

Obstacle Avoidance Methods for a Passive Haptic Display

Davin K. Swanson[†] and Wayne J. Book[‡]

George W. Woodruff School of Mechanical Engineering
Georgia Institute of Technology
Atlanta, GA 30332-0405, USA
gt6032c@prism.gatech.edu[†], wayne.book@me.gatech.edu[‡]

Abstract—

An existing two degree-of-freedom passive haptic display is used to perform obstacle avoidance tasks. Two types of controllers are examined. One attempts to control the velocity direction of the display's handle in order to guide the user around obstacles. The other controller selectively kinematically constrains the device to a single degree of freedom. The inherent passivity constraint of the haptic display imposes performance limitations on the two controllers.

I. INTRODUCTION

A. Passive Haptic Displays

The term "haptic display" is a broad one describing a class of devices that are designed to interact with a human operator's sense of touch in a tactile or kinesthetic manner. They have a wide range of uses and forms, ranging from haptic gloves used to tactically visualize virtual objects to force-feedback joysticks used for remote teleoperation.

A passive haptic display is one which cannot increase the total energy of the man-machine system. However it may generally dissipate energy, transfer energy between system components, or store energy. Any addition of energy must come externally from the human operator or from the surrounding environment. The passivity of such a device carries inherent limitations and challenges in developing control techniques since the device cannot generate arbitrary control efforts. Available efforts must obey the passivity constraints of the system.

The main advantage of using a passive interface is safety. Passivity eliminates the possibility of damage or injury due to instability or actuator malfunction. This advantage is especially important for applications where high control forces are required or where the environment is fragile (e.g., robot-assisted surgery.)

B. PTER - A Passive Haptic Testbed

This work extends from previous research performed on a robot dubbed PTER (Passive Trajectory Enhancing Robot.) PTER was designed specifically as a passive haptic interface and previous work has focused on applying it to the task of trajectory following. [1] In this application, a controller attempts to constrain the motion of PTER's handle to a specified trajectory. Given the limit on control forces due to the passivity constraint, following an arbitrary trajectory can be difficult or impossible. Another

use which may better suit PTER's capabilities is obstacle avoidance. This application allows free motion within the workspace but prevents the tip from penetrating known obstacles.

Another class of passive haptic displays that has been studied comprise steerable degree of freedom devices. [2] [3] These devices have only one degree of freedom, but the orientation of that degree of freedom is steerable. PTER has several advantages over devices of this type. First of all, PTER is holonomic. If a given situation permits free motion of the tip, the user may guide it instantaneously in any arbitrary direction without active steering by the controller. In addition, PTER is mechanically simpler than steerable DOF devices and has lower residual friction.

II. CONSTRUCTION OF PTER

PTER is a five-bar parallel linkage situated in a horizontal plane (see Figure 1.) The operator interacts with the handle attached to the tip, through which he applies motive forces and may add energy to the system. Applied tip forces are measured by a commercial force sensor. The main links, A and B, are also acted on by four electromagnetic dry friction clutches which in turn apply force to the user through the handle. Having two degrees of freedom and four actuators, PTER is overactuated. To give the reader an idea of scale, link A measures approximately 0.7 meters.




Fig. 1. Passive Trajectory Enhancing Robot (PTER)




Fig. 2. Torque Sensing Clutch — Schematic

The clutches provide controlled frictional coupling of links *A* and *B*. Clutches 1 and 2 couple arms *A* and *B*, respectively, to ground. Clutches 3 and 4 couple the arms to each other, either directly or through gearing which inverts the relative axis velocity. Given the inherent restrictions on achievable control efforts and the manifest difficulty in precisely controlling the output torque of each clutch, this overactuated configuration was chosen with the intention of broadening the set of controls available to the controller for any given state of the device. A more complete description of PTER's original construction and operation can be found in [1].

Figure 2 is a diagram of one of PTER's clutches. The clutches are commercial dry friction clutches modified to incorporate an integral torque sensor. This modification was intended to allow feedback control of the torques generated by each clutch. The friction interface consists of a dry friction material mounted on an electromagnetic coil, and a steel armature plate. The armature plate is connected through aluminum spokes to a central hub, which is taper-locked to a driveshaft. Any torque applied to the armature plate will be transmitted to the driveshaft through the spokes, and vice versa. The spokes are instrumented with strain gauges to measure their deflection. With this information, transmitted torque can be computed.

