Redirected Walking in Virtual Environments
CS4760/5760

James Walker

1. Introduction

A virtual environment is a computer-generated, usually high-fidelity landscape. What distinguishes virtual environments from other computer-generated scenery is that virtual environments strive to maximize the user's feeling of immersion within the setting. This may be accomplished by having the user wear a head-mounted display (HMD) with a visor that uses stereoscopy to present a true 3D image that fills the user's entire field of vision. Other methods may also be used to achieve the feeling of total immersion, such as a CAVE system which surrounds the user with enormous walls of screens. For the purposes of this discussion, we will limit ourselves to virtual environments rendered with an HMD system.

Although fiction is replete with examples of virtual environments that provide a level of immersion equaling or even exceeding that of the real world (such as the Holodeck in the *Star Trek* TV series), in actuality the current state of virtual reality faces many obstacles to achieving this level of immersion. One such problem is that of locomotion within virtual environments.

In order to accurately convey the virtual environment to the user, the user's position and orientation must be known. This requires the use of tracking equipment. There are many methods of tracking users in virtual reality applications, but all of them have the characteristic of limited range. Other factors also limit the user's range of motion. For example, if the user is wearing an HMD, then the user is either limited by the range of the HMD's cables (if it is wired) or by its broadcasting range (if it is wireless). The single most salient factor, however, is that virtual reality simulations are usually conducted in enclosed spaces which, by definition, have limited physical dimensions.

Because of these limitations, the radius of the virtual environment is often much larger than the user's effective range in the real world. In particular, while exploring a virtual environment, the user often must stop due to the presence of a wall in the real world even though the corresponding space in the virtual world is empty. This presents a severe limitation to users' ability to interact with larger-than-room-sized virtual spaces.

Various hardware solutions have been employed to address this problem. One such device is the omni-directional treadmill, as described by Darken et al [1]. This device is constructed from one giant treadmill with many smaller, perpendicular treadmills attached to its surface; a treadmill of treadmills. With this setup, the user can walk an unlimited distance in any direction. Another useful device which accomplishes the same objective is a human-sized hamster ball as presented by Medina et al [4]. However, these devices suffer from the problem that they do not accurately reproduce the inertial feedback experienced by humans walking in the real world. Other novel locomotion devices also exist, but to date, none has been sufficiently effective to gain widespread acceptance. Additionally, such hardware is often very expensive to purchase and maintain.

To address these drawbacks, researchers have attempted to formulate a better, more cost-effective solution to the locomotion problem. In particular, the issue of whether it is possible to address this limitation with a software-only solution has received significant attention. One of the most promising solutions is a technique called **redirected walking.** The rest of this paper will explore the development of this innovative approach in detail, as well as the related topic of **change blindness**.

2. Background

2.1. Overview

Redirected walking relies heavily on various methods of manipulating users' position and orientation in virtual environments without their noticing. This is accomplished by applying a variety of *gains*. Despite the term, gains do not necessarily indicate a rate of increase; they can also indicate a decrease. There are three kinds of gains commonly employed in virtual environments: rotational, translational, and curvature.

- **Rotational gains** cause a user's rate of rotation in virtual space to by either greater or less than the user's physical rotation. This can be applied when the user rotates his head or upper body without moving his legs, or when the user rotates by adjusting his footing. For example, if the user rotates 60 degrees in the real world, a rotational gain of 5/4 could be applied so that the user rotates 75 degrees in the virtual world.
- **Translational gains** cause the user to to move faster or slower when the user translates (i.e., changes position) in the virtual environment. If the user translates by walking (the usual form of movement in virtual environments), a translational gain causes the user to walk a farther or shorter distance in the virtual environment than in reality. For example, if the user walks 30 meters in the real world, a translational gain of 2/3 could be applied so that the user moves 20 meters in the virtual world.
- **Curvature gains** cause the virtual world to rotate while the user is walking. When the user is walking in a straight line, the virtual environment can be rotated by a small amount without the user consciously perceiving the rotation (exact figures are discussed in section 2.2); however, without realizing it, the user will walk in a curved path to compensate for this rotation. With this technique, it is possible to make the user walk in a circle in the real world while the user walks in a straight line in the virtual world.

