Abhijeet Vijayakar
School of Computing
The University of Utah
50 S. Central Campus Drive,
Rm. 3190
Salt Lake City, UT 84112

John Hollerbach
jmh@cs.utah.edu
School of Computing
The University of Utah
50 S. Central Campus Drive,
Rm. 3190
Salt Lake City, UT 84112

Effect of Turning Strategy on Maneuvering Ability Using the Treadport Locomotion Interface

Abstract

Turning strategies on the Sarcos Treadport, a linear treadmill locomotion interface, are developed and compared in a tight maneuvering task. A rate control strategy employing sidestep is compared to a proportional control strategy employing head gaze and torso trigger. The maneuvering task involves walking down a narrow corridor and avoiding obstacles that are placed at different separations to change the task difficulty. The performance metric was the number of times a subject collided with the obstacles or the corridor's walls. Traversal time and traversal distance were also characterized. The proportional control strategy was clearly found to permit more-precise maneuvering.

1 Introduction

One of the most common platforms for locomotion interfaces is a linear treadmill, such as on the Sarcos Treadport. (See figure 1.) On the Treadport, the treadmill is augmented with a mechanical tether that acts as a goniometer to measure user position and orientation. These measurements are employed for active control of forward motion and of turning. The natural display of linear motion is a strength of linear treadmills, but the unidirectional belt motion requires a somewhat artificial turning method. This paper contrasts selected examples from rate control strategies and proportional control strategies from a standpoint of one's ability to maneuver in tight spaces.

Templeman, Denbrook, and Sibert (1999) divide any control technique for virtual locomotion into two parts: the control action made by the user and the controlled effect produced by the system. The most natural is when the control action and control effect are the same, such as turning while walking on the ground: we step diagonally and turn our bodies simultaneously. Proprioception, vision, and vestibular feedback all consistently signal that a turn has taken place. On the opposite extreme is using a joystick while seated: the joystick typically controls the rate of turning in the virtual world, but only the visual display indicates that a turn has taken place. (By *turning* we mean a reorientation of the walking direction.)

Between natural walking and joysticks are many different possible turning methods that are adapted according to limitations of different locomotion interfaces or to task requirements. One-dimensional treadmills are one of the most cost effective and useful locomotion interfaces, and, although turning cannot be completely natural, our goal with the Sarcos Treadport is to find the

Presence, *Vol. 11, No. 3, June 2002,* 247–258

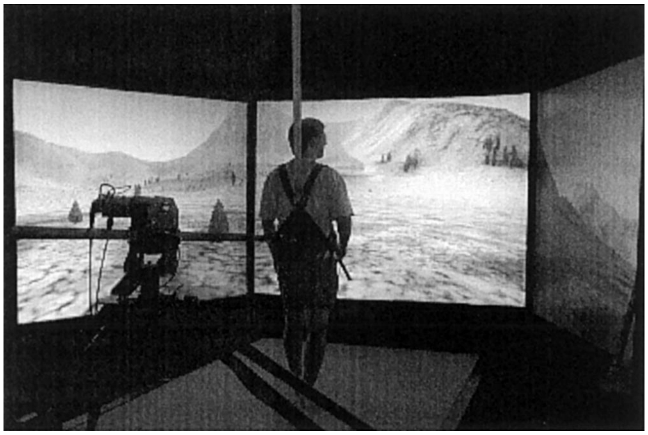

Figure 1. *The second-generation Treadport. The user walks on a 10 ft. × 6 ft. belt while viewing a three-wall, back-projected display in a flared arrangement. Depicted is a simulation of walking in Hogum Cirque, an area in the Wasatch Mountains. A mechanical tether attaches horizontally to the user's back via a worn harness. The tether's six joints measure user position and orientation, and the linear boom is motorized to exert push or pull forces. A ceiling strap acts as a safety restraint system.*

most-effective turning strategies for such devices. Other devices such as two-dimensional treadmill belts offer more-natural turning, but there can be other limitations such as cost, complexity, and motion speed, which are less favorable than for linear treadmills.

1.1 Turning on One-Dimensional Treadmills

One major strength of linear treadmills is the low cost because they are commodity items that are readily adapted into locomotion interfaces. Initial implementations employed passive treadmill belts (Brooks et al., 1992; Witmer & Kline, 1998) in which the belt moves only by virtue of users pushing with their feet. The user has to brace against bars to absorb the reaction forces; one approach (Brooks et al., 1992) is to use handlebars that can also control turning. The handlebars are typically used in proportional control mode; that is, the angle of the handlebars is the same as the angle of turning. This is in contrast to the rate control mode of most joysticks, wherein the angle of deflection of the joystick controls the angular velocity of turning.

In both cases, the control device has to be recentered, or *reindexed,* to stop turning. Reindexing is a fundamental aspect of turning on linear treadmills: any action that the user performs to effect turning must be compensated for with a similar reverse action before the user can turn in the opposite direction.

