

ABSTRACT

Sharif Razzaque

Redirected Walking

(Under the direction of Frederick P. Brooks Jr.)

There are many different techniques for allowing users to specify locomotion in human-scale, immersive virtual environments. These include flying with a hand-controller, using a treadmill, walking-in-place, and others.

Real walking, where the user actually and physically walks in the lab, and virtually moves the same distance and in the same direction in the virtual scene, is better than flying. It is more input-natural, does not require learning a new interface, results in a greater sense of presence, and theoretically results in less simulator sickness.

One serious problem with real walking, however, is that the size of the virtual scene is limited by the size of tracked area. For example, for an architect to really walk in a virtual prototype of a house, the tracked area must be as large as the house. This requirement makes real walking infeasible for many facilities and virtual scenes.

To address this limitation, I have developed Redirected Walking, which by interactively and imperceptibly rotating the virtual scene around her, makes the user turn herself. Under the right conditions, Redirected Walking would cause the user to unknowingly walk in circles in the lab, while thinking she is walking in a straight and infinitely long path in the virtual scene.

In this dissertation I develop Redirection, discuss its theoretical and physiological underpinnings, and presents results to show that it can be used: 1) to make the user turn themselves 2) without causing the user to be aware of Redirection 3) without unacceptably increasing the user's level of simulator sickness and, most importantly, 4) to useful effect:

A) In head-mounted display systems, the user can experience a virtual scene larger than the lab while also having the benefits of real walking.

B) In an open-backed, three-walled CAVE, users can have the increased presence and input-naturalness normally associated with a fully enclosed CAVE

I also present guidelines for VE practitioners wishing to use Redirection, based on the theory and observations reported in this dissertation.

DEDICATION

To my parents, Umme Salma Razzaque and Abdur Razzaque.

ACKNOWLEDGEMENTS

Over the seven years that I worked on this dissertation, many people have helped me. I have seen free sharing of original ideas, gentle criticism, equipment, code, and data among individuals, projects, and institutions; and enjoyed a passionate, warm, and supportive community of friends and researchers. For each area in which our department did not have expertise, those external researchers who did welcomed me, took interest in my work, and guided me. One could not hope for a better setting in which to work and explore.

In particular, I thank all my committee members—Fred Brooks, Mary Whitton, Gary Bishop, Don Parker, Mark Hollins, and Anthony Steed—for their interest in me and this work, and their inspiration, guidance, patience, and encouragement.

There are many others to whom I am grateful:

- Zachariah Kohn, who co-developed and conducted the RW experiments with me, and developed the 3D spatial audio for this research.
- My colleagues and friends at University College London—David Swapp, Mel Slater, and Anthony Steed—who developed and administered the RWP experiments with me and dedicated much of their facility’s resources to this work.
- At the Human Interface Technology Lab of the University of Washington-Seattle: James Lin, Henry Duh, and Don Parker, who are my guides to self-motion perception and simulator sickness; and Cameron Lee and Konrad Schroder for technical help with my videoconference proposal meetings and oral exams.
- Greg Welch, Stephen Brumback, and Kevin Arthur, who spent hours reviving dead Hiball trackers during the first experiment. I remember with particular gratitude the time that Greg Welch came in at midnight so I would not have to cancel the next day’s session.
- Geoffrey Melvill-Jones of the University of Calgary for discussion and for his videos regarding the podokinetic after-effect.

- Robert Kenney of RSK Assessments, Inc., for discussion about the SSQ and its use in flight simulators and VEs.
- Kim Swinth of UC Santa Barbara and Sarah Nichols of the University of Nottingham for their respective SSQ data and discussion.
- Denny Proffitt of the University of Virginia and Mark Hollins for help with psychophysical experimental procedures.
- Bill Chung and Sorren LaForce at NASA's Ames Vertical Motion Simulator for discussion and demonstrations of washout—and for letting me fly!
- Ian Strachan, editor of Jane's Simulation and Training and former UK Royal Air Force test pilot, for his qualitative assessments and comparisons of motion simulators.
- Montek Singh, Leandra Vicci, and Gary Bishop for helping me understand Fourier and Laplace analysis.
- Henry Fuchs for suggesting the idea of waypoints.
- Russ Taylor for help with VRPN and numerous other technical and intellectual issues.
- Eric Burns, Dorian Miller, and Luv Kohli for further development of Redirection.
- The entire Effective Virtual Environments (EVE) Team at the UNC-Chapel Hill Computer Science Department and particularly Paul Zimmons, Mark Harris, Angus Antley, Mike Meehan, Ben Lok, Paul McClaurin, and Brent Insko.
- The technical services group at the UNC-Chapel Hill Computer Science Department, including David Harrison, Kurtis Keller, John Thomas, Mike Stone, Bil Hays, David Musick, Jane Stine, Linda Houseman, Fred Jorden, Murray Anderegg, Alan Forest, John Sopko, Chester Stephen, Mike Carter, and Brian White for heroic technical support during many emergencies.
- The anonymous subjects who volunteered to participate in the experiments.
- Sally Robertson for editing help.
- Beth Nassef and Ajith Mascarenhas, my around-the-clock work partners these last six months of writing.
- My parents, for their unwavering and unconditional help, despite not understanding exactly what I was working on or why it was so exciting to me.

- Fred Brooks and Gary Bishop, for changing the way I see and think.

I am also grateful for support and funding from:

- University of North Carolina Board of Governors Fellowship
- The Ross and Charlotte Johnson Family Dissertation Fellowship
- NIH National Center for Research Resources and National Center for Biomedical Imaging and Bioengineering
- UK Equator EPSRC project
- Office of Naval Research
- The Latané Center for Human Science

CONTENTS

List of Tables	xv
List of Figures.....	xvi
Chapter 1: Overview.....	1
1.1 Goal	1
1.2 Thesis Statement and Other Results.....	2
1.3 Overview of Dissertation	3
1.4 Virtual Environment Systems.....	3
1.5 Locomotion	4
1.6 Redirected Walking	6
1.7 Redirected Walking-in-Place.....	8
1.8 Presence	9
1.9 Self-Motion Perception	9
1.9.1 Auditory.....	10
1.9.2 Tactile, Proprioceptive, and Podokinetic	10
1.9.3 Vestibular.....	11
1.9.4 Visual	12
1.9.5 Visual and Vestibular Senses Complement Each Other.....	12
1.9.6 Combining Information from Different Senses into a Coherent Self-Motion Model ...	13
1.9.7 Hypothesis of How Redirection Works	14
1.10 Simulator Sickness.....	16
1.10.1 Theory.....	16
1.10.2 How Redirection Avoids Sickness	17
1.10.3 Quantitative Measures of Sickness with Redirection.....	17

1.11	Descriptions of Experiments.....	19
1.11.1	RWp and RW.....	20
1.11.2	RWP.....	20
1.11.3	RDT.....	21
1.12	Noticing Redirection.....	21
1.12.1	Informal Assessment.....	21
1.12.2	Operational Definition of <i>Notice</i>	22
1.12.3	Experienced Users and the Lower Bound of the Detection Threshold of Rotation	22
1.13	Lab Size Required for Infinite Virtual Scenes.....	23
1.14	Steering Algorithms for Unrestricted Exploration of Arbitrary Virtual Scenes.....	24
1.15	Conclusions.....	26
Chapter 2:	Locomotion Interfaces.....	28
2.1	Locomotion.....	28
2.2	Locomotion in Virtual Scenes.....	28
2.3	Locomotion Techniques.....	29
2.3.1	Flying.....	29
2.3.2	Leaning.....	29
2.3.3	Treadmills.....	29
2.3.4	Walking-in-Place.....	30
2.3.5	Real Walking.....	30
2.3.6	Manipulating the World.....	30
2.4	The Difficulties of Comparison.....	31
2.5	Attributes Relevant to This Thesis.....	32
2.5.1	Input-Motion-Naturalness.....	32
2.5.2	Ease of Learning & Ease of Use.....	32
2.5.3	Motion Cues.....	33
2.5.4	Simulator Sickness.....	33
2.6	Comparison of VE Locomotion Techniques in Terms of Attributes Relevant to This Thesis.....	33

2.6.1	Flying	33
2.6.2	Treadmills.....	33
2.6.3	Real Walking	34
2.6.4	Redirected Walking.....	34
Chapter 3: Simulator Sickness		35
3.1	Consequences of Simulator Sickness.....	36
3.2	Difficulties in Understanding Simulator Sickness	37
3.3	Factors That Aggravate Simulator Sickness.....	37
3.4	Theories of the Mechanisms of Simulator Sickness	38
3.4.1	Cue Conflict.....	38
3.4.2	Postural Instability.....	39
3.4.3	Poison.....	40
3.4.4	Rest-frames and the Internal Mental Motion Model.....	40
3.5	Measuring Simulator Sickness	42
Chapter 4: Self-Motion Perception		44
4.1	Difficulties in Studying Self-Motion Perception	45
4.2	Overview.....	45
4.3	The Vestibular Sense	46
4.4	Auditory Self-Motion Perception	53
4.5	Proprioceptive and Tactile Self-Motion Perception.....	53
4.6	The Podokinetic system	53
4.7	Visual Self-Motion Perception	55
4.8	Visual Perceptual Stability.....	57
4.9	Integration Among the Senses.....	60
4.10	Visual-Vestibular Interaction.....	60
4.10.1	Tilt and Linear Acceleration Ambiguity.....	60
4.10.2	Washout in Flight Simulators.....	62

4.10.3	Differences between Visual and Vestibular Motion-Sensing in the Frequency Domain and in Onset Timing	63
4.10.4	The Vestibulo-Ocular Reflex	65
4.10.5	The Optokinetic Reflex.....	66
4.10.6	The OKR and VOR Complement Each Other.....	66
4.11	Efference-Copy Prediction	67
4.12	Proprioceptive-Vestibular Interaction	69
4.13	Proprioceptive-Visual Interaction	70
4.14	The Internal Mental Motion Model.....	71
4.15	Quantitative Characterizations of the Senses.....	74
Chapter 5:	How Redirection Works— Qualitative Arguments Based on Self-Motion Perception Theory.....	76
5.1	Self-Motion is the Simplest Explanation for the Sensory Cues Caused by Redirection.....	76
5.2	Non-Visual Cues	77
5.3	Algorithm Description in Terms of What the User is Doing.....	78
5.3.1	While Standing Still	78
5.3.2	While Really Turning the Head.....	78
5.3.3	While Walking.....	79
5.4	Improvements to Redirection Suggested by Self-Motion Perception Literature.....	80
5.4.1	Looking Down	81
5.4.2	Running	81
5.4.3	Faraway Virtual Objects	82
5.4.4	Taking Advantage of Podokinetic High-Pass Characteristics	82
Chapter 6:	Steering the User during Unrestricted Walking.....	86
6.1	Steer the User Toward the Center of the Lab	87
6.2	Proposed Algorithm: Steer the User Onto a Circular Orbit.....	90
6.3	Proposed Algorithm: Steer the User Toward Changing Targets.....	91
6.4	Guidelines for Designers of Steering Algorithms.....	93
Chapter 7:	The Redirected Walking Experiment: RW.....	94

7.1	Task and Virtual Scene	94
7.2	Subjects.....	95
7.3	VE System Details.....	96
7.4	Redirection Algorithm	96
7.5	Observations and Lessons Learned.....	98
7.5.1	The HMD Veil Increases User Discomfort.....	98
7.5.2	Redirection’s Sensitivity to Tracking Glitches.....	99
7.6	Spatial Audio.....	99
7.6.1	Motivation.....	99
7.6.2	Sound Cues	100
7.6.3	Earphones	101
7.6.4	Spatial Audio Algorithms	101
Chapter 8:	The Redirected Walking-in-Place Experiments: RWP	102
8.1	Overview.....	102
8.2	Motivation.....	103
8.3	Virtual Scene and User Task	104
8.4	VE System Details.....	105
8.5	Users.....	105
8.6	Experimental Measures.....	106
8.7	Experiment RWP-I	108
8.8	Problems with the RWP Implementation Revealed in RWP-I and Rectified in RWP-II .	109
8.8.1	Redirection Algorithms	109
8.8.2	Walking-in-Place Detection.....	111
8.9	Experiment RWP-II Results.....	113
8.10	Observations and Summary of Results.....	115
8.11	Comparison to Other Locomotion Techniques in CAVEs.....	116
Chapter 9:	Experiments to Determine What Level of Injected Scene Rotation Users Will Notice	119
9.1	The Lower Bound of Imperceptible Rotation Rate.....	119

9.2	A Precise Definition of <i>Notice</i> —a Review of Concepts from Psychophysics.....	119
9.2.1	Detection Thresholds.....	120
9.2.2	Signal Detection Theory: Sensitivity and Bias.....	120
9.2.3	Methods for Determining Thresholds.....	122
9.3	Experimental Designs.....	125
9.3.1	Adjustment of Visual Scene Angular Velocity While Standing Still: RDT-sev.....	126
9.3.2	Adjustment of Visual Scene Oscillation Frequency While Standing Still: RDT-ssv..	127
9.3.3	Detection of Direction of Scene Rotation While Walking.....	127
9.4	Experimental Details.....	131
9.5	Results	133
9.6	Caveats.....	134
Chapter 10:	The Simulator Sickness Questionnaire and its Bearing on Redirection	136
10.1	Background on Statistical Analysis Techniques	136
10.2	Power Analysis	137
10.3	History and Development of the SSQ	139
10.4	Diagnostic and Statistical Power of the SSQ for Flight Simulators	140
10.5	Application of SSQ to General Purpose VEs.....	140
10.6	SSQ Scores from Redirection vs. Real Walking.....	144
10.7	Redirection Induces Less Simulator Sickness than Turning Manually.....	145
Chapter 11:	Future Opportunities.....	147
11.1	Redirected Avatar Limbs.....	147
11.2	Wireless HMD VE System.....	148
11.3	The Effects of Redirection on Spatial Cognition.....	150
11.4	SSQ Scaling for VEs and General-Population Users	150
11.5	Better Walking-In-Place Implementation.....	151
11.6	Redirection of Walking Speed	152
11.7	Using Virtual Distracters.....	153
11.8	High Body Momentum.....	153

11.9	The Effects of Spatial Audio.....	153
Chapter 12: Guidelines for Developers.....		155
12.1	Guidelines for all VEs.....	155
12.2	Guidelines Specific to Redirection.....	156
12.3	Guidelines for Redirected Walking-in-Place	158
Appendix: Laplace Analysis Background		159
1	Introduction	159
2	The Semicircular Canals.....	159
3	Feedback Systems.....	159
4	Further Restrictions on Systems	160
5	Step Response	160
6	Exponential Decay.....	161
7	Linearity: Response of Filters in Series	163
8	Frequency Response	164
8.1	Signal Analysis in the Frequency Domain	164
8.2	Filter analysis in the Frequency Domain.....	164
8.3	Phase Offset.....	166
8.4	Cutoff Frequencies.....	167
9	Bodé Plots	168
10	Converting Between Exponential Decay and Frequency Response Representations.....	169
11	Transfer Functions: Laplace Domain Representations of Filters.....	169
12	Filters That Compute the Derivative and Integral of a Signal	171
13	A Single Filter can act as Both an Integrator and Differentiator.....	172
14	Converting a Filter to Operate on the Integral or Derivative of its Input	173
References		176

LIST OF TABLES

Table 1.1 - Of the four sensory channels addressed in this dissertation, only two (visual and auditory) can be directly controlled by the VE systems I used.	16
Table 1.2 - List of experiments and their abbreviations.	20
Table 3.1. Factors that correlate with decreased susceptibility (in users) to simulator sickness.....	37
Table 3.2 - Qualities of VE systems and flight simulators that increase simulator sickness.	38
Table 3.3 - The SSQ questionnaire.	43
Table 4.1 - Summary of values of the band-pass filter characteristics of three sensory modalities for inducing a sensation of rotation.....	75
Table 4.2 - Various reported rotation detection thresholds of the semicircular canals.....	75
Table 5.1 - A comparison of how each cue is stimulated to induce PKAR, for the original experiment [Gordon 1995] and my VE system proposal.	83
Table 7.1 - Description of labels, scenario-related purpose, and VE system response of each virtual wall-mounted button.....	95
Table 8.1 - Number of subjects for whom data was collected for each experiment and condition.....	106
Table 8.2 – The six questions from the presence questionnaire used in the RWP experiments.	107
Table 8.3 - A model that predicts a user’s sense of presence as a function of how much she noticed the rotations, how much she saw the back wall, and how much she turned her head.	109
Table 8.4 – The questions used to determine if the subjects noticed that the virtual scene rotated, compared to other phenomena which did not happen. The aggregate responses for each group are listed in the right-hand columns.....	114
Table 9.1 - The possible outcomes from a single signal detection trial.....	121
Table 9.2 - The chance rotation rate (CRR) and other data for the staircase sessions of experiment RDT-wcv.	133
Table 10.1 - Comparison of SSQ data from various sources.	141

LIST OF FIGURES

Figure 1.1 - Piglet and Pooh go hunting for Woozles and keep finding more and more sets of Woozle tracks, not realizing they are following their own tracks around the bush.	1
Figure 1.2 - Left: A user wearing an HMD, standing in a tracked lab. Right: A user in a CAVE.	4
Figure 1.3 - The partial floor plan of a real house and a view of the kitchen of a virtual model of the house.	6
Figure 1.4 - Left: The virtual scene used in experiment RW I. Right: Overhead views of the actual path in the virtual scene (above in blue) and in the real lab (below in red), drawn to scale.	7
Figure 1.5 - Left: A CAVE with an open back wall (with the virtual scene turned off). Right: An overhead diagram of the same CAVE.	9
Figure 1.6 - A cut-away illustration of the outer, middle, and inner ear, revealing the vestibular system	11
Figure 1.7 - A Bodé plot showing the response of the semicircular canals (SCCs).....	11
Figure 1.8 - Three optical flow patterns.	12
Figure 1.9 - The contribution of the visual and vestibular (or inertial) systems to the perception of a step function in angular velocity.....	13
Figure 1.10 - Visual and vestibular responses (compiled from several sources) as a function of frequency	13
Figure 1.11 - The virtual pit scene.....	15
Figure 1.12 - Box and Whisker plots of the SSQ scores for the Hand-Controller Turning and Redirection groups from experiment RWP-II.....	18
Figure 1.13 - The portion of the RWP questionnaire to gauge the extent to which subjects noticed the room rotation, compared to other phenomena which did not actually occur.	22
Figure 1.14 - Simulated paths a user walking an infinitely long straight line in the virtual scene under worst-case conditions.....	24
Figure 1.15 - Illustrations of sample paths of a user from three different steering algorithms.	25
Figure 1.16 - How the Steer-to-Center algorithm handles unexpected changes in the user's path.	25
Figure 2.1 - Bowman's taxonomy of flying locomotion techniques.	32
Figure 4.1 - The human inner ear labyrinth.....	46
Figure 4.2 - The macula.	47
Figure 4.3 - A single hair cell.	48
Figure 4.4 - Two views of the hollow, fluid-filled, vestibular bone structures, showing the three semicircular canals and their ampullae in relation to the cochlea.	48
Figure 4.5 -A simplified diagram of a single semicircular canal.....	49

Figure 4.6 - Cupula being distorted by motion.....	49
Figure 4.7 – A Bodé plot of cupula deflection as a function of the frequency of sinusoidal head rotational velocity.....	51
Figure 4.8 - Hydrodynamic properties of canal-cupula-endolymph system during a step up and down in rotational velocity.	52
Figure 4.9 - The rotating treadmill used by Gordon et al.	54
Figure 4.10 - The rotating turntable used by Weber et al.....	54
Figure 4.11 - Three types of optical flow patterns.....	56
Figure 4.12 - An optokinetic drum	58
Figure 4.13 - The frame and light illusion	59
Figure 4.14 - A flow diagram showing motion-state estimation from multiple sensory cues.	60
Figure 4.15 - Otolith ambiguity in sustained acceleration.....	61
Figure 4.16 – A false sensation of pitch due to forward acceleration.....	61
Figure 4.17 - A flight simulator with a motion base (the NASA Ames VMS).....	62
Figure 4.18 - Washout	63
Figure 4.19- The visual-vestibular crossover.	63
Figure 4.20 - The contribution of the visual and vestibular (inertial) senses, in the time domain, to the perception of a step in rotational velocity (about the yaw axis)	64
Figure 4.21 - Washout allows the simulator’s cab to stay within its range while making the pilot feel like she continues to accelerate.	65
Figure 4.22 - Efference-copy during rotation of the eye.....	68
Figure 4.23 - A process diagram of self-motion perception, with re-afference and efference copy prediction ...	68
Figure 4.24 - A simplified plot of PKAR velocity as a function of time.....	70
Figure 4.25 - A model of self-motion perception, showing contributions of the internal mental motion model and of efference copy and re-afference	71
Figure 4.26 - The anti-gravity room.	73
Figure 4.27 - The Ames Room illusion.....	74
Figure 5.1 - A simulated path of a user, who is walking in a straight line in the virtual scene, but due to PKAR-Redirection, is walking in a spiral in the lab.....	84
Figure 5.2 - A simulated path of a user, computed using the same simulation and PKAR-Redirection algorithm as in Figure 5.1, but where the user turns left by 90 degrees once during the simulation, and otherwise walks straight.....	85

Figure 6.1 - Steer-to-Center algorithm.....	87
Figure 6.2 - Informal testing of the Steer-to-Center algorithm.....	88
Figure 6.3 - A recorded path of a person walking a relatively straight path. The wobble is related to the person shifting weight from one foot to the other.	88
Figure 6.4 - Left: The steering rate is attenuated by multiplication by the sine of angle θ , the angle between the user's heading and the vector pointed toward the lab center. If the user is pointed perpendicular to the lab center, $\sin(90)=1$ and the steering rate is not attenuated. As the user turns past the lab center (as in Figure 6.1), the steering changes smoothly. Right: A sample path of the user steered toward and then through the lab center.....	89
Figure 6.5 – Steer-to-Center algorithm.....	89
Figure 6.6 – A problem with the Steer-to-Center algorithm.....	90
Figure 6.7 - Left: The user is steered onto a circular path orbiting the lab center. Superimposed are three hypothetical sample paths that the user could take in the virtual scene and in the lab.	90
Figure 6.8 – Steer-onto-Orbit algorithm.....	91
Figure 6.9 - Steer-to-Changing-Targets algorithm.....	92
Figure 6.10 - If the user is steered through target A and then happens to be facing directly away from both targets A and C, the system must not choose C as the next target.....	92
Figure 7.1 - Left: A user's view in the headset as she walks toward the button to sound the alarm. Right: A view of the entire virtual room (the front wall is removed for clarity).	95
Figure 7.2 - A flow diagram of the Redirection algorithm used in experiment RW.....	97
Figure 7.3 - Left: Overhead views of the path taken by the user in the virtual scene (above left, in blue) and the laboratory (below left, in red). Right: The user's path superimposed onto the virtual scene.....	98
Figure 7.4 - A user's view in the headset as she walks toward the button to close the windows. An antique radio, used for presenting pre-recorded instructions, is in the foreground.....	100
Figure 8.1 - An illustration of how RWP works.....	104
Figure 8.2 - The path in the virtual scene taken by one subject in the Redirection group.....	104
Figure 8.3 - The hand-tracking sensor attached to a hip-worn camera bag in order to track the torso orientation.	105
Figure 8.4 - Theta is the angle between the user's torso heading and the front CAVE wall.	110
Figure 8.5 - The RWP-I algorithm.....	110
Figure 8.6 - The RWP-II algorithm.....	111
Figure 8.7 - Left: The accelerometer for detecting footsteps (the black box with the white wire) was attached to the top of the blue head-tracking sensor. Right: A sample footstrike as recorded by the accelerometer.....	112

Figure 8.8 –Regression lines and actual data points, showing how much subjects saw the back wall, as a function of how much they turned their head, and which experimental group they were in.	114
Figure 9.1 – Idealized response curves resulting from the constant stimulus technique.	123
Figure 9.2 - An idealized sample progression of stimulus intensity when using the staircase method to estimate the stimulus’s detection threshold.	125
Figure 9.3 - A pilot subject manipulating the control knob in experiment RDT-scw.	126
Figure 9.4 - Photographs of a subject during trials of experiment RDT-wcv.	128
Figure 9.5 - The staircase progression of two sessions of experiment RDT-wcv.	129
Figure 9.6 - Top: The staircase progression of subject 8, session 1, but with 6 randomly interspersed trials where there was no rotation.	130
Figure 9.7 - The response curve from subject 2, session 2; performed using the constant stimulus technique.	130
Figure 9.8 - Views of the RDT virtual scene.	131
Figure 9.9 - The rotation-rate during the start-up period of each trial in experiment RDT-wcv.	132
Figure 9.10 - Response curves from the constant stimulus sessions.	134
Figure 10.1 - Measuring more peoples’ heights can uncover a significant difference between the heights of women and men.	138
Figure 10.2 - The SSQ scores from one of our VEs [Meehan 2003] have a similarly shaped distribution to that presented in Kennedy 1993. However, the scales are very different.	142
Figure 10.3. Box and Whisker plots of SSQ scores for hand-controller turning vs. Redirection.	146
Figure 11.1 - The user’s tracked virtual hand penetrates the virtual antique radio	147
Figure 11.2 - As the user lowers her hand onto a virtual tabletop, her real hand location may penetrate the virtual table. The VE system displays her virtual hand such that it stays on top of the table, while her real hand is actually beneath the virtual table.	148
Figure 11.3 - A prototype wearable image generator I built in 2001.	149
Figure 11.4 - Determining the direct path back to the starting place without revisiting the intermediate stopping points requires path integration.	150
Figure 1 – A graphical representation of a filter.	160
Figure 2 - The step response of an arbitrary low-pass filter.	161
Figure 3 - The step-response of an arbitrary low-pass filter and high-pass filter.	161
Figure.4 - The output of a low-pass filter with time-constant T_1 , when given a step up and step down in rotational velocity.	162

Figure 5 - The output of a single-pole, single-zero high-pass filter with time-constant T_2 when given a step up and step down in rotational velocity.	162
Figure 6 - Filter composition: Filter C is composed of Filter A and B. A's output is B's input.	163
Figure 7 - The output of the SCCs, computed by composing the T_1 low-pass and T_2 high-pass filters.	163
Figure 8 - A square wave decomposed as the sum of a series of sinusoids.	164
Figure 9 - Three sinusoids of the same frequency.	164
Figure 10 - The output of arbitrary idealized high- and low-pass filters when the input is a square wave (left), and the frequency response of those filters (right).	165
Figure 11 - The square-wave response of the same steep roll-off filters as in Figure 10, with idealized and realistic phase offsets.	166
Figure 12 - The phase offset and square pulse output of an arbitrary filter with no delay (upper) and a fixed-time delay (lower).	167
Figure 13 - A fixed time duration (30ms, for example) constitutes a greater phase offset for higher frequency sinusoids.	167
Figure 14 - The frequency response of low- and high-pass single-pole filters compared to that of the more complex filters from Figure 10 and Figure 11.	168
Figure 15 - A single-pole filter's frequency response plotted on a log magnitude and a log frequency scale. ...	169
Figure 16 - Bodé plot of an ideal integrator and differentiator.	171
Figure 17 - A cosine and sine wave. The integral of $\cos(t)$ is $\sin(t)$, which lags 90 degrees behind. The derivative of $\sin(t)$ is $\cos(t)$, which leads 90 degrees ahead. This explains the 90 degree phase lag and lead of the integrating and differentiating filter, respectively.	172
Figure 18 - A Bodé plot of SCC response in terms of angular velocity.	173
Figure 19 - Composing the SCC filter with an integrator turns it into one that accepts angular acceleration instead of velocity.	174
Figure 20 - A Bodé plot of SCC response in terms of angular acceleration.	174
Figure 21 - A Bodé plot of SCC response in terms of angular displacement.	175

Chapter 1: Overview

1.1 Goal

The goal of this work is to simulate walking in large, life-sized, immersive virtual environments (VEs) in a way that allows virtual scenes to be larger than the physical space available, captures the naturalness and sense of *presence*¹ associated with real walking, and does not increase the simulator sickness suffered by the user. To this end, I have developed the technique *Redirection*—making the user turn herself by interactively and imperceptibly rotating the virtual scene about her. Under the right conditions, Redirection can cause the user to unknowingly and continuously walk in circles in the lab, while *thinking* she is walking on a straight and infinitely long path in the virtual scene and real world. This is similar to a situation described by Milne in *Winnie-the-Pooh*, in which Pooh and Piglet unknowingly walk around and around in a circle while hunting Wozzles (Figure 1.1) [Milne 1926].



Figure 1.1 - Piglet and Pooh go hunting for Wozzles and keep finding more and more sets of Wozzle tracks, not realizing they are following their own tracks around the bush [from Milne 1926, copyright Penguin Group Books for Young Readers, used with permission].

¹ The user's feeling that she is really *in* the virtual scene, rather than the feeling of viewing it on a display.

1.2 Thesis Statement and Other Results

In this dissertation, I develop the technique Redirection and present results to show that it can be used:

- 1) to make the user turn herself;
- 2) to useful effect:
 - a. in head-mounted display (HMD) VE systems, the user can experience a virtual scene larger than the lab while also having the benefits of real walking;
 - b. in an open-backed three-walled CAVE®, the user can have the increased presence and input-naturalness normally associated with a fully enclosed CAVE;
- 3) without causing the user to be aware of Redirection or modify her conscious behavior because of its use in the VE system;
- 4) without unacceptably increasing the level of simulator sickness suffered by the user.

Beyond supporting the thesis statement, I present other results of my research:

- hypothesized mechanisms for why Redirection works, in terms of current self-motion perception and simulator sickness theories;
- experimental results of how fast the virtual scene can imperceptibly rotate under worst-case conditions (1 deg/s);
- an estimate of what size lab is required for the user to walk an arbitrarily long straight virtual path (30 by 30 meters);
- algorithms for steering the user in the lab while she is freely exploring arbitrary virtual scenes;
- a waypoints technique for implementing Redirection in labs which are not large enough;
- observations from several implementations of Redirection;
- guidelines for developers wishing to implement Redirection.

I give simulation results in situations where the available tracking area was large enough for experimental work.

1.3 Overview of Dissertation

This dissertation is written for a general Computer Science audience. Since this work draws on many different areas, I discuss background and relevant literature over several chapters. These background topics include virtual environments, locomotion techniques, statistical power analysis, self-motion perception, psychophysical measurement techniques, simulator sickness, and Laplace analysis. The chapters may be read independently or even skipped, depending on the background of the reader. Chapters 2-4 and the Appendix are background on virtual locomotion interfaces, simulator sickness, self-motion perception, and Laplace analysis. Chapters 5 and 6 present proposed mechanisms of Redirection and algorithms to steer users. Chapters 7-10 discuss the experimental designs and results. Chapter 10 discusses issues in designing an experiment to show that Redirection does not increase simulator sickness, and includes supporting background material on power analysis. The Laplace Analysis Background Appendix is different from much of the literature in that it does not assume knowledge of electrical engineering or make analogies to circuit design.

For the person reading this dissertation in order to implement Redirection in her own virtual environment systems, I suggest Chapters 1, 5, 6, 7, and 12.

This first chapter is an introduction and synopsis of the dissertation. Much of what I discuss here is covered in greater detail elsewhere in the dissertation.

1.4 Virtual Environment Systems

Immersive virtual environment systems attempt to give the user the impression that she is *in* a synthetic or *virtual* scene. Many such systems *track* the location of the user's head (in the real world) and present visual imagery as seen from the user's viewpoint in the virtual scene. As the user turns her head, she sees the imagery that she would see if the virtual scene were real and she were actually in it. In some systems, the visual imagery is presented on video displays directly in front of the user's eyes inside a headset or *head mounted display* (HMD) (

Figure 1.2). Other systems, such as the CAVE Automatic Virtual Environment² (CAVE), project the imagery onto the surfaces (walls, floors, etc.) of the real room [Cruz-Neira 1993]. Some VE systems and vehicle simulators forgo tracking the user's head because she is seated (approximately fixing her head's position to a known location), use specialized optics to reduce the registration errors that would otherwise result from small head movements, and simulate only the faraway virtual objects (which are less sensitive to small errors in head position values). Desktop 3D graphics are not addressed in this dissertation.




Figure 1.2 - Left: A user wearing an HMD, standing in a tracked lab. Right: A user in a CAVE.

1.5 Locomotion

As stated earlier, the goal of this work is to simulate *walking* in large virtual scenes. Walking is one form of *locomotion* (the self-movement of an organism from one place to another). The simulation of vehicles such as cars, planes, boats, etc. is beyond the scope of this dissertation. Locomotion is distinct from *navigation* or *way-finding* (finding a route between two locations), which is a cognitive task.

² CAVE® is a recursive acronym and a registered trademark of the University of Illinois Board of Trustees and FakeSpace Systems, Inc. I use the term to generically refer to CAVEs® and CAVE-like displays. Other generic terms include *Immersive Projection Technology* (IPT) and *Surround-Screen Virtual Reality* (SSVR), but neither are commonly used.

There are many techniques for allowing users to specify locomotion in human-scale, immersive virtual scenes, and these are detailed in Chapter 2. These include *flying* with a joystick or other hand-controller [Robinett 1992], using a treadmill [Brooks 1992], *walking-in-place* (where the user makes walking motions but keeps herself physically in the same spot) [Slater 1995], leaning [Peterson 1998], and others [Stoakley 1995; Miné 1997]. The choice of locomotion technique has been shown to affect the user's experience, sense of presence [Slater 1998; Usoh 1999], and, I believe, the level of simulator sickness. Presence is important for many VE applications such as training and phobia desensitization [Hodges 1994], and simulator sickness is a serious problem for many users [Kolasinski 1995].

Real walking, where the user actually and physically walks in the lab, and virtually moves the same distance and in the same direction in the VE, is better than flying with a joystick or walking-in-place. Real walking is more input-natural³ and does not require learning a new interface. It has been statistically shown to result in a greater sense of presence than flying, and there are strong arguments and some evidence that it is more presence-inducing than walking-in-place [Slater 1995; Usoh 1999]. Based on literature on simulator sickness [Kolasinski 1995; Kennedy 2003b], I also believe that real walking results in less simulator sickness than other means of locomotion.

One serious problem with real walking, however, is that the size of the virtual scene is limited by the size of tracked area or lab (whichever is smaller). For example, for an architect to really walk in a virtual prototype of a house, the tracked area must be as large as the house. This requirement makes real walking infeasible for many virtual scenes and facilities. Henceforth, I use the term *lab* to mean the physical tracked space where the user is during her VE session, regardless of whether the physical room is actually a lab or some other kind of space (e.g., an industrial design studio). If the tracked area is smaller than the physical room, *lab* refers to only the tracked part.

³ The inputs the user makes to the VE system are more like the motions a person makes to walk in the real world.

The first time I experienced real VE walking was in a detailed, realistic, and beautiful virtual scene of a particular house. I was deeply impressed with the sense of presence it invoked in me. At the same time, I was disappointed I could not explore the area beyond the virtual kitchen, because only the kitchen fit into the lab (Figure 1.3).




Figure 1.3 - The partial floor plan of a real house (right) and a view of the kitchen of a virtual model of the house (left). Only the kitchen of the house (red dashed outline) fits into our lab.

1.6 Redirected Walking

To address this limitation of real walking, Redirected Walking works by making the user turn herself by interactively and imperceptibly rotating the virtual scene around her. Under the right conditions, Redirected Walking would cause the user to unknowingly walk in circles in the lab, while thinking she is walking on a straight and infinitely long path in the virtual scene.

In 1994, Michael Moshell and Dan Mapes at the University of Central Florida, attempted to manipulate VE users into unknowingly walking along an arc while thinking they were walking straight. Simulator sickness and limitations in VE systems, particularly tracker technology, thwarted them. I independently came up with this idea in 1999, and by then general improvements in VE systems and the development of accurate, low-latency, wide-area trackers made Redirected Walking feasible.

To maximize presence, the injected rotation (the Redirection) should be imperceptible. The goal of the algorithm is to exploit the limitations of human perceptual mechanisms for sensing position, orientation, and movement, so as to minimize the intrusiveness of the injected rotation. The amount and direction of rotational distortion injected is a function of the user's real orientation and position in the lab, linear velocity, and angular velocity.

In the extreme, Redirected Walking could cause the user to walk in a large circle in the lab, while she thinks she is walking in a straight line in the virtual environment. Theoretically, if there were enough tracked

area for the complete circle, the VE system could present a virtual scene of infinite extent. Given a lab of limited size, there is a trade-off: the more rotational distortion (resulting in the user walking in tighter arcs), the larger the virtual environment one can present. However, the more rotational distortion, the more likely the rotations will intrude on the user's consciousness.

To make Redirected Walking usable for labs of limited area, one can circumvent the above trade-off by forcing the user to look around at strategically placed *waypoints* in the virtual scene. While the user is rotating herself to look around, the system can inject substantially more rotational distortion without it being perceived. The virtual scene is rotated so that a direction which was previously out of tracker range is now safely within the lab. The distance between adjacent waypoints must be less than the length of the tracking area. Figure 1.4 illustrates the use of waypoints. Whereas the need for waypoints imposes a major constraint on the virtual scene's design, I believe that many tasks, such as the fire-drill task users performed in the first user study (later referred to as RW I), naturally lend themselves to waypoints.




Figure 1.4 - Left: The virtual scene used in experiment RW I. Subjects⁴ performed a fire-drill task, pushing buttons on the wall to activate an alarm, close the windows, etc. The path of one subject is shown in blue, superimposed on the floor. Right: Overhead views of the actual path in the virtual scene (above in blue) and in the real lab (below in red), drawn to scale. As the user zigzagged through the virtual scene, she unknowingly walked back and forth between the ends of the lab instead.

⁴ In this dissertation, I use the term *subject*, rather than *participant* to refer to the human test users who volunteered for the experiments in this work. I avoid *participant* for two reasons. First, it is vague—many people participate in an experiment and not all of them are subjects. The experimenters and technicians also participate. Second, some collaborators use the term *participant* to mean the person who is experiencing the immersive virtual environment, even if this experience or session is not part of any experiment. The rationale is

1.7 Redirected Walking-in-Place

Redirected Walking requires the use of an HMD and a large tracking area (Figure 1.2). CAVEs are much more common than large-area trackers—it was estimated there were 600 CAVEs in 2001 and three to eight new ones were being installed every month [Coffin 2001]. In a CAVE, the choice of locomotion technique is also important. Redirection can also be applied to locomotion in CAVEs if combined with walking-in-place. Previous research [Slater 1995] shows that walking-in-place results in higher presence than flying with a hand-controller. However, even with walking-in-place, the CAVE user must still turn in the VE using a hand-controller.

Traditionally, if a user wishes to move toward an object in the virtual scene, she must first rotate the virtual scene using a hand-controller (e.g., joystick) so that the virtual object is in front of her. Previous research suggests that input motions that are more natural lead to a greater sense of presence [Slater 1998]. Data from one of my studies (RWP-I) show a correlation between a user's sense of presence and her physically turning the body (to face a virtual object) instead of turning the world with a joystick. This suggests that either a user who is more present would rather turn her body, or that a user who turns her body is more likely to be present. With Redirected walking-in-place (RWP), our goal is to allow the user to turn in the VE by turning her body instead of using a joystick.

The problem with turning the body, however, is that the vast majority of CAVEs have only three vertical walls [Coffin 2001] (Figure 1.5). If the user turns her body, she will eventually face the open back wall. Redirected walking-in-place slowly and imperceptibly rotates the virtual scene, while the user is walking-in-place, so that the user is made to turn toward the front wall of the CAVE without noticing. While the user is

that, in real life, a person is a participant in the world, not a user of the world. Since immersive virtual environments attempt to simulate the real world, the person experiencing it is a *participant* not a *user*. To avoid confusion, I use the terms as follows: *person*, when referring to anyone of the human race (e.g., a person must eat to survive); *user*, when referring to the subset of persons that experiences a virtual environment; and *subject*, when referring to those users who volunteered to allow the experimenters to collect data on them. I appreciate their active collaboration and contributions and do not use the term *subject* disrespectfully.

standing in one place and turning her head to look about the virtual scene, the system scales the rotation so that she can see more of the virtual scene before turning so far that she sees the open back wall.




Figure 1.5 - Left: A CAVE with an open back wall (with the virtual scene turned off). Right: An overhead diagram of the same CAVE.

1.8 Presence

The user's sense of *presence* is roughly the feeling of being *in* the virtual scene. For some applications, presence is the most important attribute [Hodges 1994]. There is debate as to the precise definition of presence and the best ways of measuring it. For the purpose of this work, I use Slater's definition [Slater 1999]: *Presence* is an internal psychological and physiological state of the user. It is distinguished from *immersion*, which refers to the sensory stimuli presented (such as display field-of-view and imagery update rate) and the virtual scene. Immersion and presence are related in that many researchers believe greater immersion of the user (e.g., by way of a wider field of view) evokes a greater sense of presence in the user (at least up to some saturation level).

1.9 Self-Motion Perception

My goal is to simulate self-motion via walking and make the injected virtual scene rotation imperceptible. This goal is aided by an understanding of the perception of self-motion.⁵ There are several

⁵ Some researchers use the term *self-motion perception* to mean the perception of one's translation only, and *self-motion cognition* to mean the perception of both one's translation and rotation. I use *self-motion perception*

sensory channels (or *modalities*) that provide information on how and where one is moving, such as auditory, visual, vestibular, and proprioceptive. Each of these contributes information to one's awareness of self-motion, and under certain circumstances each can elicit a sensation of self-motion by itself. Humans rely on these sensory cues for balance and orientation [Dichgans 1977] and to determine whether they themselves are moving (self-motion) or if the objects around them are moving (external motion).

1.9.1 Auditory

Humans have the ability to deduce qualities of their environment from the way the environment sounds (e.g., large rooms sound different than small rooms) and the ability to localize sound sources. Several mechanisms for this are discussed in Chapter 4. As a person moves, the perceived source of the sound moves appropriately (in relation to the person's head). In fact, a moving sound source alone can by itself cause a stationary person to feel as if she is moving [Lackner 1977a]. As I detail in Chapters 7 and 12, having good spatial auditory cues can greatly contribute to the effectiveness of VEs.