2.2. Redirected Walking

The first successful application of redirected walking appears in a 2001 study by Razzaque et al [6]. In this study, the researchers modeled a virtual hallway much longer than the physical laboratory where the study took place. In the study, the user was directed to press a series of buttons inside the virtual world in a simulated fire drill. Pressing these buttons brought the user from one end of the virtual hallway to the other in a zig-zag pattern.

Each time the user turned in the virtual environment, a negative rotational gain was applied. As a result, the user had to physically turn 180 degrees in the real world in order to turn a much smaller angle (slightly over 90 degrees) in the virtual world. This caused the users to walk back and forth in the real world while following a zig-zag pattern in the virtual world.
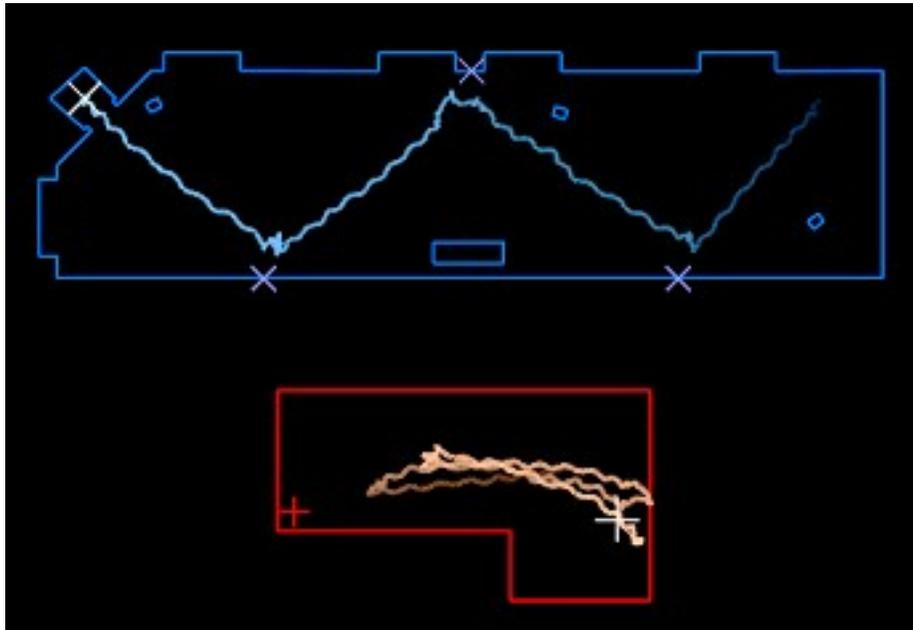
Figure 1. *To-scale diagram showing the paths walked by users in the virtual environment (top, blue) and the real world (bottom, red) in Razzaque et al's study.*

In this study, users never reported that they suspected there was any difference between their virtual movement and their real-world movement. Additionally, the researchers did not observe that manipulating the users' rotational gains caused any increase in simulator sickness. In virtual environments, users sometimes suffer from simulator sickness (with symptoms such as a headache or nausea), so it is important to test whether new techniques cause an increased incidence of this phenomenon.

This study was groundbreaking because it not only demonstrated that redirected walking is possible without alerting the user to the manipulations, but also that it can enable users to explore virtual spaces much larger than their corresponding physical space without using any special hardware.

However, the study also had a major limitation. The users in the study were not allowed to explore the virtual space at will. Instead, they were given an assigned path to follow and were not permitted to deviate from that path. This severely limits the usefulness of the technique, since in most virtual environment scenarios, the user should have some degree of autonomy to explore the space of their own volition. A technique that does not permit this kind of user autonomy has limited practical application.

Other researchers have expanded on the initial work by Razzaque et al. The rest of this section discusses several important studies that have advanced the efficacy of the redirected walking technique.