More-advanced linear treadmills are active: the belts are motorized, and the belt speed is based upon measurements of user position. For example, in the Treadport the amount by which the user's position is forward of center controls the velocity of the belt (Christensen, Hollerbach, Xu, & Meek, 2000). The resulting linear motion is more natural. The belt motors are easily sized to accomodate fast motion; on the Treadport, a user can accelerate at 1 g and reach speeds of 12 mph, which is about the best that average runners can do. A large belt size helps too, by providing safety margins on all sides, by permitting a variety of postures including crawling, and by facilitating turning through accomodation of some actual sideways motion. Two generations of Treadports have been built. In the first generation, the belt is 4 ft. × 8 ft. in width and length. In the second generation, the belt was made 6 ft. × 10 ft. specifically to address these points (Hollerbach, Xu, Christensen, & Jacobsen, 2000).

For the Treadport and in view of the high accelerations possible, forward motion is made even more natural by the application of artificial inertial forces, which are otherwise missing on locomotion interfaces because the user is stationary with respect to the ground (Christensen et al., 2000). These artificial inertial forces are applied by the active boom of the mechanical tether attachment to the user. The naturalness and high responsiveness to linear motion is another major strength of linear treadmills.

In the first-generation Treadport, a dual rate control strategy was employed (Christensen et al., 2000). At slow speeds, the amount of body twist controls the rate of turning. At fast speeds at which body twist is not as comfortable, the amount of sidestep from the treadmill center controls the rate of turning. (See figure 2). In between these speed extremes, the two strategies are

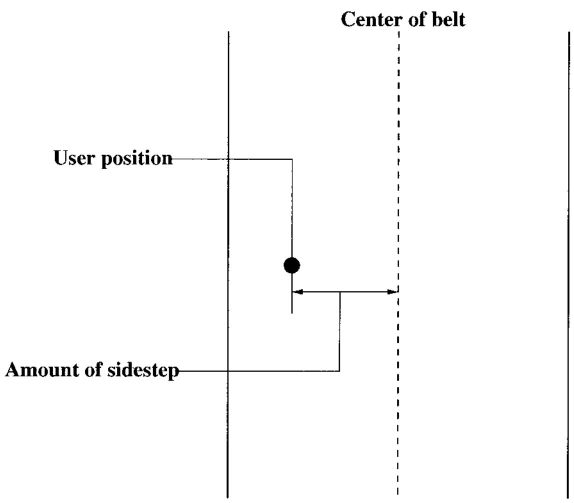

Figure 2. *Strategy 1: The amount of sidestep determines the rate of turning.*

blended. Both strategies require reindexing to stop turning: reorienting to a front-facing orientation for body twist or stepping back to the center for sidestep.

An interesting approach to make turning more natural on linear treadmills is to place the treadmill on a turntable, as was done for the ATLAS system (Noma, Sugihara, & Miyasato, 2000). Forward motion and the intent to turn are measured by camera tracking of markers on the feet. The amount of lateral motion in a step guides the swiveling of the treadmill by the turntable. By swiveling the treadmill in the direction of walking, the ATLAS system achieves a more natural gait pattern in turning. Mechanical limitations may limit the turning performance though. There will be lags in rotation due to inertia of the treadmill platform; this problem would be much worse if a large belt such as for the second-generation Treadport was used. Due to lags in rotation, it is possible that footfall occurs at a slant relative to the desired walking direction, and some correction by the user could be necessary on the next step. This problem may be compounded by the rather narrow belt of the current ATLAS, which could result in stepping off the side of the belt. The net result is that some care has to be exercised in the timing of foot placement when turning.

1.2 Turning on Two-Dimensional Treadmills

Two-dimensional treadmill belts allow for natural turning. The omni-directional treadmill (Darken, Cockayne, & Carmein, 1997) consists of a main belt made of rollers. These rollers are made to spin in an orthogonal direction to that of the main belt by another belt underneath. The Torus treadmill (Iwata & Yoshida, 1999) consists of a main belt fashioned from twelve small treadmills connected side by side. As the main belt moves, the small treadmill belts move orthogonally to create a two-dimensional motion.

In their present implementations, the naturalness of turning is somewhat compromised by limitations on forward motion. Walking on the rollers of the omni-directional treadmill is reportedly unsteady (Darken et al., 1997). Users are prone to lose balance if they turn while accelerating from rest, and especially if they decelerate to a stop and turn simultaneously, due to a misalignment between the direction of forward motion and the centering motion of the belt. The Torus treadmill suffers from a small walking area and underpowered belt motors that limit walking speeds to 1.2 m/s. This initial implementation has led to slow gait and short steps.

The drawbacks of these particular early 2-D treadmill designs could possibly be remedied by redesign. The field of locomotion interfaces is at an early stage, which can be characterized as an exploration of the design space of possible devices. The field has not yet matured to a point at which any particular design type can be said to have been optimized and its limitations fully understood. However, we can say that 2-D belt designs necessarily lead to more-complex and -expensive platforms, which may limit their proliferation.

1.3 Active Tether Design Implications

Although turning on linear treadmills detracts from the naturalness of locomotion, such treadmills do have offsetting advantages in cost and ease of forward motion. In addition, the active mechanical tether of the Treadport allows the display of wall constraints, inertial forces (Christensen et al., 2000), and even slopes

(Hollerbach et al., 2001). These tether effects support an argument that an active tether is a necessary part of any locomotion interface, but the presence of a tether has design implications.