Even though torque feedback was not used during the experiments outlined in this paper, the construction of the modified clutches impacts PTER's overall performance. The spokes were designed to be thin enough to allow satisfactory measurements from the strain gauges without requiring large amplification gains. This design choice has the unfortunate effect of introducing compliance between the driveshaft and the armature plate, causing oscillations after large or sudden changes in control forces are applied.

III. CONTROLLING PTER

A. Applying PTER to Obstacle Avoidance

Previous work on controlling PTER has concentrated on trajectory-following tasks. This involves constraining the tip of PTER to an arbitrary trajectory in opposition to operator-applied forces. Most of this work has involved using a modified impedance controller. Since PTER is passive, the set of achievable applied tip forces is limited and it is impossible to emulate arbitrary impedances in the workspace.

Obstacle avoidance is another application of PTER which had been previously considered, but not implemented. In an obstacle avoidance task, a payload or tool

would be attached to PTER's endpoint, which would then be moved or manipulated by the human operator. The controller would prevent the user from moving the tip into a restricted area of the workspace, and could also guide the user around the restricted area in the event that the user desires to traverse it. A typical obstacle avoidance task involves free movement with localized restrictions rather than precise control of the position of the endpoint. This type of task may be more easily performed by PTER given its inherent restrictions on generated forces.

Practical examples of obstacle avoidance tasks would be robotic-assisted surgery, where the controller keeps instruments away from organs and tissues that are not being operated on, or in a material handling task where an operator is manually maneuvering a heavy piece of equipment through a crowded workspace. Discussions with personnel from major airlines indicate obstacle avoidance would be very attractive in maintenance of expensive aircraft engines.

B. Requirements of Obstacle Avoidance Controllers

The purpose of an obstacle avoidance controller is twofold. Such a controller should prevent the user from entering an obstacle and should attempt to redirect the user around obstacles. Redirection is especially useful if the operator is not familiar with the configuration of the obstacles in the workspace, or if haptic feedback is the operator's only source of information about the workspace.

The controllers presented in the following section fulfill the above requirements using two different methodologies. For each controller, an "obstacle" does not necessarily represent an actual physical object, but an arbitrary region in the workspace. In a typical application, an obstacle would enclose a physical object. The controllers as presented require a priori knowledge of the location and configuration of obstacles in the workspace. A possible modification to these controllers would be a system that senses obstacles and actively updates the controller's workspace model. Such a system would be useful in dynamic environments.

IV. OBSTACLE AVOIDANCE CONTROL METHODS

A. Velocity Controller

The velocity controller attempts to redirect the direction of the tip velocity in order to prevent contact with defined obstacles in the workspace. See Figure 3. It steers the tip by attempting to drive the actual velocity of the tip v_a to a direction tangent to the obstacle surface, v_d . To do this, it commands a force F_d tangential to the current velocity direction. Applying such a force will change the direction of the velocity but not the magnitude. Since PTER is passive, it may or may not be able to apply force in the direction desired by the controller. If the appropriate force can not be applied with any combination of clutches, a force will still be applied in a direction as close to the desired force as possible.

In the event that the controller cannot prevent contact with the obstacle, it simultaneously commands large con-

trol efforts from two clutches in an attempt to totally immobilize the device. The controller releases the device only when the user applies a force in a direction that would push the tip back outside of the obstacle.

The controller only operates once the tip passes inside a boundary at a specified distance from the obstacle. If the tip is not within the boundaries of any obstacle in the workspace, the controller does nothing and the tip is free to be moved by the operator.

B. Single Degree of Freedom (SDOF) Controller

Whereas the velocity controller described above attempts to exert specific forces on the operator through PTER's tip, the single degree of freedom (SDOF) controller works by constraining the device to a single degree of freedom. When a single clutch is locked, PTER's motion is constrained to a single curve in the workspace.

Within this framework, obstacles in the workspace must be built out of SDOF curves representing the possible motion of the tip when a given clutch is locked at a given point in the workspace. For example, Figure 4 shows an obstacle made up of four SDOF lines—two representing clutch 3 being locked and the other two representing clutch 4. When the tip penetrates the obstacle through one of the clutch 3 lines, a high torque is commanded to clutch 3, locking it up, and the motion of PTER's tip is effectively restricted along that line. Likewise, if the tip penetrates through a clutch 4 line, clutch 4 is locked up, and the tip is restricted to move along that line. Figure 5 is a grid representing a set of SDOF lines for all four clutches in the workplace; this illustrates the shapes of obstacles that may be implemented with this controller. Once a clutch has been locked, the controller will free a clutch if it detects the operator applying a force to move the tip back outside of the obstacle. If the tip has not penetrated any of the ob-

stacles, the controller does nothing, allowing unrestricted motion of the tip.