1.9.2 Tactile, Proprioceptive, and Podokinetic

Humans can sense movement in their joints, muscles, and viscera, and can sense pressure and slippage on the skin. These cues are important for walking, as they tell a person where her limbs are and when her feet are touching the ground. These cues indicate the relative motion of the person's body (i.e., how the limbs move relative to the torso). Of particular interest is the *podokinetic* (motion-of-feet) system, which contributes to both the sensation and control of one's orientation while walking. As described in Chapter 4, the podokinetic sense of orientation is plastic, and experimental subjects have been induced (after several minutes of habitation to the proper stimuli) to unknowingly walk in tight arcs even in the absence of visual cues [Weber 1998; Jürgens 1999].

to refer to one's perception of all forms of self-motion (orientation, translation, limb motion, twisting the torso, etc.).

1.9.3 Vestibular

The vestibular system is able to sense motion of the head with respect to the world. Physically, the system consists of labyrinths in the temporal bones of the skull, just behind and between the ears. The vestibular organs are divided into the semicircular canals (SCCs) and the saccule and utricle (Figure 1.6). As a first-order approximation, the vestibular system senses motion by acting as a three-axis rate gyroscope (measuring angular velocity) and a three-axis linear accelerometer [Howard 1986b]. The SCCs sense rotation of the head and are more sensitive to high-frequency components of motion (above roughly 0.1 Hz) (Figure 1.7), whereas real motions are full spectrum. Because of this, it is often not possible to determine absolute orientation from vestibular cues alone. Humans use visual information to complement and disambiguate vestibular cues. Missing or misleading visual cues can lead to life-threatening motion illusions in pilots [Berthoz 2000; Cheung 2000]. On the other hand, flight simulators and Redirected Walking take advantage of the ambiguity of the vestibular cues.




Figure 1.6 - A cut-away illustration of the outer, middle, and inner ear, revealing the vestibular system [adapted from Martini 1998].




Figure 1.7 - A Bode plot showing the response of the semicircular canals (SCCs), modeled using a two-pole band-pass filter with time-constants of 3 ms and 10 sec (Howard 1986). The dashed lines denote the 0.1 to 5 Hz region, in which the SCCs are most sensitive.

1.9.4 Visual

Visual cues are the dominant modality of perceiving self-motion. Visual cues alone can induce a sense of motion—this is known as *vection*. The kinds of visual processing can be separated into landmark recognition (or *piloting*), where the person cognitively identifies objects (e.g., chairs, windows) in her visual field and so determines her location, and optical flow. Optical flow is a lower-level phenomenon wherein the movement of light patterns across the retina are sensed. In most situations, the optical flow field corresponds to the motion field. For example, if the eye is rotating in place, to the right, the optical flow pattern is a laminar translation to the left. When a person is moving forward, the optical flow pattern radiates from a center of expansion (Figure 1.8). Both optical flow and landmark recognition contribute to a person’s sense of self-motion [Warren Jr. 2001; Riecke 2002].




Figure 1.8 - Three optical flow patterns. Left: Laminar translation (which would result from turning one’s head left). Center: Radial expansion (which would result from moving forward). Right: Circular (which would result from rolling about the forward axis).

1.9.5 Visual and Vestibular Senses Complement Each Other.

As mentioned, the vestibular system is most sensitive to high-frequency motions. On the other hand, the visual system is most sensitive to *low*-frequency components of motion. The vestibular and visual systems complement each other (Figure 1.9). This is a critical concept in self-motion perception, and Redirection aims to take advantage of this. The crossover frequency of the two senses (Figure 1.10) has been reported to be about 0.07 Hz [Duh 2001b].




Figure 1.9 - The contribution of the visual and vestibular (or inertial) systems to the perception of a step function in angular velocity. The vestibular system detects the initial, high-frequency step, whereas the visual system perceives the sustained, low-frequency rotation [from Rolfe 1986].




Figure 1.10 - Visual (orange solid line) and vestibular (blue dashed line) responses (compiled from several sources) as a function of frequency [adapted from Duh 2004].

1.9.6 Combining Information from Different Senses into a Coherent Self-Motion Model

Each sensory modality provides information about different qualities of a person's motion. These pieces of information are fused to create an overall sense of self-motion. There are two challenges that must be addressed by this process. First, the information must be fused quickly so that it is up to date and relevant (e.g., the person must know that and how she has tripped in time to regain balance and footing before hitting the ground). Second, the total information, across all the sensory channels, is often incomplete. One theory is that the human processes sensory self-motion cues in a manner similar to the Kalman filter [Rolfe 1986]: At any given time, a person has a *model* or hypothesis of how she and the surrounding objects are moving through the world. This model is based on assumptions (some of which are conscious and cognitive, while others are innate or hardwired) and previous sensory information. New incoming sensory cues are evaluated *in terms of this model*, rather than new models being continuously constructed from scratch. For example, if a person is on a

stopped train, and the stopped train on the adjacent track starts to accelerate, she might have the brief sensation that *her* train has started moving instead. This model is consistent with all her sensory information thus far, perhaps until she looks out the other side of her train and notices the trees are stationary (relative to her train). She has a moment of disorientation or confusion and then, in light of this new information, she revises her motion model such that her train is now considered stationary. In short, *one perceives what one is expecting to perceive*. This is an explanation for why so many illusions work [Gregory 1966]. An illusion is simply the brain's way of making sense of the sensory information—a model of the world, based on assumptions and sensory information, that happens to be wrong [Berthoz 2000].

Perception is an active process, inseparably linked with action [Berthoz 2000]. Because sensory information is incomplete, one's motion model is constantly tested and revised via interaction with the world. The interaction among cues provides additional self-motion information. For example, if a person sees the scenery (e.g., she is standing on a dock, seeing the side of a large ship only a few feet away) shift to the left, it could be because she herself turned to her right, or because the ship actually started moving to her left. If she has concurrent proprioceptive cues that her neck and eyes are turning to the right, she is more likely to conclude that the ship is still and that the motion in her visual field was due to her actions. The active process of self-motion perception relies on prediction (of how the incoming sensory information will change because of the person's actions) and feedback.

1.9.7 Hypothesis of How Redirection Works

The most basic reason why Redirection works imperceptibly, I believe, is simply that the user is not expecting it. The world rotating about the center of the head is uncommon in the real world; the world doing so in response to the person turning her head, is even less common. While she is using the VE system, her perceptual motion model that says she is herself turning, rather than the world turning around her, is sufficient to explain her sensory cues. This is the case even for VE systems that do not employ Redirection⁶—the user

⁶ And also not using a locomotion technique such as flying or a treadmill.

perceives the virtual world as remaining stable, rather than moving in response to her movements [Jaekl 2002]. Even if the user consciously knows about Redirection, the illusion is still convincing. This is similar to the virtual pit scene (Figure 1.11), wherein the user finds herself on the edge of a virtual precipice, but consciously knows she is on real, solid ground. Still, the user often cannot make herself step out across the virtual precipice and, when she can, it requires strong willpower [Usoh 1999; Meehan 2001]. There are many illusions whose effectiveness is not reduced by the observer knowing how the illusion works [Gregory 1966].



Figure 1.11 - The virtual pit scene. Left: A photograph of the real world. The user knows there is a real floor. Center: view of the virtual scene. Right: The virtual scene from the user's viewpoint.

As mentioned earlier, humans have visual, auditory, vestibular, and proprioceptive senses of self-motion. Humans rely on these sensory feedback cues for balance, orientation, and to distinguish self-motion from external motion. Previous research suggests that keeping multiple cues consistent (with each other and with the user's internal mental motion model) increases the chance that the user will perceive rotation as self-motion as opposed to external motion [Lackner 1977a]. The goal is to maximize the probability that the user will perceive all of the movements of the virtual scene as self-motion, rather than as the world moving arbitrarily around her.

Since the VE systems used for this dissertation work can create only synthetic visual and auditory cues, the challenge is to simulate self-motion while keeping all the cues, even those that the system cannot control, consistent. Since the vestibular and proprioceptive cues are more sensitive to high-frequency motions, when the user is walking on a straight virtual path, the Redirection algorithms only inject smooth and gradual rotations of the visual and auditory scene. Because the vestibular and proprioceptive senses are not sensitive to this kind of low-frequency rotation, conflict between the visual-auditory and vestibular proprioceptive cues is minimized.

Table 1.1 - Of the four sensory channels addressed in this dissertation, only two (visual and auditory) can be directly controlled by the VE systems I used. Thus, the goal is to manipulate the visual and auditory cues in a such a way that they remain consistent with the vestibular and proprioceptive cues. If the synthetic visual cues produced by the VE system portray only low-frequency head rotations, the vestibular system will not be able to detect that those low-frequency rotations are absent in the user's real-world motion, and thus the virtual visual and real-world vestibular cues will be consistent.

<i>Sense</i>	<i>Cues come from</i>
Visual	VE System
Auditory	VE System
Vestibular	Head motion in real world
Proprioceptive	Body motion in real world

Even while standing still, the user unknowingly rotates her head and torso with the virtual scene. We hypothesize that the user's own balance mechanisms are responsible for this [Lackner 1977b]. Subjects instructed to remain balanced and standing still can be caused to sway by VE visual cues alone [Howard 1986a]. While walking, in an attempt to stay on a virtual trajectory that she perceives as straight, the user unwittingly veers in the direction of the injected rotation. At waypoints, the rapid turning while looking around causes substantial vestibular stimulation. Against this high-frequency background, an additional vestibular stimulation that would be noticed were the user not turning herself is now less noticeable. Therefore, the user does not notice the increased rotational distortion we inject while she is looking around.

For Redirected Walking to be maximally successful, the user should register and respond to the continuously updated orientation of the VE, without recognizing it as externally induced. When the technique keeps the visual, auditory, and vestibular cues consistent, the added rotation should cause users to change direction, and it should be unnoticed. Furthermore, the additional virtual scene rotations caused by Redirection should not increase the simulator sickness of the user.

1.10 Simulator Sickness

1.10.1 Theory

Simulator sickness, in which the user becomes sick during the simulation (but not in the real situation which is being simulated [Pausch 1993]) is a serious problem for VEs. Its symptoms are similar to those of motion sickness and include nausea, dizziness, blurred vision, disorientation, and vertigo [Kennedy 1995; Kolasinski 1995]. There are a myriad of different theories of the mechanisms behind motion and simulator sickness, and these are detailed in Chapter 3.

The theories have many ideas in common, and the take-away message from all of them is that sickness can arise when one's motion model (described above) is invalidated by conflicting incoming sensory information. For example, if a person is watching a wide-screen movie where the camera is moving, the visual cues indicate she is moving but her vestibular cues tell her she is still. Having an operational motion model is critical for survival. It is used not just for navigation, but also for maintaining balance and posture [Stoffregen 1988], and even for stabilizing the eyes so they can function properly [Draper 1996]. Having an invalidated motion model is serious, debilitating, and can result in sickness.

I believe simulator sickness to be a serious impediment to user acceptance of VEs. Kolasinski reports one study in which 45% of the users reported symptoms of simulator sickness after using a commercially available VE system for 20 minutes [Kolasinski 1995]. I know of no VE system that does not induce sickness for at least some users (even if the sickness has not been formally quantified).

1.10.2 How Redirection Avoids Sickness

As mentioned above, the visual and vestibular systems are sensitive to different frequencies of motion. To minimize both the user's simulator sickness (above the level caused by VE systems that do not use Redirection) and the user's conscious detection of the rotation, the Redirection algorithm keeps the rotations of the virtual scene as low-frequency as practical. This low-frequency property keeps the injected rotations below the visual-vestibular crossover frequency, thus minimizing the conflict between visual and vestibular cues. Previous research shows that differing visual and vestibular cues (i.e., the visual cue is from one motion path and the vestibular cue is from another) are more likely to cause sickness when those cues are in a frequency band where both the visual and vestibular systems are sensitive [Duh 2004]. If the cues are in a frequency range where either channel is insensitive, there is less conflict and less sickness.

1.10.3 Quantitative Measures of Sickness with Redirection

A standard measure of simulator sickness is Kennedy's Simulator Sickness Questionnaire (SSQ) [Kennedy 1993]. It is used to compute a simulator sickness score from a user's subjective self-report of the severity of various symptoms she experiences after being exposed to a VE.

The subjective reports have great statistical noise and variation from person to person, and many users are minimally susceptible to simulator sickness. The initial calibrating analysis, with more than 3600 SSQ

reports, relied on 75th percentile SSQ scores, instead of a significance test on the mean scores, to differentiate troublesome (in terms of sickness) flight simulators from acceptable ones [Kennedy 1993].

The most common alternative to using Redirection is to use a hand-controller. Using Kennedy's method, the results of experiment RWP-II suggest that walking-in-place with Redirection results in *less* simulator sickness than walking-in-place while using a hand-controller to turn (Figure 1.12).

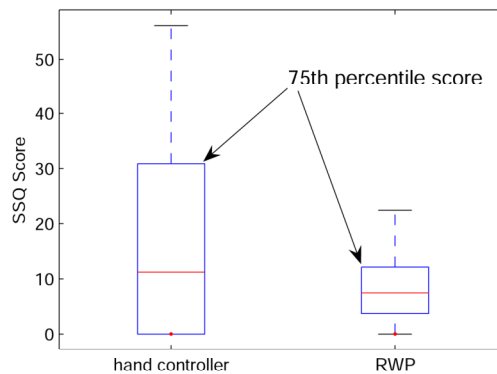


Figure 1.12 - Box and Whisker plots of the SSQ scores for the Hand-Controller Turning and Redirection groups from experiment RWP-II. The 75th percentile score (the metric used by Kennedy et al. in the original SSQ work) is lower for the group using Redirection.

A VE practitioner might like to know if Redirection causes additional simulator sickness compared to a similar VE system that uses real walking. More specifically, one would like to know if the additional sickness caused by Redirection is enough to make a previously acceptable VE system troublesome once Redirection has been added to it. I hypothesize that the increase in simulator sickness caused by Redirection, if there is any, is insignificant (not statistically but operationally). However, I cannot support this hypothesis quantitatively.

A *power analysis* (described in Chapter 10) could be used to argue that Redirection does not cause an increase in SSQ scores. Performing a power analysis requires 1) an estimate of how SSQ scores vary from user to user, for users and VE systems similar to the ones employed in this work, and 2) an estimate of how much SSQ scores would have to increase, because of Redirection, in order to be considered troublesome. These parameters are known for flight simulators and military pilots, but I argue that those values do not apply to general-population users and HMD VE systems. In my analysis, I have used SSQ scores from roughly 300 exposures to the real-walking VE system in our lab (200 of which were collected specifically for this dissertation's analysis). From this I have estimated how SSQ scores vary in our users. I do not, however, have any meaningful estimate for the second parameter, because I have not seen (and therefore have no SSQ data

from) any system in our laboratory that is considered troublesome (with respect to simulator sickness). Furthermore, no researcher I queried from the (non-flight simulator) VE community could quantify an SSQ score that would be high enough to be considered troublesome. If I arbitrarily propose the same SSQ effect that Arthur found in his simulator sickness vs. HMD field-of-view study [Arthur 2000] (his difference in mean SSQ scores was roughly 2), the power analysis indicates that an experiment would require between 900 and 1400 subjects to support the hypothesis that Redirection does not produce an SSQ effect of this size! If I assume an effect size of 5 SSQ points, the experiment would require 270 subjects. Experiments of this size are not feasible in our laboratory. But, because Redirection appears to cause less simulator sickness than using a hand-controller, *a fortiori*, Redirection does not unacceptably increase the level of simulator sickness in the user, who must locomote by some means.

1.11 Descriptions of Experiments

For the thesis work I and my colleagues at UNC and UCL⁷ conducted several informal trials and formal user studies. I list these in

1.11.1 *RWp and RW*

The purpose of this user study was to determine the viability of Redirected Walking with waypoints and spatial audio. I tested the technique on a single group of participants who were instructed to complete a fire-drill task in the virtual scene pictured in Figure 1.4. Observations from the study suggest this technique works: Redirected Walking causes users to change their walking direction without noticing and enables larger VEs while providing the benefits of real walking. The subjects did not know about Redirection, were not familiar with the size of the lab, and were led into the lab blindfolded. Subjects were surprised, after completing the task and removing the headset, to find that the real lab was much smaller than the virtual scene.

Table 1.2. Informal trials are denoted with a lowercase i.

⁷ The Effective Virtual Environments team in the Computer Science Department of The University of North Carolina at Chapel Hill and the Virtual Environments Laboratory in University College London.

1.11.2 RWp and RW

The purpose of this user study was to determine the viability of Redirected Walking with waypoints and spatial audio. I tested the technique on a single group of participants who were instructed to complete a fire-drill task in the virtual scene pictured in Figure 1.4. Observations from the study suggest this technique works: Redirected Walking causes users to change their walking direction without noticing and enables larger VEs while providing the benefits of real walking. The subjects did not know about Redirection, were not familiar with the size of the lab, and were led into the lab blindfolded. Subjects were surprised, after completing the task and removing the headset, to find that the real lab was much smaller than the virtual scene.

Table 1.2 - List of experiments and their abbreviations. Experiments i1-i6 were informal and performed on the experimenters themselves.

Abbreviation	Experiment title	Chapter
i1	Constant rotation while standing still and walking and looking down	5
i2	Rotation rate proportional to walking velocity	7
i3	Scaling of walking speed	11
i4	Steering toward center of lab	6
i5	Constant rotation while walking back and forth along a virtual line	7
RWp	Tuning of Redirection with waypoints – pilot	7
RW	Redirection with waypoints and spatial audio	7
RWP-I	Redirection with walking-in-place in a CAVE I	8
i6	Translational Redirection in a CAVE	11
RWP-IIp	Redirection with walking-in-place in a CAVE II – pilot	8
RWP-II	Redirection with walking-in-place in a CAVE II	8
RDT-scw	Redirection rate detection thresholds – while Standing with Constant Velocity scene rotation using method of adjustment	9
RDT-ssv	Redirection rate detection thresholds – while Standing with Constant-frequency Sinusoidal scene rotation using method of adjustment	9
RDT-wcv	Redirection rate detection thresholds – while Walking with Constant-Velocity scene rotation using staircase and constant stimulus	9

1.11.3 RWP

The purpose of the RWP experiments was to test Redirected Walking-In-Place. Participants carried out a task in the same virtual scene as the RW experiment. This time the task required them to freely explore instead of visiting specific places in order. The results of these user studies show that RWP is viable: Users, in a three-sided CAVE, using RWP, can freely explore the VE, do not notice the rotations, suffer less simulator sickness, and see the missing back wall of the CAVE less often than users who use a hand controller to turn in the virtual scene. One participant even reported thinking he was in a fully enclosed, four-sided CAVE. We did

not find any effects of RWP on the user's sense of presence. I suspect this is because our ability to measure presence was not statistically powerful enough. However, the data from these studies suggest a model of presence that indicates that users who more rarely see the open back wall of the CAVE feel more present.

1.11.4 RDT

The RW and RWP experiments verified that specific algorithms and turning-rate functions are effective for those users who do not know about Redirection. In order to use Redirection in real applications (instead of in lab experiments), VE practitioners would benefit from knowing the likelihood that any given rate of injected rotation will be noticed by experienced VE users (who would not be naïve about Redirection after using it many times). This information, for example, would bound the amount of tracked space required to have the user walk in a full circle while thinking she was walking in a straight line, which in turn would allow for real walking in infinitely large virtual scenes.

The Redirection Detection Threshold (RDT) series of experiments was an attempt to find conservative estimates of rotational detection thresholds. From previous research and my own observations, we expect a number of factors (unrelated to the rotation rate) to make Redirection less detectable. These include spatial audio, naïve subjects, engaging tasks (users are distracted), having the user walk at a consistent rate for several minutes, tasks that encourage the users to turn their heads and change direction often, and virtual objects that are farther away from the user. The RDT experiments were conservative in that they did not use spatial audio, told participants that the scene would rotate and that they should watch for it, forbade users from looking around or changing direction, and used a virtual scene that had virtual objects very close to the subjects. *Under these worst-case conditions, the average threshold level of rotation appears to be 1 degree per second.* This is explained in greater detail below.

1.12 Noticing Redirection

Redirection aims to rotate the virtual scene so that the user compensates by turning herself, without noticing the rotation.

1.12.1 Informal Assessment

In the first study (RW), in which subjects followed a zigzag path in the virtual scene, none of the 11 subjects (who experienced Redirection using the final version of the algorithm) seemed to notice the rotations.

These subjects were unfamiliar with the size of the lab (they entered the room while walking backward and donned the headset in the dark). Upon removing the headset, all subjects were surprised at the size of the lab, and were surprised to learn that they had been walking back and forth between the ends of the lab as they were zigzagging through the virtual scene.

1.12.2 Operational Definition of Notice

The two RWP studies attempted to investigate, in a quantitative manner, whether subjects notice the rotation. Subjects filled out a questionnaire which included a question to indicate which of several phenomena (such as the room flickering or rotating) each subject noticed while in the virtual room (Figure 1.13). Subjects did not report that the room rotated any more than they reported any of the other listed phenomena. For example, subjects were just as likely to report that the room changed size as they were to report that the room rotated. Furthermore, those subjects who used Redirection did not report having experienced rotation any more than those subjects who used the hand-controller to turn (the RWP experiments were *between-groups* studies). If *notice* is defined operationally, users do not appear to notice the rotations from Redirection.

<i>These things happened during my experience</i>	Circle Yes or No
1. the brick room became larger or smaller	YES / NO
2. objects disappeared and reappeared	YES / NO
3. parts of the brick room got brighter or dimmer	YES / NO
4. the brick room rotated	YES / NO
5. parts of the brick room flickered	YES / NO

Figure 1.13 - The portion of the RWP questionnaire to gauge the extent to which subjects noticed the room rotation, compared to other phenomena which did not actually occur.

1.12.3 Experienced Users and the Lower Bound of the Detection Threshold of Rotation

As far as I have seen, naïve users, who do not know about Redirection, do not appear to notice the rotations. But if Redirection is used in practice, then users will use the VE system several times and may learn about its operation. (I assume regular users of the VE system will eventually see the size of the tracked space.)

A separate question worth asking is: if a user is aware of Redirection *a priori*, and is looking for its effects, does it still work? Any claims I make about how much rotation a naïve user will notice say nothing about an experienced user.

Experiment RDT aimed to answer the question of how much rotation an experienced user will notice. Literature suggests that detection thresholds are higher (users notice less) when they are engaged in a task

[Rolfe 1986]. Furthermore, having a spatialized auditory virtual scene that is aligned with the virtual scene will also cause users to notice the rotation less [Lackner 1977a]. Finally, when virtual objects are closer to the user, I speculate that the rotation is more noticeable because of the errors in the motions, caused by modeling and tracking errors, of those objects that are close to the user. How much rotation can the VE system apply when users are not engaged in a task or not hearing the spatial audio? For example, how much will an experienced user notice when she is uneventfully walking through a quiet virtual hallway?

I measured, for eight subjects, how likely they were to correctly identify the direction in which the virtual scene rotated, as a function of the virtual scene's angular velocity. This was done under the above conditions (no spatial audio, no distracting task, etc.), while the subjects had just started walking after standing still. The results, overall, suggest that 1.0 deg/s is the detection threshold under these worst-case conditions.

1.13 Lab Size Required for Infinite Virtual Scenes

If the user thinks she is walking in a straight line (in the virtual scene) but is actually walking along an arc, given a large enough lab, eventually she will walk in a full circle. This would allow her to walk in an infinitely large virtual environment. How large does this tracked lab space need to be?

Under the worst conditions, such as those measured in experiment RDT (see 1.12.3 above), the turning radius is extremely large (45 meters). However, these detection thresholds are for the initial few meters, when the user first starts walking. As time in the VE progresses, the VE system should be able to slowly and imperceptibly increase the rate of rotation, because the podokinetic and vestibular systems act as high-pass filters (although I have not tested this). There are several examples in the literature wherein subjects would, after being properly adapted, unknowingly turn themselves at up to 45 deg/s in a circle with a diameter of less than 2 meters [Gordon 1995; Weber 1998; Jürgens 1999]. The data from these studies suggest that this effect exponentially approaches an asymptote (or *charges*⁸) at a rate described by a time-constant of between 6 and 12

⁸ *Charge* is a term from electrical engineering—a capacitor *charges* and *discharges* exponentially.

minutes. If the user begins walking along a straight virtual track, the VE system could, in the worst case, turn her imperceptibly at 1.0 deg/s. But after 10 minutes of walking, that rate would be roughly between 13 and 18 deg/s (assuming a 12- or 6-minute time-constant). Figure 1.14 illustrates the path required for an infinitely long virtual walk. Under the worst-case conditions (no spatial audio, user does not turn her head, etc.), and using the threshold angular velocity from the RDT study, the amount of lab space required is 30 meters by 30 meters. Under more favorable conditions, I expect the lab space requirement to be smaller. Finally, a VE system designer may find it acceptable for some users to notice the rotation injected by Redirection, in exchange for requiring even less space. It is conceivable (but untested) that noticeable rotation is still better than the unnaturalness of using some other locomotion technique for large virtual scenes.




Figure 1.14 - Simulated paths a user walking an infinitely long straight line in the virtual scene under worst-case conditions. The blue circle is for a user walking at 0.75 meters/s and imperceptibly turning at constant rate of 1.0 deg/s with a Redirection algorithm that does not take advantage of PKAR. The black spiral is for the same walking velocity and an initial turning rate of 1.0 deg/s, but with a Redirection algorithm that assumes a PKAR charging time-constant of six minutes. The lab space required when taking advantage of PKAR is roughly 30 meters by 30 meters.

1.14 Steering Algorithms for Unrestricted Exploration of Arbitrary Virtual Scenes

When a large enough tracked area is available, I expect that Redirected Walking will allow the user unrestricted exploration (without the use of waypoints) of an arbitrarily large virtual scene. For this to happen, the system must steer the user in the lab, to keep her from colliding into the lab walls, without knowing her intended path in the virtual scene. In Chapter 6, I discuss three different algorithms for doing this: steering the user to the lab center (Steer-to-Center), steering the user onto a circular orbit around the lab center (Steer-onto-

Orbit), and steering the user toward alternating targets in the lab (Steer-to-Alternating-Targets) (Figure 1.15). Informal testing revealed that Steer-to-Center is problematic in that the user often walks through the lab center and is then headed directly away from it, and it is then difficult to turn her back toward the center. I propose Steer-onto-Orbit and Steer-to-Alternating-Targets to remedy this problem, but have not tested them. The algorithm used must also be able to handle situations where the user takes unexpected turns (e.g., away from the path toward the lab center) (Figure 1.16).




Figure 1.15 - Illustrations of sample paths of a user from three different steering algorithms. The user's location and orientation are represented a black arrowhead and the path as a curved pink line with an arrowhead. Left: Steer-to-center. Center: Steer-onto-Orbit. Right: Steer-to-Alternating-Targets.




Figure 1.16 - How the Steer-to-Center algorithm handles unexpected changes in the user's path. Three hypothetical sample paths that the user could take in the virtual scene (right, in blue) and in the lab (left, in red). If the user walks straight in the virtual scene (path 1), she is steered along a smooth path (in the lab) through the lab center. If the user decides to take a 90-degree right or left turn in the virtual scene (paths 2 and 3), her 90-degree turn becomes something like a 45-degree turn in the lab. After the turn, the user is again redirected toward the lab center.

1.15 Conclusions

The idea for Redirection came from a lesson in my childhood wilderness survival course: people without a heading reference (e.g., when lost in the woods) tend to unknowingly walk in large circles. When blindfolded, people tend to walk in tighter arcs.⁹ I wondered: if the visual cues are intentionally deceptive can I cause someone to turn even faster while still not realizing it? And could this be used to allow real walking in larger virtual scenes? Informal experimentation with an early implementation was promising. I further observed that some amount of imperceptible rotation was possible while the user was standing still, though less than when the person was walking. More imperceptible rotation was possible when the user was marching on the spot, but even more when the user was actively turning. What were the mechanisms behind this illusion, and how could a VE practitioner take advantage of this? This led me to pursue the techniques used in flight simulation, and that in turn led me to vestibular motion perception.

Redirection has promise. It allows for real walking in larger virtual scenes. It is input-natural, works without the user noticing it, and results in less simulator sickness than using a hand-controller. In addition to the experimental findings, my observations attest to the potential of Redirection.

On one occasion, in the course of demonstrating the system, a visitor picked up the headset and held it in his hands while trying to observe the virtual scene rotating (by watching a copy of the HMD's video stream on a large wall-mounted projector screen). The scene did not appear to rotate, causing me to wonder if the Redirection was even turned on. It turns out the visitor was actually and unknowingly turning the HMD with his hands (presumably using the images on the projection screen to stabilize his hands)! Once the headset was placed on a chair, the scene rotation was readily observable.

In the course of developing the VE system used in the experiment RW, a colleague and I were surprised at how strong the visual illusion could be. In some cases, a user was about to walk into a lab wall

⁹ This was demonstrated by blindfolding the students and having us walk from one goalpost to the other on a football field—none of us reached it. Each student turned in an arc of differing curvature. Several students turned so sharply that they walked off the field before even reaching the zero-yard line!

because of mistuning that made the open space in the virtual scene appear in front of her. When this happened, we would instruct her to stop walking, then look left and right, and then continue walking in the original direction. During this short time, the virtual scene and the subject rotated 90 degrees in the lab and then walked away from the wall, while she thought she continued to walk in the same direction.

In another situation, we left the tracked hand-controller (normally carried in the user's hand) on top of a tripod while we were testing the drawing of the virtual scene. While the user is wearing the headset, the tracked hand-controller appears as a virtual hand. To our dismay, we saw the virtual hand was slowly moving along a large arc even though the real controller was still. We suspected the tracking system had drift or noise and spent hours trying to diagnose the problem. Then we realized that we had not turned off the virtual scene rotations—the whole scene was turning but we could not tell. Thus it appeared, even to us developing the system, that the virtual hand was moving! This was very encouraging. In informal tests, users who knew about Redirection were still not able to detect it while in the VE.

The observations from this work lead to guidelines for VE practitioners wishing to make use of Redirection; I offer them in Chapter 12

Chapter 2: Locomotion Interfaces

There are many ways to have the user specify how her viewpoint moves in virtual scenes. In this chapter, I compare several of them. To aid this comparison, I first consider my goal for the user interface—locomotion.

2.1 Locomotion

Locomotion is a special type (or subclass) of movement. *Locomotion*, as used by life scientists, refers to the act of an organism moving itself from one place to another. This includes actions such as walking, flying, and swimming. For humans, locomotion is walking, running, crawling, jumping, swimming, etc. Locomotion is not movement of one's limbs or head, nor is it swaying on the spot.

2.2 Locomotion in Virtual Scenes

Locomotion in virtual scenes is comparable to real-world locomotion. The user must be able to go from a virtual desk to a virtual bookshelf. Other authors have referred to this kind of virtual movement as *travel* [Bowman 1997], but I prefer *locomotion* because it is used in life sciences and because *travel* implies moving a significant distance—in everyday language, one travels from New York to Chicago, not from the desk to the bookshelf. *Locomotion* is distinct from *wayfinding*, which refers to the cognitive task of determining a route for reaching a destination. I avoid the term *navigation* because it is used in the literature to refer to both wayfinding and locomotion.

Some authors have a different interpretation for the term *locomotion*. Hollerback refers to a *locomotion interface* as any virtual movement technique that requires the user to make repetitive motions. Using his definition, holding a joystick in the forward position (thus keeping the virtual viewpoint moving forward) is not locomotion, while repeatedly pushing a button is. I do not use this definition.

In this chapter, I discuss the user interface techniques for virtual locomotion, or *locomotion techniques* for short. Locomotion techniques differ in many ways. For example, some are designed for non-immersive

desktop 3D graphics applications, using a mouse or keyboard, while some apply only to fully immersive VEs. Some allow the user to control the viewpoint continuously (with six degrees of freedom), while others have participants move by choosing viewpoints from a menu (e.g., VRML's bookmarks). I only consider locomotion techniques which:

- 1) apply to fully immersive VEs (e.g., CAVES and HMDs);
- 2) have the user control the viewpoint continuously;
- 3) apply to human-scale virtual scenes, where the goal is to simulate a person walking and without a vehicle. The simulation of vehicle motion is beyond the scope of this thesis.

2.3 Locomotion Techniques

2.3.1 *Flying*

The most common locomotion technique in VE systems is *flying* using a joystick or some other hand-controller. When the user pushes a joystick or presses a button, she moves forward in the virtual scene. She can still move about locally by leaning or taking a real step in any direction (if her head is tracked by the VE system). The effect is similar to that of walking about on a moving flat-bed truck or flying carpet [Robinett 1992]. When the user presses the button, the truck moves forward in the virtual scene. The user can simultaneously move about on the truck bed.

There are significant variations in how flying is implemented—*forward* has many interpretations. Some VE systems have the user move in the direction she is looking (gaze-directed). Others move the user in the direction she is pointing with her hand-controller. Still others interpret forward using a vehicle metaphor—forward is toward the front wall of the CAVE.

2.3.2 *Leaning*

Similar to flying, leaning techniques move the user in the virtual scene in the direction in which she is leaning [Peterson 1998; LaViola Jr. 2001]. Most implementations also control the rate of travel—the farther the user leans, the faster she moves. Leaning has the advantage of not requiring a hand-controller.

2.3.3 *Treadmills*

There are a number of techniques that simulate the physical act of walking. Several groups have experimented with treadmills [Brooks 1992; Hollerbach 2000]. As the user walks forward on the treadmill, she moves forward in the virtual scene. Motorized treadmills raise safety concerns, whereas, due to the friction,

passive ones require effort beyond walking in the real world. One interesting variation is the Army's use of a stair-stepper instead of a flat treadmill [Lorenzo 1995]. All of these have the limitation that the treadmill has a preferred orientation—it is difficult, disorienting, and often impossible to turn on the spot in the virtual scene. The UNC treadmill, for example, had handlebars to steer like a bicycle.

To allow turning on the spot, several groups have developed two-dimensional treadmills [Darken 1997; Iwata 1999], where the user can walk in any direction on the ground plane. Existing implementations are mechanically complex and noisy and have a small-area walking surface and limited speeds.

2.3.4 *Walking-in-Place*

When using walking-in-place, the user makes walking motions (lifting the legs) but stays on the same spot physically. The VE system detects this motion and moves her forward in the virtual scene [Slater 1995; Usoh 1999; Templeman 1999]. Like flying, it does not require a large lab or tracking space.

2.3.5 *Real Walking*

If the virtual scene is the same size as or smaller than the tracked space, then *real walking* is feasible. Here the user's movement in the virtual scene corresponds exactly to her movement in the lab. If she walks five meters in the lab, she also walks five meters in the virtual scene.

2.3.6 *Manipulating the World*

In the other extreme, there are locomotion techniques that are nothing like walking in the real world. Multigen's SmartScene® product¹⁰ and Miné [Miné 1997] have demonstrated techniques where the user can grab the virtual scene and move it toward him. By repeatedly grabbing points in the virtual scene and pulling them in, the user can locomote from one place to another. Even less like the real world, Stoakley's Worlds-in-Miniature technique has the user manipulate a hand-held, doll-house-sized model of the virtual scene. The user moves in the virtual scene by moving a doll (representing herself) to the desired location in the miniature virtual

¹⁰ As of 2004, this is owned by Digital ArtForms™.

scene. Then the miniature virtual scene grows to become human-scale and the user finds herself in the desired location in the virtual scene [Stoakley 1995].

2.4 The Difficulties of Comparison

Why are there so many varied locomotion techniques? Each must have its own advantages and shortcomings. However, it is very difficult to test them and make quantitative comparisons among them. First, many of the techniques require specialized and custom hardware, which is available only to those who developed it and only for a limited time until the equipment and lab space are reallocated to newer research. Many times, researchers present the technique without making quantitative comparisons to any other techniques. Even when the researchers do conduct user studies to compare techniques, the results are often not widely comparable. Each study has a small and different user population (i.e., expert vs. novice VE users, university students vs. military pilots vs. architects), and the task in each user test is different. Furthermore, the non-locomotion parameters of each VE system, such as frame-rate, resolution, field-of-view, latency, tracking noise, and visual fidelity are different and, in many cases, not measured.

Finally, there is trouble comparing a set of techniques because different studies use different evaluation metrics. Some investigate the time to complete a task [Bowman 1997], others examine presence [Usoh 1999], while yet others examine spatial awareness or orientation [Peterson 1998].

One particularly impressive comparison of locomotion techniques is a study by Bowman, Koller, and Hodges (Figure 2.1) [Bowman 1997]. They evaluated several variations of *flying* by re-implementing them to run on the same hardware and having each subject user perform the same task using all the flying techniques. Bowman evaluated each technique using several criteria: ease of learning, spatial awareness, speed, accuracy, and cognitive load. Instead of claiming that one technique is better than all the rest, they defined several metrics and pointed out that different applications have different needs. A phobia-treatment application may be more concerned with naturalness and presence, whereas a game may be more concerned with getting the user to the target quickly. The application designer should first decide which attributes are important, then choose a locomotion technique that optimizes the particular attributes important for that application.

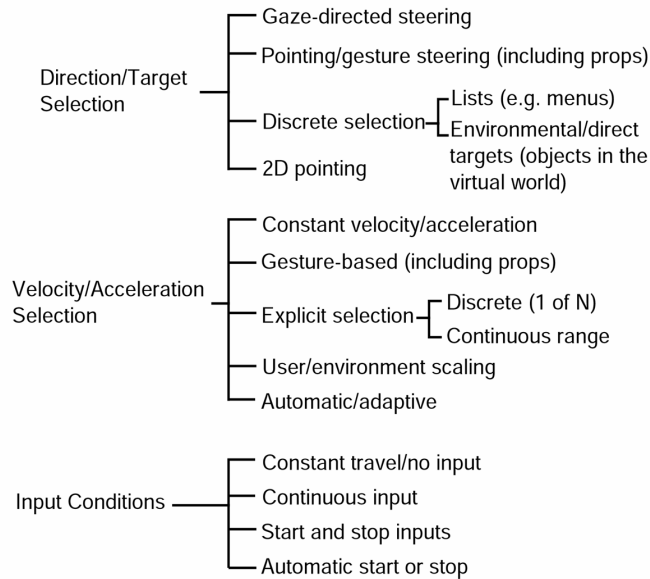


Figure 2.1 - Bowman's taxonomy of flying locomotion techniques [from Bowman 1997].

2.5 Attributes Relevant to This Thesis

The goal of this work is to simulate real-world walking in the virtual scene. As such, I am not concerned with reducing the time the user requires to get from point A to point B. To achieve the goals of this thesis work, if it takes a person 30 minutes to walk two miles in the real world, she would also take 30 minutes to move the same distance in the virtual scene (and get just as tired doing so). The qualities I am concerned with are naturalness, ease of learning, presence, portrayal of motion cues, and incidence of simulator sickness.

2.5.1 *Input-Motion-Naturalness*

Naturalness, or the lack thereof, is frequently used in everyday language to describe user interfaces. However, this quality is hard to define and measure. For the purpose of this thesis, I define *input-motion-naturalness* as the similarity of the user's physical motions, while using the interface, to the physical motions she would make doing the real-world task. For example, inputting text by writing on a tablet is more *input-motion-natural* than typing (even though typing is often faster). *Real walking* is more *input-motion-natural* than manipulating the world.

2.5.2 *Ease of Learning & Ease of Use*

Some locomotion techniques are easier to learn and/or easier to use than others. This attribute is not always identical to naturalness (or *input-motion-naturalness*). For example, walking-in-place is more natural

then flying with a joystick but, because of implementation problems or users' familiarity with joysticks, some users find it harder to walk-in-place [Whitton 2005].

2.5.3 Motion Cues

While walking about the real world and while flying a real aircraft, a person experiences many motion cues. All modern flight simulators and VE systems portray visual cues. Optical flow by itself can induce a sensation of movement [Regan 1986]. Beyond the visual cues of motion, however, there are several others [Howard 1986a; Cheung 2000], which are introduced in Chapter 1, and detailed in Chapter 4:

- 1) vestibular – inertial forces sensed in the inner ear;
- 2) proprioceptive – sense of movement in joints, muscles, and viscera;
- 3) cutaneous/tactile – sense of pressure on the skin generated by inertial forces, and by contact with objects, including the feet touching the ground;
- 4) auditory – audible spatial cues including the acoustics of the scene and wind noise.

2.5.4 Simulator Sickness

Simulator sickness is discussed in detail in Chapters 3 and 10. In brief, one popular theory of simulator sickness and motion sickness, called the *cue conflict theory*, is that it is caused by conflicting visual and vestibular motion cues. For example, if a person is watching a wide-screen movie where the camera is moving, the visual cues indicate she is moving but her vestibular cues tell her she is still. This conflict can cause motion sickness.

2.6 Comparison of VE Locomotion Techniques in Terms of Attributes Relevant to This Thesis

2.6.1 Flying

All forms of VE flying locomotion present visual cues of motion. Some also provide acoustic and wind cues. None provide proprioceptive or vestibular cues.

2.6.2 Treadmills

An ideal treadmill works by perfectly canceling the user's physical motion—as she moves forward, the treadmill rolls backward to keep her centered on the treadmill's walking surface. This has been compared to running on slippery ice [Darken 1997] and does not induce the proper vestibular motion cues—as the user accelerates forward in the virtual scene, she does not receive vestibular cues of moving forward [Templeman

1999]. She does receive, however, the vestibular cues from her head bobbing and her feet striking the ground. In addition, treadmills have the advantage (over flying) of simulating the proprioceptive cues of walking.¹¹ The stair-stepper technique and variable-incline treadmills such as the University of Utah's Treadport [Hollerbach 2000] have a significant advantage in that they can convey hilly terrain and the changes in physical effort associated with such terrain.