The Treadport's tether is mounted on the platform and has a horizontal protruding boom. (See figure 1.) The boom does not interfere with the physical environment, such as the CAVE-like visual display in the front. If the treadmill were mounted on a turntable such as for the ATR ATLAS, such a boom sticking out the back would interfere with a stationary CAVE display. A head-mounted display (HMD) instead of a CAVE would be one solution. The advantages of an HMD include portability and the ability to project an object between a user's hand and eyes, but our strong preference is for the CAVE displays because of higher resolution, safety, convenience, and the ability to see one's body as part of the virtual world.

A recent development for the ATLAS system (Noma, personal communication, 1991) is to mount a front screen and projector directly to the platform, with which it rotates on a turntable. There are limitations on screen size, especially in view of the ATLAS's ability to tilt and roll the platform for slope presentation.

Two-dimensional treadmills would require two-dimensional tethers. If a CAVE is employed, a ceiling-mounted gantry design for the tether would be required to avoid interference with the screens. The omni-directional treadmill employs such a 2-D tether, capable of applying an 89 N force in a given direction. By comparison, the Treadport's tether can apply a 315 N force. The higher force levels are necessary to properly simulate slope and inertial forces. An active, two-axis, ceiling-mounted tether with sufficient force has definite attractions but presents more of a design challenge than does a single-axis tether.

1.4 New Turning Strategy for the Treadport

The dual rate control strategy originally employed for the Treadport (figure 2) has proven to be effective. From having hundreds of people—ranging from elementary school children to adults of all ages—try this system, we have found that users adapt quite readily to this turning control mode. Often, users discover how to control turning quite rapidly without being told how. Effectiveness is not the same as naturalness, though, and we wanted to see if a better turning strategy could be devised for linear treadmills. For example, a turning approach may have implications for wayfinding and for maneuverability.

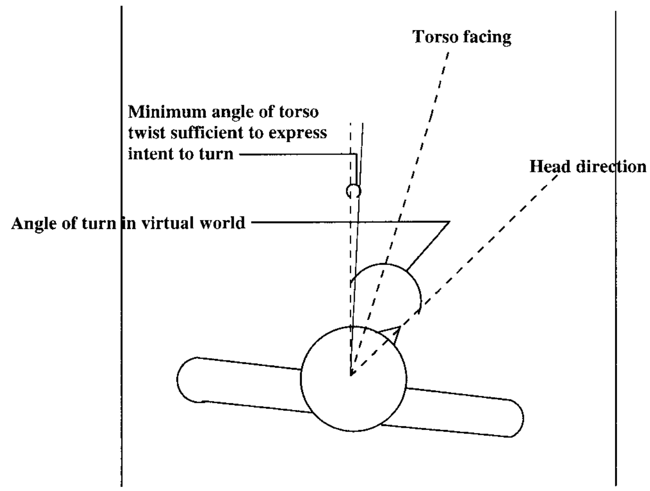

Figure 3. *Strategy 2: To initiate turning, the torso needs to be twisted enough to express intent to turn and the angle of head twist must be constant for a sufficient period of time.*

After informal experimentation, an alternative strategy using proportional control and head gaze direction with torso trigger was devised. (See figure 3.) The angle of head twist directly determines the desired angle of turn, but only after the torso has also been twisted in the same direction by a sufficient angle, ϵ_{twist}. This torso trigger is necessary to allow a user to look around freely without turning, and the threshold ϵ_{twist} prevents small body motions from inadvertently causing turning. The head gaze also has to be held constant for a small period of time.

The new strategy actually mimics real turning, because, as we step to turn, both the head and body begin facing in the new direction. This is possible on the Treadport because the belt is 6 ft. wide, which allows some amount of sideways excursion. If one imagines an extremely wide linear treadmill, then turning could be done completely naturally because the sideways excursion could be as large as necessary, but practicality will limit the width. The Treadport is much wider than other treadmills currently employed in locomotion interfaces (but not in some elite training centers), and a large width is important for this turning flexibility. One still has to reindex (which is the nonnatural part of this strategy), but probably this is the best that can be done on linear treadmills because reindexing will always be necessary.

When reindexing to the center, the visual display must be slewed in the opposite direction to the amount of turn. Rather than abruptly shifting the visual image, this slewing is done by an exponential time function. (See subsection 2.2.) The viewpoint moves quickly toward the center of the front screen at the start of the turn, but slows down near the end of the turn. The hope is that the moving display will draw the user to reindex back to center.