C. Qualitative Comparison of Controllers

The velocity controller and the SDOF controller are each suited for both guiding the user around obstacles and preventing impingement into obstacles. Each has advantages and disadvantages.

The velocity controller has the advantage of being able to model obstacles of arbitrary shape while the SDOF controller requires that obstacles comprise SDOF lines in the workspace. An arbitrary obstacle shape must be surrounded by appropriate SDOF lines as shown in Figure 4 in order for the SDOF controller to avoid it. This wastes space by including unrestricted workspace within the constraint of the obstacle. Also the velocity controller starts working once the restricted region around an obstacle is entered, before the tip enters the obstacle. The SDOF controller does not operate until the tip is already inside the obstacle. Normally the penetration distance is small, but there are some cases where it can be quite large, as will be illustrated in Section V. This fact necessitates making the SDOF obstacle larger than the physical obstacle in the workspace in order to accommodate the possibility of significant penetration.

In its favor, the SDOF controller does not have the limitations on generated forces that affect the velocity controller. The SDOF algorithm is very simple compared to the velocity controller. The velocity controller in some cases computes ideal control forces that violate the passivity constraint, so the controller must compromise and settle for an achievable control that approximates the ideal control. This results in reduced efficacy of the velocity controller due directly to the device's passivity. The SDOF controller does not suffer from this shortcoming, as it relies solely on the kinematics of the device to create constraining forces. In the direction perpendicular to the constraining




Fig. 3. Velocity Controller - Schematic and Definitions




Fig. 4. An Obstacle Defined for use with the SDOF Controller




Fig. 5. Set of SDOF Lines

forces the friction forces can be kept very low.

Unless the user has encountered a concave corner of an SDOF obstacle where two SDOF lines meet, the SDOF controller will never immobilize the tip of the device. The velocity controller does immobilize the device if the endpoint enters the obstacle. The velocity controller also typically requires more rapid action that often completely separates the surfaces of the clutch, which is then reapplied a short time later with an undesirable impact. This is partly due to poor velocity estimates obtained from position encoders at low velocities. These effects give the SDOF controller a smoother overall feel than the velocity controller, which can feel jerky at times. This effect is illustrated with experimental data in the following section.

V. EXPERIMENTAL COMPARISON OF CONTROL METHODS

A. Experimental Description

In order to compare the performance of the velocity and SDOF controllers, a set of experiments were performed. A rig that implemented a constant-force torsional spring as the forcing source was used. This was done in order to obtain as close to a repeatable test as possible. One end of the spring was fixed at a single point in the workspace $(-0.40\text{m}, 0.80\text{m})$, and the other was attached to a plastic collar which fit over PTER's handle. The plastic collar was free to rotate about the aluminum handle; this interface has relatively low friction, minimizing any unwanted frictional forces caused by the rotation of the collar about the handle as PTER is moved through the workspace. Some tests were also performed with a human operator in order to roughly determine controller validity under real-world conditions.

For each test, a physical object of circular shape is assumed to be in the workspace. A circular obstacle was thusly implemented with the velocity controller. Since the SDOF controller cannot model an exact circle in the workspace, two obstacles were implemented with SDOF lines, both enclosing the circle. Many SDOF lines may be used to enclose the circle, but obstacles with 4 and 8 lines were used in the experiments. While a greater number of lines may more closely follow the shape of the actual object in the workspace and reduce the amount of wasted space around the object, the complexity of the controller is greatly increased if a large number of facets are used to define the obstacle.

B. Obstacle Penetration

A primary performance measure for the controllers discussed above would be obstacle penetration. The amount the tip penetrates a given obstacle depends on the velocity of the tip when it first crosses the obstacle boundary. When the tip crosses a boundary and a clutch or clutches are locked, there will be a small amount of slippage at the clutch interface before the particular degree of freedom is restricted. In general, higher clutch velocities mean higher energies to be absorbed by the locking clutch, resulting in higher amounts of slip. Figures 6 and 7 show performance