In 2004, Kuhl [3] conducted a virtual reality study in which subjects' rate of rotation was manipulated. Subjects were not told about the rotational gain, but they were given tasks to acclimate them to the difference compared with their real-world rotation. After this practice, the subjects were asked to perform accurate rotations without the benefit of visual feedback. The result of the study indicated that, given a small amount of practice, users are able to adjust their expectations to accurately measure their rotation even without visual feedback when a rotational gain is applied. This result suggested that the rotational gain technique applied by Razzaque et al might be generalizable.

In 2006, Williams et al [10] explored the effects of applying a fixed translational gain in virtual environments on subjects' perceptual accuracy. The study applied translational gains of 1:1 (no gain), 2:1, and 10:1. Subjects' accuracy was measured by showing them objects, having them translate across the virtual space, and then, with graphics turned off, asked to rotate until they were facing the relevant object. The result of the study was that a translational gain of 2:1 or 10:1 only slightly affected users'

perceptual accuracy, and there was no difference between the 10:1 and 2:1 conditions. The study also examined whether there was any difference in performance between users who play many videogames and those who do not. No difference in performance was observed between the two groups, suggesting that experience with similar interfaces is not a prerequisite to be able to adjust to perceptual manipulations in virtual environments. This study further reinforced the potential utility of using perceptual manipulations to enable people to explore large virtual environments without constantly needing to manually reorient.

In 2008, Steinicke et al [7] set about quantifying the perceptual thresholds at which users can detect various visual manipulations in virtual environments. Up until that point, studies had relied on subjects self-reporting whether they detected any difference in perceptual manipulations. In this study, subjects were instructed to perform various actions (such as rotating to face a certain object), and then, in a two-alternatives-forced-choice (2AFC) answer scheme, used a remote control to indicate whether they believed the rate of change in the movement was greater or less than normal. The point where the user responded "greater" 50% of time was taken as the point where the user could not perceive the virtual movement as being different from the corresponding physical movement.

The following list summarizes specific numerical outcomes. Subjects did not detect:

- rotational gains below +68%;
- rotational gains above -10%;
- translational gains within +/-22%;
- curvature gains along a circle with a radius of at least 24m.

In sum, the study showed that significant alterations can be applied before users are able to detect the difference, which is a highly positive outcome in favor of using perceptual manipulations to facilitate exploration of large virtual environments.

In 2010, Steinicke et al [8] repeated the study and obtained similar results, still suggesting a high rate of perceptual manipulation before users notice the discrepancy. In this study, subjects did not detect:

- rotational gains below +49%;
- rotational gains above -20%;
- translational gains below +26%;
- translational gains above -14%;
- curvature gains along a circle with a radius of at least 22m.

It is worth noting that, for both rotational and translational gains, users are less sensitive to increases than decreases in the rate of change. This is important because it means that applying positive gains is to be preferred whenever possible in a redirected walking algorithm.