2.6.3 Real Walking

In real walking the vestibular and proprioceptive cues are perfect. The visual cues lag behind the other cues because the VE system must register the motion and respond with updated imagery, and this takes time (35 to 150 ms¹²). From the cue conflict theory, one would expect real walking to induce less simulator sickness than walking on a treadmill or flying. Real walking has been shown to be perceived as more natural and to result in a greater sense of presence than flying and there is evidence that it compares similarly against walking-in-place [Usoh 1999].

2.6.4 Redirected Walking

Real walking is the best technique when simulator sickness, motion cueing, and presence [Usoh 1999] are the characteristics of concern. One significant limitation of real walking, however, is that the tracked space must be as large as the virtual scene. It cannot represent entire buildings in a room-size lab. This is the problem Redirection is designed to address.

Though Redirected Walking reduces the space required for real walking in large virtual scenes, it still requires a large tracked area (though not as large as the virtual scene). The other techniques (besides real walking and Redirected Walking) are commonly implemented in very small tracked areas—often 1 to 2 squared meters.

¹¹ Even though the stride length might not be the same as in the real world.

¹² The end-to-end latency of VE systems I have measured, including the response time of the HMD.

Chapter 3: Simulator Sickness

The most frequent criticism of the concept of Redirected Walking is that, by spinning the virtual scene, Redirection will induce simulator sickness. In this chapter, I describe simulator sickness, factors that contribute to it, theories behind its causes, and some methods of measuring it.

Some users report some sickness symptoms after exposure to virtual environments: dry mouth, nausea, dizziness, visual aftereffects (flashbacks), pallor, sweating, ataxia (loss of balance), and even vomiting. VE practitioners commonly refer to this phenomenon as *simulator sickness* or *cybersickness*. The occurrence of these symptoms varies wildly from person to person and among VE systems. These symptoms are a critical problem for the use of VEs. Twenty to forty percent of military pilots suffer from simulator sickness, depending on the simulator [Kolasinski 1995].

Several other sicknesses have symptoms similar to those of simulator sickness. These include motion sickness, space sickness, and certain kinds of poisoning. Some argue that simulator sickness should be defined as only the sickness that results because the simulator does not perfectly simulate the real-world situation [Pausch 1993]. For example, if a passenger suffers motion sickness in some real airplane, then in a perfect simulator, she would suffer the exact same motion sickness. In this view, the simulator sickness is only the additional symptoms and the severity of those symptoms that she would suffer in a less-than-perfect simulator. Kennedy argues that the standard motion-sickness diagnostics are less relevant to sickness resulting from flight simulator exposure than they are to true motion sickness. For example, vomiting is an indicator of motion sickness, but does not occur regularly enough in simulator sickness to be a statistically useful indicator of it.

Kennedy provides an operational measurement of simulator sickness (the SSQ) independent of motion sickness [Kennedy 1993].¹³

3.1 Consequences of Simulator Sickness

Beyond preventing people from using VEs, simulator sickness has other repercussions. Sometimes symptoms linger for days after the exposure and can affect motor control and coordination [Draper 1998]. If the user operates machinery, drives, or pilots an aircraft after the VE exposure, the simulator sickness can affect performance in these real-world situations and put people in danger. In one case, a pilot, hours after a VE exposure, saw the real world invert while driving his car [LaViola Jr. 2000]! In fact, Kennedy et al. have proposed quantifying the severity of simulator sickness as a mapping (using measures of a pilot's ability to keep balance) to blood alcohol level, the legal metric of alcohol intoxication used to determine if a person is fit to drive [Kennedy 1995; Cobb 1998]. Under this proposal, a user, after some VE exposure, might be considered as unfit to drive as someone who has an illegal blood-alcohol level. In some cases, U.S. Marine and Navy pilots are restricted from flying for 12 to 24 hours after using a simulator [Kennedy 1992]. While this alleviates the risk to the pilot, it is very expensive in terms of the pilot's time and reduces the cost-effectiveness of the simulator.

Finally, while training in a simulator, users might adapt their behaviors to avoid becoming sick. For example, pilots in simulators may avoid looking at the outside (virtual) scenery and focus on just the cockpit instruments instead [Kennedy 1992]. Since these new behaviors are not appropriate while flying a real aircraft, simulator sickness can cause miss-training.

¹³ This operational definition of simulator sickness does not necessarily correspond to Pausch's definition of being the additional sickness, beyond true motion sickness, caused by the simulator.

3.2 Difficulties in Understanding Simulator Sickness

Users report symptoms. From these self-reports researchers and practitioners are aware there is a problem. But the study of this sickness is not straightforward. Not all users who experience the same VE get sick. Some users who get sick during one VE exposure do not get sick on the next exposure. Not all users who do report sickness report the same symptoms. The symptoms can take minutes to hours to appear, and many of the symptoms are subjective and not externally observable. To further complicate its study, there are many proposed causes for and aggravating factors of simulator sickness [Kolasinski 1995; Draper 1996].

3.3 Factors That Aggravate Simulator Sickness

Before delving into what causes simulator sickness, it is useful to consider how characteristics of the user and of the system affect the level of sickness suffered by users. The factors can be divided into those of the individual user and those of the VE system or simulator.

Table 3.1 lists some known user characteristics, in no particular order.

Table 3.1. Factors that correlate with decreased susceptibility (in users) to simulator sickness

Being in good health (illness, hangovers, stress, fatigue, and medications increase sickness)
Age
VE experience
Spatial reasoning / ability to mentally rotate 3D shapes
Male gender

Kolasinski reports that females report higher symptoms, but she thinks this might be because males tend to under-report vulnerability [Kolasinski 1995].

Table 3.2 lists the characteristics of VE systems, divided into two categories. The characteristics in the first column are technical shortcomings of the equipment. As technologies improve, one expects these to be reduced and thus decrease the resulting simulator sickness. The characteristics in the second column are qualities that are often desired in VE systems. Longer exposures are required to simulate long missions or to allow VE users to carry out meaningful tasks. Stereoscopic displays can improve task performance [Pang 2000]. But these qualities may increase sickness as a result of the shortcomings in the first column. For example, if the system lag is reduced or other parameters are tuned, users are sometimes able to stay in the

simulator longer [Strachan 2001]. Stereoscopic displays are more likely to induce simulator sickness,¹⁴ and higher field-of-view (FOV) displays result in better training and higher levels of user immersion [Arthur 2000].

Table 3.2 - *Qualities of VE systems and flight simulators that increase simulator sickness.*

<i>Equipment Shortcomings</i>	<i>Desirable Functions of System</i>
Tracker inaccuracies (temporal and spatial)	Stereo display
Low update rate	Long exposure duration
High latency/lag	Wide FOV display
Mismatches between display parameters and image generator's parameters (i.e., incorrect FOV setting)	Free head movement
Display flicker	Viewpoint motion controlled by someone other than the viewer
Headset weight	High rates of simulated linear or rotational acceleration

3.4 Theories of the Mechanisms of Simulator Sickness

There are four major theories of the causes of simulator sickness: cue conflict, postural instability, poison, and the rest-frame hypothesis. These are strongly related to the concepts of self-motion perception discussed in Chapter 4. The theories discussed here are summarized from Draper and LaViola, except where noted [Draper 1998; LaViola Jr. 2000].

3.4.1 Cue Conflict

Situations such as riding in vehicles, using VEs, and flying flight simulators can result in cue conflict between (at least) the visual and vestibular senses. For example, in the cabin of a boat on rough seas, the visual system tells the person she is staying still (because the insides of the cabin are not moving with respect to her), while the vestibular system tells her she is rolling with the waves. Similarly, in a VE where the user is moving forward (in the virtual scene) with a joystick, her visual system is telling her she is moving while her vestibular

¹⁴ But that effect might be caused by incorrect stereoscopic parameters. For example, mismatches between the inter-pupillary distance and field-of-view of the user and the corresponding values used by the image generator can increase simulator sickness [Draper 1996].

system is reporting that she is standing still. The cue conflict theory (also known as the *sensory conflict* theory) says that situations like these cause the central nervous system to have problems coordinating and integrating the different cues into a consistent mental model of self-motion, and that this is the cause of the symptoms of sickness. This cue conflict theory explains not only motion sickness and simulator sickness, but also space sickness.

3.4.2 Postural Instability

A problem with the cue conflict theory is that it does not explain why a cue conflict leads to symptoms such as nausea and dizziness. The postural instability theory [Stoffregen 1991] argues that the symptoms are a response to postural instability, not cue conflicts. Humans have mechanisms that allow us to maintain balance and posture. For example, a person standing still is not completely motionless. She is wobbling and wavering back and forth (think of an upside-down pendulum). As she drifts in one direction, her mechanisms sense this and apply the appropriate muscle controls to correct it. These balance mechanisms are adaptive. But if the mechanisms are in a state where they do not apply the correct control, postural instability results. For example, imagine a user in a VE. If the visual scene shows that she is accelerating forward, she might lean forward to compensate. But since she is really standing still, this action would make her less stable instead of improving stability. The postural instability theory says this inability to maintain balance is necessary for motion and simulator sickness and that it precedes the other symptoms of sickness. This theory explains why some users suffer from simulator sickness while others do not. It also explains how some users can adapt to VEs (such that they do not get as sick in subsequent exposures) and how sailors can eventually acquire their *sea legs* and become less prone to sea sickness. As the balance mechanism learns how to control posture and balance in a new situation, there is less postural instability in that situation.

The postural instability theory was originally proposed to refute the cue conflict theory. There are cases when cue conflict does not cause sickness and the cue conflict theory has no explanation why. But one could partially reconcile the cue conflict and postural instability theories—in the situations where cue conflict does exist, this conflict could result in postural instability, which causes the other symptoms.

The cue conflict and postural instability theories provide an explanation for the vertigo some people experience at heights. When a person is looking at the floor, the motion of the floor in her visual field is one source of postural feedback. But when the person is standing at a great height and looking down, the movement

of the ground does not correspond to her motion in the way her body is expecting it (a 1 cm movement of the head results in much less movement of the image of the ground in the eye if the ground is 20 meters away instead of 2). In the context of the cue conflict theory, this inconsistency between the vestibular cues, the visual cues, and the body's expectations results in cue conflict. In the postural instability context, the lack of the expected visual feedback degrades the person's postural stability and this results in vertigo.

3.4.3 *Poison*

Another theory of motion, simulator, and space sickness is the poison theory. This theory gives an evolutionary explanation for why cue conflict or postural instability causes symptoms such as nausea and vomiting.

Ingesting certain neurotoxins can affect the coordination of the different senses and also affect motor control. The body then responds to this poisoning with the nausea, vomiting, and fatigue. This response provided an evolutionary benefit because it served to expel the offending poison and discouraged the person or animal from moving about until the motor control and senses returned to normal. Over the course of human evolution the cause of such degradation in coordination was very likely to be poison, and so the sickness response was appropriate and advantageous. The poison theory suggests that VEs, space travel, and modern transport in planes, boats, and cars cause neurological changes similar to poisoning, and that the nausea, vomiting, and other symptoms are an inappropriate response by the body.

This theory can also be reconciled with the others. Poisoning could affect the balance mechanisms and thus cause postural instability. The motion and simulator sickness that results from postural instability, then, is the response by the body which has been fooled into "thinking" it was poisoned.

The above three explanations of simulator sickness are physiological. The next is different in that it is a mental explanation.

3.4.4 *Rest-frames and the Internal Mental Motion Model*

The rest-frame hypothesis says the brain has an internal mental model of which objects (in the world) are stationary and which are moving [Prothero 1995]. The model is formed from the sensory cues seen so far. The objects which one's brain thinks of as stationary are termed *rest-frames*. In the real world the visual background is a rest-frame—humans assume the earth and sky are not moving. To perceive motion, one must first decide which objects are stationary. Then the remaining objects' motions (and even one's self-motion) are

perceived relative to those rest-frames. Once one has a motion model, new sensory information is interpreted with the model (in other words, one perceives what one expects to perceive). This is analogous to a steady-state Kalman filter, and in fact, the processing of the sensory signals by the central nervous system has been modeled as an optimal estimator [Rolfe 1986; Berthoz 2000].

When the brain receives new sensory cues that invalidate the current mental model, simulator sickness can occur. For example, when a person is sitting in a stopped train and the train on the next track starts to move, the person is briefly disoriented because she is not sure which train just started to move. In other words, while both trains are stopped, both trains are rest-frames. When one starts moving, the brain has to re-evaluate that mental model.

This rest-frame hypothesis implies that simulator sickness is inextricably tied to one's internal mental model of self-motion. Unlike the cue conflict theory, this hypothesis claims that sickness does not result from conflict between sensory cues, but results when the sensory cues conflict with the brain's motion model.

Gregory and Berthoz claim that illusions are when the brain has a mental model that consistently explains all the sensory cues it has received, but that model is incorrect [Gregory 1966; Berthoz 2000]. Berthoz believes that the inability of the brain to come up with a consistent model to explain the sensory cues can lead to panic attacks and dizziness [Berthoz 2000]. This fits in with the postural stability theory—the lack of a satisfactory cognitive model of what is moving and what is not can degrade a person's ability to maintain balance.

The rest-frame hypothesis has important implications for VEs. Many VEs create an illusion of self-motion in the user by visual cues (this is known as *vection*) without manipulating vestibular cues. For example, a user might *fly* through a virtual building while standing still in a CAVE. For the illusion to succeed, the VE must cause the user to choose the virtual ground to be the rest-frame instead of the real floor of the CAVE. But then the vestibular cues will conflict with this mental model.

There is experimental evidence to support the rest-frame hypothesis. Factors which enhance the sensation ofvection, such as a wider field-of-view (FOV) display, result in a higher sense of presence in the virtual environment. But these very same factors also increase the level of simulator sickness [Prothero 1998; Fleming 2002]. Presenting a grid-like visual pattern, superimposed on the view of the virtual scene, that is fixed relative to the real world has reduced simulator sickness [Prothero 1999; Duh 2001a]. However, the presentation of the fixed pattern and the user's choice of it as the rest-frame, could reduce the sensation of

motion and presence that the VE was designed to induce in the first place! If a user feels more present in the virtual scene than in the real world, she presumably chooses the virtual ground as the rest-frame instead of the real ground. The implication of the rest-frame hypothesis is that a VE can have one or the other, less simulator sickness or less vection, but not both.

Whereas there are several theories of the mechanisms behind simulator sickness, they complement each other. While designing and evaluating VEs, it is useful to keep them all in mind.

3.5 Measuring Simulator Sickness

To investigate simulator sickness, many VE researchers use the tools and theories developed by the flight simulator community. As early as 1960, pilots have been reported to have, after using flight simulators, suffered many of the same motion sickness-like symptoms that VE users report. It is from this community that the term *simulator sickness* comes [Miller 1960].

One could argue that flight simulators are a subset of VEs. But one could also argue that VEs are a subset of simulators (not all of which employ computer graphics). VEs and flight simulators have different historical origins and the practitioners and users of each are from different communities. The first flight simulators were in use by 1910 and did not make use of computer graphics until the 1960s [Rolfe 1986], whereas computer graphics are inseparable from the history of VEs, the first of which was made in 1968 [Robinett 1992]. In this dissertation, I refer to flight simulators and VEs as distinct, non-overlapping, categories.

Is someone suffering from simulator sickness? This is not a yes/no question. The most common method of measuring simulator sickness (both in flight simulators and VEs) is Kennedy's Simulator Sickness Questionnaire (SSQ) [Kennedy 1993].

The SSQ characterizes the level of sickness on three linear scales, based on subjective self-reports of the symptoms. As a person feels sicker, she scores higher on one or more of these scales. The SSQ questionnaire is administered after a VE or simulator exposure and consists of 16 questions (Table 3.3). Users mark the severity of the 16 different symptoms on a scale of zero to three. The SSQ produces four scores, a total (overall) sickness score and three orthogonal subscale scores. The subscales are nausea, oculomotor discomfort, and disorientation.

Table 3.3 - The SSQ questionnaire. Each subject marks the severity of each symptom on a four-point scale [from Kennedy 1993].

1. General discomfort	none	slight	moderate	severe
2. Fatigue	none	slight	moderate	severe
3. Headache	none	slight	moderate	severe
4. Eye Strain	none	slight	moderate	severe
5. Difficulty Focusing	none	slight	moderate	severe
6. Increased Salivation	none	slight	moderate	severe
7. Sweating	none	slight	moderate	severe
8. Nausea	none	slight	moderate	severe
9. Difficulty Concentrating	none	slight	moderate	severe
10. Fullness of Head	none	slight	moderate	severe
11. Blurred Vision	none	slight	moderate	severe
12. Dizzy (with eyes open)	none	slight	moderate	severe
13. Dizzy (with eyes closed)	none	slight	moderate	severe
14. Vertigo	none	slight	moderate	severe
15. Stomach Awareness	none	slight	moderate	severe
16. Burping	none	slight	moderate	severe

Another, less common, measure of simulator sickness is the balance test, or postural stability test.

There are many variations of this test. They all attempt to quantify sickness by measuring how well a user can balance after the VE or simulator exposure [Kolasinski 1994]

Chapter 4: Self-Motion Perception

What are the mechanisms by which humans sense their own motion? What are the limitations of these mechanisms? These questions are central to Redirection, which works by exploiting these mechanisms and limitations. Beyond Redirection, understanding these mechanisms is of benefit to all VE designers, even those who do not use Redirection, because all VEs exploit perceptual limitations. For example, video displays emit only red, green, and blue light, since almost the full range of colors that humans can perceive can be represented by a combination of these. Understanding human motion perception bears on VE design questions such as: if one chooses to have users locomote by flying with a joystick, how could this be implemented so as to minimize simulator sickness and maximize a feeling of motion?

Self-motion perception¹⁵ is intertwined with motion and simulator sickness and is an important consideration when selecting or designing a VE locomotion technique. This chapter will cover self-motion perception. The next chapter will discuss how it applies to Redirection.

Self-motion perception is a popular ongoing area of research. It bears on topics much broader (and arguably more important) than VEs. These include understanding and preventing illusions that cause aircraft pilot errors, motion sickness, disorientation, loss of balance, and even panic attacks [Berthoz 2000].

¹⁵ Some researchers use the term *self-motion perception* to mean the perception of one's *translation only*, and *self-motion cognition* to mean the perception of both one's translation and rotation. I use *self-motion perception* to refer to one's perception of all forms of self-motion (orientation, translation, limb motion, twisting the torso, etc.).

4.1 Difficulties in Studying Self-Motion Perception

The mechanisms of self-motion perception, motion sickness, and balance are not completely understood. Some have asked me why I do not just “look up the numbers” to see if Redirection works. There are different values in the literature, measured under many different conditions. The psychophysical values appear to vary from person to person, with age, with how the person is moving, and whether the motion is active or passive [Howard 1986a; Schweigart 2002]. The models and experiments require assumptions, and these assumptions do not necessarily hold during normal human walking. Furthermore, it is not possible to stimulate and measure each motion-sensing organ separately, as there is crosstalk among the different perceptual systems and even between them and the person’s conscious state (e.g., what task she is attending to; what assumptions she is making about her motion; whether she is able to predict her motion). The exact conditions of Redirection and VEs have not previously been studied to my knowledge. Here I survey the relevant concepts and literature and fit Redirection into them, recognizing that the theories might change as self-motion perception research continues. In my exploration of this topic, I was surprised to find that many researchers in self-motion perception are just as interested in VEs and Redirection-like techniques as tools for investigating self-motion perception as I am in using self-motion perception research for understanding and improving Redirection [Jacobson 2001; Warren Jr. 2001; Jaekl 2002].

4.2 Overview

Humans rely on multiple cues to perceive how they are moving relative to the outside world around them:

- 1) visual;
- 2) vestibular – inertial forces sensed in the inner ear;
- 3) proprioceptive/kinesthetic – forces and tensions in joints, muscles, and motion of viscera;
- 4) cutaneous/tactile – sense of pressure on the skin generated by inertial forces, gravity, and by contact with objects;
- 5) auditory – audible cues such as localized sound sources.

Most VE systems synthesize only visual motion cues. Immersive VE systems create synthetic visual cues to match the user’s real vestibular and proprioceptive cues—as the user turns her head, the computer-generated images are updated to reflect this movement. The real vestibular, real proprioceptive, and virtual

visual cues all consistently convey the fact that she turned her head. However, this is not true when the user is flying, which is the locomotion technique that most VE systems employ. The user stands still while she flies through the virtual scene using a joystick or hand-controller, and the VE system provides only the visual cues of that flying motion. Some VE systems also convey spatial-audio cues of movement and a smaller number have the user walk-in-place on a treadmill, in order to also provide some proprioceptive cues of locomotion.

4.3 The Vestibular Sense

The vestibular sensing organs are part of the inner ear—set in the cave-like *labyrinth* beyond the ear drum (Figure 4.1). They are divided into two main components, the *semicircular canals* (SCCs) and the *otolith organs*. Roughly, the SCCs act as angular rate-gyros and sense rotation of the head, whereas the otolith organs act as linear accelerometers. Under normal conditions (when the person is walking on the earth, not in space or a vehicle), the otolith organs sense linear acceleration and the tilt of the head relative to gravity.

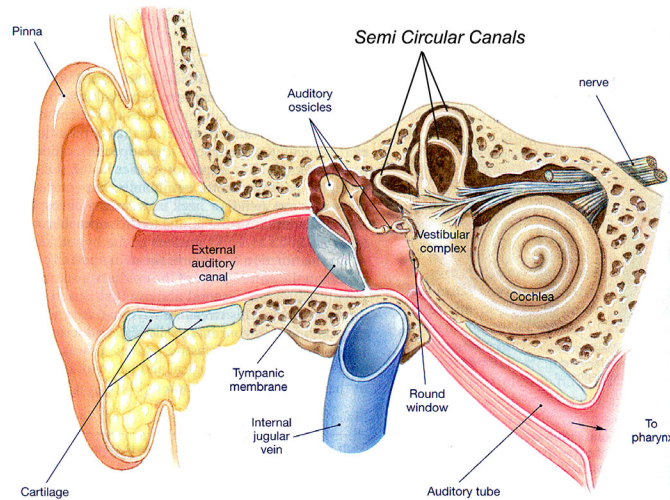


Figure 4.1 - The human inner ear labyrinth [adapted from Martini 1998].

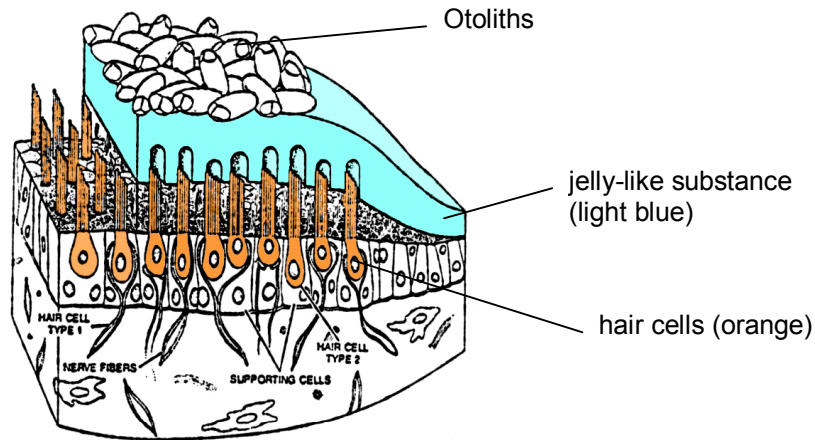


Figure 4.2 - The macula. The otoliths are shown embedded on top of the jelly-like substance. The hair cells, shown in orange, have cilia that extend upward into the jelly-like substance [adapted from Howard 1986b].

On each side of the head, there are two otolith organs, the *utricle* and the *sacculle* (Figure 4.1). Within each of these is the *macula*, which contains crystals of bone-like material—the *otoliths*. Otoliths are embedded in a jelly-like substance (Figure 4.2). As the head is linearly accelerated or tilted, the inertial forces on the otoliths cause the jelly-like substance to deform. There are also hair cells¹⁶ in the otolith organs, the tips of which are embedded in the jelly-like substance and the bases of which are anchored in stiffer supporting tissues (Figure 4.2, Figure 4.3). When the jelly-like substance deforms, the hair cells are also deformed and this deformation is encoded as nerve impulses which are sent to the brain stem. There is a very important fundamental physical limitation of the otolith organs—they cannot distinguish between tilting and acceleration.¹⁷ Furthermore, the nervous system’s interpretation of the signals from the otolith organs seems to rely almost entirely on the direction and not the magnitude of the gravitational/inertial force [Cheung 2000]. Later I describe how this is advantageous for flight simulators but can also lead to pilot disorientation and even crashes in real aircraft.

¹⁶ Despite the name, these are unrelated to the hairs and follicles found in the skin of mammals.

¹⁷ In fact, no inertial sensor can distinguish between gravity and acceleration.

The hair cells (Figure 4.3) are a common theme in the inner ear. In the cochlea, hair cells transduce sound vibrations into nerve impulses; in the SCCs, they transduce head rotation; in the otoliths, they transduce linear acceleration.

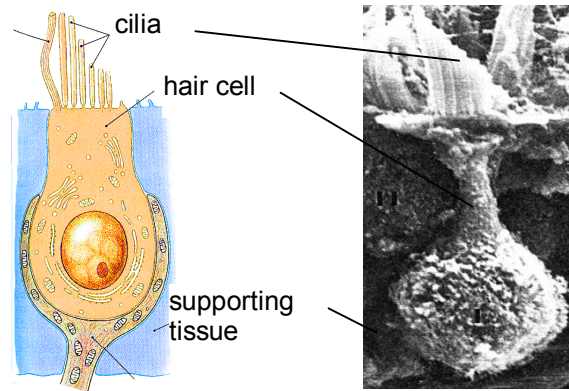


Figure 4.3 - A single hair cell. On the left is a diagram of a hair cell (orange) set in supporting tissue (blue), with the cilia at the top. On the right is an electron micrograph of an actual hair cell. See Figure 4.2 and Figure 4.5 to see the size of the hair cells compared to the otoliths and the cupula [adapted from Martini 1998; Cheung 2000].

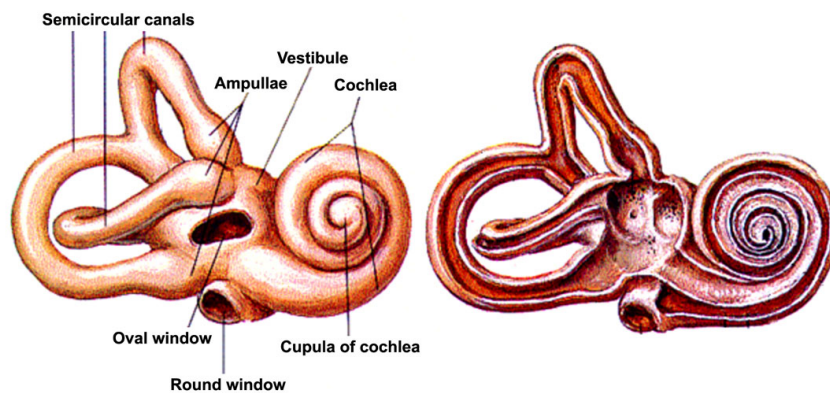


Figure 4.4 - Two views of the hollow, fluid-filled, vestibular bone structures, showing the three semicircular canals and their ampullae in relation to the cochlea. The left-side diagram shows the outer surface, while the right-side diagram is a cut-away view, showing the inner surface [adapted from Netter 1997].




Figure 4.5 -A simplified diagram of a single semicircular canal, showing the ampula, the cupula, and the hair cells' cilia embedded in the cupula [adapted from LaViola 2000; Martini 1998].




Figure 4.6 - Cupula being distorted by motion. As the head rotates (right), the semicircular canal and cupula rotate relative to the endolymph fluid (who's inertia resists the rotation). Thus, the endolymph presses against the cupula, causing it to distend.

In each side of the head, there are three SCCs lying in three mutually orthogonal planes. Each can sense rotation about one axis. Thus the set can sense rotations about all three axes (Figure 4.4). Each SCC is a roughly toroidal tube filled with *endolymph* fluid. Each SCC has a point along the toroid where it becomes wider, the *ampula* (Figure 4.4, Figure 4.5, Figure 4.6). Inside the ampula is the *cupula*—a thin flap that extends across the interior. When the head rotates, it forces the endolymph to flow, which in turn distends the cupula like a sail billowing in the wind. Hair cells embedded in the cupula (Figure 4.6) encode this distention into neural impulses.

I concentrate on the SCCs—they sense head rotation, and this is the motion of concern for Redirection. My statement above that the SCCs act as rate-gyros is overly simplistic and only true for head rotations of particular frequencies and durations. It is important to consider the frequency/phase and time response of the vestibular organs.

It is worthwhile to consider the distension of the cupula as a function of head rotation—a mechanical system with presumably no neural processing.¹⁸ If a particular kind of head rotational movement does not displace the cupula, one assumes the rotation is not detected by the vestibular system.

In order to model the function relating cupula distention to head movement, the literature commonly makes several simplifying assumptions. For example, writers assume the SCC is a perfect toroid rather than its actual irregular shape, infer¹⁹ some physical constants that cannot be measured or estimated from physical observation (such as the coefficient of elasticity of the cupula),²⁰ and ignore variation from person to person. I use the model presented by Howard [1986b].

This model relates the angle of displacement of the endolymph fluid and cupula, θ , to the angular velocity of the head, α and is expressed as:

$$\alpha H = k\theta + r \frac{d\theta}{dt} + H \frac{d^2\theta}{dt^2}$$

where:

H moment of inertia of the endolymph and cupula

αH force acting on the cupula

k elastic coefficient of the cupula

r viscous resistance coefficient

The relationship between cupula displacement and head velocity can also be expressed as a transfer function:

¹⁸This might not be a safe assumption to make, as there are “back-channel,” or efferent, nerves that transmit signals from the central nervous system to the hair cells. The function of these back-channels is unknown to me, but it is conceivable that they affect the stiffness of the cilia, and thus the sensitivity.

¹⁹The inferences are made from high-level neural impulses and human behavior during experiments.

²⁰The displacement of the cupula is less than 10 microns and difficult to observe.

$$\frac{\theta}{\alpha}(s) = \frac{1}{(T_1s + 1)(T_2s + 1)}$$

where:

T_1 the *latency* or *short time-constant*, defined as the time for the deflection of the cupula to reach 1/e of its maximum displacement after an instantaneous change in the rotational velocity of the head, cited to be in the range of three to five milliseconds

T_2 the *recovery time* or *long time-constant*, defined as the time for the cupula to return to 1/e of its central resting position after the head rotation stops. This is inferred to be between three and 16 seconds, depending on which study is cited (Table 4.1).

Those readers unfamiliar with Laplace analysis should consult the Appendix—it gives a brief introductory description of filters using time-constants and also describes how filters compute the integral and derivative of a signal, using the SCC as the driving example.

How does the cupula behave when stimulated with constant-amplitude sinusoidal rotation (imagine the person shaking her head as if to say “no”)? From this model, the gain (or sensitivity) and phase lag of the SCC, as a function of the frequency of head rotational velocity, can be illustrated as a Bodé plot (Figure 4.7). Note that the gain values are relative, because there is no obvious external reference relating the value of displacement of the cupula to the amplitude of sinusoidal head rotation.




Figure 4.7 – A Bodé plot of cupula deflection as a function of the frequency of sinusoidal head rotational velocity.

Assuming the time-constants of $T_1=3$ msec and $T_2=10$ seconds, for frequencies of sinusoidal rotation roughly between 0.1 Hz and 5.0 Hz, the cupula displacement is roughly proportional to head velocity. In this frequency range, the SCCs are most sensitive. Their gain reduces dramatically above and below this frequency range. For frequencies below 0.1 Hz, the cupula deflection is roughly proportional to angular acceleration, and for frequencies above 5.0 Hz, it is proportional to angular displacement (the Appendix describes how one can understand this from the above Bode plot). Normal human angular rotations are in a range centered at roughly 1-2 Hz when walking, and up to 3-5 Hz while running [Berthoz 2000; Jahn 2000].

The SCCs do not sense rotation at a constant velocity. If the head is first still and is then rotated at a constant velocity (in which case the rotational velocity corresponds to a step function), the elasticity of the cupula and the friction of the endolymph will cause the endolymph to eventually match the rotation of the SCC and return the cupula to its undistorted position. Then, if the head-rotational velocity is returned to zero (a step down), the SCC will, instead of reporting that the initial rotation has stopped, report a rotation in the opposite direction (Figure 4.8).

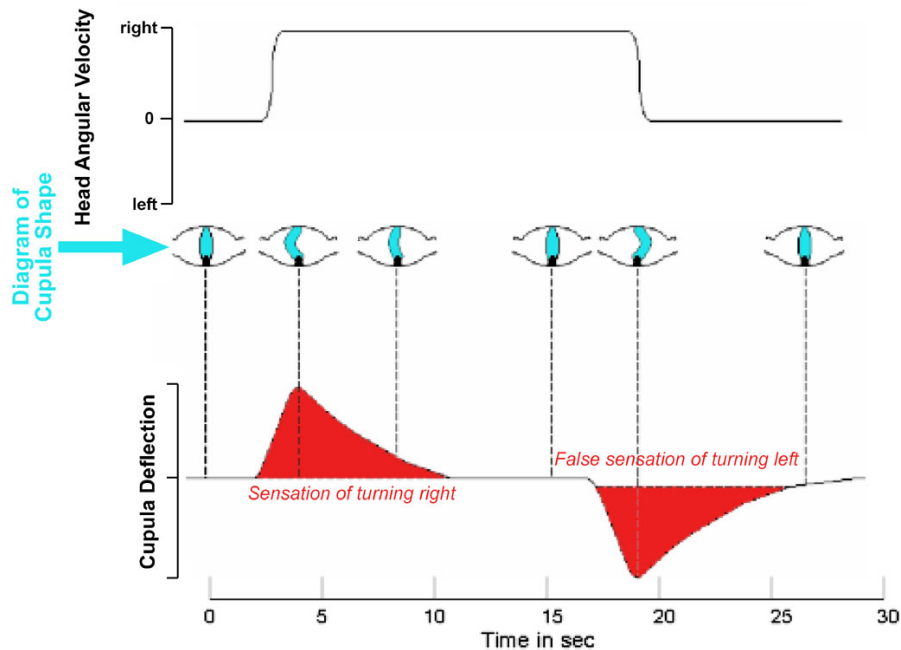


Figure 4.8 - Hydrodynamic properties of canal-cupula-endolymph system during a step up and down in rotational velocity. After 30 seconds of sustained rotation, the cupula has returned to its neutral (resting) deflection. Then, when the rotation ceases, the cupula falsely reports rotation in the opposite direction [adapted from Cheung 2000].

4.4 Auditory Self-Motion Perception

Several characteristics of audio enable humans to perceive the position of the audio source (known as *localizing* an audio source). Two such characteristics are the difference in timing and the difference in volume of the sounds reaching each ear. Additionally, the outer ear filters sound spectra differently depending on the direction the sound is coming from. This is known as the *head-related transfer-function* (HRTF) [Blauert 1996]. Audio cues, even by themselves, can create the illusion of self-motion. A slowly moving audio source (orbiting about the person), presented in darkness, can make a person report that she is rotating [Lackner 1977a].

4.5 Proprioceptive and Tactile Self-Motion Perception

The proprioceptive sense conveys the orientation and motion of the body's muscles and joints. This sense, for example, is what allows a person to successfully move her hand and arm from the outstretched position to touching her nose, even with her eyes closed. Mechanoreceptors are embedded in the muscles, tendons, and joints; they sense muscle length, muscle velocity, muscle force, and joint angles. The tactile sensors are embedded in the skin and detect pressure, texture/friction, vibration, pain, and heat flow [LaMotte 1991; Cheung 2000].

4.6 The Podokinetic system

The *podokinetic* system (a subset of proprioception) is involved in controlling and sensing walking. While *walking*, a person has at least one foot planted on the ground at any given time (as opposed to *running*, where there are moments when neither foot is touching the ground). During a step, the foot planted on the ground is the *stance foot*; the foot in the air is the *stride foot*.

A person can consciously sense the angular rotation of the foot (as it twists about the vertical axis) relative to her trunk [Mergner 1993]. Biomechanically, this angle is limited to 35-45 degrees [Weber 1998]. If a person is walking on a curved path, the podokinetic system combines the angular deflection of the stance foot for the successive steps and estimates the person's direction change. The podokinetic sense, however, has limitations, as evidenced by several symptoms. First, humans cannot maintain a constant heading while walking without other cues. In one study [Gordon 1995], healthy subjects were asked to take 50 steps in darkness and with ear plugs. Subjects drifted by as much as 72 degrees in those 50 steps. Second, humans who

are made to walk in a circular path for an extended time will then, when they are told to walk in a straight line in darkness, continue to turn without realizing that they are turning. This is *podokinetic after-rotation* (PKAR) [Weber 1998].



Figure 4.9 - The rotating treadmill used by Gordon et al. The disk was 5 ft in diameter and spun at 22 deg/s. Subjects walked around the disk such that they stayed in the same spot while the disk turned underneath them [adapted from Gordon 1995].

Gordon et al. [Gordon 1995] had subjects walk along the periphery of a rotating treadmill (a spinning disk—Figure 4.9). Subjects had the visual and vestibular cues of a constant heading, but podokinetic cues of turning at a constant speed. After some time, subjects were removed from the rotating treadmill and asked to walk in a straight line in a dark room. Subjects turned in the direction of the treadmill path without realizing— all subjects turned but thought they walked straight. Jürgens et al. concluded that PKAR is due to adaptation of the podokinetic system to constant turning [Jürgens 1999].



Figure 4.10 - The rotating turntable used by Weber et al. Subjects walked in place while the turntable spun beneath them [adapted from Weber 1998].

Weber et al. [1998] further investigated PKAR, but had subjects walk-in-place on the center of a rotating turntable instead of walking on a treadmill (the vertical axis of rotation went through the head and between the feet) (Figure 4.10). Each subject was first asked to maintain her orientation, while stepping, as the disk spun beneath her. Then the disk was stopped, the subject was transported to a dark room, and then the

subject's orientation was measured as she stepped-in-place and unknowingly turned herself. In addition to replicating the PKAR found in the Gordon et al. study, Weber et al. were able to accurately measure the angular velocity of subjects during PKAR, and quantified the velocity and time-course of PKAR. In some stimulus conditions, subjects rotated themselves at up to 22 deg/s while thinking that they were maintaining a constant heading! The angular velocity of PKAR decays exponentially (*discharges*²¹) with a time-constant on the order of 6 minutes (except for a brief start-up transient which I discuss later), and a small long-term effect with a time-constant above 60 minutes). PKAR velocity charges with exponential decay with a similar time-constant (six minutes).

Finally, PKAR velocity (when time-course effects are corrected for) is proportional to the turntable velocity, except when the velocity approaches 90 deg/s. At this velocity and step rate (2 Hz), the stance foot reaches maximum deflection (35-45 degrees) within a single step.

4.7 Visual Self-Motion Perception

Many argue that visual cues alone can completely convey a person's self-motion ([Gibson 1966], summarized in [Bridgeman 1994]). For example, if the visual scene is rich enough in details, a viewer of a movie can understand how the camera moved through the scene. In fact, there are computer vision techniques to track a camera's motion just from a video recording [Pollefeys 1998; Gibson 2003].

²¹ *charge* is a term from electrical engineering—a capacitor *charges* and *discharges* exponentially. These terms are also used to describe the *optokinetic after nystagmus* (OKAN) velocity storage system.




Figure 4.11 - Three types of optical flow patterns. Left: Laminar translation (which would result from turning one's head left). Center: Radial expansion (which would result from moving forward) Right: Circular (which would result from rolling about the forward axis).

Optical flow is a feature of imagery moving across the retina. An optical flow field contains a motion vector for each position in the visual field. There are several categories of optical flow (Figure 4.11). The first is *translation* or *laminar*, wherein the motion vectors, at each point in the visual field, are parallel and of the same magnitude. For example, if a person is rotating about the vertical axis, then the optical flow field will be translation.²² In *radial expansion* or *radiating*, the motion appears to emanate from a single point called the *expansion point*. For any point in the optical flow field, the motion vector points away from the expansion point, and the magnitude of the motion vector is greater the further away it is from the expansion point. When one is moving forward through a scene, an expansion optical flow field results, with the expansion point located in the direction that the person is headed. In *rotation*, the motion vectors are tangent to circles around a center point. The magnitude of the motion vectors is greater farther from the center point. Looking directly down at the ground while rotating about the vertical axis would result in this optical flow pattern.

For there to be visual cues of self-motion, there must be visual *structure* in the scene. When one is looking at a featureless, evenly lit wall, the image on the retina is also featureless, and thus any motion is visually undetectable. In natural scenes, textures such as grass on the ground and trees in a forest provide the visual detail for generating optical flow.

²² But the person is translating sideways; the vector length of each point depends on its distance from the person—closer objects have a greater velocity across the retina.

Visual cues alone can induce a sense of self-motion. This phenomenon is *vection*, and is the means by which many VEs (and even large-screen movie theaters) induce a feeling of self-motion. Many factors limit the perception of self-motion from visual cues alone. First, the visual scene must contain sufficient detail. Second, the retina acts as a band-pass filter (in the temporal and spatial domains). Fast, high-frequency motions cause the images to move so quickly across the retina that they cannot be faithfully transduced due to the retina's relatively slow response time [Bridgeman 1994]. Using the movie camera analogy, a film in which the camera moves too quickly or jerkily is disorienting and blurry. Professional camera operators are trained to make smooth motions and transitions, and amateur-grade video equipment has features to reduce jerky motion.

Eye motion is divided into three types: *fixation*, *pursuit*, and *saccade*. Saccadic movements are very fast (up to 1000 degrees per second) and ballistic (once they start, the destination of the eyes cannot be changed). Human eyes periodically and unconsciously perform saccadic movements roughly three times a second, separated by periods of fixation. During saccades, the eye's angular velocity is too fast for the visual system to track the outside world. Because of this, visual-only cues of self-motion cannot completely capture the full range of human self-motions [Bridgeman 1994].