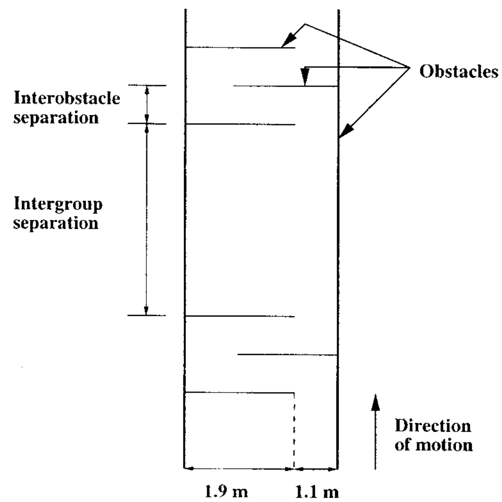

Figure 4. *Overhead view of a sample path for the experiment.*

1.5 Obstacle Avoidance Task

We have found that, if a task is sufficiently "easy," then the particular turning strategy doesn't make much difference in performance. Where we noted a difference was in tasks involving tight maneuvering, such as mazes. Based on experience in teleoperation, one would expect proportional control strategies to be better than rate control strategies for fine positioning (Sheridan, 1992). Rate control strategies are more appropriate for large excursion control. Hybrid strategies involving proportional control in central regions of control parameter ranges and rate control at the range extremes are also possible (Salcudean, Wong, & Hollis, 1995), but were not investigated here. The combinatorial possibilities for turning control strategies are large and impractical to test exhaustively. We are really looking for what after some reasonable effort seems to be the best turning strategy, rather than, say, to compare proportional control to rate control using only sidestep.

Consequently, in this paper, we seek to validate this new strategy by comparison to a previous Treadport strategy (rate control with sidestep), which is a restricted version of the full strategy of figure 2. The sidestep-based rate control strategy was chosen because it has been the predominant turning method on the Treadport for several years and has been found to be quite effective. We also conducted a number of informal studies with various rate-variable alternatives but could not come up with a better one. An alternative that was discarded, for example, was using body twist as the rate variable.

2 Methods

The virtual environment involves a narrow corridor with partitions jutting out to create a maze-like maneuvering task. The path schematic shown in figure 4 shows that obstacles (excluding the side walls) are grouped in threes, with a relatively large distance between groups. The difficulty level was then adjusted by varying the separation between the obstacles within a group. (See table 1.) Each path measured 60 m in length and 3 m in width. All obstacles (including the side walls) were 3 m high. The walls perpendicular to the user's direction of motion were 1.9 m wide; thus, there was a gap of approximately 1.1 m between such a wall and the side wall. As explained later, however, the user was represented in the program by a rectangular

Table 1. *Variation of Interobstacle and Intergroup Separation with Path Difficulty*

Difficulty level	Interobstacle separation	Intergroup separation
Low	4.5 m	6 m
Moderate	4 m	7 m
High	3.5 m	8 m

Figure 5. *User's perspective of a sample path for the experiment.*

parallelopiped 0.1 m in width; thus, the effective room for the user to pass through was 1 m.

The same path from the user's perspective is shown in figure 5. The side walls along each path were colored red. No texture was applied to the side walls, but a texture was applied to the walls perpendicular to the direction of user motion. Different textures were applied to walls flush with the left and the right edges of the path, as well as to walls occurring in paths of different difficulty. The ground resembled a typical road, with a double yellow line running along the center of the path. It is likely that the presence of a texture on the side walls would have provided the user with additional visual cues for turning. However, because these cues were absent for both strategies, it is not expected that adding such textures would have affected the observed result.

The task is not just a pure turning task because users may also shift sideways to help them get around obstacles. That is to say, stepping to the side by 1 m shifts the user's position in the virtual world by 1 m; this is true for both strategies being tested. This means that lateral shift does not by itself change the visual display: although for the sidestep-based rate control, the world would start rotating. The relatively wide belt width of the Treadport allows this flexibility. As explained later, users were required to keep a steady walking pace so that a pure lateral shift strategy could not be used.

For the sidestep-based rate control strategy, a user cannot actually step to the side without turning, whereas under the head-twist proportional control strategy a user could sidestep without turning. Whether that is an important factor is not known, but it is unlikely because of an informal study in which a torso-twist rate control strategy was employed instead. The torso-twist rate control strategy does not confound turning with stepping to the side, but results indicated that this strategy was not as effective as the sidestep-based rate control strategy.

We do not distinguish to what extent the avoidance is due to moving to the side versus heading in a new direction. An implication may be that the width of the belt is important in the performance of this turning method, but in the present studies we have not investigated that effect.

The graphics were rendered on an SGI Onyx2 running IRIX 6.5 with four R12000 CPUs, two InfiniteReality Engines and 512 MB of RAM, and displayed on three CAVE-style, back-projected screens with a total field of view of approximately 180 deg. Stereo projection was not used, and the subjects viewed the graphics display with both eyes. The controller for the Treadport ran on a Motorola 68000 processor board with 4 MB of RAM and using VxWorks v.5.2 as the operating system and ControlShell as the application program. The user's head position was tracked for Strategy II using the InterSense IS-600 head tracker. The user's torso position and orientation are measured by the six joints of the mechanical tether.

2.1 Experimental Procedure

Subjects were required to travel from one end of the path to the other with as few collisions as possible. No tether force feedback to simulate reaction forces was implemented; thus, subjects could "pass through" any of the obstacles. The primary performance measure was the number of obstacles that subjects collided with when traversing the virtual path for strategies 1 (figure 2) and 2 (figure 3). Of secondary importance were the time taken for the traversal of each path and the average distance traversed per path.

Twelve subjects (seven males and five females) were used to conduct the experiment. Two of the subjects were expert users; the others were first-time users (or relatively inexperienced users) of the Treadport. Subjects were first asked to read a written description of the experiment. They were then allowed to practice walking on the Treadport for approximately 5 min., then they were allowed to practice on a test virtual path to get used to the turning strategy for 2–3 min. Finally, they were given three practice runs on the experimental task.