In 2008, Engel et al [2] explored an algorithm for dynamically computing rotational gains. Previous studies that used rotational gains to manipulate the user's path relied on the user following a strict predetermined path with precomputed rotational gains. Engel et al developed an algorithm that, although it still required the user to follow an approximate path, allowed for some amount of deviation and adjusted the rotational gains dynamically to prevent the user from running into walls in the physical space. The algorithm also attempted to minimize the gains required to steer the user to reduce the perceptibility of the manipulations.
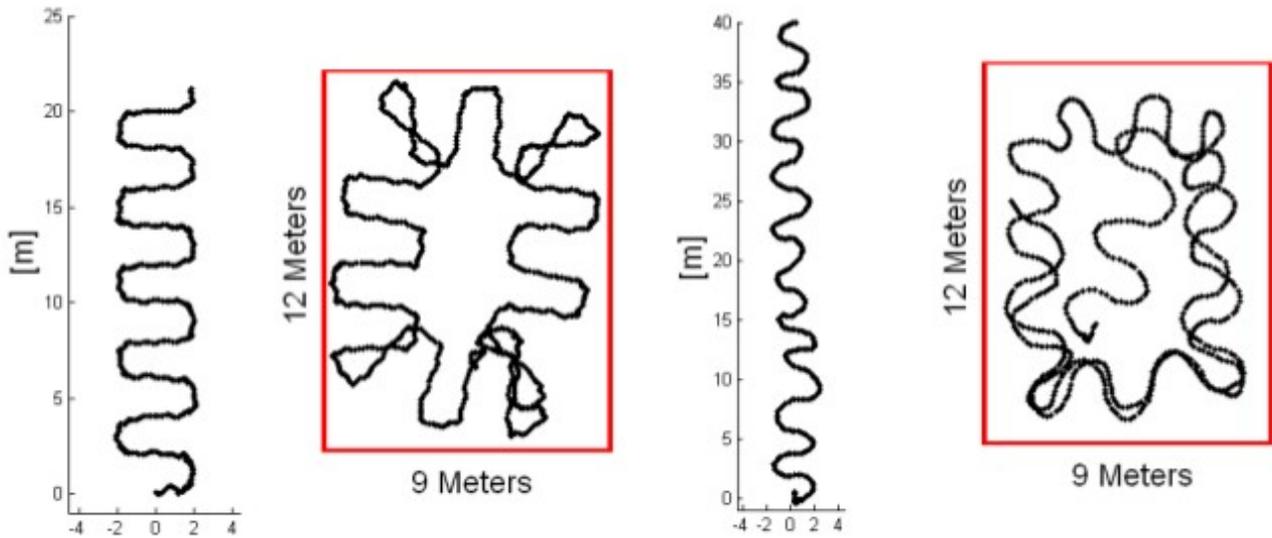
Figure 2. *The left diagram shows the user following a predetermined path in the virtual environment (left) and the path they followed in the physical lab (right) due to the algorithm's dynamically calculated rotational gains. The right diagram shows the same information, except this user deviated slightly from the prescribed path. As the diagram shows, the algorithm was still able to keep the user within the confines of the physical space.*

This work was a step towards generalized redirected walking, but it still required the user to approximately follow a predetermined path. As such, it still did not constitute a complete redirected walking solution.

In 2012, Neth et al [5] explored the implementation of a dynamic curvature gain algorithm. In this study, they discovered that users' perceptual threshold for curvature gain varies based on their locomotive velocity. When subjects walked at 0.75m/s, they could detect curvature gains of $0.1m^{-1}$ (which would result in a circular path with a radius of 10m). Conversely, when subjects walked at 1.0m/s or more, they detected a curvature gain of $0.036m^{-1}$ (circular radius of 27m). In other words, they found that users are more sensitive to curvature gains when walking faster.
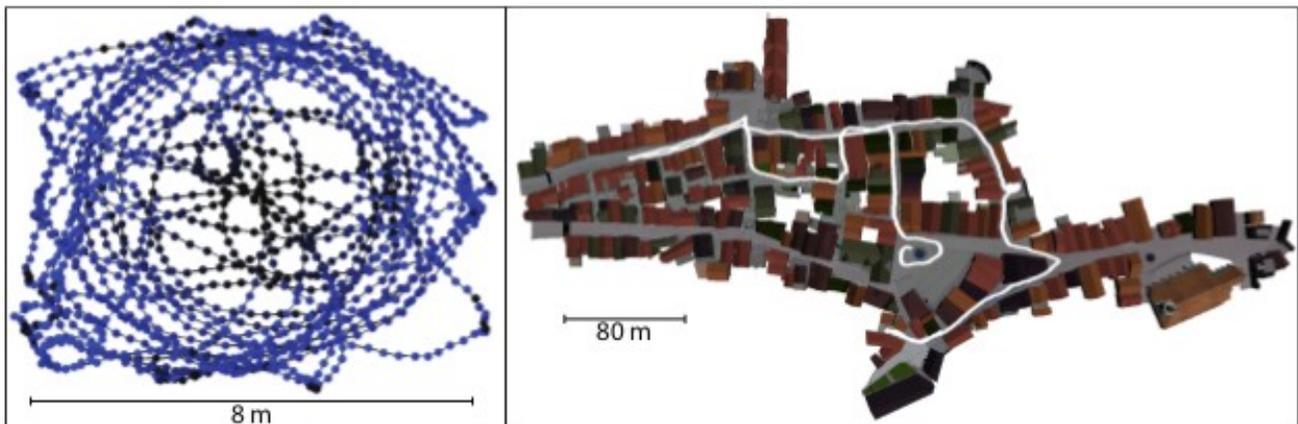