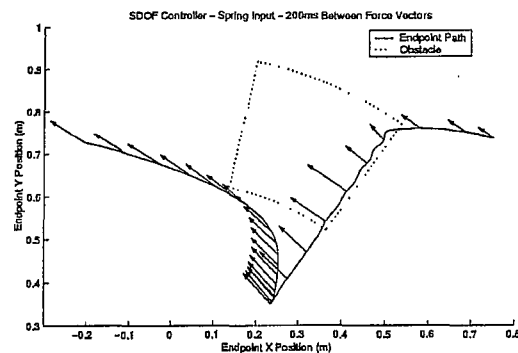


Fig. 6. SDOF Control; 4 SDOF Obstacle; High Penetration Speed

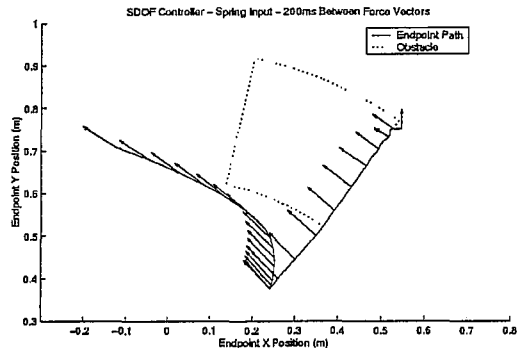


Fig. 7. SDOF Control; 4 SDOF Obstacle; Low Penetration Speed

of the SDOF controller implementing the 4-SDOF obstacle with high and low penetration velocities, respectively. The endpoint is initially held stationary at the right side of the plot. When the tip is released, the spring draws the endpoint towards the left. The arrows on the plots represent the measured endpoint forces applied by the spring and are spaced equidistant in time.

Note in the low-velocity case that the amount of penetration is small (2 mm) compared to the high velocity case (11 mm). This is because clutch four slips more in the high-velocity case before reaching a static state. Figure 8 shows results from the velocity controller. In this case, the controller is unable to redirect the tip around the obstacle and is forced to lock the tip once it enters the obstacle. This case can be compared with the high-velocity SDOF case, as in both cases the tip starts roughly the same distance from the obstacles. Maximum penetration for the velocity controller is 4 mm, smaller than the 11 mm measured for the similar SDOF test. The primary reason for this is that due to the buffer region defined around the obstacle by the velocity controller, control efforts slow the tip down before it encounters the obstacle.

Figures 9 and 10 show high entrance velocity performance of the SDOF controller when modeling an obstacle with 8 facets. These plots show the variability of obstacle penetration with the SDOF controller depending on the clutch being used and the location of the tip in the workspace. When navigating below the obstacle, the tip




Fig. 8. Velocity Controller; Tip Locking Condition

penetrates 8 mm. When navigating above it, the tip penetrates the obstacle 44 mm. The large disparity is due to the different levels of mechanical advantage of each clutch and on the location of the tip in the workspace. The line that is crossed in Figure 9 is a clutch 2 line. Clutch 2 acts on arm B, which is the shortest of PTER's arms, giving it the weakest mechanical advantage.

These tests show that both the velocity and SDOF controllers effectively minimize penetration of the obstacle, though to different degrees. The velocity controller in general works better than the SDOF controller in this respect due to the fact that it acts on the device before the tip actually contacts the obstacle.

C. Redirection of Tip

In addition to minimizing penetration of the obstacle as discussed above, the velocity controller and the SDOF controller both have the ability to redirect PTER's tip around an obstacle in the workspace. Looking at Figures 9, 10, 11, and 12, it is clear that both controllers can guide the tip around the obstacle in both directions. Looking just at the endpoint paths, it appears that the velocity controller does a better job at redirecting the tip; the path gradually changes direction and does not actually penetrate the obstacle, unlike in the SDOF cases. However, there is a factor other than endpoint path which affects controller performance. "Feel" or "smoothness" is a characteristic of any haptic display that is difficult to measure. Within the framework of obstacle avoidance, the primary goal is to keep the endpoint of PTER away from obstacles in the workspace, but a secondary goal is to be comfortable to the user. A smoother controller will likely result in less fatigue to the user and will be gentler for the equipment being positioned by the device. It is easy for a human operator to make judgements of smoothness, but a quantitative measurement is not evident. To that end, several human operators have professed that the SDOF controller feels smoother than the velocity controller.

To explore this, the actuator efforts for the experimental trials shown in Figures 9 and 11 were investigated. Figure 13 shows a time history of the voltage commands sent to the active clutches' power supplies by the controller for each experiment. The SDOF control efforts are shown on the left and the velocity control efforts on the right. It is clear that there is more on-off action with the velocity




Fig. 9. SDOF Control; 8 SDOF Obstacle; High Penetration Speed




Fig. 10. SDOF Control; 8 SDOF Obstacle; High Penetration Speed

controller due to the controller switching between active clutches. This especially happens at low velocities due to both the oscillatory motions caused by the clutch spoke compliance and numerical jitter in the clutch velocities. The numerical jitter is due to the fact that the velocity is estimated from digital encoder measurements of the arm positions.