During the actual experiment, three paths were presented in ascending order of difficulty, and subjects were allowed to walk along each path twice. The order of strategies was shuffled between subjects so that half the subjects were presented strategy 1 first and half were presented strategy 2 first. The average time for the entire experiment was approximately 40 min. The reason for always presenting paths in increasing order of difficulty was to allow subjects further practice. Although there may be some effect of not randomizing difficulty level, this effect would be expected to be the same between the two strategies.

Subjects were asked to maintain a comfortable walking pace without slowing down or stopping. During the experiment, the experimenter directed subjects to maintain their walking speed if necessary. This means that subjects could never just use a sideways shuffle to get around objects; turning to walk in a new direction was always being employed. At the same time, we did not restrict the proportion of sidestep versus turning to get around obstacles, and we did not quantify the relative contribution of each.

For detecting collisions with any wall, the user was represented as a rectangular parallelopiped 2 m high, 0.1 m wide, and 0.05 m thick. Subjects were given no visual clue about this representation. On the Treadport, it was observed that people experienced a subjective sense of collision with an object when their eyepoint passed through, or sufficiently close to, the object. Presumably the "looming effect" (the tendency of any object through which the viewpoint passes to visually enlarge until it fills the entire display area) triggered this subjective sensation of collision. For the software to record a collision precisely when the user also subjectively experienced a sense of collision, the width of the parallelopiped would have to be twice the "sufficiently close" distance (because the width of the parallelopiped was symmetric about the user's eyepoint). It was found through trial and error that 0.1 m was the largest width at which there was no noticeable disparity between when the user felt that he had collided with an object and when the software recorded a collision. The width of the parallelopiped was therefore taken to be 0.1 m in the experiment.

2.2 Preexperiment Tuning

Before the experiment was conducted the parameters for the control strategies were tuned. In strategy 1, the amount of sidestep, s, controls the rate of turning, R, as $R = K_{sidestep} \cdot s$, where $K_{sidestep}$ is a proportionality constant.

In strategy 2, a user who has indicated a turn of ϕ degrees will need to reindex to center. The visual display is slewed in the opposite direction by $-\phi$ degrees so that the new viewpoint is at the center. This slewing is done by an exponential time function, $\phi(t) = 1 - e^{-K_{twist}t}$, where K_{twist} is an exponential decay constant. The viewpoint will appear to move quickly toward the center of the front screen at the start of the turn, but will slow as the turn nears completion. Presumably, the user will be drawn to turn back to center as the object that he wants to turn to moves in that direction, thus effecting a (somewhat) natural reindexing. The total parameters to be tuned for strategy 2 are as follows.

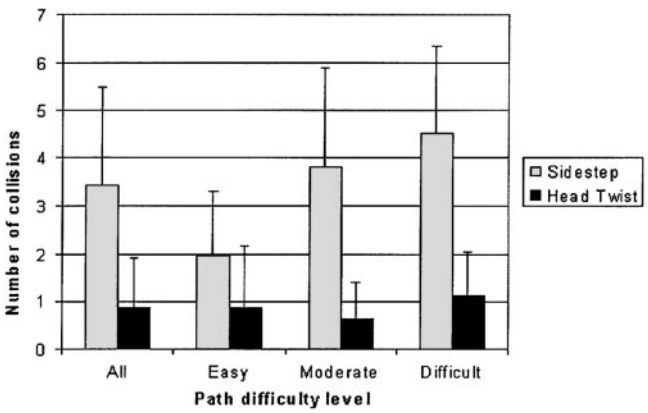

Figure 6. *Average number of obstacles that users collided with on paths of different difficulty levels. Error bars are shown.*

- The angle, ϵ_{twist}, beyond which a twist of the torso indicates intent to turn.
- The time, $t_{constant}$, for which the angle of head or torso twist needs to remain constant to indicate desired angle of turn.
- The exponential decay constant, K_{twist}, that determines the average rate of turning.

Two expert users were put on the Treadport, and the values of parameters were changed until a particular set of values felt "right" for both users. The parameter values obtained from this procedure were as follows.

- $K_{sidestep}$ = 1.5 rad/m second
- ϵ_{twist} = 5 deg.
- $t_{constant} \approx 1/15$ sec.
- K_{twist} = 1.5 /sec.

3 Results

Figure 6 shows the average number of obstacles that users collided with on paths of different difficulty levels considered separately, as well as for all paths combined. Because each subject was presented three different paths, each at one level of difficulty, and allowed to traverse each path twice, there was a total of six data points per subject. The graph for all paths combined is thus drawn from a total of 72 data points, and the graphs for paths of each difficulty level are drawn from 24 data points each.

The head-twist strategy performs much better than the sidestep strategy, both when all the data is taken together and when the data for each difficulty level is considered separately. As the difficulty level of the path increases, the performance of the sidestep strategy worsens appreciably; the head-twist strategy performs worst on the "difficult" path and (counterintuitively) better on the "moderate" path than the "easy" path, but the difference between its performance on the moderate path and the difficult path is not as pronounced as for the sidestep strategy.