4.8 Visual Perceptual Stability

Perceptual stability is the phenomenon of perceiving that the outside world is stable and still. Given that the eyes dart about quickly and unconsciously during saccades, how is it that humans experience visual perceptual stability instead of the world jerking about as the eyes saccade and the head rotates? How does a person know if the movement of the images across the retina is due to self-motion (of the eye or person) or is a result of motion of the external world? Visual cues alone cannot resolve this. I present three situations that demonstrate this, and later describe other perceptual mechanisms to explain them.



Figure 4.12 - An optokinetic drum, where the person is seated on a stationary chair while the surface of the drum, which has alternating vertical stripes painted on it, rotates about the person [from Hain 2005]

1) Consider a person seated in an *optokinetic drum* (Figure 4.12), which is an upright cylinder with vertical black and white stripes painted on the inside [Mach 1875]. The room is initially dark and the drum is rotating at a constant velocity about the gravity axis. The chair is fixed to the ground and the drum is spinning around the chair. Then the lights are turned on, and the person sees the vertical stripes inside the drum moving in the direction of the drum rotation. Initially, she will correctly report that she is still and the drum is spinning. But after several minutes, she will report that the drum has slowly stopped spinning, and now she and her chair are spinning in the opposite direction! One might argue that a uniform sideways optical flow pattern occurs in nature only when the person is rotating.²³ However, this does not explain why she initially reports that the drum is rotating.

2) If a person closes one eye, and rotates the other by gently pushing her finger against the upper or lower eyelid, she will see the world appear to momentarily rotate sideways. Optical flow does explain this situation, but does not account for why a person does not have a similar experience during normal eye motions.

²³ If a person were translating sideways, as if looking out of a train, the optical flow pattern would not be uniform—the closer objects in the scene would have a faster optical velocity than the faraway objects.

3) If one projects a single spot of light onto a sheet of cardboard, and then moves the cardboard sideways as another person observes it, the observer perceives the cardboard as stationary while the light spot appears to move (Figure 4.13).




Figure 4.13 - The frame and light illusion [from Gregory 1966].

The phenomenon of perceptual stability is important for theoretical and clinical reasons. *Oscillopsia* is an illness where afflicted persons perceive the world oscillating about them during head motions. Perceptual stability is also critical for Redirection—the goal is to make the virtual scene appear stable (fixed in space) when it is, in fact, rotating.

Just as research on self-motion perception is ongoing, so is that on perceptual stability. I have found many papers disagreeing with each other. From my layman's viewpoint, there appear to be three complementary categories of mechanisms to account for perceptual stability:

- 1) humans integrate cues from many senses;
- 2) humans predict changes in sensed cues caused by their own actions;
- 3) humans also integrate a mental model of the world and their self-motion, based on previous experience and expectations of world consistency.

In the next sections, I explore these mechanisms and use them to explain the above three puzzling phenomena.

4.9 Integration Among the Senses

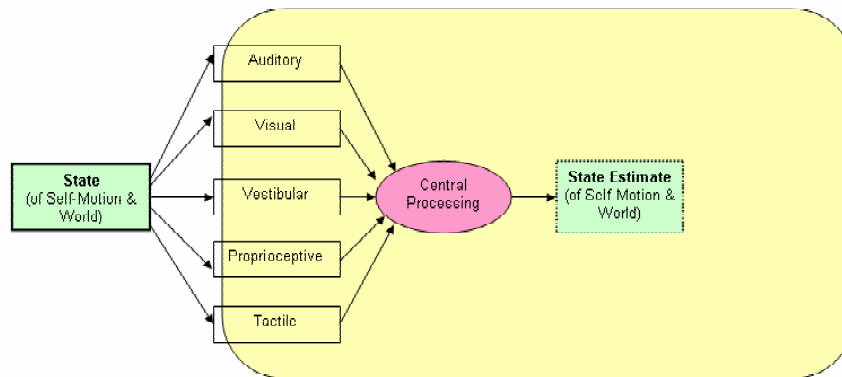


Figure 4.14 - A flow diagram showing motion-state estimation from multiple sensory cues. The area inside the rounded rectangle represents the internal state of the person [adapted from Rolfe 1986; Cheung 2000].

Humans combine information from the senses to perceive their self-motion. When information from one sense is incomplete or ambiguous, another sense can often, but not always, provide information that fills in the gap. The cues from various senses are often redundant. It is useful to think about the mechanisms of self-motion perception not as ways of *knowing* how the person is moving through the world, but as ways of *estimating* the state of self-motion, sometimes with incomplete or ambiguous sensory information [Rolfe 1986].

4.10 Visual-Vestibular Interaction

Integration of visual and vestibular cues is perhaps the most studied of the sensory integrations.

4.10.1 Tilt and Linear Acceleration Ambiguity

The visual and vestibular senses turn out to be very complementary. Recall that the otolith organs (or indeed, any accelerometers) cannot distinguish between linear acceleration and tilt (Figure 4.15).




Figure 4.15 - Otolith ambiguity in sustained acceleration. When the head tilts backward, the otoliths move back (relative to the head). But when the person is linearly accelerated forward (green arrow) the otoliths move in the same way. Thus, the macula is unable to distinguish between tilting backward and accelerating forward. The red arrows show the sum of forces (gravity and inertial) acting on the otoliths [adapted from Martini 1998; Cheung 2000].

This ambiguity between tilt and linear acceleration is resolved from visual cues. When visual cues were missing, this ambiguity has led to aircraft accidents. For example, during low visibility conditions (i.e., darkness or clouds) pilots have misinterpreted a high linear acceleration (such as those encountered in take-off from carriers or aborted landings) as a pitch upwards (Figure 4.16). If the pilot compensates for this illusory pitch up, by directing the aircraft to pitch down, the aircraft will actually be pointed downwards while she thinks it is level. This phenomenon has on several occasions led to the loss of aircraft and life [Berthoz 2000; Cheung 2000]. On the other hand, flight simulators with motion bases take advantage of this ambiguity.




Figure 4.16 – A false sensation of pitch due to forward acceleration. During high forward accelerations and without visual cues, pilots have confused the resultant force vector (the red vector marked “R”) with the normal gravity vector (the blue vector marked “G”) and thus perceived the aircraft pitch upward when it was really flying level [from Cheung 2000].

4.10.2 Washout in Flight Simulators

In a flight simulator with a motion platform, as the pilot accelerates the virtual aircraft, the simulator will move the cab (which contains the pilot and a mockup cockpit) forward to provide the cues of linear acceleration. But after continued acceleration, there will be no more room for the cab to continue to move forward—it will run out of room on its track (Figure 4.17).




Figure 4.17 - A flight simulator with a motion base (the NASA Ames VMS). Left: The cab, shown in yellow, contains the pilot and mockup of the plane's cockpit. It can tilt (yellow arrows) and translate in three directions (the green, red, and blue arrows). The distance it can move is limited by the length of the red and blue tracks, and the distance it can move up and down is limited by the height of the building. Right: A photograph of the flight simulator [from "The VMS Motion Base" 2005].

To address this, the simulator will slowly tilt the cab back (so gravity instead of linear acceleration pushes the pilot back into her seat) while the cab's linear motion is also stopped gradually (Figure 4.18). The tilting and deceleration is below the rotational motion threshold of the SCCs, so the pilot is unaware of it. Despite the fact that the cab is stopped, the pilot still feels like she is accelerating forward. This technique, known as *washout*, is performed while the view "out the window" (presented in video-screens on the cab) shows visual cues of the plane accelerating linearly (and not slowing down or tilting) and results in a very convincing illusion [Strachan 2001].




Figure 4.18 - Washout: As the cab of the simulator nears the end of its travel, it is slowed down and tilted back. The pilot perceives that she continues to accelerate.

4.10.3 Differences between Visual and Vestibular Motion-Sensing in the Frequency Domain and in Onset Timing

The visual system is better at capturing cues of lower-frequency motions, whereas the vestibular system is better at detecting higher-frequency motions. In the middle frequencies, both senses contribute to our perception of motion (Figure 4.19). Consequently, the vestibular system is initially more sensitive to a sudden onset of velocity. But after some time of sustained velocity, the vestibular cues subside while the visual sensitivity to this motion increases and takes over (Figure 4.20).




Figure 4.19- The visual-vestibular crossover. This graph shows, in the frequency domain, the relative contributions of visual and linear vestibular cues to postural stability ([adapted from Duh 2004]).




Figure 4.20 - The contribution of the visual and vestibular (inertial) senses, in the time domain, to the perception of a step in rotational velocity (about the yaw axis) [from Rolfe 1986].

This crossover further explains why washout in flight simulators works.²⁴ After sustained linear acceleration, the vestibular system is less sensitive to that constant acceleration. Thus, the cab can gently slow down without the pilot noticing (Figure 4.21).

This also explains the situation of the person in the rotating optokinetic drum presented above (Section 4.8). When the person first opens her eyes and sees the drum rotating, her visual system is presented with a step function in rotational optical flow (it went from zero to constant velocity immediately). Her vestibular sense tells her that she *could not* have just started rotating (otherwise it would have detected it). Thus, she perceives correctly that the drum is rotating. However, as she continues to rotate, the vestibular contribution is reduced. Were she really rotating, her vestibular system would report the initial step in rotational velocity, and then decrease after roughly 30 seconds. Therefore, after 30 seconds, she is not expecting any vestibular cues of motion, and there is no other sense to correct her visual system from telling her she is moving.

²⁴ Washout can be implemented as a high-pass filter between the signal that represents the simulated plane's acceleration and the signal that controls the acceleration of the simulator's cab [Rolfe 1986].

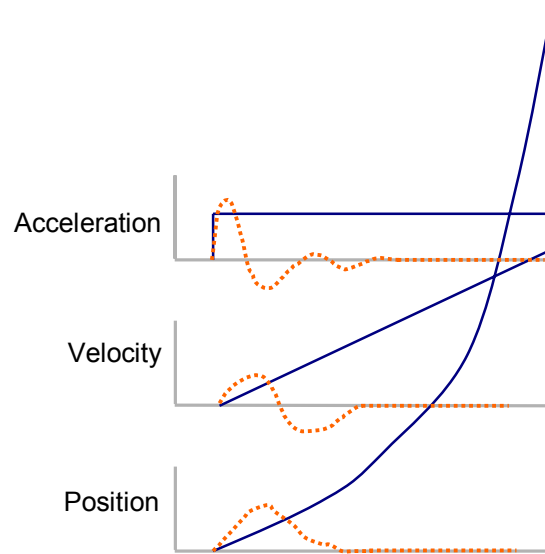


Figure 4.21 - Washout allows the simulator's cab to stay within its range while making the pilot feel like she continues to accelerate. The solid blue lines show the acceleration, velocity, and position of the virtual plane, whereas the dashed orange lines show the corresponding values of the cab. Because the vestibular system is not sensitive to low frequencies of motion, the pilot does not notice that the acceleration has ceased [Chung 2000].

4.10.4 The Vestibulo-Ocular Reflex

The vestibulo-ocular reflex (VOR) is a basic visual-vestibular interaction and is thus worth studying just for this reason. It is also employed to study the question: how much does each sense contribute to sensing motion in each frequency band?

The VOR acts to stabilize the eyes as the head moves. For example, as the head moves quickly to the right, the VOR rotates the eyes to the left to compensate. This eye stabilization serves to keep the optical flow on the retina in the low-frequency range to which the retina is sensitive. In other words, it allows for stable vision during movements of the head. Without this basic stabilizing reflex, head motion would be the greatest source of optic flow [Draper 1998].

There are several reasons to explore the VOR in the context of visual-vestibular integration. Its effects can be measured more easily than can purely visual or vestibular cues. Its understanding touches on concepts that underlie many mechanisms of self-motion perception.

If a seated person is rotated in a dark room, the eyes will still compensate (at least initially) for head motion. The mechanism behind this is the *dark VOR* (dVOR). The eye can only move a certain angle in its socket before reaching the end of its travel. If the head is subjected to sustained rotation, then the eyes will exhibit *nystagmus*. Nystagmus is a repeating sawtooth-like pattern where the eye slowly moves in one direction

(the *slow phase* of nystagmus), then very quickly saccades back to the center of its orbit (the *quick phase* of nystagmus). Nystagmus can be induced by both vestibular cues (vestibular nystagmus) and visual cues (visual or optokinetic nystagmus). To study the VOR under different conditions, researchers measure the gain and phase of the nystagmus's slow phase in response to head movements. A gain of 1 and phase angle of 0 degrees would be perfect compensation of the head movement.

4.10.5 The Optokinetic Reflex

If a person is still and presented with a uniformly translating visual field (as in the optokinetic drum described above in Section 4.8), the eyes will attempt to stabilize relative to the moving visual pattern (i.e., the eyes will track a point on the inside of the rotating drum). The mechanism behind this is the *optokinetic reflex* (OKR). The person perceives she is rotating and the drum is still—for if this were really the case, then the eyes would be similarly stabilized on a drum stripe during slow nystagmus. When a person is rotated in lighted conditions, the VOR and OKR work together to stabilize the eyes.

4.10.6 The OKR and VOR Complement Each Other

When a still person is then rotated at a constant velocity in the dark, the eyes exhibit VOR-induced vestibular nystagmus, as mentioned above in section 4.10.4. This begins within 4 to 14 ms of the onset of the rotation. However, the eyes do not maintain this nystagmus indefinitely. Its gain decays with a time-constant of approximately 25 seconds.²⁵ On the other hand, the OKR has a longer start-up latency (roughly 20-30 seconds [Draper 1998], but does not decay with constant optical flow. As the VOR becomes less compensatory (of the head movement), the OKR “kicks in” for seamless stabilization. The VOR is most effective for head rotation in the 1-7 Hz frequency range, and less effective at lower frequencies, particularly those below 0.1 Hz. On the other hand, the OKR is most effective at frequencies below 0.1 Hz and has decreasing effectiveness in

²⁵ Despite the fact that the long time-constant of the SCCs is about 3 to 16 seconds. The extra time is due to higher-level processing in the central nervous system [Draper 1998].

the 0.1 Hz to 1 Hz frequency range. Thus the VOR and OKR complement each other in the frequency range and onset timings of normal head movements [Draper 1998].

It should be noted that the dark VOR never completely compensates for head motion even in ideal conditions. The gain averages 0.95 instead of 1.0. This suggests that the OKR corrects the residual error left by the VOR.

4.11 Efference-Copy Prediction

Another important property of VOR performance is that it depends on whether the rotation is active or passive. Active rotations, where the person moves her head herself, result in more effective eye stabilization (higher VOR gains and less phase lag) than passive rotations, where the person's head is moved by something else (in many experiments, a motor moves the chair). One mechanism responsible for the increased effectiveness of the VOR during active head movements is *efference-copy prediction*.

Efference refers to a nerve signal that goes from the central nervous system (CNS) to some peripheral effector, such as a muscle. An *afference* is a nerve signal that goes from a peripheral sensor to the CNS. Changes in sensory cues can be caused by the person sensing the changes (e.g., when one turns her head, it normally results in a change in the image projected on her retina); these changes in cues are *re-afference*. Changes not directly caused by motion of person are *ex-afference* (e.g., trees swaying in the wind). When one commands her neck muscles to move her head, the CNS copies this motor command or efference—the *efference-copy*. The CNS uses the efference-copy to predict the re-afferences that will result from that motor command. By predicting the resulting re-afference, the person's CNS can initiate responses sooner (in the case of the VOR resulting from active head rotation) or account for them in the perceptual cues (Figure 4.22). There are other reflexes similar to the VOR that use neck [McCrea 1999], trunk, and even leg motion cues instead of vestibular cues to stabilize the eyes [Howard 1986a], and efference-copy mechanisms have been hypothesized in many of them.




Figure 4.22 - Efference-copy during rotation of the eye [adapted from Gregory 1966].

This explains the second puzzling phenomenon above in Section 4.8—why does one see the world shift when one pushes the finger against one’s eyelid? When eye movements are made with the eye muscles, an efference copy of those movements is used to compensate for this in the visual perception of motion. The CNS is predicting that the image on the retina will shift due to this eye movement and compensates for it in the re-afference. However, there is no such compensation when the eye movement is made with the finger pressing on the eyelid, thus the world appears to shift.

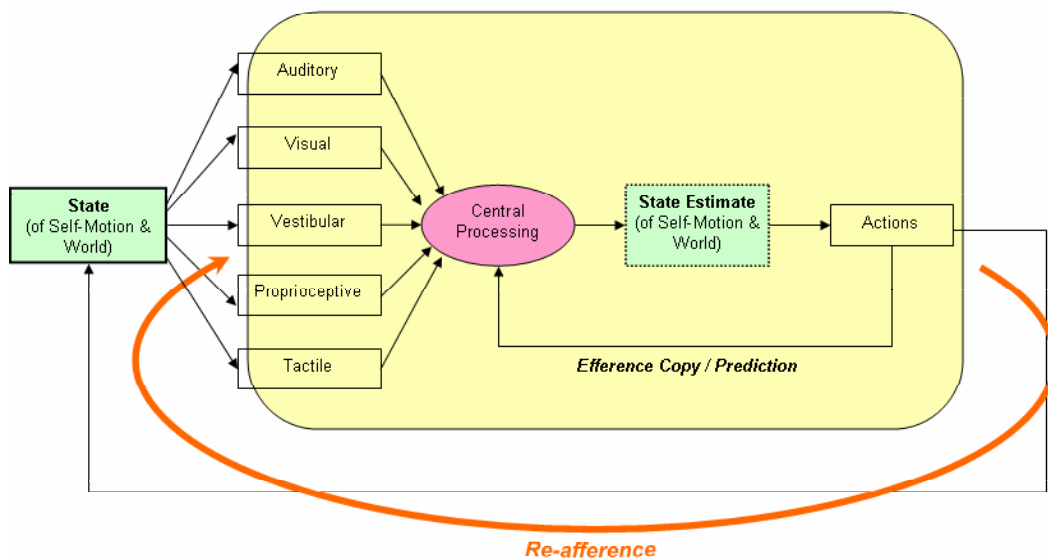


Figure 4.23 - A process diagram of self-motion perception, with re-afference and efference copy prediction [adapted from Howard 1986; Cheung 2000; Rolfe 1984].

However, efference-copy does not completely explain perceptual stability. The motion of limbs and the eyes are not completely specified by the efference command. There is noise in the muscles. External forces prevent the limbs from going to the exact position commanded. The pose of the eye at the end of a saccade cannot be precisely predicted by the efference-copy. One theory of perceptual stability during eye saccades says that the CNS’s estimate of the eye’s pose is recalibrated by matching the new image on the retina to what

was predicted by the person's internal mental model of the world [MacKay 1966]. The residual errors left by the efference-copy-prediction are corrected using the mental model. This is somewhat analogous to how, during head movements, the OKN uses visual cues to correct residual errors left by the VOR.

4.12 Proprioceptive-Vestibular Interaction

Recall the podokinetic after-rotation experiments I describe in section 4.6. Subjects rotate themselves when removed from the turntable and asked to step-in-place in the dark. The PKAR velocity is initially zero, quickly increases to its maximal value, and then decays with a time-constant of six minutes. Weber et al. [1998] hypothesized that the rate of PKAR velocity increase was due to the vestibular cues. The SCCs are able to sense sudden changes in rotational velocity. They act as high-pass filters with a time-constant between 4 and 16 ms. If PKAR peaked immediately after subjects began stepping in the dark, the SCCs would detect it, and subjects would not believe they are keeping a constant heading. Instead, the PKAR velocity is initially inhibited by the vestibular contribution to the person's sense of orientation. PKAR increases at a rate that is just on the edge of that which can be sensed by the vestibular system. This hypothesis was confirmed by Earhart et al. [2004], who measured the PKAR on subjects with non-functioning vestibular systems and healthy subjects. They found subjects with a non-functioning vestibular sense increased to peak PKAR almost instantaneously, due to the lack of vestibular inhibition of PKAR.




Figure 4.24 - A simplified plot of PKAR velocity as a function of time. The dark dashed curve shows the PKAR for subjects without functioning vestibular organs. In these subjects, PKAR instantly reaches maximum velocity, then slowly decays. The lighter curve shows PKAR in normal subjects, where the vestibular cues inhibit PKAR initially, but have no effect later. The initial inhibition in PKAR keeps the vestibular system from detecting the PKAR [from Earhart 2004].

While a person is walking on a straight path, the vestibular system cannot be used to keep her directional heading constant, because of its high-pass quality. But it is well suited to detecting the high frequency sideways jolts that result from the stride foot landing on the rotating surface of the treadmill or turntable. But subjects did *not* notice their PKAR or the resulting jolts. Weber et al. [1998] hypothesized that efference-copy mechanisms are responsible for this. Because the body is actively turning the feet, this turning is subtracted out of the vestibular signal. This is similar to how people do not see the world move when the eyes rotate—because the eye movement is active, the motion is subtracted out from the visual cues of motion.

4.13 Proprioceptive-Visual Interaction

There is other evidence that suggests that visual cues of a limb's position dominate the proprioceptive cues of that same limb's position [Welch 1986]. In fact, another VE interaction technique related to Redirection takes advantage of this [Burns 2005].

The interactions between the podokinetic and visual systems are the most relevant to Redirection. Jürgens et al. [1999] studied the PKAR effect under many different stimulus conditions, each some combination

of visual, podokinetic, and vestibular cues, and used the results to model how each contributes to PKAR. Regarding the optical cues, they experimentally discovered that, even while a subject is keeping her feet still (not stepping), visual and optical cues can also induce (change) PKAR. The optical cue contribution to PKAR has band-pass characteristics, with a high-pass time-constant of 600 seconds and a low-pass time-constant of 90 seconds. This is consistent with the motions of optical flow that result invection (see Table 4.1 for comparison).

During PKAR, if any optical cues become available (e.g., the lights are turned on), they heavily dominate the habituated podokinetic, vestibular senses (that result in the person turning unknowingly), and all turning ceases. Future versions of Redirection should take advantage of this by using visual cues to induce PKAR.

4.14 The Internal Mental Motion Model

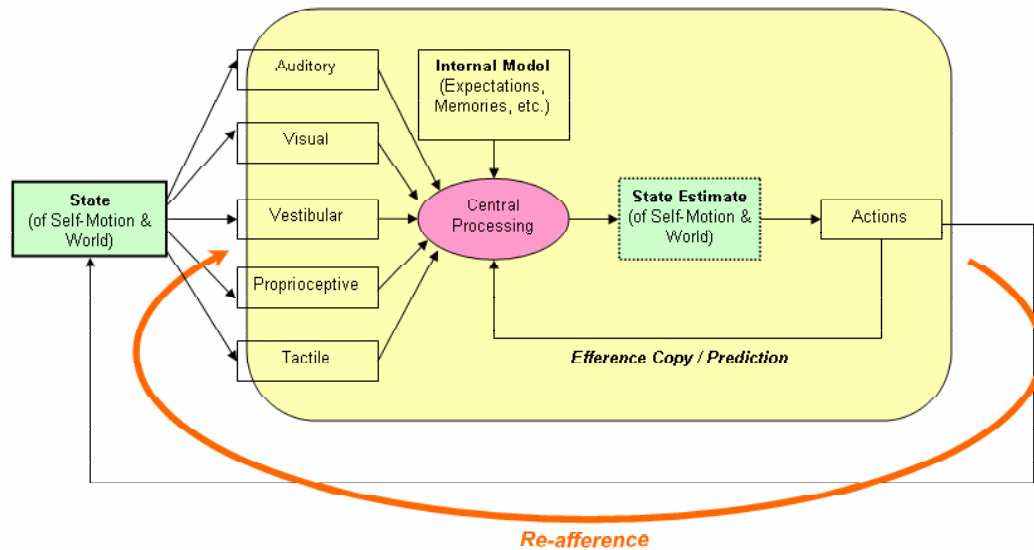


Figure 4.25 - A model of self-motion perception, showing contributions of the internal mental motion model and of efference copy and re-afference [adapted from Howard 1986a; Rolfe 1986; Cheung 2000].

The different sensory cues often reinforce each other, but even all together they can be incomplete. In addition to using the sensory cues, humans also rely on an internal mental motion model of the scene and their self-motion through it. To quote Gregory:

“It is not difficult to guess why the visual system has developed the ability to use non-visual information and to go beyond the immediate evidence of the senses. By building and testing hypotheses, action is directed not only to what is sensed but to what is likely to happen, and it is this that matters. The brain is in large part a probability computer, and our actions are based on the best bet in a given situation. The human brain makes efficient use of its rather limited sensory information” [Gregory 1966 p225].

This internal mental motion model is based on previous experience and on future expectations (Figure 4.25). Imagine a blind person feeling her way through a familiar room. Her vestibular, proprioceptive, and tactile cues tell her how she is moving in the room. But this information is incomplete—she may not know exactly how far she has walked into the room. When she comes across a familiar object whose location she knows (e.g., a particular sofa), she then knows exactly where she is in the room. It turns out this internal mental motion model is a very powerful part of human perception—humans often perceive what they expect to perceive based on previous experience. Under the term *internal mental motion model*, I group many different persistencies of expectations and memories, from short-lived to permanent. One can imagine that a person’s mental model of the position of objects in a room is transient—often re-evaluated and re-learned. On the other hand, there are other more permanent (or “hardwired”) assumptions. Breaking those assumptions can lead to strong illusions, even if one consciously knows exactly which assumptions are false. For example, the mental model assumes that certain features in the scene are vertical (i.e., trees, walls, etc.), even when they are not. This assumption leads to one’s unsettling experiences in the “anti-gravity” houses (Figure 4.26) and the Ames room (Figure 4.27). Again, Gregory expresses this elegantly:

“The perceptual system has been of biological significance for far longer than the calculating intellect. The regions of the cerebral cortex concerned with thought are comparatively juvenile. They are self-opinionated by comparison with the ancient striate area responsible for seeing. The perceptual system does not always agree with the rational thinking cortex” [Gregory 1966 p224].

Another assumption is that of consistency—one assumes that the external scene does not change (in the short term).²⁶ At the end of a saccade, one assumes the external world has not significantly changed since the beginning of the saccade (which began just a few milliseconds earlier), and thus the measurement of the position of the eyes is recalibrated to the expectation of where the visual features of the scene should be, as determined by the mental model. In the rotating optokinetic drum, the person (who is not moving) eventually

²⁶ This does not imply that one assumes the world has no motion. The state of the world could include the paths of objects moving relative to the background.

experiences the illusion that she is rotating because her internal motion model assumes the outside scene is not moving. The optical flow (of the scene uniformly moving to the right or left) is consistent with this assumption, and the vestibular cues do not contradict it (because the vestibular system cannot sense constant rotation).

Yet another mental-model assumption is that larger objects and farther-away objects are more likely to be perceived as staying still, despite any optical flow they might cause [Gibson 1966; Brandt 1975]. This explains the final situation described above in Section 4.8—the projected light spot appearing to move when the cardboard it is projected on is really moving. Here, the cardboard’s larger size makes one perceive it as still, and thus the brain decides that the light is moving to achieve consistency with the light’s motion relative to the cardboard.

The internal mental motion model of the environment and self-motion carries expectations for how one’s own actions will affect it. When one acts, the new perceptual cues are compared against the internal motion model, thereby refining, or invalidating it—this is re-afference. The process of estimating self-motion using the mental motion model and limited and ambiguous sensory cues is consistent with the *rest-frame hypothesis* of motion sickness (discussed in Chapter 3)—when the mental model of self-motion is invalidated by new and inconsistent sensory cues, motion sickness results.




Figure 4.26 - The anti-gravity room. “Nevertheless, these tourist attractions contain some of the strongest visual illusions known. Familiarity with how they are constructed will not break the illusion. ...When you enter the house, you will notice that it has a strange tilt. All references to the true horizontal are removed from your sight. This is always true whether you are just outside the house or inside it. For example, there is always a wooden fence around the house to remove any significant comparisons to the true horizontal...The anti-gravity house is actually built at an angle of 25° off the true horizontal. This will explain every effect seen. Once in the area of an anti-gravity house you are always comparing the effects to what you are used to—normal-level floors and walls that are perpendicular to the ground” [from “Mystery Spot” 1997].




Figure 4.27 - The Ames Room illusion. Left: The two women are of equal size, but the non-rectangular shape of the room makes one of them (who is farther away from the viewer) appear much smaller. Right: An overhead diagram showing the actual shape of the Ames Room. The circles represent the locations of the women [from Gregory 1966].

4.15 Quantitative Characterizations of the Senses

Psychophysical values, such as the thresholds for detecting a motion, and the cutoff frequencies and time-constants²⁷ that determine the sensitivity of sensory systems vary from experiment to experiment, person to person, and even situation to situation. For example, even without being aware of it, a person is much more likely to detect and respond to self-motions that are active (caused by the person—e.g., when she turns her neck) rather than passive (when someone else turns her neck) [Howard 1986a; Draper 1998]. Also, the likelihood of a person detecting a rotation depends not just on the magnitude (i.e., acceleration or velocity) but also on the duration. The smaller the magnitude, the longer the duration required to detect it [Howard 1986b]. These values are summarized in Table 4.1 and Table 4.2.

²⁷ The Appendix provides an explanation of how a filter's time-constants and cutoff frequencies (or corner frequencies) are related.

Table 4.1 - Summary of values of the band-pass filter characteristics of three sensory modalities for inducing a sensation of rotation. *Italicized values have been computed from the non-italicized values (to aid comparison). Each source is itself a summary of other research results.*

vestibular				
<i>high-pass</i>		<i>low-pass</i>		<i>source</i>
<i>time-constant</i>	<i>cutoff/corner frequency</i>	<i>time-constant</i>	<i>cutoff/corner frequency</i>	
10 s	.016 Hz			Jürgens 1999
5-7 s	.023-.032 Hz			Draper 1998
10 s	.016 Hz	0.003 s	53 Hz	Howard 1986
3.8 s	.042 Hz	0.005 s	32 Hz	Howard 1986
16 s (yaw)	.01 Hz			Cheung 2000
7 s (pitch)	.023 Hz			Cheung 2000
4 s (roll)	.04 Hz			Cheung 2000
0.16 - 3.2 s	.05 - 1 Hz	.023 s	7 Hz	Draper 1998

visual (optical flow)				
<i>high-pass</i>		<i>low-pass</i>		<i>source</i>
<i>time-constant</i>	<i>cutoff/corner frequency</i>	<i>time-constant</i>	<i>cutoff/corner frequency</i>	
600 s	.0003 Hz	90 s	.0018 Hz	Jürgens 1999

podokinetic				
<i>high-pass</i>		<i>low-pass</i>		<i>source</i>
<i>time-constant</i>	<i>cutoff/corner frequency</i>	<i>time-constant</i>	<i>cutoff/corner frequency</i>	
		360-720 s	.0002 - .0004 Hz	Weber 1998; Weber 2000
		360 s	.0004 Hz	Jürgens 1999

Table 4.2 - Various reported rotation detection thresholds of the semicircular canals.

<i>vestibular detection thresholds</i>	<i>source</i>
0.1 deg/s ²	Draper 1986
0.2 deg/s ²	Howard 1986
.25-3.0 deg/s ²	Howard 1986
0.2- 0.7 deg/s ²	Howard 1986
.05 – 2.2 deg/s ²	Howard 1986
0.3 deg/s ²	Howard 1986
2.5 – 4.5 deg/s (velocity)	Howard 1986
0.14 deg/s ² (yaw)	Cheung 2000
0.5 deg/s ² (pitch)	Cheung 2000
0.5 deg/s ² (roll)	Cheung 2000

Chapter 5: How Redirection Works— Qualitative Arguments Based on Self-Motion Perception Theory

This chapter describes self-motion perception as it applies to Redirection. It covers both how the Redirection technique currently works and also how it might be improved.

The overall impression I intend that the reader get from the previous chapter is that humans produce and maintain a model of their self-motion with incomplete and noisy sensory information, from many separate senses. After the first informal trial session of Redirected Walking, I was amazed and delighted at how convincingly users were fooled into changing direction. Given what I now know of self-motion perception, I am no longer amazed that Redirection works—I am amazed that we humans can do the daily activities that are taken for granted—standing without falling over, climbing down a crowded subway staircase surrounded by people moving at different speeds, and walking across a street while avoiding cars.

5.1 Self-Motion is the Simplest Explanation for the Sensory Cues Caused by Redirection

I believe the overriding explanation for why Redirection works is that *human motion perception machinery does not expect it*. The world almost never rotates about the center of one's head. The simplest explanation, based on real-world experience, is that the world is fixed and stable, and one is moving her own head instead. This is a strong illusion, much like the Ames room illusion (Figure 4.27), where one is so accustomed to seeing a rectangular room that, when faced with an Ames room, the perceptually (even if not cognitively) likeliest explanation is that the person in it is changing sizes as she walks from corner to corner. Redirection attempts to encourage the illusion of a stable world in several ways.

Berholz and Gregory convincingly argue that perception is not passive observation, but is intertwined with action [Gregory 1966; Berthoz 2000]. In Redirection, the virtual objects respond to translational head movements as real-world, ground-fixed objects would. If a virtual object floats in front of a user, the user can

move about and inspect it from many angles. In one study, Jaekl et al. varied the gain between real-world motion and virtual motion (e.g., for every 1 cm the user moved forward in the laboratory, her position in the virtual scene moved forward by 2 cm) and discovered that even when there is significant mismatch in the gain, the user still has the illusion of perceptual stability. She perceives that she is moving correctly in the virtual scene, and the virtual objects that make up the scene are fixed in space [Jaekl 2002]. The mind tenaciously holds on to the assumption and perception of a stable world. Even low-level reflexes such as the vestibular ocular reflex (VOR) adapt to correct for visual-vestibular mismatches in order to bring about perceptual stability [Draper 1998]. The virtual scene in the Jaekl study consisted of a simple, unnatural, and nondescript sphere, lit from the inside, with no cues of scale. Despite this, the subjects perceived a stable virtual world.

Redirected Walking was developed for more complex and familiar human-scale virtual scenes such as the interiors of buildings. A user can walk closer to a virtual window frame, inspect it from many angles and stick her head through it just like a real window frame. This supports her assumption that the window is not some floating and moving object, but firmly attached to the wall and anchored to the ground. When Redirection rotates the window frame, it rotates it in perfect unison with the wall and floor. The optical flow pattern is an even laminar translation, which is normally encountered only during self-motion. All of this encourages the perceptual system to believe that the most likely explanation for the cues it is receiving is that the world is not rotating about the user.

5.2 Non-Visual Cues

There are non-visual sensory cues that could betray the trickery of rotating the virtual world about the user. The strategy with Redirection is to minimize these potentially conflicting cues. To my advantage, the non-visual senses detect different kinds of motion than the eyes (e.g., the SCC sense higher-frequency motions), so the non-visual cues may not conflict with the visual ones. Several studies show that when sensory cues do conflict, the visual cues usually dominate [Jürgens 1999]. The expectation of a stable world is so strong, and our self-motion perception mechanism so plastic, that even when there is conflicting sensory information, research shows that it is sometimes suppressed by the perceptual mechanisms. For example, in the studies of the podokinetic effect [Weber 1998], subjects who were unknowingly turning themselves, even quickly (22 deg/s), did not notice the sideways jolts that resulted from their stride foot striking the rotating disc, despite the fact the jolts are the kind of motion that the vestibular system easily detects under normal conditions.

To encourage the illusion of world stability, the Redirection algorithm used in the RW and RWP experiments presented a 3D spatial audio scene as well as the visual one. The scenes were rotated in unison. The virtual audio also helped to mask the real-world noises, which might interfere with the illusion or betray the users' real orientation in the laboratory.

5.3 Algorithm Description in Terms of What the User is Doing

The Redirection algorithms take advantage of several special conditions.

5.3.1 *While Standing Still*

While the user is standing still, the system rotates the scene in a slow, smooth and low-frequency manner. The goal is to keep this motion undetectable by the vestibular system, which is more sensitive to higher-frequency motions (whereas the visual system is sensitive to lower-frequency motions). The user rotates herself with the virtual scene. If she is standing still with her feet fixed to the floor, there is a theoretical limit to how far the system can rotate the virtual scene—the maximum angle a foot can turn, relative to the trunk about the gravity axis, is 35-45 degrees [Weber 1998]. There is also a limit to how much the torso can twist comfortably. Before experiment RPW-I, I assumed that at some point, with continued rotation, the user must become aware of her unnatural body position (via proprioception). In experiment RWP-II, one subject stood in place and rotated with the virtual scene for almost 100 degrees, adjusting his feet as necessary, without noticing. It appears that people do not stand for long without moving their feet. The person shifts weight from foot to foot and makes small adjustments as needed, seemingly without noticing. This happens even in the real world, for example when people are waiting in line at the bank.

5.3.2 *While Really Turning the Head*

When the user is turning her head, Redirection can imperceptibly rotate the virtual scene much faster. In experiment RWP-II, the system magnified the user's head's angular velocity, such that the user could see more of the virtual scene before seeing the missing back wall of the cave. This is similar to Jaekl's experiment where the illusion of perceptual stability was maintained despite the virtual head motion being mismatched to the real-world head motion [Jaekl 2002]. In experiment RW, the system also magnified the user's head angular velocity, but the gain was continually chosen so as to steer the user toward the target waypoint in the real laboratory.

The world rotating about the center of the head is rare in the real world; the world doing so *in response to the person turning her head*, is singular. The perceived rotation of the world across the eye happens exactly when the user turns her head. The body assumes that this optical flow is due to the user's own self-motion. I hypothesize that this is why the system can inject faster virtual scene rotation in response to the user really rotating her head and that the underlying mechanism for this is an efference-copy. In real life, when a person moves her head, the body makes a copy of the motor command to the neck (the efference-copy), and uses it to predict that there will be a resulting change in the sensory cues as a result of the motor command (re-efference). The body accounts for the changes in efference due to efferences and suppresses them from the consequence perception. The external world has not changed state due to the head movement, so the resulting change in sensory cues is accounted for at a subconscious level and ignored by higher conscious perception. There is margin for error in the process, as the both the efferences and afferences have noise, and the muscles do not carry out their commands exactly. The efference-copy mechanisms are tolerant to errors. When a person turns her head left, the body expects the world to rotate right by the corresponding amount, and suppresses this perception at the higher level—the person does not see the world rotate across the retina. And if the external world does not rotate across the eye by the exact amount the neck was instructed to move, the differences are written off as perceptual and motor error²⁸ [MacKay 1966]. Thus, with Redirection, the system can sneak in faster virtual scene rotation during active head movements.

5.3.3 While Walking

Researchers have confirmed that a person walking will veer in the direction in which the visual scene is rotated [Warren Jr. 2001]. As the person is walking, the VE system can get away with slightly more rotation than while the user is standing still. (In experiment RWP-II, the virtual scene rotated at a maximal rate of 19

²⁸ This is the evaluation or comparison model, as opposed to the cancellation model. Both are described in MacKay [1966].

deg/s when the user was walking-in-place, compared to 5.6 deg/s while standing still—subjects did report noticing the rotation.)

Given some arbitrary, constant injected scene-rotation rate, I hypothesize that the faster a person is stepping, the less noticeable that rotation rate. The Redirection algorithm (in experiment RW) used linear velocity as an approximation of step rate (the faster one is walking, the greater the step rate), but I believe that step rate, rather than linear velocity, is what determines the rotation threshold.

While a person is standing still and viewing a rotating visual scene, I hypothesize that the more the ankles, legs, and hips twist, the more likely she is to notice the twist. There is also a maximal angle to which these joints can twist. But when the person is marching on the spot she can be made to turn continually without noticing. For any fixed rotation rate, the stance foot is less likely to reach its maximal angle when the step rate is higher, and so the person is less likely to detect the rotation.

Also I hypothesize that the podokinetic system estimates total change in orientation by summation over successive steps. For any fixed rotation rate, the podokinetic system must sum over more successive steps when the step rate is higher, and so the person is less likely to detect the rotation.

Finally, with each step, the high-frequency jolt that results from a foot striking the ground briefly blurs the vision. During this time, vision is suppressed [Grossman 1989] and when the suppression ends, the brain re-fixates or reacquires whichever object it was previously looking at. Any unpredicted motion of the target between fixations is chalked up to noise. The greater the step rate, the more often vision is suppressed, resulting in greater opportunity for imperceptible visual scene rotation.

I hypothesize that the faster the user is stepping, the more unnoticed rotation can be injected. Thus, the system can rotate the scene faster when the user is walking (or stepping in place) than when she stands still. If this is correct, future systems would benefit from the use of sensors to measure directly when the feet lift from the floor and when they strike it, as these signals would not be affected by changing stride lengths.

5.4 Improvements to Redirection Suggested by Self-Motion Perception Literature.

The previous sections describe my best current hypothesis of how Redirection works. The mechanisms of self-motion perception have further implications for Redirection that future developers should consider.

5.4.1 *Looking Down*

While the user is looking down, she should be able to detect the rotation of the virtual scene more easily. This is because, with Redirection, the scene is rotated only about the vertical axis (the virtual floor stays in the same plane as the real floor). When the user is looking down, the rotations form a rotational optical flow pattern (where the top of the visual field has a different direction and speed than the bottom) rather than a laminar/translation shear pattern, as when the person is looking straight ahead (Figure 4.11). This rotational flow pattern does not occur in response to any natural head or body rotations, so it is less likely to be perceived as self-motion. I have experienced this myself, when testing Redirection in virtual pit scene (Figure 1.4). When I looked down (to keep myself from walking off the ledge), the virtual scene rotations were much more noticeable. Future implementations of Redirection should rotate the scene less, the more the user's head is pointing down. Some people tend to look at the floor while walking—users who did this in the virtual scene would be troublesome, as the system would not be able to steer them as effectively. This could be remedied by giving the user a task to force her to keep her gaze level.

5.4.2 *Running*

I have argued that the faster a person is walking, the more rotation the VE system should be able to get away with. Running may be an exception and hence may require a different algorithm. A running stride is different from walking and may have different neuro-motor control mechanisms. Jahn et al. showed that blindfolded subjects whose vestibular organs are disrupted by electrical stimulation are *less* affected during running than during walking [Jahn 2000].²⁹ This implies that some other sense, besides vestibular and visual, dominates the control and sensing of orientation while running. Because of this, one cannot simply assume that the Redirection algorithm for walking applies equally well to running.³⁰

²⁹ Though it may be that the specific electrical vestibular stimulus they used only disrupts walking, and a different pattern of stimulus would be required to interfere with running.