Figure 3. *Diagram showing the path walked by a user in the real world (left) and the virtual environment (right) under curvature gain.*

The study also evaluated the effectiveness of a curvature gain algorithm by allowing users to freely explore a virtual city environment and measuring how far users were able to walk before reaching a physical boundary in the real world and undergoing manual reorientation. This distance was found to be 15m on average when using a static curvature gain and 22m when using a curvature gain that dynamically adjusts based on the user's walking speed.

This study was an important landmark because it implemented a redirected walking technique in an environment that permitted autonomous exploration by the users and significantly increased how far users could walk before reaching a physical boundary.

2.3. Related Techniques

There are also other important techniques related to redirected walking. One of these is manual reorientation, also known as *resetting.* When the user is about to collide with a physical obstacle in the real world, the program informs the user to stop and reorientation occurs.

In 2007, Williams et al [11] investigated three different reset algorithms, *freeze-back up*, *freeze-turn*, and *turn.* In all cases, when the user reaches a boundary in the real world, the program instructs the user to stop. In *freeze-back up,* the program then instructs the user to step backwards a certain number of steps. In *freeze-turn,* the program instructs the user to turn around in physical space while their orientation in the virtual environment is kept the same. In *turn,* the program instructs the user to turn 180 degrees, and as the user does so, the user's orientation in the virtual environment is turned 360 degrees (a form of rotational gain). User error as a result of cumulative resets was found to be the lowest in the *freeze-backup* condition. Most users also reported that they subjectively preferred *freeze-backup* over the other reset methods.

An algorithm that relies completely on resets to prevent users from colliding with the boundaries of their physical space is primitive and does not provide an ideal experience. However, the ability to reset is a necessary component of exploring virtual environments, because even an excellent redirected walking algorithm will sometimes fail to prevent users from approaching real-world obstacles.

A related technique which could potentially be combined with redirected walking to further reduce the number of manual reorientations is *change blindness.* Change blindness is the phenomenon wherein people fail to notice changes to an object or scene; for example, if someone enters a room, and the door behind the person changes position, and when the person turns around, he fails to notice that the door is in a new location.

Suma et al [9] exploited the change blindness phenomenon to achieve the same result as redirected walking; that is, to enable users to explore a large virtual space without reaching the boundary of their corresponding physical space. In this study, users entered a series of rooms where the position of the door was changed while the user's back is turned, thereby altering the user's course and enabling the user to explore a much larger virtual space.

With this technique, users were able to explore a virtual space more than 4 times as large in both dimensions as the physical space, without requiring any manual reorientations. A post-experiment questionnaire that embedded two real questions (whether users detected anything changing position and that they were walking in circles) within several decoy questions revealed that only 1 out of 77 participants noticed the trick. Additionally, despite the changing environment, users were able to maintain their orientation within the environment as demonstrated by their ability to draw accurate maps of the virtual space.
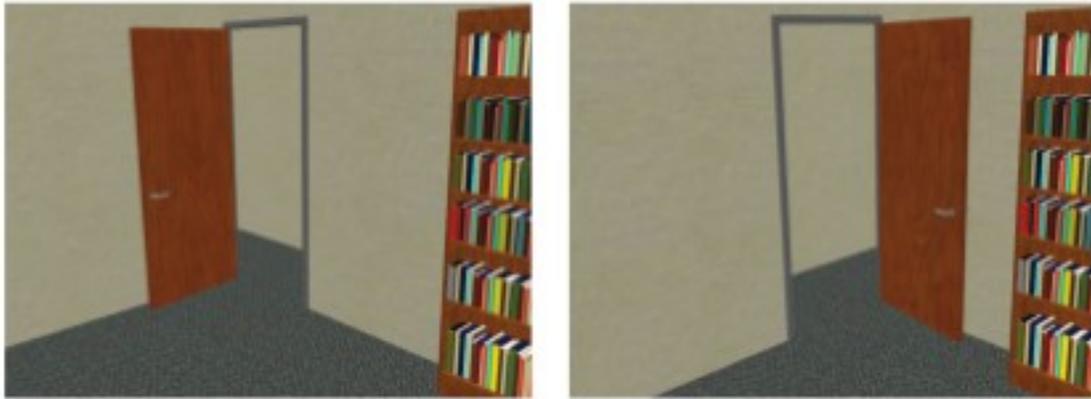
Figure 4. *Exploitation of the change blindness phenomenon by Suma et al [9].*
Left: The room prior to the scene change.
Right: The room after the scene change.