A 3 (strategy) × 3 (difficulty) × 2 (order) mixed-design ANOVA was performed on the mean number of collisions with strategy and difficulty as within-subjects factors and order as a between-subject factor. The analysis revealed a significant effect of strategy, $F(1, 10) = 75.42$, $p < .001$, difficulty, $F(2, 20) = 12.83$, $p < .001$, and a strategy × difficulty interaction, $F(2, 20) = 10.04$, $p < .001$.

Overall, subjects collided with fewer obstacles using the head-twist strategy ($M = 0.88$) compared to the sidestep strategy ($M = 3.44$). There was no effect of order of strategies ($p = .28$) and no interactions with order.

3.1 Effect of Difficulty Level

Across both strategies, planned contrasts showed that performance was better on the low-difficulty path ($M = 1.42$) compared to the moderate path (($M = 2.23$), $F(1, 10) = 7.36$, and $p < .01$) and the difficult path (($M = 2.83$), $F(1, 10) = 27.39$, and $p < .001$). Separate one-way ANOVAs (three levels of difficulty) for each strategy condition examined the strategy × difficulty interaction. This analysis indicated that the difficulty effect was attributed to performance with the sidestep strategy, ($F(2, 22) = 17.59$, $p < .001$) but not the head-twist strategy ($p = .28$). Planned contrasts indicated that, for the sidestep strategy, performance on the easy path ($M = 1.96$) was better than on the difficult path (($M = 4.54$), $F(1, 11) = 41.13$, and $p <$

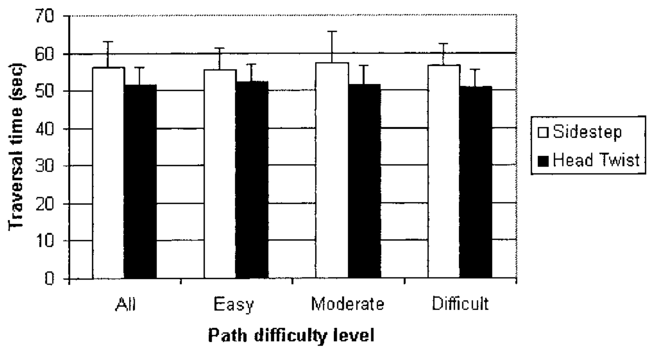

Figure 7. *Average path traversal times for all subjects at different difficulty levels for the sidestep versus head-twist strategies. Error bars are shown.*

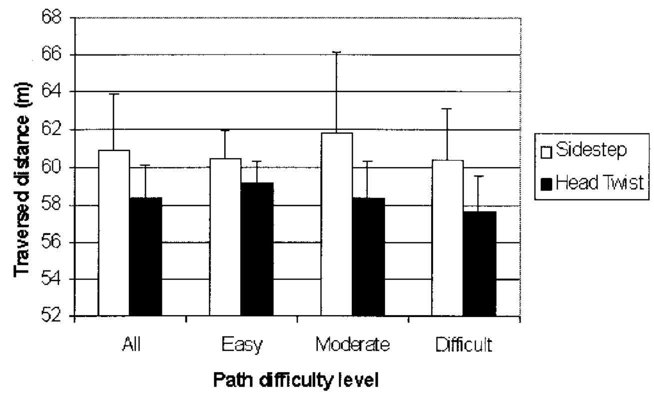

Figure 8. *Average path traversal distances for all subjects at different difficulty levels for the sidestep versus head-twist strategies. Error bars are shown.*

.001) and the moderate path (($M = 3.83$), $F(1, 11) = 17.81$, and $p < .001$).

It is concluded that the difficulty level of the paths did affect performance, but less so for the head-twist strategy than the sidestep strategy. When the paths were relatively easy, users seemed to be able to "get by" with even the sidestep strategy, but, as the difficulty level went up, their performance degraded rapidly. There was more of a "performance buffer" with the head-twist strategy: increasing level of path difficulty did not degrade the performance quite as much.

3.2 Traversal Time and Traversed Distance

Due to some missing data points, analysis of users' traversal time could be performed with only 64 (rather than 72) samples. (See figure 7.) The average time a user took to traverse a path using the sidestep strategy was 56.38 sec.; with the head-twist strategy, this was 51.54 sec. This difference proved to be highly significant.

A 3 (strategy) × 3 (difficulty) × 2 (order) mixed-design ANOVA was performed on the mean traversal time with strategy and difficulty as within-subjects factors and order as a between-subject factor. There was a significant effect of strategy ($F(1, 10) = 14.4$, $p < .01$). Users took less time to traverse a path using the head-twist strategy ($M = 51.32$ sec.) compared to the sidestep strategy ($M = 56.38$ sec.). There were no other effects or interactions.

Because users were asked to maintain a constant pace throughout the experiment (and this would likely be the same for both strategies), the significant difference observed could be because subjects in general took a longer route with the sidestep strategy than the head-twist strategy. If subjects were unable to control their motion precisely with the sidestep strategy (as indicated by the significant difference in the number of collisions), they would be more likely to veer off the optimum path through the virtual world, and this would increase their traversal time as well. Subjects could have walked straight through all the obstacles to minimize the traversal distance, but that was not observed because the task was one of obstacle avoidance.