³⁰ I have not explored Redirection for users who are running, because the VE system equipment is delicate and its cables have the potential to trip the user.

5.4.3 *Faraway Virtual Objects*

In experiment RDT-wcv, several subjects reported that they were most likely to detect the virtual scene rotations as they approached and passed through the doorway, which was in the middle of their path through the virtual scene. I hypothesize that this is because virtual scene rotation is more noticeable when parts of the virtual scene are close to the user.

When a person's head is rotating, the optical flow is not exactly laminar/translation, because the two eyes are offset from the center point of head rotation—each eye translates as it rotates during head rotations. A 45-degree head yaw results in a 10 to 15 cm translation of the eyes [Jaekl 2002]. For objects in the scene that are far away, the effect of this translation on the optical flow pattern is negligible. But for objects that are nearby, the effect is more pronounced and detectable. The farther away an object is, the less its position on the retina changes as the eye translates. As an illustration of this, consider a passenger viewing the scenery from a moving train—the nearby trees zoom past the observer while the moon appears stationary. Furthermore, a moving background is more likely to be perceived as being still and result in illusory self-motion than a moving foreground [Brandt 1975].

In the VE systems I have used, the position of each eye and the position of the head's center of rotation are estimated from the tracker's head-sensor position and orientation. These values vary from person to person and are difficult to model for each individual. Each time any particular user dons the headset, it can rest on the head differently. The errors introduced by this are negligible for virtual scenery far away, but the presence of nearby virtual objects coupled with the inaccuracies of eye and center-point positions might cause the virtual scene rotations to become more noticeable.

VE systems should model the centers of rotation of the eyes and head as accurately as possible. One way to find the head's center of rotation, for each user, would be to have the user stand still and turn her head left and right several times. If the head-tracking sensor's position is not at the center of the head, then it will move in a circular arc. The center-of-radius of the arc would be the head's center of rotation.

5.4.4 *Taking Advantage of Podokinetic High-Pass Characteristics*

One should not be surprised that the podokinetic system for sensing and controlling orientation is not sensitive to slow gradual changes in direction. As described by Earhart et al., “When walking in everyday

environments, one must change walking direction frequently to round corners and avoid obstacles. In fact, walking a straight line is the exception, rather than the norm” [Earhart 2004]. The Redirection algorithms I have implemented as part of this dissertation do not take advantage of PKAR, but I consider how future implementations could do so.

In the original experiment [Weber 1998], PKAR was induced by having subjects walk along the periphery of a rotating treadmill (a spinning disk). Subjects had real-world visual and vestibular cues of a constant heading, but podokinetic cues of turning at a constant speed. I propose that a VE system can induce PKAR by slowly increasing the visual scene rotation rate in accordance with PKAR’s charging time-constant of 6-12 minutes (Table 5.1). Since the user (unknowingly) turns herself with the visual scene, this should cause the podokinetic system to habituate to the rotation. A user who is attempting to walk straight (in the virtual scene) at a constant speed would gradually spiral inward (in the real world) without being aware of this, as illustrated in Figure 5.1.

Table 5.1 - A comparison of how each cue is stimulated to induce PKAR, for the original experiment [Gordon 1995] and my VE system proposal.

	<i>Gordon 1995</i>	<i>Proposed VE System</i>
<i>Visual</i>	Straight: Real-world	Straight: Synthetic (via HMD)
<i>Vestibular</i>	Straight: Real-world	Straight ³¹ : Real-world
<i>Podokinetic</i>	Turning: Synthetic (via treadmill)	Turning: Real-world

³¹ The user is turning in the real-world, but the increase in rotation is so slow that the vestibular system cannot detect the rotation and reports that the person is walking straight. The time-constant of the vestibular system is much smaller than that of the podokinetic system.




Figure 5.1 - A simulated path of a user, who is walking in a straight line in the virtual scene, but due to PKAR-Redirection, is walking in a spiral in the lab. The simulation assumes that the user walks at a constant 1.4 meters/s, that the PKAR charging time-constant is six minutes, and that the maximum PKAR velocity is 22 deg/s. Weber et al. report [Gordon 1995] that some subjects turned at that rate without being aware they were doing so. Other forms of Redirection could be used with PKAR-Redirection to result in a path that requires less tracked area (Figure 1.14).

This new component of the Redirection algorithm (henceforth called *PKAR-Redirection*) would be in addition to the previous components based on the speed at which the user is turning her head, her walking speed, and the direction in which the computer would like to steer her. Because PKAR has a large time-constant, PKAR-Redirection could not be used to actively steer the user as she consciously changes direction in the VE. Figure 5.2 shows the simulated path of a user who walked straight, then turned left 90 degrees, and continued walking straight in the virtual scene. While the user is turning herself in the virtual scene, the PKAR Redirection cannot quickly change. Instead, PKAR-Redirection could be used to establish the user on a circular orbit in the laboratory, and then the other components of Redirection could return her to the circular path when she deviates from it. I develop this idea further in Chapter 6.




Figure 5.2 - A simulated path of a user, computed using the same simulation and PKAR-Redirection algorithm as in Figure 5.1, but where the user turns left by 90 degrees once during the simulation, and otherwise walks straight. Other forms of Redirection would be required to return the user to the original orbit point.

Chapter 6: Steering the User during Unrestricted Walking

When a large enough tracked area is available, I expect that Redirected Walking will allow the user unrestricted exploration of arbitrarily large virtual scenes without the use of waypoints. For this to happen, the system must steer the user in the lab³² to keep her from running into the lab boundaries, without knowing her intended path in the virtual scene. I present several algorithms for steering the user in these situations. I offer these algorithms and observations to system designers as a starting point, as this work is incomplete. I have only implemented one of the algorithms, and even that was not tested in a sufficiently large tracked area.

I assume that people tend to continue on their current (torso) heading. This is not to say that people do not change directions or walk in curved paths. Rather, at any given time, the best prediction one can make (without any special knowledge of the virtual scene or her task in it) is that the person will continue on her current heading.³³ I also assume the lab shape is roughly square, not a skinny corridor or complex multi-room shape.

³² As mentioned in Chapter 1, I use the term *lab* to mean the physical tracked space in which the user walks. If the tracked space is smaller than the physical room, then the goal is to prevent the user from walking out of the tracked space. If the tracked space is larger than the physical room available (e.g., using some GPS-like tracking technology), then the goal is to prevent the user from colliding with the walls.

³³ It would be worthwhile to test this assumption, and also to characterize the paths humans take while they are exploring arbitrary settings. With what frequency do humans turn their heads, bodies, and change walking directions? What percentage of the time, and under what circumstances, do they side-step or walk backwards?

6.1 Steer the User Toward the Center of the Lab

The first strategy I explored to keep the user away from the lab boundaries was to try to keep her in the center of the lab, by continually steering her toward the lab center. The assumption underlying this is that the lab center is the optimal place to be (the user has the least chance of exceeding the lab boundaries while there). If the lab center is to the left of the user's heading vector, the virtual scene rotates leftward in order to steer her to her left. If the lab center is to the right of her heading, she is steered right (Figure 6.1).

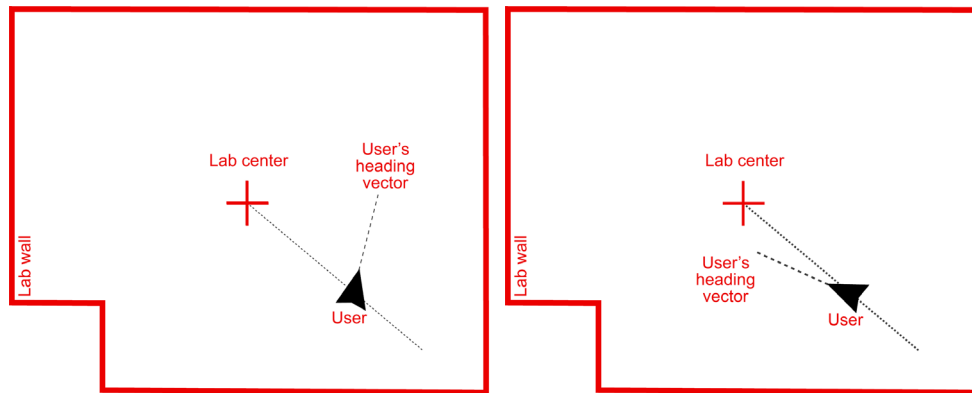


Figure 6.1 - Steer-to-Center algorithm: The user's position and heading in relation to the lab center. Left: If the system is trying to steer the user toward the lab center, the center is on her left and the virtual scene rotates left in order to steer her toward the center. Right: If she turns left, past the lab center, the system then switches the virtual scene rotation toward the right.

Informal testing revealed (Figure 6.2):

- 1) When the user turns past the center of the room (Figure 6.1), the direction of the virtual scene rotation changes abruptly (because the lab center was on her left but is now on her right). This abrupt change is very noticeable to the user.
- 2) The user's head position and orientation wobble. Even when walking on a direct path, the head bobs from side to side and the orientation sways left and right (Figure 6.3). The tracking system introduces additional noise. This interacts with the observation above and results in the virtual scene appearing to vibrate about the user's head.




Figure 6.2 - Informal testing of the Steer-to-Center algorithm. Top: Views of the virtual scene used in testing. Bottom: An overhead view of the user's path in the virtual scene (blue) and lab (red). The lab center is shown with a red cross and the user's heading vector as a pink line.



Figure 6.3 - A recorded path (of head position, projected onto the ground plane) of a person walking a relatively straight path. The wobble is related to the person shifting weight from one foot to the other.

I addressed these problems by having the steering rate change smoothly. This was done by attenuating the steering rate by multiplying it by the sine of the angle between the user's heading vector and the vector from the user's position to the lab center (Figure 6.4). When the user is headed perpendicular to the lab center (pointed neither toward nor away from it), the steering rate is unmodified ($\sin(90)=1$), but when then the user is pointed directly toward the lab center, the steering rate is completely attenuated ($\sin(0)=0$). As she turns past the lab center (Figure 6.1), the steering attenuation changes smoothly. This technique seems to be less intrusive and does not suffer from the virtual scene vibration.




Figure 6.4 - Left: The steering rate is attenuated by multiplication by the sine of angle θ , the angle between the user's heading and the vector pointed toward the lab center. If the user is pointed perpendicular to the lab center, $\sin(90)=1$ and the steering rate is not attenuated. As the user turns past the lab center (as in Figure 6.1), the steering changes smoothly. Right: A sample path of the user steered toward and then through the lab center.

Although I assume that the best prediction is for the user to continue walking along her current heading in the virtual scene, this steering algorithm does not depend on it. If the user suddenly changes direction in the virtual scene, her route in the lab will momentarily divert from the path toward the lab center. But since the system is continually trying to steer her toward the center, this diversion will be corrected shortly afterward (Figure 6.5).




Figure 6.5 – Steer-to-Center algorithm: Three hypothetical sample paths that the user could take in the virtual scene (right, in blue) and the corresponding paths she would take in the lab (left, in red). If the user walks straight in the virtual scene (path 1), she is steered along a smooth path (in the lab) through the lab center. If the user decides to take a 90-degree right or left turn in the virtual scene (paths 2 and 3), her 90-degree turn becomes something like a 45-degree turn in the lab. After the turn, the user is again redirected toward the lab center.

A significant problem with this algorithm is that when the user is pointed directly away from the lab center, the system does not steer her back toward it (Figure 6.6). In this configuration, the tracking noise and normal head wobble cause the system to switch between steering her right and left. But if the Steer-to-Center algorithm is successful in steering the user toward the center, she will then walk through the lab center and then

be walking directly away from it—this is exactly the problem situation! This invalidates one of the assumptions on which this strategy is based—the lab center is not the safest place for the user, if she is heading away from the center. One potential way, which I have not implemented, to address this is to bias the steering in one direction when the user is near the lab center. The following two algorithms are designed to address this problem situation in different ways.

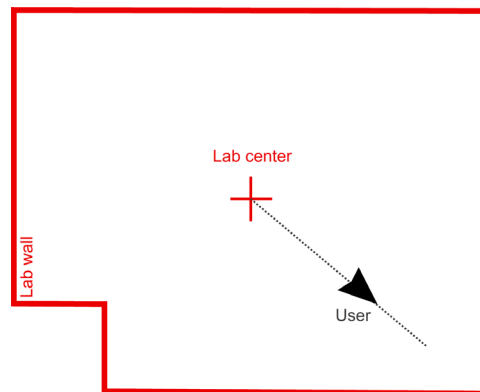


Figure 6.6 – A problem with the Steer-to-Center algorithm: When the user is pointed directly away from the lab center, steering her toward the center again is problematic. The system shifts back and forth between steering the user right and left due to head wobble.

6.2 Proposed Algorithm: Steer the User Onto a Circular Orbit

Instead of steering the user toward the lab center, the Steer-onto-Orbit algorithm tries to steer her onto a circular path which orbits the lab center. Once she is on this path, she can continue walking in a straight path in the virtual scene while staying on the circular path in the lab (Figure 6.7). If she takes a turn in the virtual scene, her lab path momentarily deviates from the circular orbit, but then the system steers her back onto it (Figure 6.7).




Figure 6.7 - Left: The user is steered onto a circular path orbiting the lab center. Superimposed are three hypothetical sample paths that the user could take in the virtual scene (right, in blue) and in the lab (left, in red). If the user walks straight in the virtual scene (path 1), she is steered along a smooth path onto the circular orbit. If the user decides to take a 90-degree right or left turn in the virtual scene (paths 2 and 3), her 90-degree turn becomes something like a 45-degree turn in the lab. After the turn, the user is again redirected onto the circular orbit.

It is conceivable that the user could take a turn in the virtual scene that happens to lead her away from the lab center (Figure 6.8), resulting in a problem situation similar to that with the Steer-to-Center algorithm. However, I expect this situation would occur less often in the Steer-onto-Orbit algorithm. In the Steer-to-Center algorithm, this problem situation is a direct result of the algorithm being successful, whereas in the Steer-onto-Orbit algorithm, it happens only by chance. With Steer-to-Center, it happens as the user is walking in a straight line through the center. But with Steer-onto-Orbit, it happens when the person is turning away from the orbit path. While the user is turning away, the system can amplify her angular velocity in the virtual scene (by rotating the virtual scene in the direction opposite to that in which she is turning) to reduce her turning away from the lab center. When she is pointed away from the lab center, the situation can be addressed in the same manner as with the Steer-to-Center algorithm (by biasing the steering in one direction). In fact, given that the podokinetic system (Chapter 4) can be fooled by having the user turn consistently in one direction for several minutes, it may be useful to guide the user onto a circular orbit in a particular and consistent direction (e.g., always clockwise or always counter-clockwise).




Figure 6.8 – Steer-onto-Orbit algorithm: The user could happen to take a turn such that her path in the lab has her pointing directly away from the lab center. I expect this to happen less often than with the Steer-to-Center algorithm.

6.3 Proposed Algorithm: Steer the User Toward Changing Targets

Another approach to avoiding the situation where the user is made to walk through the lab center and then heads directly away from it is to steer the user toward changing targets in the lab instead (Figure 6.9).




Figure 6.9 - Steer-to-Changing-Targets algorithm: Left: The system is steering the user toward target A. Right: Once she has reached it, the system then selects target C (because the user's heading is pointed closer to target C than A or B) and steers the user toward it.

At any given time, the system tries to steer the user toward a particular fixed target (all of which are centrally located but also spaced sufficiently apart). Once the user walks through that target and is pointed directly away from it, the system chooses another target and then steers the user toward that one. The system ensures that the user is never pointed directly away from the target to which the system is trying to steer her. For example, the new target which the system chooses must not be collinear with the user and the immediately previous target (Figure 6.10). If the user unexpectedly turns in the virtual scene such that she is now facing away from the current lab target, the system then selects a different, more convenient, target and steers her toward that one.




Figure 6.10 - If the user is steered through target A and then happens to be facing directly away from both targets A and C, the system must not choose C as the next target, as it would be just as problematic as steering her toward target A (which is the very problem Steer-to-Changing-Targets was designed to solve).

This Steer-to-Changing-Targets algorithm bears some resemblance to algorithm used for Redirected Walking with waypoints (Chapter 7). The critical distinction is that the targets exist only in the physical lab

space, and there are no waypoints in the virtual scene—the user is free to walk along any arbitrary path or direction she chooses in the virtual scene.

6.4 Guidelines for Designers of Steering Algorithms

In summary, any algorithm to steer the user during unrestricted exploration of the virtual scene should not assume that the lab center is the optimal place in the lab. On the other hand, the algorithm should be able to accommodate head wobble and tracker noise. Additionally, I believe it to be useful to assume the user will continue along her current heading, but the algorithm must be able to handle her unexpectedly changing heading. Finally, in real life, people often sidestep obstacles (e.g., a fire hydrant on the sidewalk). In the RWP experiments, I observed several subjects having difficulty getting around obstacles (they could only walk in the direction in which they were looking). I have not attempted to determine how often people sidestep or what the steering algorithm should do during sidestepping. Nevertheless, the system designer should consider this if the virtual scene is to have many obstacles (e.g., a restaurant crowded with tables and chairs).

Chapter 7: The Redirected Walking Experiment: RW

This chapter contains the details of the experimental design, methods, and observations from both experiment RW, which was an institutional-review-board-approved study on naïve subjects, and the various pilot sessions, in which colleagues participated as test users (i1, i2, i5 and RWp, from

7.1.1 *RWp and RW*

The purpose of this user study was to determine the viability of Redirected Walking with waypoints and spatial audio. I tested the technique on a single group of participants who were instructed to complete a fire-drill task in the virtual scene pictured in Figure 1.4. Observations from the study suggest this technique works: Redirected Walking causes users to change their walking direction without noticing and enables larger VEs while providing the benefits of real walking. The subjects did not know about Redirection, were not familiar with the size of the lab, and were led into the lab blindfolded. Subjects were surprised, after completing the task and removing the headset, to find that the real lab was much smaller than the virtual scene.

Table 1.2). This chapter also discusses the use of spatialized audio. This was the first experiment in our laboratory to do so, and the only such experiment included in this dissertation work.

The purpose of experiment RW, conducted by Zachariah Kohn and me, was to investigate the viability of Redirected Walking with waypoints. Can users carry out a task that requires them to walk a path that does not fit in the laboratory?

7.2 Task and Virtual Scene

The task subjects performed was a simulated fire drill. Subjects were immersed in a virtual brick room approximately twice the length and width of the 10- by 4-meter tracked area. Four buttons mounted on the virtual walls served as waypoints (Figure 7.1). Subjects were asked to visit and, using a tracked hand-controller, push the virtual buttons, in a particular order. Each button had a label, a purpose in the fire-drill scenario, and an auditory and/or visual response (Table 7.1).



Figure 7.1 - Left: A user's view in the headset as she walks toward the button to sound the alarm. Right: A view of the entire virtual room (the front wall is removed for clarity).

Table 7.1 - Description of labels, scenario-related purpose, and VE system response of each virtual wall-mounted button..

<i>Sequence</i>	<i>Label</i>	<i>Purpose in scenario</i>	<i>Response</i>
1	"Practice"		Made a clicking sound (which emanated from the button) to confirm to the user that it was pushed
2	"Alarm"	Sound the Alarm	Started a loud, mechanical, ringing-bell sound emanating from the ceiling, whose intensity faded away after several seconds so that the subject could hear subsequent noises
3	"Window"	Close Windows	Moved the window glass and frames down to the closed position, while playing a motor and gears whirring sound, followed by the sound of latch closing when the window glass reached the closed position
4	"Halon"	Activate the fire suppression system	Started a hissing noise emanating from the ceiling

The buttons were located eight meters apart in both the virtual scene and in the lab; subjects had to really walk in order to virtually locomote from one to another. After pushing all four buttons, subjects were instructed to leave the virtual room through the doorway. The path took the subject through the virtual room in a zigzag pattern. The subjects had to stop at each waypoint to push the buttons and were instructed to walk calmly. They were instructed not to wander aimlessly about the room but to look around to locate the next button before walking toward it.

7.3 Subjects

Eleven subjects participated in experiment RW. Subjects were at least 18 years old, in their normal condition of good health, without having consumed alcohol or cold medicines, without a history of epilepsy,

able to communicate in English, able to walk without assistance, and with normal vision and hearing (in both eyes and both ears). Most importantly, subjects were selected such that none were familiar with our laboratory or had even visited the building in which it was housed, so that they would not know (nor be able to infer) the size of our laboratory.

Subjects were paid \$10 per hour (each subject participated for roughly one hour) and were told they could withdraw from the experiment at any time and would still receive the \$10.

7.4 VE System Details

For this study, subjects wore a Virtual-Research™ V8 HMD, with a 60-degree diagonal field of view and a 4:3 aspect ratio. My colleague and I added a black cloth veil to the HMD to prevent the subjects from seeing the laboratory. Stereo visual imagery was generated at 30 frames per second using one graphics pipe and one processor of an SGI Onyx2 Reality Monster™ computer. A wide-area optical tracker provided position and orientation of the user's head and right hand. This tracker was a custom-built predecessor of the 3rd Tech Hiball 3000 system. Ceiling-mounted LED fiducials were sited at roughly 400 Hz per sensor, and Kalman-filtered position and orientation reports were generated at 70 Hz. The end-to-end latency, including tracker filtering, network delays, and image generation, was measured to be between 50 to 115 ms (the average was roughly 80 ms). Spatialized audio was generated by an Aureal AU8830A2 A3D 2.0 processor sound card in a Dell PC. The audio was presented through Sennheiser HD250 II sealed, circumaural headphones.

7.5 Redirection Algorithm

The Redirection algorithm used in this experiment used the user's linear and angular velocity as input parameters. The tracking system reported only position and orientation, so a first difference (velocity = (most recent position sample – previous position sample) / time difference between samples) was used to compute the velocities. However, differencing doubles the noise already present in the tracker's position and orientation reports, resulting in a velocity measures that were too noisy to use. To address this problem, position and orientation reports were box-filtered over the four most recent reports.

The algorithm employed three separate components of rotational distortion. During development (experiments i1, i2, and i5), Kohn and I discovered that even while the user was standing still, the system could slowly rotate the virtual scene and the user unwittingly turned in the same direction without noticing. To

exploit this, the Redirection algorithm used in this study injected a small, baseline amount of constant rotational distortion, even when the subject was standing still. Second, it used a component of rotation related to the user's walking speed. Third, when the user turned herself, a higher-frequency motion, it injected additional rotation proportional to the user's angular velocity. The rotational distortion injected in any frame was the maximum of the three components: constant rotation, rotation proportional to the user's linear velocity, and rotation proportional to the user's angular velocity.

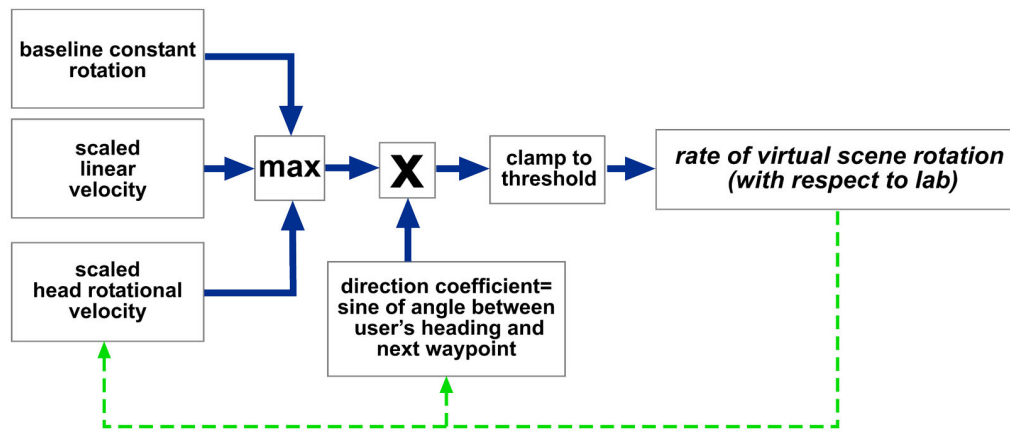


Figure 7.2 - A flow diagram of the Redirection algorithm used in experiment RW. The dashed green line represents the feedback that occurs via the user turning herself in response to the rotation of the virtual scene.

The system scaled this distortion rate by a *direction coefficient*. The direction coefficient was a measure of how much and in which direction the system needed to steer the user. This coefficient was dynamically calculated by computing the sine of the angle between the user's direction in the VE and the direction the system desired her to take in the lab.³⁴ As implemented in the study, the desired direction was the direct path toward the next waypoint (the virtual button toward which the subject was currently walking). Finally, the system compared the scaled rate to a threshold for imperceptible angular distortion. If the distortion

³⁴ The sine function was arbitrarily chosen, and I do not claim it is the ideal function. It varies between 0 and 1, and changes smoothly as the user changes heading. This results in something similar to a *proportional control system*. Prior to using this function, I experimented with a *bang-bang control system*, and those results are described in Chapter 6.

rate exceeded the threshold, it was clipped to the threshold value. The threshold was set to the rate that seemed imperceptible to all of the algorithm testers in the pilot experiment RWp.

The user zigzagged through the virtual scene, walking from one wall-mounted button to the next. After pushing a button, the subject turned to see the next button. As she turned herself to see the next button, the system injected larger amounts of distortion by scaling the user's rotation rate. After the subject turned, the next virtual button was almost lined up with the farthest wall of the lab. Any small misalignment that remained was then made up once the subject started walking, by applying rotational distortion proportional to the subject's walking speed. This yielded the arced real paths seen in the lower left portion of Figure 7.3.




Figure 7.3 - Left: Overhead views of the path taken by the user in the virtual scene (above left, in blue) and the laboratory (below left, in red). The user walked in a zigzag pattern through the virtual scene while walking back and forth within the tracker space. The tracker space and virtual scene are to scale. Crosses denote waypoints. Right: The user's path superimposed onto the virtual scene.

7.6 Observations and Lessons Learned

Users were able to complete the fire-drill task. All of them were surprised, upon removing the headset, to learn they were walking back and forth between the ends of the lab rather than zigzagging thru it, as they had in the virtual scene. This result demonstrates the viability of Redirected Walking.

Beyond this primary observation, some others merit discussion.

7.6.1 *The HMD Veil Increases User Discomfort*

A veil was hung from the cowl of the HMD, in order to prevent the user from seeing the laboratory. This covered the user's face and was made of a heavy, black, velvet-like material. One user became uncomfortable and asked to stop the session while walking from the virtual "practice" button toward the "alarm" button, even before the virtual scene began to rotate. During the debrief interview, she revealed that she began to feel ill even before the experimental session began. We recalled that the previous subject was

wearing a strong perfume. Upon donning the headset ourselves, we observed it was very stuffy, smelled strongly of perfume, and made breathing difficult. We suspect this is the cause of this subject's sickness. Other subjects also reported the headset being stuffy.

7.6.2 *Redirection's Sensitivity to Tracking Glitches*

In several sessions during the pilot experiment (RWp), the tracker lost acquisition of the user's position. When this happens, this particular tracker model continues to report the last known position of the user. When the tracker reacquires the user's position, it then begins to report the user's current position. Thus, when the tracker loses acquisition, the system believes the user has stopped moving and the visual imagery becomes still. When the tracker reacquires, the system believes the user has instantaneously jumped to a new position (from the old position where the user was when the tracking was first lost).

Much to our surprise, many users continued walking for several seconds when the tracker stopped updating, despite the fact that the visuals were no longer updated. For example, if the user was three meters from a virtual painting when the tracker failed, she sometimes continued to walk toward the painting, even though the painting did not appear to get any closer. (I have observed this phenomenon in subsequent VE systems that do not use Redirection.) During this time, the Redirection algorithm cannot steer the user, and she risks walking into a position from which the algorithm can no longer recover (e.g., to the lab boundary).

7.7 Spatial Audio

7.7.1 *Motivation*

In previous VE systems in our laboratory, users reported breaks-in-presence from hearing noises from the real laboratory (e.g., footsteps, people talking, etc.) while they were in the virtual scene. Even worse, when experimenters spoke to the user, the user would often turn to the source, a disembodied voice speaking to them [Usuh 1999]. Kohn and I worried that, during Redirection, the laboratory noises and experimenter voices would not only disrupt the user's sense of presence, but also provide her cues about her real orientation.

The motivation for implementing a spatial-audio virtual scene in addition to the visual virtual scene was to 1) shut out real-world sounds and 2) increase the user's immersion in the virtual scene and strengthen Redirection's ability to fool her, by having an additional consistent, controlled cue of the user's orientation in the virtual scene.

7.7.2 *Sound Cues*

In addition to sounds that were triggered by the wall-mounted buttons, we also included other background noises. These included traffic and bird noises from outside the virtual windows, and fan noises from air vents in the room.

Rather than allowing experimenters to speak directly to the subject while she was in the virtual scene, we prerecorded instructions and commonly needed phrases such as “Please do not run.” These recordings appeared to emanate from antique radios placed throughout the virtual scene (Figure 7.4), when an experimenter triggered them via a wireless keypad worn on the experimenter’s torso.



Figure 7.4 - A user’s view in the headset as she walks toward the button to close the windows. An antique radio, used for presenting pre-recorded instructions, is in the foreground.

From post-session interviews, we found that spatial audio masks real-world noise very successfully—none of the subjects reported hearing noises from the laboratory. Our observations of the benefits of the sounds came not from the users’ comments, but rather from the lack of them. We believe that, when supporting audio cues are designed properly, the user often does not notice them. But if the sound cues are made louder or otherwise more obvious, they become unrealistic and the user notices this. On the other hand, if the sound cues are missing, the user notices their absence or reports hearing distracting noises from the real world. Our speculation is based on observations from pilot testing (RWp) and from the development of subsequent VE systems with spatial audio.

7.7.3 Earphones

To block laboratory noises, we chose circumaural earphones, which completely cover the user's ears. At the time of the experiment and of this writing, circumaural earphones attenuate outside noise better than noise-canceling headphones, particularly for non-periodic noises such as door noises and speech.

7.7.4 Spatial Audio Algorithms

The Aural spatial audio hardware implements wave-tracing of sound and the head-related transfer function (HRTF). Once our implementation was working properly, we found the combination of this system and the acoustic model of the virtual scene to be very convincing. At one point during development, Kohn secretly changed the source of a recorded instruction from a table-top radio to a ceiling vent. Upon hearing this recorded instruction from within the virtual scene, I immediately perceived it was coming from the ceiling air vent and exclaimed, "Someone's stuck in the vent!"

There are numerous spatial audio products and, in my experience, they each have vastly different technique and quality. For example, our audio implementation (in later VE systems) using a Creative Audio EAX™ product (which does not use wave-tracing) was not able to reproduce my experience with the voice coming from the ceiling vent. The implications of the fidelity of the spatial audio system are discussed in Chapter 12.

Chapter 8: The Redirected Walking-in-Place Experiments: RWP

8.1 Overview

The results of the RWP experiments are summarized in Chapter 1. This chapter contains the details of the experimental design, methods, results, and observations from experiments RWP-I, RWP-II-p, and RWP-II from

8.1.1 *RWp and RW*

The purpose of this user study was to determine the viability of Redirected Walking with waypoints and spatial audio. I tested the technique on a single group of participants who were instructed to complete a fire-drill task in the virtual scene pictured in Figure 1.4. Observations from the study suggest this technique works: Redirected Walking causes users to change their walking direction without noticing and enables larger VEs while providing the benefits of real walking. The subjects did not know about Redirection, were not familiar with the size of the lab, and were led into the lab blindfolded. Subjects were surprised, after completing the task and removing the headset, to find that the real lab was much smaller than the virtual scene.

Table 1.2.

My colleagues at UCL and I carried out experiments in order to test a hypothesis on several variants of RWP: RWP results in a lower frequency of the open back CAVE wall coming into the subject's field-of-view than turning with a hand-controller, and users do not notice the injected virtual scene rotation.

From the results of experiment RWP-I, it is clear that RWP (as implemented in RWP-I) did not meet its objective of having users see the back wall less while not noticing the virtual scene rotations. Furthermore, the implementation of walking-in-place was also troublesome. Based on our observations and subjects' comments about RWP-I, the RWP technique was revised. We tested this new RWP implementation with a six-subject pilot study (RWP-IIp) and then conducted another full experiment (RWP-II) to verify RWP efficacy. I explain both the original and revised algorithms in this chapter.

8.2 Motivation

The most common method of locomotion in CAVEs is to fly using a hand-controller (i.e., joystick or wand). Many users have trouble adapting to this interface and find it distracting [Usuh 1999]. Flying with a joystick results in a lower sense of presence than walking-in-place [Slater 1995]. Holding and manipulating the joystick is also an encumbrance, since the user can no longer use that hand for other tasks.

When the user walks-in-place she moves in the direction in which her head or torso (depending on the implementation) is pointing. Walking-in-place allows a user in a virtual scene to move through the virtual world, including turning in any desired direction using her body. This is highly problematic in a three-walled CAVE, because users will invariably turn such that they notice the blank wall. The goal of RWP is to allow her to virtually walk in any direction, even in complete circles, in the virtual scene and never see the blank wall.

Traditionally, if a user wishes to move toward an object in a virtual scene, she must first rotate the virtual scene using a joystick or other hand-controller so that the virtual object is in front of her. This is unnatural (one never rotates the world in real life) and causes a mismatch between proprioceptive and visual cues; previous research shows a positive correlation between appropriate body movement and increased presence [Usuh 1999] (i.e., a user who turns her body in order to rotate herself in a virtual scene is more likely to feel present than one who uses a joystick to rotate the world). With RWP, the goal is to enable the user to turn in the virtual scene by turning her body instead of using a joystick, while also reducing the proportion of time she sees the open back wall.

RWP works by interactively and imperceptibly rotating the virtual scene about the user. This rotation causes the user to continually turn toward the front wall of the CAVE (Figure 8.1).

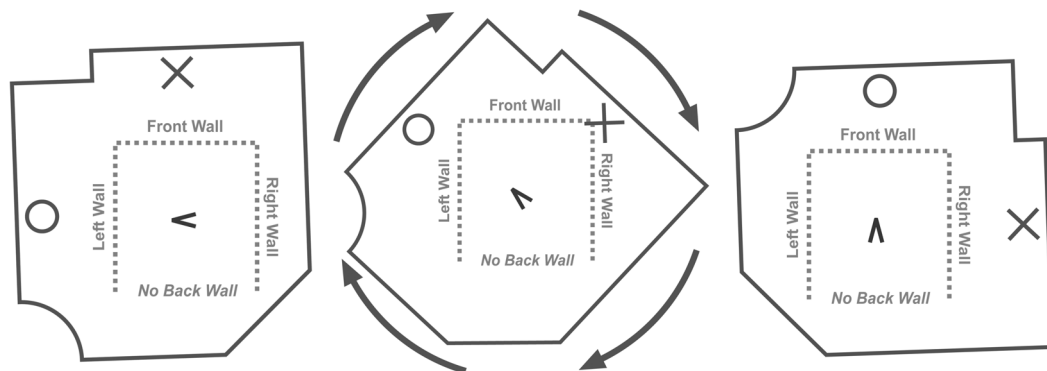


Figure 8.1 - An illustration of how RWP works. Left: The user (the arrowhead in the center) turns to face the circle. Center: The system responds by slowly turning the virtual scene to the right. Right: Such that the circle is now behind the front wall.

8.3 Virtual Scene and User Task

The virtual scene in these experiments was the same brick room as the one as used in the RW experiments, but the task was different. Instead of visiting each wall-mounted virtual button in a particular order, the task required the subject to explore the room. It involved the four yellow signs: “Alarm”, “Halon”, “Practice” and “Window,” Subjects were asked to find, approach, and read all four signs, then to revisit³⁵ them in alphabetical order. This task forced the subject to walk about and explore the large virtual room and was specifically designed to involve many substantial changes of direction (Figure 8.2). Before beginning the task, subjects were familiarized with the VE equipment and practiced walking-in-place. The total virtual scene exposure was approximately 10 minutes.




Figure 8.2 - The path in the virtual scene taken by one subject in the Redirection group.

³⁵ Subjects were instructed to stand in front of the sign, facing towards it, instead of pushing the wall-mounted button underneath the sign, as subjects did in the RW experiments. Subjects’ hands were not tracked in the RWP experiments.

8.4 VE System Details

The CAVE-like system used in this experiment was a Trimension ReaCToR™ with four projection surfaces (three vertical walls and the floor). A SGI Onyx2, using four graphics pipes (one per screen) and 5 processors, generated imagery at 22.5 frames per second for each eye. Subjects wore Crystal Eyes™ shutter glasses to view sequential stereo imagery. The refresh rate of the four cathode ray tube (CRT) projectors was 90 Hz (45 Hz in each eye). An Intersense™ IS-900 tracker provided the position and orientation of the subject's head and torso at 180 Hz. The IS-900 wand, which is normally held in the user's hand, was attached to the subject's waist with a hip-worn camera bag to track the subject's torso orientation (Figure 8.3). For a hand-controller, subjects held a Logitech™ wireless computer mouse. In the control condition, where the user turned using the mouse, pushing the right button rotated the virtual scene to the right. Similarly, pushing the left button rotated the virtual scene to the left. Both groups of subjects wore the same torso tracking equipment, and both groups carried the hand-controller, even though it was not used by the subjects who used Redirection.



Figure 8.3 - The hand-tracking sensor attached to a hip-worn camera bag in order to track the torso orientation.

Both groups moved forward in the virtual scene by walking-in-place. Experiments RWP-I and RWP-II used different techniques for detecting when the user was stepping. I describe both techniques in detail below. Regardless of the detection technique, when the system detected that the user was stepping, it moved the user's viewpoint in the virtual scene in the direction in which her *torso* was pointing.

8.5 Users

We recruited 44, 6, and 30 people for experiments RWP-I, RWP-Iip, and RWP-II respectively, from around the UCL campus by advertisement, and paid them \$7.50. They were randomly assigned to the control or experimental group. Subjects were asked to carry out a task in a virtual scene. The control group turned in the virtual scene using a hand-controller and the experimental group used RWP. Apart from the turning

method, the task and equipment were the same—both groups completed the same task in the same virtual scene and both used walking-in-place to move. The final allocation of subjects to conditions is reported in

Table 8.1.

Table 8.1 - Number of subjects for whom data was collected for each experiment and condition.

Experiment	Hand-controller Control Group	Redirection Experimental Group
RWP-I	13	15
RWP-IIp	3	3
RWP-II	12	14

Due to loss of data (from equipment failures), the final allocations were 13 (RWP-I) and 12 (RWP-II) people to the control group, and 15 (RWP-I) and 14 (RWP-II) people to the Redirection (experimental) group. These experiments were approved by both the UNC Institutional Review Board and the Joint UCL/UCLH Committees on the Ethics of Human Research.

8.6 Experimental Measures

The experimental variables used in the analysis were as follows:

i) *Saw_back_wall*: This is a measure of how often during the session the open back wall was within a 40-degree field-of-view of the subject. Similar measures were taken at varying fields-of-view (2, 20, 65, 90 and 106 degrees). *Saw_back_wall* was computed on a frame-by-frame basis and is reported as a percentage—the number of frames where the back wall was in the field-of-view divided by the total number of frames during the task.

ii) *rotate*: In addition to this objective measure of rotation, we included a question that assessed the extent to which people actually noticed whether the room was rotating. In order not to alert subjects to this possibility, the question of whether they had noticed the room unexpectedly rotating was embedded among a series of similar questions, such as whether they noticed the virtual scene flickering, getting brighter or darker, or changing size.

iii) *assq*: Each subject filled out the Simulator Sickness Questionnaire (SSQ) [Kennedy 1993] immediately before and after her experimental session. The SSQ is designed for use only after the VE

exposure. We administered the SSQ before the VE exposure only to detect subjects who did not meet the requirements for participating in the experiment.

iv) *pres*: Self-reported presence—this was assessed by six questions in the post-session questionnaire, following exactly the format used on several previous and subsequent occasions [Slater 1999; Usoh 1999; Slater 2000; Meehan 2003; Zimmons 2004]. The six questions are listed in Table 8.2. A higher score indicates greater reported presence. The overall score for a subject is the number of high scores among the six questions, where a high scoring question is any question to which the subject answered with a 6 or 7. Hence the overall score is a count variable (ranging from 0 to 6) and is treated as a binomial response variable in a logistic regression.

Table 8.2 – The six questions from the presence questionnaire used in the RWP experiments.

I had a sense of "being there" in the brick room	[1. not at all ... 7. very much]
There were times during the experience when the brick room was the reality for me	[1. at no time ... 7. almost all the time]
The brick room seems to me to be more like	[1. images that I saw ... 7. somewhere that I visited]
I had a stronger sense of	[1. being in the lab... 7. being in the brick room]
I think of the brick room as a place in a way similar to other places that I've been today	[1. not at all ... 7. very much so]
During the experience I often thought that I was really standing in the brick room...	[1. not very often ... 7. very often]

v) *sdhead*: We also measured how much a subject turned her head and torso while carrying out the task (standard deviation of head and torso orientation over the course of the virtual scene exposure). Previous studies [Slater 1998] have found that presence is positively correlated with such body movement. In this situation, though, when users turn their heads or their bodies, the virtual scene (ideally without the user noticing) rotates to compensate. However, if this rotation is noticed, one would expect that this would decrease presence, since it conflicts with everyday experience. Therefore, there is potentially a complex relationship between *saw_back_wall* and the amount of head and torso rotation (*sdhead*). It turned out, as expected, that head and torso movement are almost perfectly correlated ($R^2= 0.98$ in RDW-I and $R^2= 0.95$ in RDW-II), so in subsequent discussion I refer only to head rotation.

8.7 Experiment RWP-I

The major hypothesis for which this first RWP experiment was designed was:

(a) *saw_back_wall* is lower for the RWP condition than for the control condition.

The secondary hypotheses which we hoped this experiment would illuminate were:

(b) Users in the RWP condition would not report noticing the virtual scene rotations any more than in the control condition. Also, users in the RWP condition would not report noticing rotations more than other (non-existent) phenomena.

(c) RWP does not significantly increase the users' level of simulator sickness (*assq*) compared to turning with a hand-controller.