This study proved as an effective proof-of-concept of the change blindness technique. However, like many redirected walking studies, it tightly restricted users' actions, forcing them to enter and exit rooms in a strict sequence. Additionally, change blindness seems difficult to generalize. The most effective method for exploiting change blindness would be to custom-tailor each virtual environment, and that would not always be possible. For example, it is difficult to envision how change blindness could be applied to a wide open outdoor environment. Nonetheless, change blindness may have significant utility as a software solution to enable users to explore large virtual spaces.
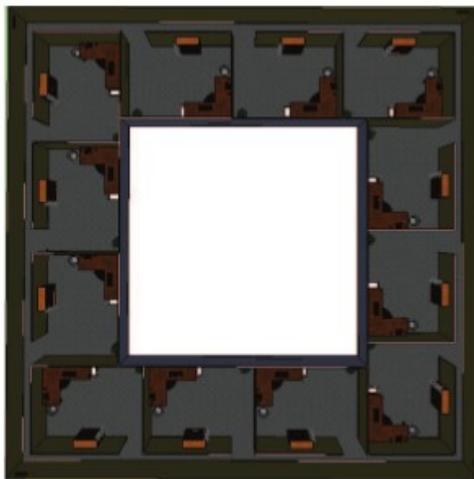

Figure 5. *With change blindness, users were able to explore this*
60x60' virtual environment from within the confines of a
14x14' physical laboratory.

3. Towards a Complete Software Solution

3.1. Current Research

The virtual reality research group at Michigan Technological University is currently conducting research to further generalize redirected walking (R. Zhang, personal communication, February 5, 2013). They have developed an algorithm that dynamically applies translational and rotational gains during undirected user exploration for the purpose of minimizing manual reorientations.

For translational gains, while the user is translating, the algorithm assumes that the user will continue at the current direction and speed and applies a gain that will cause the nearest virtual obstacle in front of the user to coincide with the nearest physical obstacle in front of the user. The algorithm allows a user-specified threshold and will clamp the rate of gain within that threshold to avoid applying an excessive amount of gain.

For rotational gains, the algorithm assumes that the user rotates in place and evaluates 360 possible directions (one per degree). It applies three heuristics with configurable weights to decide which of these directions is best and applies a rotational gain to achieve that direction. The heuristics are:

- **Preference for using tracked space.** Real-world space that does not overlap with physical space is not "reachable" from the perspective of the virtual environment and is effectively wasted. Therefore, the algorithm prefers rotations that maximize the overlap between physical and virtual space.
- **Preference for small rotations.** Large rotations require higher gains which would be more noticeable and potentially more disorienting to the user. Therefore, the algorithm prefers directions which require small rotations.
- **Preference for longest unobstructed path.** The algorithm prefers directions which, assuming the user translates straight forward after the rotation is completed, result in the longest unobstructed path. In other words, the algorithm tries to maximize the amount of linear space in front of the user.

Sometimes, the algorithm still fails to prevent an impending user collision with a real-world obstacle. When this occurs, manual reorientation occurs using the *turn* procedure as described in the reset study by Williams et al [11].