This hypothesis was borne out by an analysis of the total distance traversed by subjects with the two strategies. (See figure 8.) A 3 (strategy) × 3 (difficulty) × 2 (order) mixed-design ANOVA was performed on the mean distance traversed with strategy and difficulty as within-subjects factors and order as a between-subject factor. The analysis indicated an effect of strategy ($F(1, 10) = 50.54$, $p < .001$) and a strategy × difficulty interaction ($F(2, 20) = 3.89$, $p < .05$). Overall, the effect of strategy showed that greater distance was traversed using the sidestep strategy ($M = 60.82$ m) compared to the head-twist strategy ($M = 58.25$ m). Separate 3 (difficulty) × 2 (order) ANOVAs were performed on traversed distance for each strategy to assess the interaction. There was an effect of difficulty of the path for the head-twist strategy ($F(2, 20) = 6.50$, $p < .01$), but not for the sidestep strategy ($p = .14$). For the head-twist strategy, planned contrasts indicated that less distance was traversed for the high-difficulty path compared to the easy path ($F(1, 10) = 22.67$, $p < .001$). There was no difference between the moderate and easy paths ($p = .16$).

Many subjects stopped immediately after clearing the last obstacle, several meters short of the end of the path; this probably accounts for the otherwise surprising fact that these lengths are approximately the same as the straight-line distance from one end of the path to the other. The conclusion is that subjects did indeed travel a significantly shorter distance with the head-twist strategy than with the sidestep strategy, at least partially explaining why there was a significant difference in traversal times.

3.3 User Preference

All subjects except one preferred the head-twist strategy to the sidestep strategy; the one dissenter—a novice user—felt that both head-twist and sidestep strategies were equally good (although this subject's performance in the experiment, as measured by the number of collisions, was worse with the sidestep strategy than the head-twist strategy). Most subjects felt that having the head angle determine the angle of turn allowed them to simply "walk towards" an object, which is especially important to make the sharp turns required in the experiment. This was a definite advantage over the sidestep strategy, in which many subjects tended to turn too far when trying to avoid an obstacle, colliding with the side wall as a result. One subject lost her bearings completely while navigating one of the virtual paths using the sidestep strategy: she was unable to estimate when to reindex in order to stop turning, and kept going around in circles. Most people, on the other hand, adapted quickly and easily to the head-twist strategy.

4 Discussion

The head-twist strategy employing proportional control with torso trigger significantly outperforms the sidestep strategy employing rate control. Fewer collisions were recorded at all levels of path difficulty, and the degradation in performance as the level of path difficulty increased was much less. Users also took significantly less time to traverse a given path, and the path followed was significantly shorter.

User feedback strongly indicated that the ability to "look where you were going" (a consequence of using proportional control coupled to head twist) was an important benefit of the head twist strategy. With the sidestep strategy, users had to essentially perform a mental integration to determine how much of a sidestep, if maintained for the time they had going at their current pace, would be sufficient to allow them to pass through the available gap between two obstacles. Because the head-twist strategy requires a twist of both head and torso as when stepping to turn, more appropriate kinesthetic and vestibular feedback than for the sidestep strategy would be provided to the user. Users could maintain a strictly forward motion when twisting the head and torso, or they could in fact actually step diagonally into the turn, which is permitted because of the large belt width of the Treadport. In the latter case, turning is in fact completely natural. What makes it unnatural is the need to reindex, but reindexing will be a feature of any turning control method on a linear treadmill. The head-twist strategy may be about the best that can be done.

This paper has presented an initial exploration of a new turning method for the Treadport, and more research is required on the head-twist strategy. The head-twist strategy does not use "pure" proportional control. Although the extent of the turn in the virtual world is proportional to the extent of the control action (the user's head twist), the rate of turning is determined by a decaying exponential with decay factor K_{twist}. A possible refinement of the strategy would be to make the factor K_{twist} dependent on some characteristic of the user's motion on the Treadport (for example, proportional to the user's velocity), but this has not been done in the current implementation. In addition, there is a delay $t_{constant} = 1/15$ sec. between when the user performs the control action and when the controlled effect takes place. This delay might be reduced but not eliminated because of the need to avoid inadvertent turns due to body jitter.

Only certain types of turning can be performed with this strategy. In the real world, it is perfectly possible to turn while the head and torso are facing straight ahead, or while they are twisted in opposite directions. Such turns are not possible with this strategy: the torso has to be twisted enough to express intent to turn, in the direction of desired turn, and the head has to be twisted in the same direction by the amount of turn desired. We also do not consider large-angle turns, such as completely turning around. A turn greater than 90 deg. to the left or right would remove the user from the workspace of the screens. One possibility that was mentioned earlier is to implement a hybrid control scheme, in which smaller turn angles are done by proportional control and larger turn angles are done by rate control (Salcudean et al., 1995).