(d) There is a relationship between presence and locomotion technique.

The results for each of the above were:

(a) There was no significant difference in the mean values for *saw_back_wall* between the two conditions. In the control condition the mean value was $8.4\% \pm 13.7\%$ and in the Redirection condition it was $11.2\% \pm 5.1\%$. In other words, the implementation of RWP used in this experiment did not result in a decreased frequency of looking toward the blank wall. However, the variance for the RWP condition was significantly lower than for the control condition ($p < 0.0005$).

(b) Subjectively, the number of subjects in the RWP condition who noticed that the world was unexpectedly rotating was much higher (7/15) than for the control group (1/13). I have two suspicions for why this happened, and these are detailed below in Section 8.8.1.1. For all the other such variables (virtual scene flickering, changing size, etc.) the results were evenly distributed between the two conditions.

(c) There was no significant difference between the conditions regarding simulator sickness (*assq*). The means of the SSQ scores are 11.8 ± 13.2 and 10.2 ± 8.5 , respectively. The SSQ produces scores between 0 and 100. As described in Chapter 10, Kennedy suggests that a SSQ score above 15 is cause for concern, and a score above 20 indicates a problem simulator [Kennedy 2000].

(d) Presence did not significantly vary between conditions. This is no surprise, given that the subjects in the RWP group noticed the back wall as much as the control group and noticed the virtual scene rotations.

We also found also a relationship between reported presence, *pres*, and *saw_back_wall* when we took into account the amount of head rotation. This result is consistent with earlier ones from Slater’s laboratory—the more a subject turned her head or torso, the higher her sense of presence (other things being equal). As expected, the more a subject noticed that the virtual scene rotated, the lower her reported presence. Finally, the more the open back wall came into her (40-degree) field-of-view, the lower her sense of presence. Table 8.3 summarizes the resulting model that predicts a user’s sense of presence as a function of how much she noticed the rotations, how much she saw the back wall, and how much she turned her head.

Table 8.3 - A model that predicts a user’s sense of presence as a function of how much she noticed the rotations, how much she saw the back wall, and how much she turned her head. The coefficient column shows the parameter estimate for the corresponding variable in the logistic regression analysis, and the S.E. column shows the standard error of the estimate. The χ^2 column shows the chi-squared value for deletion of the corresponding variate from the model. This should be compared with the tabulated 5% value of 3.841 on 1 degree of freedom. In other words, no variable can be removed from the model without significantly worsening the overall fit.

<i>Variable</i>	<i>Coefficient</i>	<i>S.E.</i>	χ^2
rotate	-1.6	0.68	6.4
saw_back_wall	-0.17	0.061	8.6
sdhead	0.029	0.0093	11.1

8.8 Problems with the RWP Implementation Revealed in RWP-I and Rectified in RWP-II

From the results of experiment RWP-I, it is clear that RWP did not meet its objectives. Based on our observations and subject reports from RWP-I, the RWP technique was revised.

8.8.1 Redirection Algorithms

8.8.1.1 ...as used in RWP-I

The Redirection algorithm for these experiments was very similar to that used in the RW experiments. Redirection works by continuously injecting rotational distortion. As illustrated in Figure 8.5, there were three components that contributed to the virtual scene’s rotation rate: a baseline rotation rate that dominated when the user was standing still and not turning her head, a higher rate that dominated when the user was walking-in-place (and not turning her head), and a rotation rate proportional to her head’s rotational velocity.

The system took the maximum of the above three components, and then scaled it by a *directional coefficient*, to rotate the virtual room such that the subject was made to turn smoothly toward the front wall. This coefficient was calculated by computing the sine of half the angle Θ between the subject's torso orientation in the CAVE and the front wall of the CAVE (Figure 8.4). I chose the sine function because its value smoothly changes as it crosses zero when the user is directly facing the front wall. Half the angle Θ was used so that the further the user turned from the front wall, the greater the directional coefficient. As described in Chapter 6, this proportional control system prevents the virtual scene from appearing to vibrate (which was a problem in the first implementations of Redirected Walking).




Figure 8.4 - *Theta* is the angle between the user's torso heading and the front CAVE wall.

We observed two problems with this algorithm. First, the rotation rates were not high enough to prevent the user from turning toward the back wall. Once she saw the back wall for some period of time, it became obvious that the virtual scene was rotating (as virtual objects scrolled past the edge of the CAVE walls into the darkness). Second, when a user turned her head to look over her shoulder and then turned back, while keeping her feet stationary, the virtual object that was previously directly in front of her would have rotated significantly. Since the user's feet stayed firmly affixed to the floor, she was able to detect this rotation.

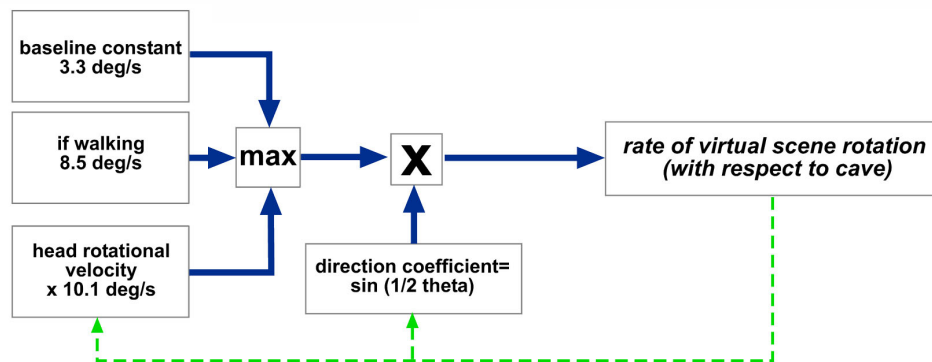


Figure 8.5 - The RWP-I algorithm. The green dashed line represents feedback via the user turning herself in response to the virtual scene rotating.

8.8.1.2 ...as used in RWP-IIp and RWP-II

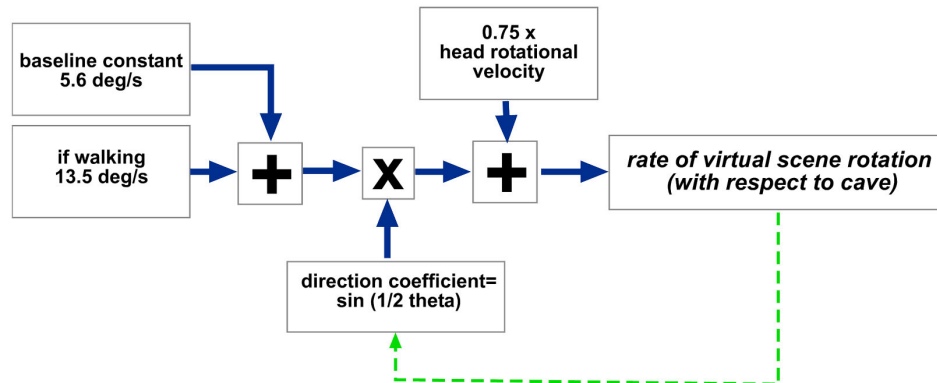


Figure 8.6 - The RWP-II algorithm.

Several changes were made to the algorithm. The baseline and walking rotation rates were increased. The system added these two components (instead of taking the maximum of them) before scaling by the directional coefficient. Finally, when the user turned herself (a higher-frequency motion), the system amplified her virtual angular head velocity. When the user slewed her head (i.e., she quickly turned her head to look over her shoulder and then quickly turned her head back), objects in the virtual scene did not rotate any more than if she had not turned her head (since the scene rotation caused by her turning her head one way canceled out the scene rotations caused by her turning her head back to its original position). This last mode of rotation is similar to LaViola's Auto Rotation technique [LaViola Jr. 2001].

8.8.2 Walking-in-Place Detection

8.8.2.1 Neural-Network Detection from Head Position

In earlier studies and in experiment RWP-I, a neural network detected when the user was walking-in-place, from head-position data. This technique has the great advantage of not requiring any hardware beyond what is normally used in VE systems. However, the neural network requires training, and its success depends entirely on how well the training data (usually the manually annotated tracking data from one person) matches the style of walking of each particular user. In experiment RWP-I, there was one subject for whom the neural network did not work at all. Other subjects found it very difficult to use. Users would try to increase the correct detection probability by stepping harder, but this rarely worked. Other subjects found that it occasionally reported stepping when the user did not step.

Furthermore, the neural network required a full stepping cycle to determine if the user had just started or stopped walking-in-place. If, as required in the experimental task, the user walked up to a virtual wall and then tried to stop just in front of it, the neural network would not detect that she had stopped until roughly half a second later, and by then she would have penetrated the virtual wall. Walking through a virtual wall is a break-in-presence event and, in the case of this experiment, even prevented the subject from completing the task (because she was stuck outside the virtual room, hanging in empty virtual space).

8.8.2.2 Detection of Foot-Strikes from Accelerometer Signal

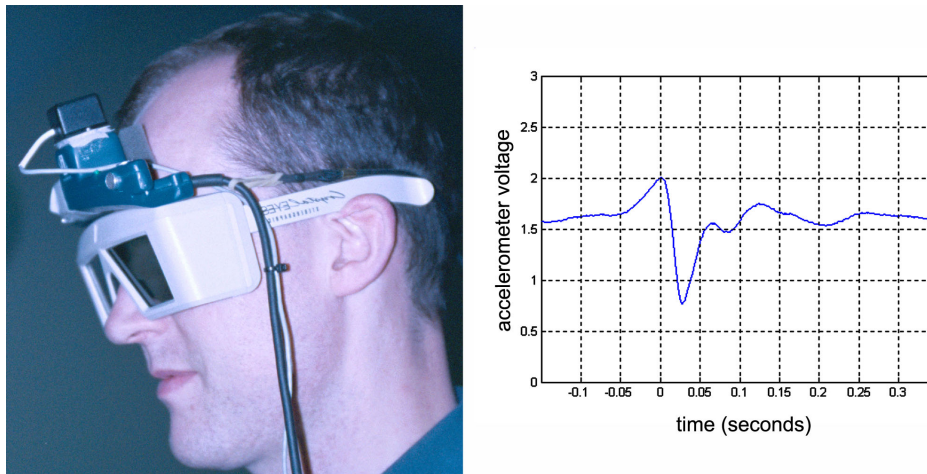


Figure 8.7 - Left: The accelerometer for detecting footsteps (the black box with the white wire) was attached to the top of the blue head-tracking sensor. Right: A sample footstrike as recorded by the accelerometer. The vertical axis is voltage, where 1.5 v corresponds to 1 g.

An alternative approach was used in experiment RWP-II. For this technique, a Crossbow™ solid-state accelerometer was mounted to the tracker's head-sensor (Figure 8.7) [Kohn 1999]. When the user's foot struck the CAVE floor, the vibration was detected by the accelerometer, with only a few milliseconds of latency. This technique, although it required additional hardware and an extra cable, worked much better. It eliminated the latency and false-detection problems associated with the neural network technique. Additionally, when a user found that it missed detecting a step, she would step harder, and this actually improved detection. Furthermore, the user could increase her virtual walking speed by increasing her step rate. This was not possible with the neural network, as it reported only the binary presence or absence of head-bobs. One flaw with the accelerometer technique as we implemented it was that it did not work when the user looked down at her feet,

because the accelerometers axis was not aligned with the direction of the footstrike-induced vibration. This problem can be addressed by looking for characteristic vibrations in all three axes of the accelerometer.

8.8.2.3 Problems with the User Creeping Forward

Both the neural network and the accelerometer-based techniques for walking-in-place had the drawback that, occasionally, users would unknowingly creep forward when walking-in-place. This would cause some of them to eventually run into the front wall of the CAVE. Other users unknowingly crept forward without actually hitting the front wall, but then reported that the imagery became very blurry. This was because they were very close to the front wall, and the number of video pixels in their field-of-view decreased dramatically.

8.9 Experiment RWP-II Results

After making substantial revisions to the RWP algorithm, we conducted experiment RWP-II to test the hypotheses:

(a) *Saw_back_wall* is lower for the RWP condition than for the control condition.

(b) Users in the RWP condition would not report noticing the virtual scene rotations any more than the control condition. Also, users in the RWP condition would not report noticing rotations more than other (non-existent) phenomena.

(c) RWP does not significantly increase the users' level of simulator sickness (*assq*) compared to turning with a hand-controller.

(d) Presence decreases with higher *saw_back_wall*.

(a) In the control condition the mean value for *saw_back_wall* is 4.1 ± 4.6 %, and in the RWP condition it is 1.7 ± 3.0 %. These values do not differ significantly. However, the degree to which people rotated their heads varied considerably within each of the two turning methods. Some users simply turn their heads more than others, even without the use of Redirection. The mean standard deviation in head rotation in the control group was 50 ± 30 degrees and in the RWP condition 50 ± 16 degrees. One would expect that the more that users rotated their heads, the greater the chance of seeing the back wall. We therefore used the amount of head rotation, *sdhead*, as a covariate in order to take into account this confounding factor.

The result is shown in Figure 8.8. This shows the scatter plot of *saw_back_wall* by *sdhead*, and the regression line shows a significant difference in intercept ($t = -2.6$ on 23 d.f., $p < 0.05$), and the overall fit of the model is high ($R^2 = 0.68$). From the figure, one sees that there is one outlying point where the subject rotated his head far in excess of anyone else. (This data point is identified as a formal outlier using a statistical technique called *leverage analysis* [Pregibon 1981]). Eliminating this point improves the fit (although the fit is good even with this point). The results suggest that for any given level of head rotation, RWP does, on the average, result in fewer turns to the open back wall than the traditional turning with a hand-controller.




Figure 8.8 –Regression lines and actual data points, showing how much subjects saw the back wall, as a function of how much they turned their head, and which experimental group they were in. The Redirection group (magenta) saw the back wall less than the mouse-turning group (black). Data points are shown as black diamonds (for the hand-controller turning group) and magenta triangles (for the Redirection group). The outlying data point is excluded from the trend line.

(b) Did those subjects who used RWP notice that the virtual scene was rotating? In the post-experimental questionnaire, we asked this as a sub-question embedded in an overall question: “During the time of your experience, which of the following happened unexpectedly? Circle yes or no for each item.” The items and results are listed in Table 8.4.

Table 8.4 – The questions used to determine if the subjects noticed that the virtual scene rotated, compared to other phenomena which did not happen. The aggregate responses for each group are listed in the right-hand columns.

<i>These things happened during my experience</i>	<i># Yes: Control</i>	<i># Yes: RWP</i>
The brick room became larger or smaller	3/12	4/14
Objects disappeared and reappeared	2/12	1/14
Parts of the brick room got brighter or dimmer	6/12	5/14
The brick room rotated	3/12	5/14
Parts of the brick room flickered	2/12	3/14

There is no significant difference between the control and RWP groups in how many subjects reported unexpected rotation. There was also no significant difference between the ratio of subjects reporting rotation and those reporting other phenomena which did not happen (such as the room changing size). This suggests that the subjects did not notice the rotational distortion induced by Redirection.

(c) There is no significant difference between the conditions with respect to simulator sickness. The mean SSQ scores are 18.0 ± 21.7 for the control group and 9.5 ± 6.9 for the RWP group. However, the mean scores for the mouse-turning group are *higher* than for the Redirection group. Furthermore, the 75th percentile scores (the metric Kennedy proposed in the original SSQ paper [Kennedy 1992]) for the mouse-turning group are higher. This suggests the Redirection results in less sickness than using a hand-controller to turn.

(d) We expected that presence would be negatively correlated with sightings of the open back wall. On the other hand, previous studies have shown that reported presence is positively correlated with the extent to which people carry out appropriate head and body turns. Presence is also sensitive to movement technique [Usoh 1999]. In experiment RWP-I, we found that presence was indeed positively correlated with head movement and negatively correlated with seeing the open back wall.

In this experiment there was no significant difference in presence between conditions. We also did not find any significant relationship between seeing the open back wall and presence. However, the main objective of RWP was to reduce the occurrence of seeing the open back wall. Because we were successful in this, we do not have as many data points where the user saw the back wall as in RWP-I. To further investigate the relationship between seeing the back wall and presence, we would need to conduct a different study where users see the open back wall more often.

8.10 Observations and Summary of Results

The results suggest that RWP has the following properties:

Independent of the amount of head rotation, RWP reduces the frequency of the open back wall coming into a 40-degree field-of-view of the user, compared to turning with a hand-controller.

Users do not notice the rotations of the virtual scene.

RWP does not measurably increase the user's level of simulator sickness.

There is some evidence to suggest that the open back wall coming into the subjects' field-of-view, even for a short time, decreases the users' sense of presence.

Some users find it cumbersome and distracting to turn with a joystick or hand-controller. During the post-session interview, one subject commented about how he used the hand-controller:

SC7 – “When I got stuck. When it would take too much turning around. I think that it was very unrealistic. Um – a very still traversing.”

Experimenter – “So you preferred to turn with your body unless you got stuck?”

SC7 – “Oh ya – uh huh.”

RWP frees the user from needing a hand-controller for movement in the virtual scene—she specifies the direction of movement with her torso.

Although the Redirection algorithm used in experiment RWP-I did not meet our goals across all users, we do have anecdotal evidence that it worked for some users. One subject (SB9, in the Redirection group), when asked how much she saw the open back wall, became visibly confused and reported:

“No – I didn’t think I noticed it all, don’t think. I don’t know... I don’t know if I ever turned around that far. But I supposed I must have because I was walking in all sort of directions but I don’t remember seeing it – no.”

RWP is a technique that is very simple to implement and does not require expensive or obtrusive additional equipment. As far as I know, it is the only technique that allows users to walk-in-place about a virtual scene within a CAVE while reducing the chance of seeing the open CAVE wall. During the post-session interview, one subject (from RWP-II) remarked:

Experimenter – “How often did you notice the black wall or curtain?”

CB4 – “I never noticed. I assumed—well I thought there was a white wall behind me. I was surprised when I looked over and saw it was open there.”

Experimenter – “When did you notice that the white wall wasn’t there?”

CB4 – “When the experiment finished—I turned around and saw it.”

8.11 Comparison to Other Locomotion Techniques in CAVEs

In addition to hand-controller-specified flying and RWP, there are several other techniques that allow users to explore large virtual scenes in open-backed CAVEs. Among these are flying specified by leaning [LaViola Jr. 2001] and treadmills [Hollerbach 2000].

Each of these methods has its own advantages and disadvantages. Treadmills provide realistic proprioceptive cues of walking. Single-axis treadmills have a preferred direction of travel, and it is impossible for the user to turn on the spot. I know of no VE system that combines multi-axis treadmills with a CAVE-like display, though it is possible to build one. Multi-axis treadmills are loud and mechanically complex. RWP, on the other hand, requires only hardware common to CAVEs.³⁶ Leaning gestures are mechanically simple but do not provide the proprioceptive cues of walking.

RWP is most similar to LaViola's Auto Rotation technique. Both RWP and Auto Rotation allow the user to turn with her body, and both respond by automatically rotating the virtual scene to keep her from seeing the open back wall. Both techniques also free the user's hands. Auto Rotation magnifies the user's orientation so she can see in all virtual directions. For example, if the user is standing in the center of the three-walled CAVE, the 270° physical field-of-view that is covered by the CAVE's walls are mapped to a 360-degree virtual field-of-view.

Despite the similarities, RWP and Auto Rotation have different objectives. RWP aims to rotate the virtual scene in a manner that is not noticeable and does not increase simulator sickness, by accounting for the visual and vestibular responses to the rotation. Also, RWP causes a subject to unwittingly turn toward the front wall, even if she is not actively turning in the virtual scene. RWP aims to improve presence and naturalness by mimicking the way a person would move through the real world—this is why it is used with walking-in-place. For example, if a person becomes tired by walking five kilometers in the real world, she will also become tired when she moves five kilometers in the VE with walking-in-place.

On the other hand, Auto Rotation aims to improve ease-of-use, and is used in conjunction with leaning for locomotion in the virtual scene. A user who specifies a five-kilometer virtual movement by leaning will not get as tired as the user who walks-in-place for that virtual distance. I know of no studies investigating if Auto

³⁶ Except for an accelerometer, which improves walking-in-place in any VE with or without Redirection. Some common models of head trackers, such as those from Intersense™, have an accelerometer built in.

Rotation is noticeable or how it affects presence or simulator sickness. I did not experimentally compare Auto Rotation to RWP.

Chapter 9: Experiments to Determine What Level of Injected Scene Rotation Users Will Notice

The experiments, presented in Chapters 7 and 8, demonstrate that Redirection is effective. This chapter presents experiments I undertook to determine how much rotation can be used in a VE system before users notice. Answering this question also provides insight into how much physical space is required to have a user walk in a full circle, thus enabling infinitely extended virtual scenes.

9.1 The Lower Bound of Imperceptible Rotation Rate

How much one can detect the injected virtual scene rotation depends on many factors, not all of which are practical to control when Redirection is used in a real VE system. In particular, the literature suggests that a person is less likely to detect a given rate of rotation when:

- 1) she is cognitively engaged in other tasks [Rolfe 1986];
- 2) she is not expecting the visual scene to rotate [Gregory 1966] ;
- 3) there is a 3D spatial audio scene that is consistent with the visual scene [Lackner 1977a];
- 4) when she is turning her head (Chapter 5);
- 5) when the objects in the visual scene are far away (Chapter 5).

The first two experiments used naïve subjects who did not know about Redirection, were distracted by a task which forced them to turn their heads, and—in the first experiment (RW)—were presented with a spatial audio scene. The studies in this chapter explore the lower bound of the detection threshold: what rate of visual scene rotation will a user detect when she is actively looking for the scene rotation, not distracted by other tasks, not experiencing a spatial audio scene, and has nearby visual objects?

9.2 A Precise Definition of *Notice*—a Review of Concepts from Psychophysics

In order to undertake these experiments, one must have a precise definition of what it means *to notice* the rotation. Psychophysics is the study of the relationship between the physical stimulation of a person's sense

organ and the resulting interpretations (or perceptions) by the person. Psychophysical methods are formalized techniques for experimentally answering the question: given a certain *stimulus*, how well can a person consciously detect its presence? The stimulus can be of any sensory modality, for example, a pure 500 Hz sound tone, a visual pattern on a piece of paper, the 30 Hz vibration of a cell phone, or in the case of these experiments, the rotation of the virtual scene.

9.2.1 Detection Thresholds

The *detection threshold* for any given stimulus signal is the minimal intensity at which the person can detect it. Ideally, below this intensity, the signal is undetectable, and above it the signal is detectable. Using the previous example of a 500 Hz sound tone, the detection threshold would be the loudness (or, more specifically, the sound pressure level measured in, say, decibels) at which a person can detect the sound.

Ideally, at a loudness level infinitesimally greater than the detection threshold, the person would detect the tone every time it is played. At a loudness infinitesimally smaller than the threshold, the person would never detect the presence of the tone when it is presented. Unfortunately, signal detection does not behave this way. For a fixed signal of fixed intensity, a person may, for example, only detect the signal 60% of the time. One assumes that the greater the intensity of the signal, the greater the probability that the person will detect the stimulus. In certain experimental situations (where two-alternative *forced-choice* is used), the detection threshold is considered to be the intensity at which the person correctly detects the signal 75% of the time [Snodgrass 1985]. In other situations, the threshold is considered to be at the 50% level.

9.2.2 Signal Detection Theory: Sensitivity and Bias

The problem with the concept of detection thresholds is that it assumes that the only factors associated with detecting a signal are the person's acuity (sensitivity) and the intensity of the signal. It does not account for the person's bias nor the existence of background noise. In any signal detection situation, there is unavoidable background noise. In the sound tone example, the loudspeaker will emit some noise even when it is not playing the 500 Hz pure tone stimulus. And beyond that, there is some level of unavoidable noise in the laboratory in which the experiment is conducted. Finally, there is even noise in the perceptual mechanisms *within* the person.

Given this noise, it is impossible for a person to always correctly detect the stimulus when it is present, and to correctly reject the presence of the stimulus when it is absent. Either the person hears the noise and

incorrectly interprets it to be the signal, or the person misses the signal because it is obscured by the noise. If she is forced to choose between “the signal is present” and “the signal is not present,” the probability that she reports detecting the signal depends in part on her bias. This is the basis of *Signal Detection Theory* (SDT).

A realistic example of bias is a pair of radiologists trying to determine the presence of a tumor from a noisy ultrasound image of a breast. Even if the two radiologists have the same acuity, given the same noisy image, they may report differently. One might err on the side of reporting a non-existent tumor (thinking it is in the patient’s best interest to have a biopsy rather than miss the potential tumor and thus delay treatment). The other radiologist might err on the side of giving the patient a clean bill of health (thinking it is in the patient’s best interest to not suffer from unneeded procedures, and that any missed tumor will be detected during the next routine exam). Both radiologists have the same ability to detect the tumor, but their biases lead to different detection thresholds.

Table 9.1 - The possible outcomes from a single signal detection trial.

		Signal really present?	
		Yes	No
Reported by subject	“Yes”	Hit	False Alarm
	“No”	Miss	Correct Rejection

In order to account for the bias, one measures the person’s ability to discriminate the signal from the background noise—her *discriminability* of the signal, instead of her detection threshold for the signal. This is done by purposely manipulating the bias of the observer. For example, to measure the discriminability of the 500 Hz pure tone stimulus, the experimenter might have the subject do many trials where she tries to detect the tone (Table 9.1). The subject would be paid each time she correctly detected the signal (a *hit*), but *charged* each time she incorrectly reported the signal when none was present (a *false alarm*). By varying the ratio of payments and charges, the experimenter can vary the subject’s bias. For example, if paid \$1.00 for each hit and charging nothing for each false alarm, the subject would always report detecting a signal. Similarly, if the experimenter paid nothing for a hit but charged \$1.00 for each false alarm, the subject would never report detecting a signal. By varying the pay-off ratio, and by having many trials for each particular pay-off ratio, the

experimenter can determine a probability of detection for each pay-off ratio, and from many such pay-off ratio probabilities, determine the function relating the hit rate to the false alarm rate (*the receiver operating characteristic curve*). The shape of this curve determines the discriminability of the person for that stimulus. (For details see Corso [1967] and Heeger [1977]). The discriminability of a stimulus is particular to the stimulus intensity, so the experimenter must experimentally determine the discriminability for many different stimulus intensities to understand how the intensity of the stimulus affects its discriminability.

Whereas discriminability addresses the person's bias, it has several problems. One problem is that it assumes (and depends upon) the subject being an ideal observer who is making decisions solely so as to maximize her payment (so that the experimenter can control her bias). Another problem with discriminability studies is that they require huge numbers of trials (though this number can be reduced with a *rating* procedure) [Snodgrass 1985]. This makes them less practical for experiments conducted in immersive virtual environments. In audio stimulus detection experiments, a trial can be conducted in a few seconds, and it is practical to have a subject spend a few hours in the laboratory, whereas subjects in our VE studies are standing and wearing a heavy headset, and suffer increasing symptoms of fatigue and simulator sickness as their exposure time grows. Our studies are often limited to 5- to 10-minute virtual environment exposures. Even though a study of the discriminability of the virtual scene rotation would have the most internal validity, it is impractical for this research. Howard mentions that SDT techniques do not appear to have been used in vestibular research [Howard 1986b]. I know of no study using SDT in a VE or for self-motion studies.

9.2.3 Methods for Determining Thresholds

The experiments presented in this chapter are all detection threshold experiments. There are many standard techniques for measuring detection thresholds. I describe here only those that I use.

9.2.3.1 Method of Adjustment

With the *method of adjustment* technique, the subject is given direct control of the stimulus intensity. She adjusts the intensity such that she can just barely detect it. Using the audio tone example, the subject would be given a volume knob and she would use it to adjust the volume so that she can just barely hear it. This method's advantage is that is very quick but it is also very susceptible to manipulation by the subject and subject bias. For example, if the subject wants the experimenter to believe that her hearing is more sensitive

than it really is, she could adjust the knob to her threshold setting, and then turn the volume down before reporting that she has finished the trial.

9.2.3.2 Constant Stimulus

The *constant stimulus* technique consists of many trials. In its most basic form, the experimenter presents the stimulus at some intensity level and the subject reports whether or not she can detect it. Several different intensities of the stimulus are tested, and for each particular level of intensity, many trials are conducted to determine the probability of correct detection for that intensity. This results in a *response curve*, which is usually sigmoidal (Figure 9.1). The threshold is considered to be the intensity at which the subject detects the signal 50% of the time.




Figure 9.1 – Idealized response curves resulting from the constant stimulus technique. Left: The curve from the standard constant stimulus technique. The detection threshold is the stimulus intensity at which the person claims to detect the stimulus 50% of the time. Right: The curve from the forced-choice constant stimulus technique. Here, the detection threshold is the stimulus intensity at which the person correctly detects the stimulus 75% of the time.

A more rigorous version of the constant stimulus technique presents the stimulus in only 50% of the trials. The subject chooses between one of two possible responses (for example “I detect the signal” or “I do not detect any signal”). This is a two-alternative *forced-choice*. The subject does not have the option to respond “I can’t tell.” In this version, when the subject cannot detect the signal, she must guess, and will be correct on average in 50% of the trials. The detection threshold is the lowest-intensity level (out of those tested) for which the probability of correct detection is 75%. I use this definition of threshold.

The constant stimulus technique allows the experimenter to measure the response curve of the stimulus intensity and is robust to sampling noise (sampling noise will not produce a large error in the resulting detection threshold). Furthermore, if the response curve is *non-monotonic* (not steadily increasing or decreasing, but with local dips and bulges), this will be revealed. The problem with constant stimulus is that it requires many trials

(though not nearly so many as a signal discriminability experiment), and those trials are evenly distributed about all levels of the stimulus intensities—just as many trials are used to probe some intensity level where the subject always detects the stimulus (far above the detection threshold) as are used to probe at the detection threshold. In this sense, it is inefficient.

9.2.3.3 Staircase

The staircase or *tracking* technique is similar to the constant stimulus technique, except that it is more efficient. The trials are concentrated near the detection threshold. The intensity of the stimulus at each trial depends on the intensity of the stimulus of the previous trial: each time the subject correctly detects the stimulus, the intensity is reduced for the next trial, and each time she does not detect the presence of the stimulus, the intensity is increased for the next trial. Thus, the staircase method quickly finds the detection threshold (if one defines the threshold as being at the 50% detection probability), and then oscillates the intensity just above and below it (Figure 9.2). The disadvantage of the staircase method is that it is vulnerable to noise in the signal. If the subject responds “yes” at some intensity level below the actual detection threshold, then the next stimulus will be of even lower intensity, and the experiment will take more trials to converge onto the detection threshold. Even worse, if the response curve is non-monotonic, the staircase method might converge onto a local minimum (which is not the detection threshold, since the detection threshold is the minimal intensity at which the stimulus can be detected), without giving the researcher any indication that the response curve is non-monotonic.

If one defines the threshold as the intensity at which the subject has a 75% probability of correct detection (such as in this dissertation), the above staircase technique converges to an intensity level *below the detection threshold*. There are several other adaptive techniques which estimate the 75% detection threshold, such as PEST [Taylor 1967].




Figure 9.2 - An idealized sample progression of stimulus intensity when using the staircase method to estimate the stimulus's detection threshold.

9.3 Experimental Designs

To determine virtual-scene rotation detection-thresholds, I explored three different experimental designs, most of which were unsuccessful. The three experiments are RDT-scv (while user is Still, with Constant Velocity rotation), RDT-ssv (while user is Still, with Sinusoidal Velocity rotation), and RDT-wcv (while user is Walking, with Constant Velocity rotation). I explored several minor variations of each of the three experimental designs, and abandoned all but RDT-wcv.

The semicircular canals (SCCs), detect head angular velocity. However, they cannot continue to detect an indefinitely sustained, constant-velocity rotation. For stimuli below 5 Hz, the lower the frequency of the velocity stimulus (Figure 4.7), the less sensitive the SCCs (semicircular canals) are to that stimulus. Furthermore, for stimuli of constant frequency and constant amplitude, stimuli presented for a greater duration are more likely to be detected [Howard 1986b]. The vestibular system's latency (the time required in order to detect the stimulus) is greater for lower intensity stimuli. This interaction between angular velocity, frequency, and time complicates these experimental studies. For example, whether a person can detect 0.1 deg/s at time t depends on the rotation rate before time t .

9.3.1 Adjustment of Visual Scene Angular Velocity While Standing Still: RDT-scv



Figure 9.3 - A pilot subject manipulating the control knob in experiment RDT-scv.

The study's goal was to determine what constant rotation rate a subject could detect while standing still. The subject wore an HMD and observed the scene rotating toward the right about the center of her head. This study used the method of adjustment described above—the subject used a rotary knob (Figure 9.3) to control the angular velocity of the virtual scene. For each trial, the starting velocity was randomly chosen from the range $0 - 1\text{deg/s}$ and the gain of the control knob was randomly chosen to prevent the subject from learning the relationship between the knob angle and the controlled velocity of the scene.

The problem with RDT-scv was that the rotation was more noticeable while the pilot subjects were turning the knob (in either direction—increasing or decreasing velocity) than when the knob was held still. The faster the knob was turned, the more detectable the rotation. The detection threshold for changing velocity is much smaller than for constant velocity. For example, starting at undetectable velocity, then turning the knob quickly to the right and then back to its original position creates a step-like function in the orientation of the virtual scene, which is quite noticeable. To encourage subjects to turn the knob very slowly, I modified the software such that the knob would only respond when it was turned in one direction. Thus, subjects could only increase the velocity and thus subjects had to be sure that they could not detect the rotation before adjusting with the knob. This did not appear to improve anything, and I abandoned this experiment.

9.3.2 Adjustment of Visual Scene Oscillation Frequency While Standing Still: RDT-ssv

This study was similar to the RDT-scv, but was to determine the *frequency* detection threshold at which the virtual scene could oscillate without being detected by a subject who is standing still. Subjects used the knob to control the frequency of the virtual scene rotational oscillation. The knob had no endpoint, so that the subject could turn it any number of revolutions and could increase the frequency. Each angular unit of knob rotation controlled the log of the frequency, such that there was no minimal or maximal frequency value. For each trial, the starting frequency was randomly chosen from the range 0 to 0.5 Hz, and the gain of the control knob was randomly adjusted.

I hypothesized that the detection threshold is at a frequency below 0.1 Hz, because the semicircular canals are less sensitive to velocity below 0.1 Hz and because Duh et al. propose that the visual-vestibular crossover is below 0.1 Hz [Duh 2001b; Duh 2004]. The frequency 0.1 Hz corresponds to a peak-to-peak duration of 10 seconds. At these low frequencies, the threshold set by the subject appeared to depend on the instantaneous angular velocity (which depends on the phase of the sinusoidal oscillation) rather than the frequency. The problem in RDT-scv was also present here. The rotations were more noticeable while the subjects were manipulating the knob. Thus, the threshold set by the subjects depended on how quickly they were turning the knob, and what the phase was at the time of the detection. Because of these problems, I abandoned this experiment as well.

9.3.3 Detection of Direction of Scene Rotation While Walking

In this experiment, I investigated the detection threshold of constant-velocity rotation of the virtual scene while the user is walking in a straight line in the virtual scene. In each trial, the scene rotated to the left or right at a constant rate as the subject walked from one end of the room to the other. Once she reached the other side of the room, she told the experimenter if the room rotated right or left.



Figure 9.4 - Photographs of a subject during trials of experiment RDT-wcv. Top: The subject starting at the lab location where the trial begins. Bottom left: The subject after walking to the destination in the virtual scene. In this trial, the virtual scene did not rotate, thus the subject walked straight in the lab. Bottom right: In this trial, the scene rotated to the subject's right, thus the subject walked to her right in the lab as well. While the subject was walking, I carried the cables behind her.

The first iteration of this experiment used a modified constant stimulus technique. The computer randomly selected a rotation rate and direction from a list. These ranged from 0.5 to 5 deg/s. However, during piloting I found I could only run 20 or so trials per session. (The consent forms, training, debriefing, etc took the rest of the 60-minute session.) This only left two trials per rotation rate and direction—not enough, so I switched to a staircase technique. If the subject correctly detected the direction of rotation, on the next trial the system would reduce the rotation rate (it divided the rate by an arbitrarily chosen factor of 1.6). Similarly, if the subject was incorrect, the computer increased the rotation rate by a factor of 1.6. The system used separate staircases for the rightward and leftward rotations and randomly chose between them to keep the subject from becoming habituated. I hoped this would converge with fewer trials.




Figure 9.5 - The staircase progression of two sessions of experiment RDT-wcv. Angular velocity of each trial stimulus is plotted as a function of the trial number. The resulting detection thresholds are shown with red lines. There is one for each direction and each session.

This staircase technique converges to the rotation rate at which the subject can correctly detect the rotation 50% of the time (henceforth called the *chance rotation rate*—CCR), which is below the detection threshold rate (at which the subject must correctly detect the rotation 75% of the time). I conducted six sessions with six different subjects using this experimental design. Data from two sample sessions are shown in Figure 9.5. Many subjects had very different CCRs depending on whether the rotations were toward the right or left. Were the CCRs actually different for each direction? Or is the difference in measured thresholds due to a bias on the subjects part (when they cannot determine the direction of rotation, they guess the answer with a bias toward one side) or due to the effects of noise (there are two staircases in each session—one for each direction—and the noise randomly affects the adaptive staircase for one direction differently than for the other direction). To investigate this, I added trials where there was no rotation (subjects were unaware of this), but the subject still had to choose between right and left. The results do not suggest any obvious bias, as the 95% confidence interval includes the 50%-50% probability for each subject thus measured (Figure 9.6). However, the confidence intervals are very wide.




Figure 9.6 - Top: The staircase progression of subject 8, session 1, but with 6 randomly interspersed trials where there was no rotation. These trials are shown on the green zero velocity line. Bottom: The subject answered “left” on 4 out of the 6 no-rotation trials. If he had answered 3/6, the red line would be at the 0.5 mark on the chart. The gray area represents the 95% confidence interval.

Another possibility is that the response-curves for the virtual scene rotation are non-monotonic. To investigate this, six of the subjects who performed the staircase trials were invited to return of an additional session, and three of them agreed to do so. In these sessions, I used the constant stimulus technique. These subjects had already been trained and had a previous session of practice, and thus were able to perform many more trials, making the constant stimulus technique slightly more practical than in the first pilot sessions (5-7 trials per rotation rate and direction, rather than the original 2 trials per rotation rate and direction). Even with this balanced sampling per rate, I cannot rule out or confirm that the response curves are non-monotonic (Figure 9.7).




Figure 9.7 - The response curve from subject 2, session 2; performed using the constant stimulus technique. At +3 deg/s the threshold appears to be lower than for +1.5 deg/s.

9.4 Experimental Details



Figure 9.8 - Views of the RDT virtual scene. Subjects were instructed to walk through the doorway (left) to the painting of flowers (right).

Subjects wore an HMD which was blank most of the time the subject was wearing it. For each trial, I would present the virtual scene, and the subject would turn toward (but not move toward) a virtual painting on the wall of the far side of the virtual room. Then I would blank out the HMD again. Then I would present the virtual scene once again, and the subject was instructed to walk, at an even pace, toward the painting. In order to reach the painting, the subject had to walk through a doorway in the middle of the virtual room (Figure 9.8). The subject started walking immediately when the virtual room was presented—she did not wait for me to tell her to go. While she was walking, the virtual room rotated about her head, at a fixed angular velocity. The subject would walk a straight line in the virtual room, but in a curved path in the real world. The headset was adjusted such that the subject could not see the lab floor, and she was told to look straight ahead and walk at a constant speed, without slowing down or speeding up. When she approached the virtual painting (5 meters from the starting point) the HMD again went blank. The subject would call out “left” or “right” depending on if she thought she had veered left or right. I recorded her response by pushing the appropriate button on a wrist-mounted computer. After this, I would lead her back to the starting point in the lab, with the HMD still blanked. I took a non-direct path back to the starting point, to reduce her ability to estimate, from her path back to the starting point, in which direction she had veered in the previous trial. This was repeated until either the one-hour session was over, the subject accrued 20 minutes of non-blanked time in the HMD, or the subject asked to stop.




Figure 9.9 - The rotation-rate during the start-up period of each trial in experiment RDT-wcv. The angular velocity became maximal in 2.5 seconds. The HMD video faded in over the first second.

The video imagery of the virtual room presented in the HMD took one second to fade in from the blank screen. While the video was fading in, the scene would increase its rotation rate from 0 to the constant rate for that particular trial. The rotation rate was not a step function, but rather a sigmoidal such that velocity reached its final value in 2.5 seconds. I did this to keep the subject from noticing the rotation as it first started. Subjects were told ahead of time that the virtual room would rotate left or right and that they would veer in the same direction (in reference to the real world). I told them the room would always rotate right or left and that they had to choose one or the other. If they did not know, they should guess. After piloting on a few subjects, I decided to demonstrate this to them with a few training trials before beginning the data collection trials. While subjects were walking, headphones played white noise to mask sounds from the laboratory.

The HMD was a Virtual Research model VR8 with full color and 640 by 480 resolution. Video was monoscopic and locked to 60 Hz refresh rate and 60 frames per second. Stereoscopic video was not used because it resulted in varying frame-rates, which might have confounded the results. The tracker was a single 3rd Tech Hiball Series 3000 which reported position and orientation of the head at 250 reports per second. The tracker's multi-modal filter was turned on, such that the system automatically switched between different Kalman filters depending on whether the subject's head was still or moving. This is the default setting for this tracking system and results in less apparent tracking noise. Though it complicated the interpretation of the results (what is the effect of tracker noise on the rotational detection-threshold?) I felt the increase in the experiment's validity by using real-world VE system settings was a worthwhile tradeoff.

None of the subjects had a history of severe motion sickness or epilepsy. Each was paid \$10 for participation.