To test the efficacy of the algorithm, the researchers conducted a simulation of a user randomly exploring a virtual space with five different simulated real-world spaces and ten different simulated virtual spaces, where most of the virtual spaces were at least twice as large as the real-world spaces. Every possible combination of real and virtual spaces was tested with different clamping values for the dynamic translational and rotational gains, and the number of manual reorientations that occurred was recorded for each simulation. The following table summarizes the results of the simulation:

Rotation limit	Translation limit	Avg. Reorientations	Avg. Reduction
1.00-1.00	1.00-1.00	25.8	
0.76-1.19	0.78-1.22	17.8	31.5%
0.50-2.00	0.50-2.00	11.5	56.2%
0.76-1.19	0.78-10.0	4.2	84.1%

Table 1. *Summary of the results from simulations conducted using R. Zhang's generalized redirected walking algorithm.*

These results indicate that the algorithm has promise as a fully generalizable technique for reducing the number of manual reorientations when exploring virtual environments. This is especially true when the algorithm is allowed to apply gains which are consciously noticeable to users. Depending on how inconvenient the user finds manual reorientations, noticeable gains may be considered an acceptable measure. This threshold is likely to vary between users. The algorithm accounts for this by allowing the clamp values to be user-configurable.

However, the algorithm still makes unrealistic simplifying assumptions. Most notably, it assumes that user movement consists exclusively of walking in a straight line or rotating in place. However, users often rotate and translate simultaneously. Additional research is ongoing to further refine the algorithm's heuristics, to allow for more natural user movement and possibly incorporate curvature gain, and to test the algorithm's performance against peer algorithms.

3.2. An Ideal Solution

Since the seminal work by Razzaque in 2001 [6], redirected walking has made great strides. However, there are still many unsolved problems. It is still impossible to avoid the reset procedure altogether. This will likely remain true in the general case, so the objective of the research is to find a combination of techniques that minimizes the number of resets to the greatest extent possible.

Many current techniques still rely on fixed-rate gains, which have been repeatedly shown to have inferior performance to solutions that use dynamic gains, such as the 2012 study by Neth et al [5] and the work currently ongoing at Michigan Technological University. Further utilization and refinement of dynamic gain algorithms will be necessary to bring out the full effectiveness of this technique. As yet, there are still no existing algorithms that apply all three major types of gains in combination with each other.

Additionally, redirected walking is not the only software technique available for facilitating the exploration of large virtual spaces. The change blindness work conducted by Suma et al [9] showed highly impressive results in its ability to allow exploration of a relatively enormous space with zero reorientations, while preserving users' internal sense of orientation, and only 1 out of 77 participants noticing that objects in the virtual space were changing places. Although this technique would be difficult to generalize and is not applicable to wide open spaces, it should be possible to create an algorithm to dynamically calculate changes to indoor environments to reduce manual reorientations.

An ideal software solution would combine all of these techniques. It would use dynamic gains for translation, rotation, and curvature, with values optimized from many experiments to achieve an ideal combination of user comfort (i.e., not noticing the perceptual manipulations) and fewer reorientations. The algorithm would also exploit change blindness by dynamically calculating environmental changes to facilitate user exploration. Furthermore, because these techniques interact with each other, they could not be calculated in isolation; instead, they would have to be aware of each other and perform their calculations with awareness of their disparate methods and overlapping objectives. For example, the change blindness portion of the algorithm would take into account potential reorientations due to rotational gain and change the location of geographical features appropriately, and vice-versa.

Designing, testing, and refining such a general and robust algorithm would obviously be an immense undertaking. However, by building on the existing research, creating the ultimate general algorithm for exploration of large virtual spaces should be fully achievable in practice. All that is required is to combine all of the existing techniques in a single algorithm with mutual awareness of each other's capabilities.

4. Conclusion

One of the greatest obstacles facing the proliferation of virtual reality technology is the difficulty of exploring large virtual environments with a limited physical space. Hardware solutions tend to be prohibitively expensive and difficult to maintain, and most of them do not allow for fully natural-feeling locomotion. Thus, significant research is ongoing into flexible software solutions, which virtual reality users could utilize at no additional cost.