Another potential problem with the proposed turning strategy is that certain actions may be misinterpreted by the software, and a turn may occur where none was desired. One of the ways this could happen is if the user tries to look back over his shoulder, so that he twists his body around. For example, the user may be alerted by visual or auditory events to look back. The software will take the user through a nearly 180 deg. turn in this case, although the user did not intend to turn. The ability to use guidance cues behind the user to navigate forward may sometimes be important; hand controls for turning have been advocated for this purpose (Bowman, 2002). Again, our visual display does not currently present a surround view, so looking backward will not reveal any visual display. Some dual turning control scheme may have to be devised to cover circumstances such as this, perhaps involving some threshold to switch strategies such as looking back.

It is perhaps not surprising that a proportional control strategy works better than a rate control strategy for tight maneuvering because analogous results have been found in teleoperation (Sheridan, 1992). The more particular issue may be what the control variables are for each strategy being compared. Our goal in this study was not to compare all possible turning approaches because that would be combinatorially impractical. It is also not a pure study because we did not compare, for example, rate control using sidestep to proportional control using sidestep. Instead, after a series of informal studies examining different possibilities, we fixed on the two strategies of this paper as effective strategies for their domain. One of the strategies discarded was rate control using body twist only, which was a strategy briefly mentioned in the introduction. The key issue here does seem to be rate control versus proportional control.

The goal of this research was to find the best turning strategy for 1-D treadmills. Because of their commodity nature, 1-D treadmills hold the promise of proliferating as locomotion interface bases, and therefore the question of how to best utilize them is important. At the same time, it may be the case that the width of the belt is particularly important as well. Narrow belts such as are commonly used for exercise will not permit the sideways excursion that could be helping to make a turning strategy effective in the context of obstacle avoidance. In the future, it would be of interest to quantify how much the width of the belt affects the obstacle avoidance ability.

Comparisons of different locomotion interface types, such as 1-D versus 2-D treadmills or programmable foot platforms (Hollerbach, 2002), cannot be conducted conclusively until each type of interface has been optimized. Because our studies compared strategies for the same device, the results are more clearly interpreted than they would be in a comparison of vastly different devices such as a joystick versus a treadmill. Whether our particular turning strategy generalizes to other user interface devices is not known and was not considered.

Acknowledgments

This research was supported by NSF grant IIS-9908675 and by NSF grant CDA-96-2361. We thank Ryan Hayward for experimental assistance, Sarah Creem for assistance with statistical analysis, and Bill Thompson and Pete Shirley for advice.

References

Bowman, D. (2002). Principles for the design of performance-oriented interaction techniques. In K. Stanney (Ed.), *Handbook of virtual environments technology,* (pp. 277–300). Mahwaw, NJ: Lawrence Erlbaum Associates.

Brooks, F. P., Airey, J., Alspaugh, J., Bell, A., Brown, R., Hill, C., Nimscheck, U., Rheingans, P., Rohlf, J., Smith, D., Turner, D., Varshney, A., Wang, Y., Weber, H., & Yuan, X. (1992). Six generations of building walkthroughs: Final report to the National Science Foundation. Tech. Rep. TR92-026, Department of Computer Science, University of North Carolina.

Christensen, R., Hollerbach, J. M., Xu, Y., & Meek, S. (2000). Inertial force feedback for a locomotion interface. *Presence: Teleoperators and Virtual Environments 9*(1), 1–14.

Darken, R. P., Cockayne, W. R., & Carmein, D. (1997). The omni-directional treadmill: A locomotion device for virtual worlds. *Proc. of ACM User Interface Software and Technology*, 213–221.

Hollerbach, J. M. (2002). Locomotion interfaces. In K. Stanney (Ed.), *Handbook of virtual environments technology* (pp. 239–254). Mahwah, NJ: Lawrence Erlbaum Associates.

Hollerbach, J. M., Mills, R., Tristano, D., Christensen, R. R., Thompson, W. B., & Xu, Y. (2001). Whole-body force feedback realistically simulates slope on treadmill-style locomotion interfaces. *Intl. J. Robotics Research 20*(12), 939–952.

Hollerbach, J. M., Xu, Y., Christensen, R., & Jacobsen, S. C. (2000). Design specifications for the second generation Sarcos Treadport locomotion interface. *Haptics Symposium, Proc. ASME Dynamic Systems and Control Division, DSC-Vol. 69–2,* 1293–1298.

Iwata, H., & Yoshida, Y. (1999). Path reproduction tests using a torus treadmill. *Presence: Teleoperators and Virtual Environments 8*(6), 587–597.

Noma, H., Sugihara, T., & Miyasato, T. (2000). Development of ground surface simulator for tel-E-merge system. *Proc. IEEE Virtual Reality 2000*, 217–224.

Salcudean, S. E., Wong, N. M., & Hollis, R. L. (1995). Design and control of a force-reflecting teleoperation system with magnetically levitated master and wrist. *IEEE Trans. Robotics and Automation, 11,* 844–858.

Sheridan, T. B. (1992). *Telerobotics, automation, and human supervisory control.* Cambridge, MA: The MIT Press.

Templeman, J. N., Denbrook, P. S., & Sibert, L. E. (1999). Virtual locomotion: Walking in place through virtual environments. *Presence: Teleoperators and Virtual Environments 8*(6), 598–617.

Witmer, B. G., & Kline, P. B. (1998). Judging perceived and traversed distance in virtual environments. *Presence: Teleoperators and Virtual Environments 7*(2), 144–167.