9.5 Results

For each staircase session and direction of scene rotation (right or left), the estimated chance rotation rate (CRR) was computed as the geometric mean of the rotation rate from the last four trials. These are listed in Table 9.2. The geometric mean was used because the rotation rate of each successive trial was computed by *multiplying or dividing* the rate of the previous trial by 1.6. For sessions with fewer than nine trials per direction, the CRR was not computed. This is because the rates of those sessions did not seem to have enough trials to converge. *The average CRR for all staircase trials was 0.66 deg/s.*

Table 9.2 - The chance rotation rate (CRR) and other data for the staircase sessions of experiment RDT-wcv. In staircase sessions, the rotation rate of each trial was adaptive and converged (if enough trials were presented) to CRR. CRR is computed as the geometric mean of the last four trials.

subject #	3		4		5		1		6		7		8	
Ave speed (m/s)	.76		.62		.75		.67		.62		.72		.73	
dir.	right	left	right	left	right	left	right	left	right	left	right	left	right	left
# of trials	19	15	18	14	17	14	19	15	22	18	11	8	8	5
CRR (deg/s)	-0.73	1.11	-0.31	0.15	-2.60	4.42	-0.20	1.11	-1.00	0.12	-0.90	-	-	-

Figure 9.10 (left) shows the response curves from all three of the constant stimulus sessions. If one assumes that all users have the same detection threshold, or that a practical VE system must use an average value for all users, then all the trials from all the constant stimulus sessions can be aggregated, resulting in the overall response curve shown in Figure 9.10 (right). This curve suggests that the detection threshold hold is 1 deg/s. The detection threshold rotation rate should be greater than the chance rotation rate, and these data are consistent with this expectation. The CRR from the staircase sessions is 0.66 deg/s, whereas the detection threshold from the constant stimulus sessions is 1.0 deg/s (Figure 9.10).

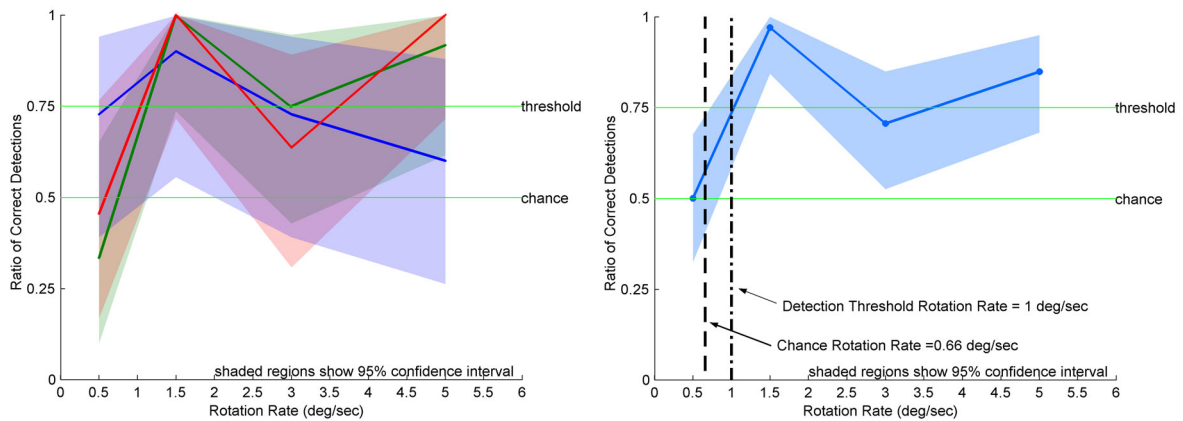


Figure 9.10 - Response curves from the constant stimulus sessions. Left: Response curves from all three sessions overlaid. Right: Response curve from all trials from all three sessions aggregated. From this curve, the detection threshold is estimated at 1 deg/s (dashed-dotted line). As expected, it is greater than the average chance rotation rate, which is estimated from the staircase session at 0.66 deg/s (dashed line).

Note that these thresholds are for walking durations of 10 seconds. As described in Chapter 4 and Chapter 5, the rotation rates can be slowly increased as governed by the high-pass characteristics of the podokinetic system.

9.6 Caveats

The margins of error in this study are quite large. This is due to the nature of the binomial distribution of trials that have only two possible outcomes.

There could also be significant learning effects during these trials. In fact, pilot testing determined that subjects required a few training trials before they could understand that they were veering to one side, even on the trials with maximum rotation rates. Some subjects reported that they were confident they were rotating but were unable to determine in which direction they turned.

Because of the above, the data from this experiment do not conclusively demonstrate each subject's detection threshold. Instead, the data provide meaningful insight about how much rotation can be present before subjects notice. For reasons described at the start of this chapter, the estimates of the detection thresholds are quite conservative. In all likelihood, a user who is engaged in a task (even free exploration) in a virtual scene with spatial audio will not detect rotation rates higher than the detection thresholds estimated here. And even if they do notice occasionally, I argue that the benefits of using real walking with Redirection (ease of

use, naturalness, and simulator sickness) make it a better choice than using a hand-controller to fly through a virtual scene.

Chapter 10: The Simulator Sickness Questionnaire and its Bearing on Redirection

Does Redirection affect the user's level of simulator sickness? The tool I used to investigate this is Kennedy's simulator sickness questionnaire (SSQ). This was presented in a milestone paper [Kennedy 1993], which I hereafter refer to simply as Kennedy. The thesis-statement result, that Redirection results in lower SSQ scores than turning with a hand-controller, is briefly discussed in Chapter 8. This chapter presents other SSQ-related results:

- 1) an argument that SSQ scores from a VE system cannot be directly compared to the standard SSQ benchmark-scores (which were derived on flight simulators);
- 2) a survey of SSQ scores from VE systems;
- 3) a power analysis to determine how many subjects would be required to experimentally determine if Redirected Walking causes greater SSQ scores than real walking without Redirection.

To support the above, this chapter also presents the origins and use of the SSQ, and a background on statistical analysis techniques.

10.1 Background on Statistical Analysis Techniques

The *test of significance* is the standard statistical analysis technique to find out if an effect exists; and *power analysis* is the technique used to show that an effect does not exist or is so small it is not a concern. In summary, it is more straightforward to show the existence of an effect than to show non-existence. More information is needed to do this, and we do not yet have this information for the SSQ scores or other measures of simulator sickness.

When conducting an experiment to determine if some measurable characteristic of one group of people is different than another group of people, (i.e., are men taller than women?) it is not acceptable just to compare the averages of each group. There is variation from person to person and error in the experimental process, and the analysis must deal with it. Each person is different and those particular persons that are measured are

randomly selected from the group. What if the persons measured happen to be the tallest ones in the group, and therefore the measurements do not faithfully represent the group? Furthermore, there is some random error in each measurement (or sample). If one measures the same person twice, the two height measurements will not be exactly the same. Standard statistical analysis procedures take these variations and errors into account.

In the above example, to determine if men are on the average taller than women, one might measure some number of men and women and then perform a significance test³⁷ on the measurements. The results of the significance test would be a p-value—the probability that the difference between the heights of the men and women who were measured is due to chance.³⁸ In general, the more people who are measured, the lower the p-value will be (if a difference really exists), hence the more confident one can be that the difference between men and women is reflective of the population and not due to chance. For example, if one measures a thousand men and women, one can be more confident in the results than if one measured just five men and women. A conclusion that two groups are different, drawn from a test result of $p=0.001$ is more credible than the same conclusion drawn from a test result of $p=0.1$. In the community of VE researchers, $p=0.05$ is the commonly used threshold. If the p-value is below 0.05, the effect is considered statistically significant. Showing that an effect exists is much more straightforward than showing that an effect does not exist.

10.2 Power Analysis

In the measurements of a subset of people's heights, one can either find or not find a statistically significant difference. If the effect is statistically significant, then one can claim it exists in the population. However, the inverse is not true. If there is not a statistically significant effect, one cannot conclude the effect *does not* exist in the population. It could well be that the effect does exist, but that too few people were

³⁷ A t-test is one example of many different significance tests. The details about each significance test or how one chooses between them are beyond this dissertation's scope.

³⁸ There is always some chance that the people who happened to be measured are not representative of the larger population. Therefore, one might observe a difference between the measured groups of people when, in reality, none exists in the larger population. The chance of this happening is the p-value.

measured to find it. If one were to measure more people, one might then uncover a statistically significant difference between men and women that was previously not apparent (Figure 10.1).




Figure 10.1 - Measuring more peoples' heights can uncover a significant difference between the heights of women and men. The leftmost plot shows, for a fictional population, the distribution of women's and men's heights. In this fictional population, men are taller on average. The middle plot shows the distribution of a random sample of 100 men and 100 women from this population. From this sample, it is not clear whether men are significantly taller. But it would be wrong to claim, from this sample, that men are not taller than women. The rightmost plot shows the distribution if we increase the sample size to 1000 men and 1000 women. This bigger sample clearly shows that men are taller.

To show that the difference between two groups is bounded by some value, one performs a *power analysis*. One specifies what minimum size effect one is looking for (e.g., the height difference between men and women) and the result of the power analysis tells how likely it is to see that result in the data. For example, one might discover that some study has an 80% chance of finding a 2 cm or greater difference in heights between men and women, if such a difference really exists in the population. After collecting the data, if there is no significant difference in the measurements, one can conclude that there is an 80% probability³⁹ that the difference between the male and female populations is not greater than 2 cm. In other words, men are not, on average, 2 cm taller than women.

³⁹ In computer science and psychology, the generally accepted value P (power) is 0.80, even though the accepted value of p (significance) is 0.05.

To perform such a power analysis, one must specify the number of measurements, the size of the effect one is looking for, and (assuming a normal distribution) the standard deviation of the measure in the population [Hays 1963].

10.3 History and Development of the SSQ

The maladies affecting pilots in early flight simulators, as far back as the 1950s, were thought to be motion sickness [Kennedy 2003b]. Before the development of the SSQ, practitioners measured this flight-simulator-induced sickness with the Pensacola Motion Sickness Questionnaire (MSQ). Kennedy reports both subjective and quantitative reasons why the MSQ is not appropriate. First, Kennedy claims sickness affects a much smaller percentage of people in flight simulators than in motion-sickness-inducing situations such as sailing in rough seas, and the symptoms they suffer are much less severe. Kennedy also argues that, in order to continuously monitor and diagnose problems with individual simulators, one needs a questionnaire that has more statistical power (is more sensitive) and is more straightforward to administer than the MSQ.

The SSQ was developed from a factorial analysis of 1,119 MSQ questionnaire pairs (pre- and post-exposure). These were filled out by U.S. Navy and Marine pilots before and after flying one of ten flight simulators used for actual and regular flight training. Of the ten simulators, five turned out to be highly symptom-inducing and the other five benign. Symptoms which did not show statistical power; either because they did not change frequency or severity between before and after the exposure, or because they occurred too infrequently; were eliminated. For example, vomiting is a definite sign of sickness, but it only occurred twice in the 1,119 simulator exposures. The remaining symptoms sorted into three distinct clusters from which Kennedy derived the nausea, ocular-motor discomfort, and disorientation scales of the SSQ. These scales were given weighting such that, in the final SSQ scores, each subscale has a minimum score (absolutely no symptoms) equal to 0 and a standard deviation of 15. “Because the number of observations was so large, the sample was treated as if it were a population that could be used as a baseline against which future simulation evaluation data could be compared.” In addition, the SSQ, as Kennedy defined it, does not include any pre-exposure questions in the scoring. Once pilots who were not in their “usual state of good health” were excluded from the data, Kennedy found the pre-exposure MSQ scores did not have useful standard deviations. Because difference measures (i.e., post-exposure – pre-exposure score) are less reliable, Kennedy decided that only post-exposure data should be used in scoring.

10.4 Diagnostic and Statistical Power of the SSQ for Flight Simulators

The SSQ is used for monitoring and diagnosing flight simulators. For example, using SSQ data collected from 3691 simulator sessions, Kennedy reported that some simulators have a shake-down period—after first operation it takes time for the systems to be properly tuned. SSQ scores drop after the first few months of operation. Similarly, when their SSQ scores were sorted by the number of days each pilot rested between simulator exposures, they found that rests of 2-5 days produced the lowest SSQ scores.

But to discover these important trends, Kennedy proposed using the 75th percentile of the SSQ scores instead of the means. The reason for this is that 40-75% of all the scores in their data were zero. Many (if not most) of the pilots did not report any sickness at all! Because of this heavily skewed distribution (Figure 10.2), mean scores, even with $n=3691$, were not as revealing. Kennedy argued that the 75th percentile score is a stable statistic and is roughly the median score for the half of the pilots who got sick.

10.5 Application of SSQ to General Purpose VEs

The SSQ has demonstrated utility for diagnosing problems in military flight simulators. Some have argued that it should also apply to VEs [Kennedy 1992; Kolasinski 1995]. Users of the HMD-VEs in our laboratory report higher SSQ scores than Kennedy's pilots. Independently, Kennedy reports similar observations (he compares simulator scores to HMD-VEs from other laboratories). He found VE systems have scores ranging from 19-55, whereas flight simulators' scores range from 8-20; and the median score (50th percentile) from VEs is higher than the 90th percentile of flight-simulator scores [Kennedy 2003b]!

Does this mean that VEs are necessarily more sickness-inducing than flight simulators? I argue not. The data on which the SSQ scoring procedure was calibrated is from male Navy and Marine pilots, whereas VEs have a very different user base. The SSQ is designed such that the standard deviation of the population's score is 15. Kennedy assumed the original SSQ data describes a population to which other SSQ scores could be compared. SSQ data from VEs in our laboratory and others, whose users are university students and the general population, have a greater standard deviation and mean than Kennedy's (Table 10.1). In addition, there are significant pre-exposure symptoms and variation. In one study (RWP-II), healthy subjects had pre-exposure scores with mean = 12 and std. dev = 21. In Kennedy's pilot data, he found that pre-scores had "very little variance." Our pre-score standard deviation is greater than that for Kennedy's pilots report after the simulator exposure!

Kennedy suggested that users with post-exposure SSQ scores above 15 consult a doctor, and that those users with scores above 20 not be allowed to leave until the sickness subsides [Kennedy 2003b]. Nichols' pre-exposure data (mean=14.17) are close to the score at which Kennedy would be concerned after the VE exposure [Nichols 2000].

Table 10.1 - Comparison of SSQ data from various sources. The top six are from VEs in which vection was not used to simulate motion (users really walked to locomote in the virtual scene). The bottom two are from the flight simulators used to develop and validate the SSQ.

study	n	Pre SSQ		Post SSQ		exposure time	source	VE details
		ave	S.D	ave	S.D			
Meehan 2003	193	7.4	11.1	14.2	21.6	5 min	my data	V8 HMD, UNC Hiball tracker, low latency (50-90 ms). Real walking near virtual pit. Data only for subjects in state of normal health
Usoh 1999				12		5 min	[Arthur 2000]	V8 HMD, UNC Hiball tracker, Real walking near virtual pit.
Arthur	43			11.6	10.9		[Arthur 2000]	Data for V8 HMD only. UNC Hiball tracker, Real walking in maze.
Lok	14	4.2	4.8	15.7	20.0		personal comm.	V8 HMD, UNC Hiball tracker. "Purely Virtual" condition only. Block manipulation task, no locomotion [Lok 2002].
Zimmons II	12	5.3		22.4		5 min	personal comm.	V8 HMD, UNC Hiball tracker. Real walking, VE. (post scores are after the first exposure) [Zimmons 2004].
Zimmons I	54	11.8	12.2	16.9	24.5		personal comm.	V8 HMD, UNC Hiball tracker, Real walking VE.) [Zimmons 2004].
Stanney				30			[Arthur 2000]	
Kimberly Swinth	206			14.5	18.5		personal comm.	V8 HMD, 6DOF tracker, "Virtual Casino", no locomotion.
Nichols 2000	19	14.2	11.6	29.9	21.6	30 min	[Nichols 2000]	V8 HMD. Locomotion via flying.
RWP-I	28	8.7	6.2	11.1	11.0	< 10 min	my data	In UCL's CAVE. Walking-in-place.
RWP-II	22	11.7	20.9	13.6	16.5	<10 min	my data	In UCL's CAVE. Walking-in-place.
Kennedy SSQ calibration	1119			9.8	15		[Kennedy 1993]	10 military flight simulators.
Kennedy	3691						personal comm. & [Kennedy 1993]	TH-57 Helicopter trainer. 3200 out of the 3691 are <=14 (est. from histogram in paper).




Figure 10.2 - The SSQ scores from one of our VEs [Meehan 2003] (right) have a similarly shaped distribution to that presented in Kennedy 1993 (left). However, the scales are very different. Roughly 3200 out of 3691 (82%) of Kennedy's scores are ≤ 14 , while 113/192 (58%) of our scores are ≤ 14 [left plot from Kennedy 1993].

Given the list of VE factors that aggravate simulator sickness, one would expect that the VEs used in the first six studies above would induce less simulator sickness than a flight simulator, not more. For example, these VEs had less vection (users really walked) and had much shorter exposures (five minutes vs. entire 1- to 4- hour flights⁴⁰). In addition, the 60-degree field-of-view of the HMD was less (the flight simulators presumably used projection domes). Why are our SSQ scores not lower than Kennedy's?

The SSQ scoring weights were scaled using self-reported data from male military pilots. I propose that the general population rates the severities of symptoms differently than military pilots. Military pilots are a different population. I suspect there are several components of this difference:

- 1) Military pilots have already been naturally selected to be less susceptible to motion sickness (if someone readily becomes motion sick, they presumably would not last very long as a navy pilot trainee).
- 2) Military pilots are much more physically fit than the general population.
- 3) Military pilots are exposed to, or trained for, situations that are more stressful than the general population experiences. A university student might report nausea resulting from pre-exam anxiety as 3 out of 5,

⁴⁰ The exposures were all greater than one hour, and the longest was four hours [Kennedy 1993]. While not part of Kennedy's dataset, some extreme cases of flight simulator exposures are 38 hours long [Strachan 2001]!

whereas a military pilot, who is prepared for combat, might consider the same level of nausea as not even worth reporting.

4) Military pilots can be grounded for having symptoms, and are thus under (unintentional) pressure to not report symptoms [Parker 2003].

5) The general population can have a bias in the other direction—they are expecting to feel simulator sickness and thus are more likely to feel symptoms.⁴¹ Some researchers prefer not using the freshmen psychology course subject pool for this reason [Hollins 2001].

6) Military pilots are mostly male, whereas university test subjects are gender-balanced. Kolasinski reports that females have higher SSQ scores than males but believes that this result is because males tend to underreport symptoms [Kolasinski 1995].

Kennedy himself notes that military pilots “are self-selected, have more experience with novel motion environments, and may be more likely to underreport symptoms” [Kennedy 2003b]. Because of this, it is not reasonable to compare the absolute score of a navy pilot flying a flight simulator to that of a university student in a VE of a living room. If the living room VE scores higher, that hardly supports a claim that the virtual living room is more sickness-inducing than the flight simulator.⁴²

⁴¹ In a week-long experiment, our research team accidentally labeled the questionnaire given to the subjects as the “Simulator Sickness Questionnaire” (instead of leaving it blank as we should have). Steve Ellis at NASA Ames joked that if we changed the title to “Sexual Dysfunction Questionnaire” in the middle of the week, we would see a statistically significantly lower score after the change.

⁴² Kennedy points out that VEs not only have higher total SSQ scores, but that the profile of sub-scores (i.e., the ratio of nausea to disorientation scores) is different from that of flight simulators [Kennedy 2003]. It may be that VE causes simulator sickness in a different way than flight simulators, or that the general-population users report different SSQ profiles even when they are exposed to the same stimulus (flight simulator or VE) as military pilots.

10.6 SSQ Scores from Redirection vs. Real Walking

Ideally, a VE practitioner would like to know if Redirection causes additional simulator sickness compared to a similar VE system that uses real walking.⁴³ More specifically, one would like to know if the additional sickness caused by Redirection is enough to make a previously acceptable (without Redirection) VE system troublesome once Redirection has been added to it. I hypothesize that the increase in simulator sickness caused by Redirection, if any, is insignificant (not statistically insignificant, but operationally insignificant). However, I am unable to quantitatively support this hypothesis.

A power analysis requires that the distribution of the user population and an effect size be known. Since these parameters from military flight simulators are not appropriate, I collected SSQ data from roughly 200 general-population users of our real-walking VE system [Meehan 2003] and aggregated them with other recent studies from our laboratory (Table 10.1) [Lok 2002; Zimmons 2004]. From this, I estimate the population mean (=14) and distribution shape (exponential⁴⁴) for non-Redirection, real-walking VEs. However, I have no meaningful estimate for an effect size. None of the real-walking VE systems on which I have data could be considered particularly sickness-causing. Furthermore, no researcher I queried from the (non-flight simulator) community could quantify a threshold SSQ score (or increase in SSQ score) above which the VE system would be considered unacceptable.

If I arbitrarily propose an effect size of 2, which is the (not significant) difference that Arthur [2000] found in his study of HMD field-of-view, then a power analysis⁴⁵ results in a sample size of roughly 900 subjects (assuming the SSQ score would increase from 11 to 13, as it did in Arthur's study) to 1400 subjects

⁴³ From a practitioner's point of view, this is an odd question to ask—if one has enough tracking area to allow for real walking in the virtual scene, why consider using Redirected Walking?

⁴⁴ I modeled the SSQ scores with an exponential distribution because they are obviously not normal distributions (figure 2{xref}). However, using a normal distribution gives similar results.

⁴⁵ Using a two-group, 1-tailed, exponential distribution power analysis, with significance level of $p=0.05$ and power of $P=0.8$. If one Assumes a normal distribution and standard deviation of 22 in both groups, as in Meehan [2003], the resulting sample sizes are similar.

(assuming an increase from 14, my population mean, to 16). If I choose an effect size of 5 (the difference between “significant sickness” and “a problem simulator” in Kennedy’s military pilot population), then the power analysis results in 266 subjects. This many subjects would be required to claim, with 80% certainty, that Redirection does not increase SSQ scores by 2 (or 5) points or more. Experiments of this size were not feasible with my resources.

10.7 Redirection Induces Less Simulator Sickness than Turning Manually

Another way to examine the effects of Redirection on simulator sickness is to compare it to the alternatives.

In the RWP-II experiment (Chapter 8) I obtained SSQ scores of two groups of users, one using Redirected walking-in-place, and the other using a hand-controller to turn (while still walking in place). Using Kennedy’s measure (75th percentile scores), I found that simulator sickness was less for the group using Redirection (Figure 10.3). My statistics consultant could not point to a power analysis for 75th percentile scores. The mean and standard deviation were 18.0 ± 21.7 for the hand-controller group and 9.5 ± 6.9 for the Redirection group. It is also interesting to consider the SSQ subscale scores. The Oculomotor subscale relates to eye strain and similar symptoms caused by the optics of the VE system. Since subjects in both conditions (hand-controller and Redirection) used the same VE system, one would expect only the nausea and disorientation subscales to show differences between conditions. Indeed, the 75th percentile scores for nausea and disorientation subscales are lower for the Redirection group.

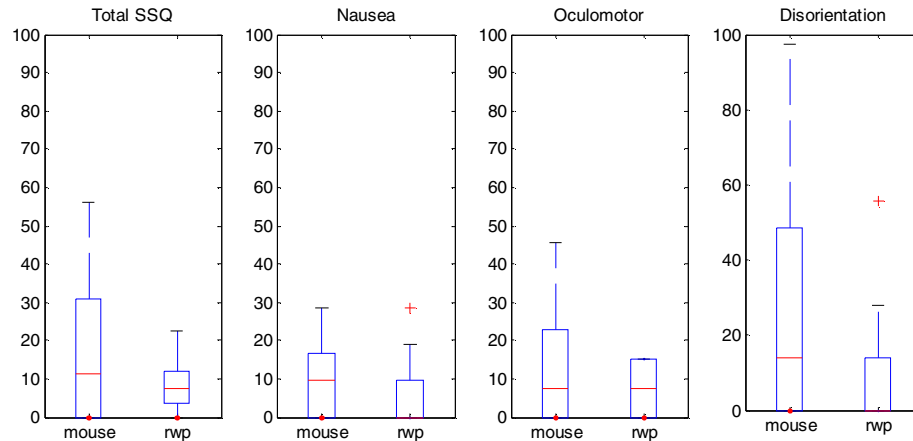


Figure 10.3. Box and Whisker plots of SSQ scores for hand-controller turning vs. Redirection. The leftmost plot shows the distribution of total SSQ scores. The 75th percentile of the total SSQ score for Redirection is less than that for turning with a hand-controller. The other three plots show the distribution of scores for the SSQ subscales—nausea, oculomotor, and disorientation. The oculomotor subscale relates to eyestrain and similar symptoms caused by the optics of the VE system. Since both groups (hand-controller and Redirection) used the same VE system, one would expect only the nausea and disorientation subscales to show differences. These plots are consistent with this expectation.

In conclusion, I cannot use the SSQ to show that Redirection does not increase simulator sickness compared to real walking (without using Redirection or any other form of virtual rotation), because I do not have a reasonable estimate of standard deviation in SSQ levels. But compared to the alternative (virtual rotation controlled by the user), Redirection results in less simulator sickness. *A fortiori*, Redirection does not unacceptably increase the level of simulator sickness in the user, who must locomote by some means.

Chapter 11: Future Opportunities

In the course of this dissertation research, I have encountered several topics that merit further investigation. In Chapter 5 and Chapter 6 I proposed several enhancements to Redirection and I believe they merit being tested. I discuss several more ideas in this chapter, in descending order of promise.

11.1 Redirected Avatar Limbs



Figure 11.1 - The user's tracked virtual hand penetrates the virtual antique radio [from Burns 2003].

Immersive VEs in which the position and orientation of the user's arms, relative to her head and the rest of her body, are tracked, have a particular problem representing virtual objects and their physical interactions with the user's (virtual) avatar. When her real hand passes into the space occupied by a virtual object, the virtual hand passes through the virtual object. This can be disruptive to the user's experience—a *break-in-presence*, to use Slater's term [Slater 2000]. The solution in desktop 3D graphics applications, where the user controls the virtual hand with a keyboard or joystick, is to simply prevent the virtual hand from penetrating the virtual object. However, the immersive VE system cannot prevent the real hand from entering the virtual object's boundary. In small spaces, robotic force-feedback devices can apply real force to the user's hand. However, the working volume of force feedback devices is limited. It ranges from about 2500 cubic cm (that of the user's hand moving at the stationary wrist) in the case of the Sensable Phantom Desktop™, to a

cubic meter in the case of the Argonne Remote Manipulator [Brooks 1990]. This approach of using an active haptic device is impractical in large VE scenes. Another approach is to build approximate real objects to match the virtual objects. This technique, known as *passive haptics* [Insko 2001] very effectively enforces the boundaries of static virtual objects and increases the user's level of presence. However, it does not lend itself to dynamic virtual objects. Examples include virtual basketballs and computer-generated virtual actors. With Redirected Walking, the entire virtual scene is rotated by the VE system, and thus even the walls and furniture are unsuitable for passive haptics.

A Redirected Avatar Limbs technique would address the above situation by allowing the virtual avatar hand to drift imperceptibly from the position of the user's real hand. As the user's virtual hand bumps into the virtual countertop, the virtual hand is kept from penetrating the object, while her real hand enters it (but the real hand is not displayed to the user inside the HMD). The real and virtual hands separate (Figure 11.2). When the user's hand changes direction and starts to recede from the virtual countertop, the virtual hand, which is on top of the counter, must lift off the counter immediately, maintaining the separation between the virtual and real hand. Then the separation is slowly and imperceptibly reduced until it is gone. Eric Burns is actively pursuing this idea [Burns 2003; Burns 2005].




Figure 11.2 - As the user lowers her hand onto a virtual tabletop, her real hand location may penetrate the virtual table. The VE system displays her virtual hand such that it stays on top of the table, while her real hand is actually beneath the virtual table. Note that the transparent hand is shown for purposes of illustration only—it is not visible to the user [from Burns 2005].

11.2 Wireless HMD VE System

In our real-walking VEs, the tethers are an impediment. They force the VE scene designers to plan for the user's path ahead of time (allowing room for the cables, for example), they require two people to carry and manage the cables during the VE session, and still tug at the user and restrict her motion. When the user is

freely exploring an arbitrarily large virtual scene (Chapter 6) and rotating multiple revolutions in the real world, the tethers will be an even more severe impediment. The cables might wrap around the user, and it will be harder for the cable-carriers to predict where the user will move. Components required to build (at reasonable effort and cost) a wireless VE system are just now becoming available. Instead of sending the high-bandwidth video (for 60 frame/s VGA streams, I estimate the bandwidth required is roughly 20 times that of a NTSC television channel) across a wireless link from the image generator to the HMD, it may be more practical to have the user wear the image generator. As part of the dissertation work, I built a prototype mobile image-generator to demonstrate the feasibility of this approach (Figure 11.3). In the next year or two, I expect fully wireless VE systems, small enough to wear around the waist, to be commercial available.⁴⁶




Figure 11.3 - A prototype wearable image generator I built in 2001. If built using newer components, an equivalent image generator would fit into a belt-worn camera bag.

⁴⁶ As of this writing, the Quantum3d Thernite™ wearable computer and Intersense IS-1200™ tracker are promising candidates.

11.3 The Effects of Redirection on Spatial Cognition

Spatial cognition refers to the abilities of a person (or animal) to understand or use spatial relations—the position and orientation of objects relative to each other and to oneself. Spatial cognition relates to navigation and path integration. *Path integration* refers to the ability of an animal to keep track of its location as it locomotes from point to point, and to determine the direct path back to the starting location without having to revisit the intermediate places (Figure 11.4).

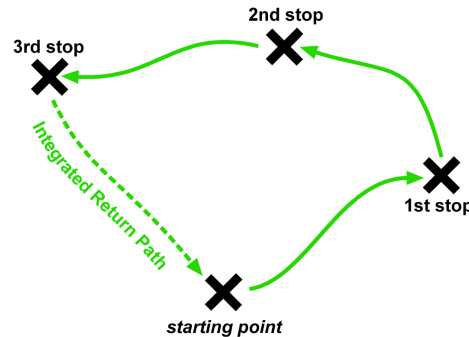


Figure 11.4 - Determining the direct path back to the starting place without revisiting the intermediate stopping points requires path integration [adapted from Breed 2001].

For some applications (such as training a person to find the exit in a burning building), it is important to consider how the VE system might affect the user's spatial cognition. I have not examined how Redirection affects spatial cognition, nor how its effects are different from the effects of non-Redirection VE systems on spatial cognition.

If one's lab is large enough to allow the user to perform the task in the virtual scene with real walking, there is little reason to use Redirection! If the lab is not large enough, then one must consider what the practical alternative techniques of locomotion are (e.g., flying, treadmill, turning with a hand-controller). Redirection's effects on spatial cognition should be compared to those alternatives.

11.4 SSQ Scaling for VEs and General-Population Users

SSQ scores from VEs are typically higher than those from military flight simulators. But this does not necessarily mean VEs induce more sickness than flight simulators. The SSQ was designed and validated on data from military pilots exposed to military flight simulators. Furthermore, 95% of these military pilots were male [Kennedy 2003b].

As suggested in Chapter 10, there are two plausible explanations for why VEs have higher SSQ scores:

- 1) VEs are more sickness-inducing than military flight simulators or;

- 2) male military pilots report fewer symptoms, and a different profile of symptoms, than general population or university students report;

I personally believe the latter—subjects exposed to VEs which should, theoretically, induce very little sickness, have higher post-exposure SSQ scores than even the troublesome flight simulators, and even have high pre-exposure SSQ scores! Furthermore, pilots are in better physical condition and are motivated to under-report symptoms.

It would be useful to resolve this. First, it would be interesting to be able to compare sickness from flight simulators to sickness from VEs. Second, absolute SSQ scores from flight simulators have diagnostic value. For example, Kennedy suggests that simulators whose users have SSQ scores of 20 are “problem simulators” [Kennedy 2003a] whereas individual users who have SSQ scores above 15 should consult a physician and those with scores above 20 should not be allowed to leave the facility until they recover [Kennedy 1992]. But these numbers may not apply to VEs. It would be useful to have similar criteria for VE, both for scientific purposes and for establishing safety guidelines.

One way to investigate this issue would be to conduct a within-subject study between VEs and military flight simulators. The same users would be exposed to VEs and military flight simulators. The subjects would ideally consist of military pilots, university students, and representatives of the general population; though it may be unacceptable to allocate expensive pilot time or risk pilot readiness for this research question.

A cheaper alternative would be to investigate if military pilots and general population users give the same subjective ratings to their symptoms. Even with no sickness, there is some baseline rate at which a person will report symptoms [Kennedy 1993]. One could compare the pre-exposure SSQ scores collected from military pilots with the pre-exposure scores of general population or university student VE users.

11.5 Better Walking-In-Place Implementation

In the first version of RWP, the system used a neural net to detect when the user was stepping. This did not work for some users. In addition, the neural net required roughly half a second to detect that a person had stopped walking. This sometimes caused users to walk through the virtual walls! RWP-II used a head-mounted accelerometer. This eliminated the false steps, but some users found they had to step vigorously in order for the step to be detected. I propose that accelerometers closer to the user’s feet (on the foot, leg, or hips) would work better, because there would be less damping of the foot-strike by the user’s body. There exist

commercially available wireless accelerometers that could mitigate the user's extra encumbrance. Another enhancement to Walking-in-place would be to have the length of the virtual stride increase as the user steps more vigorously, allowing the user to control virtual stride-length.

A bigger problem with walking-in-place, though, is that it is cumbersome to avoid obstacles. Imagine a situation where there is a chair in the middle of a person's path to a window. In the real world and when really walking in a VE, the person can quickly step sideways around the chair. However, our implementation of walking-in-place moves the user in the direction she is facing. To avoid the chair, she must turn perpendicular to the direction of her target, the window, take steps, and then, having cleared the chair, turn back toward the window. Some users in the RWP experiments found this unintuitive and cumbersome and remarked it was hard to avoid obstacles. Some walking-in-place implementations allow the user to make gestures to indicate she wants to side-step [Templeman 1999]. Another solution would be to allow users to really side-step virtual objects, letting them drift from the position in the real world where they started (i.e., the center of the CAVE). In the RWP experiments, some users actually did this (without being instructed), and it seemed to work well except that a few users approached the walls of the CAVE.

The final problem with walking-in-place is that some users unknowingly move forward in the real world when taking virtual steps—with each foot-strike, they creep forward. This seemed to be more common in users who reported higher presence. To combat this, and to allow users to really side-step around obstacles, a way of imperceptibly translating users would be useful.

Can the user be made to translate in the real world without noticing? I informally investigated slowly translating the virtual scene, while the user was stepping, to nudge her toward the center of the CAVE floor. When the virtual scene moved sideways (with respect to the user's orientation), it was very noticeable, even at quite low velocities.

11.6 Redirection of Walking Speed

With Redirected Walking, it would be useful to change the user's walking speed imperceptibly, as well as the user's orientation. This could allow for even larger virtual scenes. I implemented velocity scaling, where the user's virtual velocity was proportionally larger in the virtual environment than the real velocity (experiment i3). In informal testing, users found this noticeable and very strange. Some described it as if they sped up too quickly when they started walking, and then slowed down too quickly as they stopped walking. They did not

complain about the mismatch in virtual and real speed when walking at constant velocity. Jaekl [2002] reports that when the head's forward and backward translational movements are scaled (e.g., when the user moves forward 1 cm in the lab, the virtual viewpoint moves 2 cm), the user still feels as if the virtual scene is stationary.

A better technique might be to scale the user's acceleration (perhaps only in certain frequency bands). Acceleration changes very quickly compared to position and velocity. I did not implement this idea, as the positional reports from our head-tracker were too noisy to be used to compute acceleration, and the VE system we used was not fast enough (72 report/s head tracking rate and 20 frame/s video rate) for me to expect this to work.

11.7 Using Virtual Distracters

Dorian Miller recently investigated adding things inside the virtual scene to distract the user so that the virtual world can be rotated faster without the user noticing. In his VE, the user's task was to chase a dragon. While the user was busy with the dragon, the virtual scene rotated faster. People who are focused on a task have higher rotation detection thresholds [Rolfe 1986]. In some settings this is not appropriate—an architect's client, evaluating a proposed building interior, should not be distracted from examining the building. On the other hand, in entertainment applications, the task could be continually modified such that the system creates engaging situations just when higher levels of virtual scene rotation are needed to keep the user from exceeding the lab boundaries.

11.8 High Body Momentum

When a person is running or carrying a heavy backpack, her momentum is greater than in the situations I have thus far simulated. As a person turns while running, her body compensates for her momentum, and she leans into the turn [Imai 2001]. What are the implications of unknowingly turning a user when they have great momentum? This merits further research.

11.9 The Effects of Spatial Audio

In experiment RW, the implementation included 3D spatial audio. I firmly believe this increased the user's sense of presence and allowed for greater undetected rotation rates. Audio cues alone can cause a person to feel as if she is moving [Lackner 1977a]. It would be interesting to measure the contribution of the spatial

audio relative to that of non-spatial audio or no audio at all. However, such an experiment will be complicated by several factors.

Virtual acoustic models are different from visual models. In order to speed up acoustic simulation, the acoustic virtual model often must approximate the scene with simpler geometry than does the visual model. Acoustic models must include surface properties. For example, a glass surface reflects sound very differently from a plush sofa. Problems such as modeling mistakes and lack of fidelity in the acoustic model are, in my experience, much harder for VE practitioners to detect than problems in the visual model.

Finally, the acoustic and visual models must be registered in time and space with each other. How close to each other do the scale and orientation of the two models need to be? What if the level of registration varies from person to person? (For example, the headset's headphones in our system were not rigidly attached to the head-tracking sensor, in order to accommodate different users' head shapes and sizes). Also, what is the difference in latency between the VE's acoustic and visual systems?⁴⁷

⁴⁷ In flight simulators certified for training commercial pilots (FAA Level D), the latencies of the different sensory cues must be matched (e.g., the visual feedback of a maneuver must be presented at the same time as the motion feedback of the maneuver) [Strachan 2001].

Chapter 12: Guidelines for Developers

The observations from this work and principles from self-motion perception literature lead to guidelines for VE practitioners wishing to make use of Redirection. Many of these guidelines are collected (and summarized) from other chapters, which are referenced.

The first group of guidelines would not only make Redirection more effective, but also improve any VE, not just one using Redirection.

12.1 Guidelines for all VEs

1) **Avoid color-sequential displays.** There are many displays (both HMDs and projectors) that cycle among presenting the red, blue, and green video channels, at some output rate higher than the input video's refresh rate. For example, an HMD might accept 60 frame-per-second video, and display cycling frames at 180 Hz in this sequence: read input frame, display red channel for 5 ms, display green channel for 5 ms, display blue channel for 5ms, read next input frame, etc. In this example, the blue channel would have 10 ms more latency than the red channel. When the user turns her head quickly, the red, green, and blue channels will appear to separate in a distracting manner.

2) **Give preference to HMDs with a wider field-of-view.** In the seven years since I started this work, HMDs appear to be increasing in resolution, but decreasing in field-of-view (FOV). The Virtual Research V8™ has a 60-degree diagonal FOV, whereas the Sony Glasstron's™ is 12 degrees. This is troubling, as a wider FOV increases the user's sense of presence and spatial orientation [Arthur 2000]. A wider FOV is also more likely to evoke a sensation of self-motion, which is essential for Redirection to work [Prothero 1998]. If necessary, I would select an HMD with a lower resolution and wider FOV over one with a higher resolution but narrower FOV.

3) **Furnish the virtual scene with rich spatial detail** (Chapter 4). A virtual scene with large areas without detail is problematic. Optical flow requires structure in the visual field. If the user's view is such that only a featureless portion of an evenly lit virtual wall is visible, she will not have any visual cues of motion

when she or the virtual scene is moving. I have experienced this myself in early versions of the virtual scenes used in this research. In this situation, Redirection will not work. Furthermore, in a CAVE, featureless virtual objects are visually overpowered by the blemishes and seams in the CAVE's projection surfaces, and the real world intrudes into the user's consciousness. But even for VEs that do not use Redirection, a featureless view is disorienting to a user. I have observed this to be the case especially with inexperienced VE users. All surfaces that could fill the user's view should be textured or otherwise detailed.

4) **Present a spatial audio scene** (Chapter 7). The use of spatial audio dramatically improves the immersion of any VE. Furthermore, it makes Redirection more likely to work and masks presence-breaking noises from the real world.

5) **Use circumaural headphones.** Having the user wear sealed, circumaural headphones, even if the VE system does not present any audio cues, also helps block real-world noises. I believe they work better than noise-canceling headphones, because the real-world noises are usually aperiodic.

6) **Use a wireless VE system** (Chapter 11). While the user is walking about, the cables tug at her, cause breaks in presence, and restrict her motion. With Redirection, the cables become even harder to manage and more dangerous.

7) **Avoid speaking directly to the user** (Chapter 7). Real-world voices are a break-in-presence. With Redirection, the voice cues the user as to her spatial orientation in the lab, and can reduce the effectiveness of Redirection. When it is necessary to instruct the user, consider having voice recordings play in the headphones instead.

12.2 Guidelines Specific to Redirection

8) **Accurately model the center of the user's head** (Chapter 5). The system must rotate the virtual scene about the center of the user's head; otherwise the optical flow pattern will not fool the user into believing that she is turning instead of the virtual scene turning around her. In a CAVE, this means the rotation must be about the tracked head position, and not some fixed center point in the CAVE (as the user might drift away from the CAVE center). In an HMD system, the transformation between the head tracking sensor and the head's center of yaw rotation must be calibrated accurately and potentially for each user. Chapter 5 explains one technique for doing this.

9) **Use a veil to block out the real world** (Chapter 7). In HMD VE systems, the user must not be able to see the real lab if Redirection is to work. If the HMD, by itself, cannot ensure this, consider adding a veil. The veil must not restrict the user's fresh air, however, because that would risk increasing the symptoms of simulator sickness.

10) **Rotate the visual scene less when the user is looking down** (Chapter 5). When the user is looking down, the visual scene rotations are much more noticeable.

11) **Adapt the virtual scene and task for use with Redirection.** In applications where it is acceptable to modify the virtual scene and user's task, the use of virtual objects that come close to the user should be minimized. For example, pillars in the middle of the virtual space should be avoided. Rotations are more noticeable when virtual objects are near the user (Chapter 5), and the user might have difficulty avoiding obstacles (Chapter 11). Just as with a real building, the user is less likely to feel disoriented if the virtual scene has strong orientating cues (e.g., not having four identical walls in the virtual room). Finally, the user should have a task to perform in the virtual scene. Users engaged in a task are less likely to notice the rotations [Rolfe 1986].

