , limited only by the visibility of the OptoTrak. The maximum peak force we experienced during trials was $\sim 7.58\text{ N}$, and the maximum continuous force was calculated as $F = 3.16\text{ N}$. This is slightly underpowered for a large workspace device, but it was a known limitation of our feasibility study, due to the compressor's low mass-flow-rate capabilities. The AirWand's inertia is equivalent to the mass of the device, 70 g , which is less than the effective mass of a Phantom Premium 1.0. There is no friction in the device, since it makes no mechanical connection to ground, although there is a rotational stiffness due to the air hoses connected at the tool midpoint. In practice the device has proven itself safe. While it is startling to be exposed to a 130 psi air jet, it is not physically painful.

Fig. 9. A user interacting with the AirWand demo program. The haptic device is modeled as a 3-DOF point and can interact with virtual surfaces of varying stiffness.

Our final conclusion is that optical tracking and air jet actuation have strong potential to improve large-workspace haptic systems, but pose problems that differ from those encountered in traditional rigid-link robot design. Future efforts should focus on obtaining a larger air supply source to increase the maximum continuous force of the device, and investigating portable air supplies [4]. Since the force response of the AirWand is dominated by the amount of air flow through the system, inline flow sensors are an important component for consideration in future designs. A means of actively damping the noise should be investigated, as the air jet exhaust is quite loud, also noted by Gurocak et al. [5]. Additional degrees of force feedback will help to enhance the immersiveness of the virtual environments, but implementing these additional degrees is a challenging problem due to space constraints on the handheld tool. The AirWand technology presents an alternative means of displaying ungrounded haptic interaction, and hopefully will prove useful in a wide array of large workspace applications. We are excited by the possibility of this new haptic feedback strategy and look forward to working on these problems with future designs.

REFERENCES

- [1] F. Barbagli, A. Formaglio, M. Franzini, A. Giannitrapani, and D. Prattichizzo. An experimental study of the limitations of mobile haptic interfaces. In *Proceedings of International Symposium of Experimental Robotics (ISER2004)*, 2004.
- [2] L. Buoguilu, M. Ishii, and M. Sato. *Haptic Human-Computer Interaction*, chapter Scaleable SPIDAR: a haptic interface for Human-Scale Virtual Environments, pages 182–193. Lecture Notes in Computer Science. Springer Berlin / Heidelberg, 2001.
- [3] R. Fox, A. McDonald, and P. Pritchard. *Introduction to Fluid Mechanics*. John Wiley and Sons Inc., 2003.
- [4] M. Gogola, E. J. Barth, and M. Goldfarb. Monopropellant powered actuators for use in autonomous human-scaled robotics. In *Proceedings of the 2002 IEEE International Conference on Robotics & Automation (ICRA 2002)*, 2002.
- [5] H. Gurocak, S. Jayaram, B. Parrish, and U. Jayaram. Weight sensation in virtual environments using a haptic device with air jets. *ASME Journal of Computing and Information Science in Engineering*, 3:130–135, June 2003.
- [6] NASA Glenn Research Center. Mass flow choking: grc.nasa.gov/www/k-12/airplane/mfchck.html, 2008.
- [7] N. Nitzsche, U. Hanebeck, and G. Schmidt. Mobile haptic interaction with extended real or virtual environments. *Proceedings of the 10th IEEE International Workshop on Robot and Human Interactive Communication*, pages 313–318, 2001.
- [8] Northern Digital Inc. Optotrak technical specifications. <http://www.ndigital.com/optotrak-techspecs.php>.
- [9] S. S. Snibbe, K. E. MacLean, R. Shaw, J. Roderick, W. L. Verplank, and M. Scheeff. Haptic techniques for media control. In *UIST '01: Proceedings of the 14th annual ACM symposium on User interface software and technology*, pages 199–208. ACM, 2001.
- [10] Y. Suzuki and M. Kobayashi. Air jet driven force feedback in virtual reality. *IEEE Computer Graphics and Applications*, 25:44–47, 2005.
- [11] R. Q. Van der Linde, P. Lammertse, E. Frederiksen, and B. Ruiter. The HapticMaster: a new high-performance haptic interface. In *Proceedings of EuroHaptics*, pages 1–5, 2002.
- [12] Y. Xu, I. W. Hunter, J. M. Hollerbach, and D. J. Bennett. An airjet actuator system for identification of the human arm joint mechanical properties. *IEEE Transactions on Biomedical Engineering*, 38:1111–1122, 1991.
- [13] M. Zinn, O. Khatib, B. Roth, and J. K. Salisbury. Playing it safe: human-friendly robots. *IEEE Robotics and Automation Magazine*, 11(2):12–21, June 2004.