The results in Figure 13 agree with the human operators' judgement that the SDOF controller feels smoother than the velocity controller. When the SDOF controller constrains the tip to a single degree of freedom, a hard constraint is imposed and the tip feels as though it is gliding along a smooth surface. The continued on-off action of the velocity controller causes a jerky motion of the tip, and elicits vibratory forces on the user's hand.

D. Human Input Tests

Experiments were performed with both controllers with a human user moving the tip of PTER in lieu of the spring setup. The operator attempted to move the tip through the area containing the obstacle. He knew the general location and configuration of the obstacle, but had no visual feedback informing him of the exact location and shape.

Figure 14 shows the response of the system using each of the controllers. Both controllers satisfactorily redirect the user's motion. The velocity controller guides the user around the obstacle more quickly, but the SDOF controller allows for in-contact exploration of the shape and location of the obstacle.

VI. CONCLUSIONS

The two controllers presented in this paper meet the basic objectives of obstacle avoidance and implement the original algorithms as envisioned. The controllers both limit

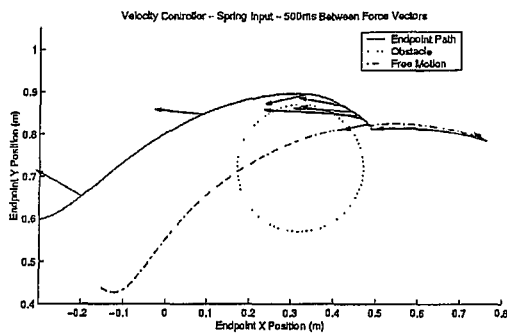


Fig. 11. Velocity Control; Path Redirection

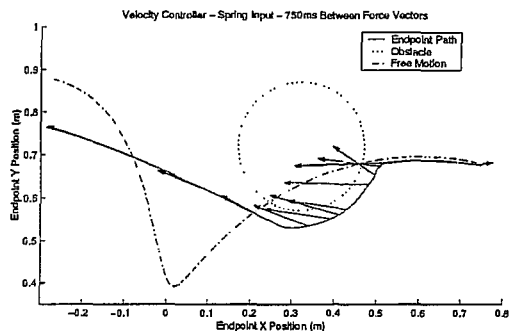


Fig. 12. Velocity Control; Path Redirection

the depth of intrusion into obstacles and attempt to redirect the user around them. Each controller does, however, have its limitations.

The SDOF controller is very smooth, simple, and allows a good level of active searching by the user. With the obstacle shapes investigated above, it will never totally immobilize the device. However, arbitrary shapes cannot be modeled. By using many SDOF lines close approximations of any shape can be realized, but using too many small segments may introduce the type of on-off actuator commands exhibited by the velocity controller. In addition, the SDOF controller at times will allow relatively large penetration of the obstacle due to high clutch velocities. Penetration can be reduced by using the controller to slow the component of tip speed normal to the obstacle so that the velocity on reaching the obstacle is small in the direction into the object. The effect would be a gradual application of one clutch instead of the sudden application used in this paper.




Fig. 13. Clutch Commands





Fig. 14. Performance of Controllers with Human Input

This is envisioned for future work.

The velocity controller is able to model any arbitrary obstacle shape, and does a good job of quickly redirecting the path around obstacles in the presence of moderate applied tip forces. If large forces are present or if the tip is in a portion of the workspace with little mechanical advantage, contact with the obstacle can occur. If this happens penetration depth is typically smaller than in a roughly equivalent SDOF task due to the fact that the controller has already slowed down the tip, but the controller also totally immobilizes the device.

These two controllers present a good example of the limitations of working with passive haptics. Each controller is limited in some way by the passivity of the device. It is difficult to make a judgement of which is more suited for obstacle avoidance. The selection of one of the two controllers would have to be made while considering the specific application at hand.

ACKNOWLEDGMENTS

This work was completed with partial support from the National Science Foundation, Grant IIS-9700528.

REFERENCES

- [1] Hurley T. Davis and Wayne J. Book. Passive torque control of a redundantly actuated manipulator. In *American Control Conference*, Albuquerque, NM, June 1997.
- [2] M. A. Peshkin, J. E. Colgate, and C. Moore. Passive robots and haptic displays based on nonholonomic elements. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 551-556, April 1996.
- [3] W. Wannasupphasit, P. Akella, M. Peshkin, and J. E. Colgate. Cobots: A novel material handling technology. In *ASME International Mechanical Engineering Congress and Exposition*, Anaheim, CA, November 1998.