Chief among these solutions is the technique called redirected walking. Redirected walking exploits humans' inability to detect perceptual manipulations within a certain threshold to steer users away from physical obstacles without the users' knowledge. Redirected walking techniques consist of translational, rotational, and curvature gains. The efficacy of all of these techniques has been demonstrated in numerous studies, but only in the last few years have general algorithms that do not assume the user follows a preset path been implemented. The current state of the research is to refine these algorithms and make them flexible enough to account for all possible types of movement by the user.

A closely related technique to redirected walking is change blindness. Change blindness exploits the fact that humans are often poor at noticing changes to their environment when they are not directly observing it. By changing the locations of portals and the directions of hallways, a change blindness algorithm can steer users away from physical world obstacles while allowing the exploration of large virtual environments.

However, even the best redirection techniques sometimes fail to steer the user away from physical obstacles. When this happens, a manual reorientation is required by telling the user to stop and either translate or rotate to a new position and orientation better suited to further exploration.

In theory, a fully general and flexible software solution to the problem of exploring large virtual environments should be achievable. This would be accomplished by combining all existing techniques and continuing to expand their generality. Additionally, further data is needed by running many experiments to see what values for these algorithms yield the best results. With this goal in mind, the ideal of facilitating the exploration of virtual environments of limitless size without requiring any special hardware can be realized.

References

[1] Darken, R. P., Cockayne, W. R., & Carmein, D. (1997, October). The omni-directional treadmill: a locomotion device for virtual worlds. In P*roceedings of the 10th annual ACM symposium on User interface software and technology* (pp. 213-221).

[2] Engel, D., Curio, C., Tcheang, L., Mohler, B., & Bülthoff, H. H. (2008, October). A psychophysically calibrated controller for navigating through large environments in a limited free-walking space. In *Proceedings of the 2008 ACM symposium on Virtual reality software and technology* (pp. 157-164).

[3] Kuhl, S. (2004, August). Recalibration of Rotational Locomotion in Immersive Virtual Environments. In *Proceedings of First ACM SIGGRAPH Symposium on Applied Perception in Graphics and Visualization (APGV) (*pp. 23-26).

[4] Medina, E., Fruland, R., & Weghorst, S. (2008, September). VIRTUSPHERE: Walking in a human size VR "hamster ball". In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* (Vol. 52, No. 27, pp. 2102-2106). SAGE Publications.

[5] Neth, C. T., Souman, J. L., Engel, D., Kloos, U., Bulthoff, H. H., & Mohler, B. J. (2012). Velocity-dependent dynamic curvature gain for redirected walking. *Visualization and Computer Graphics, IEEE Transactions on*, *18*(7), 1041-1052.

[6] Razzaque, S., Kohn, Z., & Whitton, M. C. (2001, September). Redirected walking. In *Proceedings of EUROGRAPHICS* (Vol. 9, pp. 105-106).

[7] Steinicke, F., Bruder, G., Jerald, J., Frenz, H., & Lappe, M. (2008, October). Analyses of human sensitivity to redirected walking. In *Proceedings of the 2008 ACM symposium on virtual reality software and technology* (pp. 149-156).

[8] Steinicke, F., Bruder, G., Jerald, J., Frenz, H., & Lappe, M. (2010). Estimation of detection thresholds for redirected walking techniques. *Visualization and Computer Graphics, IEEE Transactions on*, *16*(1), 17-27.

[9] Suma, E. A., Clark, S., Krum, D., Finkelstein, S., Bolas, M., & Warte, Z. (2011, March). Leveraging change blindness for redirection in virtual environments. In *Virtual Reality Conference (VR), 2011 IEEE* (pp. 159-166).

[10] Williams, B., Narasimham, G., McNamara, T. P., Carr, T. H., Rieser, J. J., & Bodenheimer, B. (2006, July). Updating orientation in large virtual environments using scaled translational gain. In *Proceedings of the 3rd Symposium on Applied Perception in Graphics and Visualization* (pp. 21-28).

[11] Williams, B., Narasimham, G., Rump, B., McNamara, T. P., Carr, T. H., Rieser, J., & Bodenheimer, B. (2007, July). Exploring large virtual environments with an HMD when physical space is limited. In *Proceedings of the 4th symposium on Applied perception in graphics and visualization* (pp. 41-48).