12) **Blank the display during lapses in tracking** (Chapter 7). Occasionally, tracking systems temporarily lose the position of the user's head. The VE system should detect this and blank the display. Otherwise, the user might continue walking and may end up in a position from which it is hard for Redirection to recover.

13) **Use every tracker report to compute head velocity.** Redirection requires the user's angular head velocity, but most tracking systems do not report it. Instead, the velocity must be computed from the tracker's orientation reports, and this amplifies any tracking noise. When the tracking rate is greater than the graphics frame-rate, most VE systems discard the tracker samples that arrive while the system is computing the next graphics frame. To minimize noise, error, and response time, all tracker reports should be used in computing head velocity.

14) **Consider using a foot-strike sensor** (Chapter 8). Higher rotation rates are possible the faster the user is stepping.

15) **Detect when the user is about to exceed the lab boundary.** As a fail-safe, the system can detect when the user is about to leave the tracked area or lab boundary. When this happens, the system could

automatically react and blank the display or otherwise warn the user. Alternatively, the system could instead instruct the user (using prerecorded verbal messages) to stop walking and then to turn in one direction, and then turn back. As described in Chapter 7, as the user does this, the system can change her real walking direction in the lab such that she heads away from the lab boundary. Meanwhile, she believes that she stopped, turned, and then continued walking in the same direction as before.

12.3 Guidelines for Redirected Walking-in-Place

16) **Track the user's torso.** When walking-in-place, the viewpoint should move in the direction in which the torso is pointing, rather than in the direction in which the head is pointing. The former is much more natural.

17) **Consider training the user to walk-in-place without creeping forward** (Chapter 11). Otherwise, they might run into a cave wall. In my experience, making the user aware they may creep forward seems to reduce its occurrence. The effect of this awareness on presence is unknown.

18) **The open back wall of the CAVE should be dark and featureless.** During experiment RWP-I, several pilot subjects reported being distracted by the red LED on the video camera that was used to record the session. When the opening comes into the user's field-of-view, she is more likely to be distracted and to remember it if there are visible features there. Heavy black curtains are one way to make the opening dark and featureless.

APPENDIX: LAPLACE ANALYSIS BACKGROUND

1 Introduction

This appendix briefly introduces the Laplace analysis techniques used in this dissertation. I expect that most of my readers are familiar with describing digital filters by their frequency characteristic (e.g., a low-pass filter) but are less familiar with time-constant description of analog filters. Laplace analysis and time-constants are particularly important in modeling the semicircular canals (SCCs) and the podokinetic system. This appendix uses the SCCs as an illustrative example. This does not cover Laplace analysis in much detail or precision. Complex numbers, the s-plane, impulse response, digital filters, and analogies to electrical circuits are omitted.

2 The Semicircular Canals

Each semicircular canal senses rotation of the head about a single axis. I discuss the arrangement, structure, and operation of the SCCs in detail in Chapter 5. In brief, each SCC is a hollow toroid of bone filled with fluid. Each SCC has a single flap extending across the fluid-filled interior of the toroid. As the head begins turning, the fluid's inertia causes it to exert a force on the flap, causing the flap to distort and bend. Sensory cells in the base of the flap transduce the distension of the flap. The inertia of the fluid and the spring-like nature of the flap cause it to return to its resting position when under constant angular velocity.

3 Feedback Systems

Laplace techniques are a way of analyzing and designing systems governed by differential equations. The rate at which the output of the system changes is (at least partially) determined by the output level. For example, the quantity of bacteria in a colony that is reproducing (without any limits on food or other resources), grows exponentially (a differential function):

$$\frac{dx}{dt} = \lambda x$$

where x is the quantity of bacteria as a function of time and dx/dt is the growth rate of the colony. Each individual bacterium produces offspring at a fixed rate, so the number of new offspring produced at any point in time, in the entire colony, is determined by the number of bacteria at that time.

Many kinds of systems can be modeled with differential equations. Examples include predator-prey relationships, analog electrical circuits, radioactive decay, economics, and mechanical systems of springs and pendulums. The semicircular canals can be modeled as pendulums.

4 Further Restrictions on Systems

In order to use Laplace analysis, not only must the systems be differential, but they must also satisfy several other constraints. First, a system must have an input and output, both of which are functions of time (Figure 1). Also, the system must be *linear* and *time-invariant*. I explain these two terms below. A linear *filter* is a system that satisfies these constraints.⁴⁸




Figure 1 – A graphical representation of a filter.

Two common ways of visualizing and analyzing a filter are to examine its *step response* and its *frequency response*.

5 Step Response

The *step response* is the output of a filter to an input which suddenly increases from zero to one and then stays at one (Figure 2). The property of being time-invariant dictates that, no matter when in time the input step occurs, a filter's step response will always be the same (Figure 2).

⁴⁸ Here, I use the term *filter* to describe an analog system. This is different from a digital finite impulse response (FIR) filter that is commonly used in computer graphics, where one convolves a sampled signal with a sampled kernel




Figure 2 - The step response of an arbitrary low-pass filter. The input is shown in blue, and the filter's output in green. Because the filter is time-invariant, shifting the time of the input has no effect on the shape of the step response.

Figure 2 and Figure 3 show the step responses of an arbitrary *low-pass filter*, while Figure 3 also shows an arbitrary *high-pass filter*. The rationale for these filter names is described below.




Figure 3 - The step-response of an arbitrary low-pass filter (left) and high-pass filter (right).

6 Exponential Decay

The step response of many filters can be modeled as exponential decay. In the filters presented above, the step-response can be fully described with just a single parameter—the *time-constant*.

The formula for exponential decay is:

$$\text{output} = \frac{1}{e^{t/\tau}}$$

where τ is the time-constant, and t is the time since the input's step. In the step-response of a high-pass filter, the time-constant is the time required for the output to drop from any value to $1/e$ (~37%) of that value.

In the step-response of the low-pass filter, the output increases from 0 to 1, but the difference between the output and 1 decays exponentially. More specifically, the output of the low-pass filter is:

$$1 - \left(\frac{1}{e^{t/\tau}} \right)$$

Simple filters, fully described by a single time-constant are *single-pole* filters. Below I will discuss more complex filters, governed by multiple time-constants and multiple poles. In the case of single-pole filters, the time-constant can be measured from a graph of the step response.

The SCCs sense the head's angular velocity. The SCCs' step response is a low-pass filter exponential curve as illustrated in Figure.4. The time-constant that describes the SCC's initial response is called the *response time* or *short time-constant*.




Figure.4 - The output of a low-pass filter with time-constant T_1 , when given a step up and step down in rotational velocity. This illustrates the onset time (or short time-constant or latency) of the SCCs. Note T_1 is greatly exaggerated for purposes of illustration.

When presented with sustained constant angular velocity, the output of the SCCs exponentially decays toward zero. It behaves like a high-pass filter (Figure 5). The time-constant describing this decay is called the *recovery time* or *long time-constant*.




Figure 5 - The output of a single-pole, single-zero high-pass filter with time-constant T_2 when given a step up and step down in rotational velocity. This illustrates the recovery time (or long time-constant) of the SCCs.

The exact value of the short and long time-constants varies from publication to publication. I summarize these values in Chapter 4. For the examples in this appendix, I use the values from Howard [1986b]: T_1 (short time-constant) = 3 1/3 milliseconds and T_2 (long time-constant) = 10 seconds. .

7 Linearity: Response of Filters in Series




Figure 6 - Filter composition: Filter C is composed of Filter A and B. A's output is B's input.

Given two independent filters A and B, one can create a new filter C in which the output of a new filter is made by feeding the output of A into the input of B. For any given input signal, $i(t)$, the output of the new filter C is simply the output of A on input $i(t)$ multiplied by the output of B on input $i(t)$. For example, the step response of filter C is the step response of filter A multiplied by the step response of B. The step response of the complete SCC is the product of the short and long time-constant step responses. This is illustrated in Figure 7. The product of two single-pole filters is a *two-pole* filter.




Figure 7 - The output of the SCCs, computed by composing the T_1 low-pass and T_2 high-pass filters. Note that T_1 is greatly exaggerated for purposes of illustration.

The linear property dictates that the order of composing the filters does not matter. The output of A fed to B produces the same response as the output of B fed to A.⁴⁹

⁴⁹ There is a practical exception to this rule, which I discuss later in a footnote to section 14 of this appendix.

8 Frequency Response

8.1 Signal Analysis in the Frequency Domain

All repeating functions can be decomposed into the sum of sinusoids of varying frequency. Figure 8 shows a square wave decomposed into sinusoids. The sinusoid at any particular frequency has a specific amplitude and phase (Figure 9). A *Fourier spectrum plot* for a given repeating signal is a plot of the magnitude and phase of the sinusoid at each frequency.




Figure 8 - A square wave decomposed as the sum of a series of sinusoids. As higher-frequency and lower-amplitude sinusoids are added to the series, the sum of the series becomes a better approximation of the square wave.




Figure 9 - Three sinusoids of the same frequency. The green has an amplitude of 2, and the red has a phase lead of 90 degrees.

8.2 Filter analysis in the Frequency Domain

Not only can signals be analyzed in the frequency domain, but so can filters. Filters have an input and output signal. The *frequency response* of a filter is, for every frequency, the output sinusoid's amplitude divided by the input amplitude. This is the filter's *magnitude*. Figure 10 shows the frequency response for an arbitrary low-pass filter. A magnitude greater than 1 indicates that the filter is amplifying certain frequencies such that those frequencies have greater amplitude in the output than in the input. A magnitude of less than 1

indicates that the filter is blocking (or attenuating) input frequencies from the output. This is where the terms *high-pass* and *low-pass* filter originate—a high-pass filter allows the higher-frequency content to pass while *attenuating* the lower frequency content, whereas a low-pass filter allows the lower-frequency content to pass while attenuating the high-frequency content.




Figure 10 - The output of arbitrary idealized high- and low-pass filters when the input is a square wave (left), and the frequency response of those filters (right). The cutoff frequency is labeled *cf*.

Figure 10 shows the frequency response of idealized high- and low-pass filters and shows the output (or response) of those same filters when fed a square wave. Note that the low-pass filter “rounds off” the vertical edges of the square wave. The vertical edges are high frequency content and do not pass through the filter, whereas the high-pass filter preserves the vertical edges but not the straight horizontal parts of the square wave (where the input value is constant). A constant value has a frequency of 0 cycles/second,⁵⁰ thus a high-pass filter does not allow it through.

⁵⁰ A signal with a constant value (or a constant offset in addition to other, higher-frequency content) is referred to as having *DC* content—content at the zero frequency. DC is an abbreviation for *direct current*. This terminology comes from electrical engineering.

The filters in Figure 10 are idealized in that they require the filter to know what the input signal does in the future. For example, in Figure 10, the low-pass output starts increasing *before* the vertical edge (the step up). This cannot be; it defies *causality*. Figure 11 shows the square-wave response of more realistic filters (meaning more like the filters one finds in physical, real-time, systems), but with the same frequency responses as those in Figure 10.




Figure 11 - The square-wave response of the same steep roll-off filters as in Figure 10, with idealized and realistic phase offsets.

8.3 Phase Offset

In reality, the output of the filter is always delayed (as shown in Figure 11). Evidence of this delay is visible in the phase response of the filter. The *phase offset* of a filter shows, for each frequency, how the phase of the input is altered. A positive phase offset is *leading*, and a negative phase offset is *lagging*. Thus, a -90 degree phase means the sinusoids are delayed in time by 90 degrees. A filter's phase offset and frequency response are often plotted together.

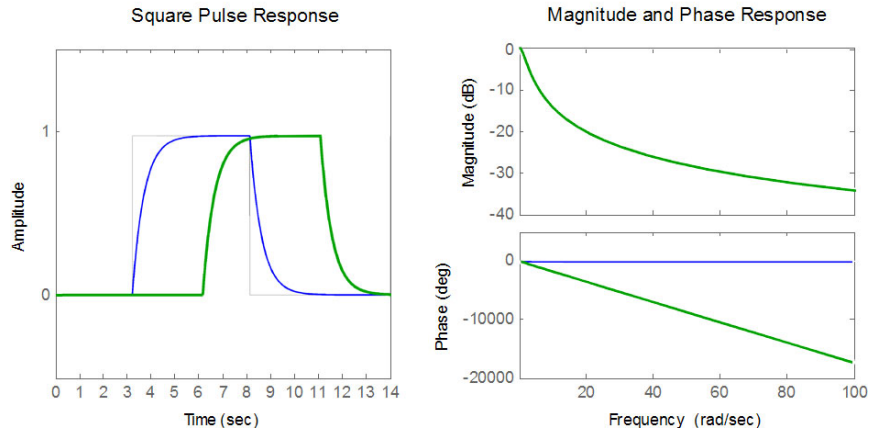


Figure 12 - The phase offset and square pulse output of an arbitrary filter with no delay (upper) and a fixed-time delay (lower). The phase offset of the fixed-time delay filter is a straight line, with higher frequencies having a greater phase lag (measured in degrees, not time). The greater the delay, the steeper the phase offset line.

A filter which only delays the input signal (and does not change any frequency's magnitudes) has a lagging phase (<0). Its phase offset plot is a straight line (when plotted against the linear phase and linear phase offset angle axes), as in Figure 12. This is because the higher frequencies have a shorter period, and thus the constant time delay will be a greater proportion of a high-frequency sinusoid's period than that of a lower-frequency sinusoid (Figure 13).




Figure 13 - A fixed time duration (30ms, for example) constitutes a greater phase offset for higher frequency sinusoids.

8.4 Cutoff Frequencies

Low- and high-pass filters allow some frequencies to pass and attenuate others. The demarcation between which frequencies pass through and which are blocked is the *cutoff frequency*. *Roll-off* refers to the

steepness of the transition in magnitude between the frequencies that pass through and the frequencies that are blocked by the filter. The idealized filters depicted in Figure 10 exhibit a sharp roll-off in frequency response. A steep roll-off is desired in some engineering applications; complex filters are required to accomplish this. For example, the filters above in Figure 10 have ten poles each. In contrast, a single-pole filter (as in Figure 2 - Figure 5) has a much more gradual roll-off. Figure 14 shows the frequency response for the single-pole high- and low-pass filters. The term *cutoff frequency* is inappropriate. Some authors [Stark 1998] prefer the term *corner frequency*.




Figure 14 - The frequency response of low- and high-pass single-pole (right) filters compared to that of the more complex filters from Figure 10 and Figure 11 (left). Note that the single-pole filters (right) have a much more gradual magnitude roll-off.

9 Bodé Plots

A filter's frequency and phase response are often visualized with a Bodé plot. A Bodé plot shows the frequency response plotted against log magnitude and log frequency axes. In this space, the asymptotes for the filter's roll-off are straight lines. A single-pole filter frequency response can be approximated with two straight lines, with a vertex at the cutoff or corner frequency. This is why it is called the *corner frequency*. At the corner frequency, the magnitude of a single-pole filter is -3 dB. The more poles a filter has, the more complex it is. For each pole, there is another corner in the Bodé diagram.




Figure 15 - A single-pole filter's frequency response plotted on a log magnitude and a log frequency scale. At the cutoff or corner frequency (point B), the magnitude is -3 dB [from Stark 1998].

One can compute the output of two composed outputs by multiplying the individual step responses.

One can compute the Bode plot of the combined frequency response of two composed filters by simply adding (instead of multiplying) the Bode plots of the individual filters, because of the log magnitude scale. Composing a high-pass and low-pass filter results in a *band-pass filter*—content within a frequency band is allowed to pass, and content and frequencies above and below this band are attenuated.

Bode plots are convenient for filter analysis because of the fact that the filter's frequency roll-off can be approximated as straight lines and because it is easy to visualize the frequency response of multiple filters composed with each other.

10 Converting Between Exponential Decay and Frequency Response Representations

The first section of this appendix describes filters as exponential decay curves. The time-constant determines the rate at which the step-response decays. The time-constant also determines the frequency response of the filter. The longer the time-constant, the more slowly the output decays, and the lower the cutoff frequency. The exact relationship is:

$$\text{cutoff frequency} = \frac{1}{2\pi\tau}$$

If τ has units of seconds, then the frequency will be in hertz (cycles per second)

11 Transfer Functions: Laplace Domain Representations of Filters

A transfer function is a mathematical function that relates the output of a filter to its input. Transfer functions are often expressed in terms of the Laplace domain, which is represented as the variable s , instead of in terms of the time domain. In this dissertation, the reader should assume transfer functions to be in the Laplace domain unless otherwise specified.

Equation 1 - The poles and zeros form of expressing a transfer-function.

$$H(s) = \text{gain} * \frac{(s - z_1) * (s - z_2) * \dots * (s - z_n)}{(s - p_1) * (s - p_2) * \dots * (s - p_m)}$$

Equation 1 is a generalized form of a transfer function expressed in poles and zeros. The terms z_1 to z_n are called the *zeros*, and the terms p_1 to p_m are called the *poles*. There can be any number of zeros and poles. All filters described in this dissertation have their zeros equal to zero (if they have any zeros at all). Thus, the transfer function can be simplified to:

$$H(s) = \text{gain} * \frac{s^n}{(s - p_1) * (s - p_2) * \dots * (s - p_m)}$$

Earlier in the appendix, I referred to a single-pole filter as being a simple filter. A single-pole low-pass filter has the transfer function:

$$H(s) = \text{gain} * \frac{1}{(s - p)}$$

Given the time-constant of the low-pass filter, τ , the single pole, p , is $= -1/\tau$ and the cutoff frequency is $p/(2*\Pi)$. A single-pole high-pass filter actually has a single pole and a single zero:

$$H(s) = \text{gain} * \frac{s}{(s - p)}$$

The zero is actually equal to 0 ($z=0$), thus it is dropped from the formula above. The pole is the same value as for the low-pass filter ($p=-1/\tau$).

Another common form of a transfer function is:

$$H(s) = \text{gain} * \frac{s^n}{(\tau_1 - s + 1) * (\tau_2 - s + 1) * \dots * (\tau_m - s + 1)}$$

This is algebraically equivalent to the previous form (Equation 1), since for each pole, there is a corresponding time-constant, determined by the above equation. This is the form used in Howard’s text on the SCCs [Howard 1986b] which I reference heavily in Chapter 4.

The Laplace domain form of the transfer function is convenient for several reasons. First, it is a complete and succinct mathematical description of the filter. Second, filters can be composed simply by multiplying their Laplace-domain transfer functions. For example, the composition of the high- and low-pass filters described in Figure.4, Figure 5, and Figure 7 is:

$$H(s) = \frac{1}{(\tau_1 + s)} * \frac{s}{(\tau_2 + s)} = \frac{s}{(\tau_1 + s)(\tau_2 + s)}$$

Finally, the Laplace-domain transfer function allows one to integrate and differentiate signals very easily.

12 Filters That Compute the Derivative and Integral of a Signal

Integration can be modeled by a low-pass filter with an infinite time-constant and a phase lag of 90 degrees. The transfer function is $H(s)=1/s$. Figure 16 shows the Bodé plot and step response of such an integrator. The integral of $\cos(t)$ is $\sin(t)$. Figure 17 shows the sine wave follows the cosine wave by 90 degrees. This is why the integrator filter has a 90-degree phase lag.

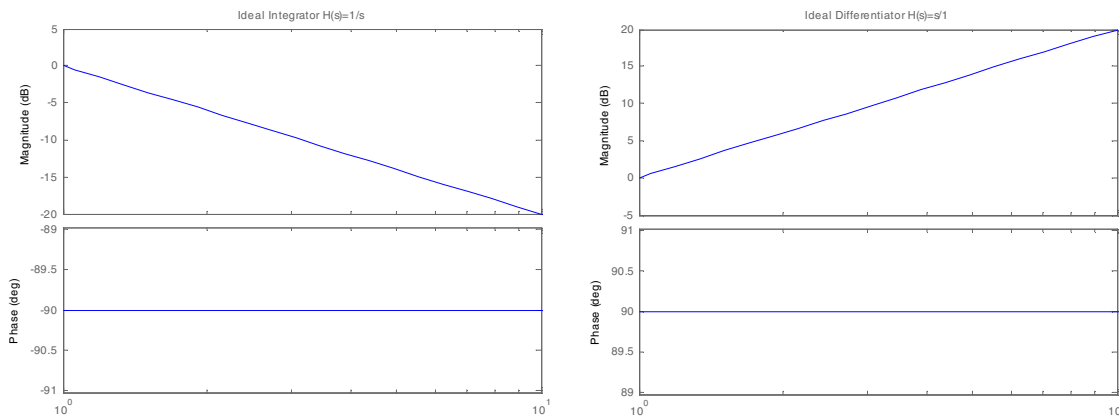


Figure 16 - Bodé plot of an ideal integrator (left) and differentiator (right). The integrator’s phase lags by 90 degrees, whereas the differentiator’s phase leads by 90 degrees.




Figure 17 - A cosine (green) and sine (blue) wave. The integral of $\cos(t)$ is $\sin(t)$, which lags 90 degrees behind. The derivative of $\sin(t)$ is $\cos(t)$, which leads 90 degrees ahead. This explains the 90 degree phase lag and lead of the integrating and differentiating filter, respectively.

Similarly, a differentiator is a high-pass filter with a phase lead of 90 degrees (a -90 phase offset). Its transfer function is $H(s)=s/1$. Figure 16 shows the Bode plot of a differentiator. The step-response is undefined because the derivative at the rising edge of the step is not continuous. The high-pass quality of the differentiator makes sense, as high-pass filters respond to quick changes in the input signal (which is high-frequency content) whereas high-pass filters output a value of zero to input signals that hold a constant value (and have no high-frequency content). The derivative of $\sin(t)$ is $\cos(t)$. Figure 17 shows the cosine wave leads the sine wave by 90 degrees, and thus the differentiator has a 90-degree phase lead.

Because integrators and differentiators are so easy to represent as transfer functions, one can intuitively detect the presence of them in the systems one is analyzing. For example, a high-pass filter has the transfer function $H(s) = s/(\tau +s)$ and the low-pass filter with the same time-constant has the transfer function $H(s)=1/(\tau +s)$. Since a differentiator is $H(s)=s/1$, the high-pass filter is simply the low-pass filter composed with a differentiator.

13 A Single Filter can act as Both an Integrator and Differentiator

Previously I discussed how the SCCs act as band-pass filters on the head's angular velocity. Howard models it with time-constants of 3 milliseconds and 10 seconds [Howard 1986b].




Figure 18 - A Bode plot of SCC response in terms of angular velocity.

For very low frequencies of angular velocity (below 0.1 Hz), the SCCs' Bode plot has 90-degree phase lead and high-pass filter characteristics. In this frequency range, they are acting as differentiators, and their outputs are proportional to the angular *accelerations* of the head.

For very high frequencies of angular rotation (above 100 Hz), the SCCs' Bode plot has a 90-degree phase lag and low-pass filter characteristics. In this frequency range, they are acting as integrators, and the outputs are proportional to the angular *displacements* of the head.

14 Converting a Filter to Operate on the Integral or Derivative of its Input

The model of SCCs presented thus far assumes that the input signal is the head's angular velocity. Since the SCCs are physical systems that respond to head rotational motion, one can just as easily model them as having angular displacement or acceleration as the input. Figure 19 shows the algebraic manipulation required to convert the SCC transfer function presented above (that works on angular velocity) to one that works on angular acceleration.

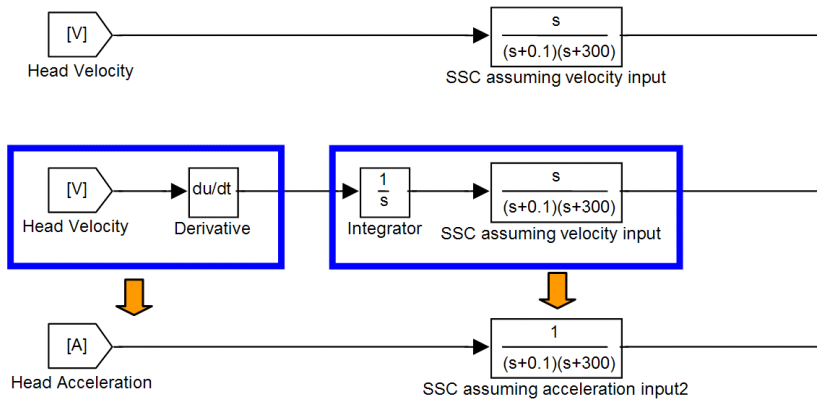


Figure 19 - Composing the SCC filter with an integrator turns it into one that accepts angular acceleration instead of velocity. The original transfer function accepts velocity (upper flow diagram). One can add a differentiator ($s/1$) and integrator ($1/s$) one after another without any effects, since they cancel each other out⁵¹ (middle flow diagram). Then one composes the differentiator with the head velocity input to get head acceleration, and compose the integrator with the original SCC filter to get a new SCC filter that accepts head acceleration instead.




Figure 20 - A Bode plot of SCC response in terms of angular acceleration.

⁵¹ By virtue of being linear, the order in which one composes filters should not matter. Earlier, in footnote 49, I mention that there is a practical exception to this rule. If the input has DC content and is fed to an integrator followed by a differentiator, it will not produce the same output as a differentiator followed by an integrator. If the input is a constant value, then the differentiator filter will output zero. Then the integrator filter will not be able to restore it back to the original constant value. Does this not violate the linearity constraint? It turns out an input signal with DC content is not a valid input signal. All inputs and outputs are assumed to be zero initially. Thus, under these constraints, there is no such thing as a non-zero constant-valued input.

In the acceleration SCC Bodé plot (Figure 20) one can see that, for the lowest frequencies, the SCC reports acceleration faithfully. For the middle frequencies, it integrates acceleration to output head velocity, and for the highest frequencies, it doubly integrates acceleration (since the phase is lagging 180 degrees and the magnitude slope is twice as steep as a single integrator) to output angular displacement.

Figure 21 shows the Bodé plot of an SCC filter that accepts orientation as input (angular displacement). In the very high frequencies, the SCCs report displacement faithfully, and for the very low frequencies, they doubly differentiate to report acceleration.




Figure 21 - A Bodé plot of SCC response in terms of angular displacement.

Nominal head movements are in the 0.1 to 5 Hz range [Howard 1986b]. Whereas one can describe the SCC response in terms of angular acceleration, velocity, or displacement, velocity is the most convenient. When described in terms of velocity, the head movements in the nominal range have no phase offset and a flat frequency response.

REFERENCES

- Arthur, Kevin (2000). Effects of Field of View on Performance with Head-Mounted Displays (Doctoral Dissertation), Computer Science. Chapel Hill
- Berthoz, Alain (2000). The Brain's Sense of Movement. Cambridge, Massachusetts, Harvard University Press.
- Blauert, Jens (1996). Spatial Hearing. The Psychophysics of Human Sound Localization. Cambridge, MA, MIT Press.
- Bowman, D. A., D. Koller and L. F. Hodges (1997). Travel in Immersive Virtual Environments: an Evaluation of Viewpoint Motion Control Techniques. Proceedings of IEEE VRAIS'97.
- Brandt, Thomas, E. R. Wist and Johannes Dichgans (1975). Foreground and background in dynamic spatial orientation. Perception and Psychophysics. **17**: 497-503.
- Breed, Michael D. (2001). "Path Integration." Retrieved May, 2005, from <http://www.animalbehavioronline.com/pathintegration.html>.
- Bridgeman, Bruce, A.H.C Van der Heijden and Boris M. Velichkovsky (1994). "A Theory of Visual Stability Across Saccadic Eye Movements." Behavioral and Brain Sciences **17**(2): 247-292.
- Brooks, Frederick P, John Airey, John Alspaugh, Andrew Bell, Randolph Brown, Curtis Hill, Uwe Nimscheck, Penny Rheingans, John Rohlf, Dana Smith, Douglass Turner, Amitabh Varshney, Yulan Wang, Hans Weber and Xialin Yuan (1992). Six Generations of Building Walkthrough: Final Technical Report to the National Science Foundation. Chapel Hill, Department of Computer Science, University of North Carolina.
- Brooks, Frederick P. Jr., Ming Ouh-Young, James J. Batter and P. J. Kilpatrick (1990). "Project GROPE-Haptics displays for Scientific Visualization." Computer Graphics **24**(4): 177-185.
- Burns, Eric and Sharif Razzaque (2003). "Avatar Redirection: An Exploration of Sensory Conflict." Presence-Connect Online Journal.
- Burns, Eric, Sharif Razzaque, A. Panter, Mary C. Whitton, M. R. McCallus and Frederick P. Brooks (2005). The Hand is Slower than the Eye: A quantitative exploration of visual dominance over proprioception. Proceedings of IEEE Virtual Reality 2005, Bonn, Germany, IEEE Computer Society.
- Cheung, Bob. (2000). "Non-visual Spatial Orientation Mechanisms." U.S. Air Force Research Laboratory for Spatial Countermeasures. Retrieved Oct 2002 from <http://www.spatiald.wpafb.af.mil/MechanismsStudies/cheung.ppt>.
- Chung, Bill (2000). (Personal Communication) Washout for motion-platform flight simulators. March 2002. Mountain View.
- Cobb, Susan V. G. and Sarah C. Nichols (1998). "Static posture tests for the assessment of postural instability after virtual environment use." Brain Research Bulletin **47**(5): 459-64.
- Coffin, Tom (2001). (Personal Communication) Number of IPTs. Oct. 3, 2001. London.
- Corso, John F. (1967). The Theory of Signal Detection. The Experimental Psychology of Sensory Behavior, Holt, Rinehart and Winston, Inc.: 445-460.

- Cruz-Neira, Carolina, D. J. Sandin and T. A. De Fanti (1993). Surround-screen projection-based virtual reality: The design and implementation of the CAVE. Proceedings of ACM SIGGRAPH '93.
- Darken, Rudolph P., William R. Cockayne and David Carmein (1997). The Omni-Directional Treadmill: A Locomotion Device for Virtual Worlds. Proceedings of ACM User Interface Software and Technology, Banff, Canada.
- Dichgans, Johannes and Thomas Brandt (1977). Visual-Vestibular Interaction: Effects on Self-Motion Perception and Postural Control. Perception. Held, R., H. Leibowitz and H.-L. Teuber. New York, Springer-Verlag. **8**: 755-804.
- Draper, Mark (1996). Can your eyes make you sick?: Investigating the Relationship between the Vestibulo-ocular Reflex and Virtual Reality. Seattle, Univ. of Washington, HIT Lab.
- Draper, Mark (1998). The Adaptive Effects of Virtual Interfaces: Vestibulo-Ocular Reflex and Simulator Sickness (Doctoral Dissertation), Department of Industrial Engineering, University of Washington, Seattle
- Duh, Henry Been-Lirn, Donald E. Parker and Thomas A. Furness (2001a). An "Independent Visual Background" Reduced Balance Disturbance Evoked by Visual Scene Motion: Implication for Alleviating Simulator Sickness. Proceedings of CHI '01, Seattle, ACM.
- Duh, Henry Been-Lirn, Donald E. Parker, J. Phillips and Thomas A. Furness (2004). ""Conflicting" motion cues at the frequency of crossover between the visual and vestibular self-motion systems evoke simulator sickness." Human Factors **46**: 142-153.
- Duh, Henry, D. Parker, J. Phillips and T. Furness (2001b). ""Conflicting" Motion Cues in the Frequency Range of Cross-Over between the Visual and Vestibular Self-Motion Systems Evoke Simulator Sickness (submitted)." Aviation Space Environmental Medicine.
- Earhart, Gammon M., Kathryn M. Sibley and Fay B. Horak (2004). "Effects of bilateral vestibular loss on podokinetic after-rotation." Experimental Brain Research **155**: 251-256.
- Fleming, Seay A., David M. Krum, Larry F Hodges and William Ribarsky (2002). Simulator sickness and presence in a high field-of-view virtual environment. CHI '02 extended abstracts on Human factors in computing systems, Minneapolis, ACM.
- Gibson, J. J. (1966). The sense considered as perceptual systems. Boston, Houghton Mifflin.
- Gibson, Simon, Roger J. Hubbard, J. Cook and T. L. J. Howard (2003). "Interactive reconstruction of virtual environments from video sequences." Computers and Graphics **27**: 293-301.
- Gordon, Carlos. R., William. A. Fletcher, Geoffrey Melvill Jones and Edward. W. Block (1995). "Is the stepping test a specific indicator of vestibulospinal function?" Neurology **45**(11): 2035-2037.
- Gregory, Richard L. (1966). The Eye and Brain: The Psychology of Seeing. London, Weidenfeld and Nicolson.
- Grossman, Gerald E., R. John Leigh, Eugene N. Bruce, William P. Huebner and Douglas J. Lanska (1989). "Performance of the Human Vestibuloocular Reflex During Locomotion." Journal of Neurophysiology **62**(1): 273-287.
- Hain, Timothy C. (2005). "Nystagmus." Retrieved July 2005, from <http://www.dizziness-and-balance.com/practice/nystagmus.html>.

- Hays, William L. (1963). Power of a Statistical Test. Statistics for Psychologists, Holt, Rinehart and Winston, Inc.: 269-277.
- Heeger, David. (1997). "Signal Detection Theory." Retrieved Jan 2003 from <http://white.stanford.edu/~heeger/sdt/sdt.html>.
- Hodges, L. F., Barbara O. Rothbaum, Rob Kooper, Dan Opdyke, Thomas Meyer, Johannes J. de Graaff and J. Williford (1994). Presence as The Defining Factor in a VR Application: Virtual Reality Graded Exposure in the Treatment of Acrophobia. Atlanta, GA, USA, Graphics, Visualization & Usability Center, Georgia Tech.
- Hollerbach, John M., Robert R. Christensen, Yangming Xu and Stephen C. Jacobsen (2000). Design Specifications for the Second Generation Sarcos Treadport Locomotion Interface. Haptics Symposium, Proc. ASME Dynamic Systems and Control Division, Orlando.
- Hollins, Mark (2001). (Personal Communication). Chapel Hill.
- Howard, Ian P. (1986a). The Perception of Posture, Self Motion, and the Visual Vertical (Chapter 18). Handbook of Perception and Human Performance. Boff, K. R., L. Kaufman and J. P. Thomas. New York, Wiley-Interscience. **1**: 18.1-18.50.
- Howard, Ian P. (1986b). The Vestibular System (Chapter 11). Handbook of Perception and Human Performance. Boff, K. R., L. Kaufman and J. P. Thomas. New York, Wiley-Interscience. **1**: 11.1-11.26.
- Imai, Takao, Steven T. Moore, Theodore Raphan and Bernard Cohen (2001). "Interaction of the body, head and eyes during walking and turning." Experimental Brain Research **136**(1): 1-18.
- Insko, Brent (2001). Passive Haptics Significantly Enhances Virtual Environments (Doctorial Dissertation), Department of Computer Science, University of North Carolina at Chapel Hill. Chapel Hill
- Iwata, H. and Y. Yoshida (1999). "Path reproduction tests using a torus treadmill." Presence: Teleoperators and Virtual Environments **8**(6): 587-597.
- Jacobson, Jeffrey, Mark S. Redfern, Joseph M. Furman, Susan Whitney, L., Patrick J. Sparto, Jeffrey B. Wilson and Larry F. Hodges (2001). Balance NAVE: A Virtual Reality Facility for Research and Rehabilitation of Balance Disorders. Proceedings of the ACM Symposium on Virtual Reality Software & Technology, Banff.
- Jaekl, P. M., R. S. Allison, L. R. Harris, U. T. Jasiobedzka, H. L. Jenkin, M. R. Jenkin, J. E. Zacher and D. C. Zikovitz (2002). "Perceptual stability during head movement in virtual reality."
- Jahn, Klaus, Michael Strupp, Erich Schneider, Marianne Dieterich and Thomas Brandt (2000). "Differential effects of vestibular stimulation on walking and running." NeuroReport **11**(8): 1745-1748.
- Jürgens, R., T. Boß and W. Becker (1999). "Podokinetic after-rotation does not depend on sensory conflict." Experimental Brain Research **128**: 563-567.
- Kennedy, Robert S. (2003a). (Personal Communication) SSQ scores. Oct 16, 2003. Chapel Hill.
- Kennedy, Robert S., Julie M. Drexler, Daniel E. Compton, Kay M. Stanney, D. Susan Lanham and Deborah L. Harm (2003b). Configural Scoring of Simulator Sickness, Cybersickness and Space Adaptation Syndrome: Similarities and Differences. Virtual and adaptive environments: Applications, implications, and human performance. Hettinger, L. J. and M. W. Haas. Mahwah, NJ, Lawrence Erlbaum Associates, Inc.: 247-278.

- Kennedy, Robert S., Lane, Berbaum and Michael G. Lilienthal (1993). "Simulator sickness questionnaire (SSQ): A new method for quantifying simulator sickness." International Journal of Aviation Psychology **3**(3): 203-220.
- Kennedy, Robert S., N. E. Lane, Michael G. Lilienthal, K. S. Berbaum and L. J. Hettinger (1992). "Profile analysis of simulator sickness symptoms: Application to virtual environment systems." Presence: Teleoperators and Virtual Environments **1**(3): 295-301.
- Kennedy, Robert S. and Michael G. Lilienthal (1995). Implications of Balance Disturbances Following Exposure to Virtual Reality Systems. Proceedings of the IEEE VRAIS'95 Conference, Research Triangle Park, NC, IEEE Press.
- Kennedy, Robert S., Kay M. Stanney and William P. Dunlap (2000). "Duration and exposure to virtual environments: Sickness curves during and across sessions." Presence: Teleoperators and Virtual Environments **9**(5): 466-475.
- Kohn, Zachariah (1999). Using Inertial Data for Locomotion in Virtual Environments. Chapel Hill, Dept. of Comp. Sci., Univ. of North Carolina at Chapel Hill: 6.
- Kolasinski, Eugenia (1995). Simulator Sickness in Virtual Environments. Alexandria, Virginia, U.S. Army Research Institute for the Behavioral and Social Sciences.
- Kolasinski, Eugenia, S.A. Jones, Robert S. Kennedy and R. D. Gilson (1994). Postural stability and its relation to simulator sickness. Proceedings of the 38st Annual Meeting of the Human Factors and Ergonomics Society. **2**: 980.
- Lackner, James R. (1977a). "Induction of Illusory Self-Rotation and Nystagmus by a Rotating Sound-Field." Aviation, Space, and Environmental Medicine.
- Lackner, James R. (1977b). Some Mechanisms Underlying Sensory and Postural Stability in Man. Perception. Held, R., H. Leibowitz and H.-L. Teuber. New York, Springer-Verlag. **8**: 805-845.
- LaMotte, Robert H. and Mandayam A. Srinivasan (1991). Surface Microgeometry: Tactile Perception and Neural Encoding. Information Processing in the Somatosensory System. Wenner-Gren Intl. Symposium Series, MacMillan Press.
- LaViola Jr., Joseph J. (2000). "A Discussion of Cybersickness in Virtual Environments." SIGCHI Bulletin **32**(1): 47-55.
- LaViola Jr., Joseph J., Daniel Acevedo Feliz, Daniel F. Keefe and Robert C. Zeleznik (2001). Hands-Free Multi-Scale Navigation in Virtual Environments. ACM Symposium on Interactive 3D Graphics, Research Triangle Park, NC.
- Lok, Benjamin Chak Lum (2002). Interacting With Real Objects Significantly Enhances Cognitive Virtual Environment Tasks (Doctoral Dissertation), Department of Computer Science, University of North Carolina at Chapel Hill. Chapel Hill
- Lorenzo, Max (1995). OSIRIS - Immersive virtual environment for night-vision battle simulation. SIGGRAPH 95 Visual Proceedings. Los Angeles.
- Mach, Ernst W. J. W. (1875). Grundlinien der Lehre von der Bewegungsempfindungen. Leipzig, Wilhelm Engelmann.

- MacKay, D. M. (1966). Cerebral Organization and the Conscious Control of Action. Brain and Conscious Experience: Study Week September 28 to October 4, 1964, of the Pontificia Academia Scientiarum. New York, Springer-Verlag: 422-445.
- Martini (1998). Fundamentals of Anatomy and Physiology. Upper Saddle River, Prentice Hall.
- McCrea, Robert A., Greg T. Gdowski, Richard Boyle and Timothy Belton (1999). "Firing Behavior of Vestibular Neurons During Active and Passive Head Movements: Vestibulo-Spinal and Other Non-Eye-Movement Related Neurons." J Neurophysiol **82**(1): 416-428.
- Meehan, Michael (2001). Physiological Reaction as an Objective Measure of Presence in Virtual Environments (Doctoral Dissertation), Department of Computer Science, University of North Carolina at Chapel Hill, Chapel Hill
- Meehan, Michael, Sharif Razzaque, Mary C. Whitton and Frederick P. Brooks (2003). Effects of Latency on Presence in Stressful Virtual Environments. Proceedings of IEEE Virtual Reality, Los Angeles, IEEE Computer Society.
- Mergner, T., F. Hlavacka and Schweigart G. (1993). "Interaction of vestibular and proprioceptive inputs." Journal of Vestibular Research **3**(1): 41-57.
- Miller, J. W. and J. E. Goodson (1960). "Motion sickness in a helicopter simulator." Aerospace Medicine **31**: 204-212.
- Milne, Alan A. (1926). ...In Which Pooh and Piglet Go Hunting and Nearly Catch a Woozle. Winnie-the-Pooh; with decorations by Ernest H. Shepard. New York, E.P. Dutton & Company: 158.
- Miné, Mark Raymond (1997). Exploiting Prioception in Virtual-Environment Interaction (Doctoral Dissertation), Computer Science, University of North Carolina at Chapel Hill. Chapel Hill
- "Mystery Spot." (1997). IllusionWorks, L.L.C. Retrieved June 25, 2005, from http://psylux.psych.tu-dresden.de/i1/kaw/diverses%20Material/www.illusionworks.com/html/mystery_spot.html.
- Netter, Frank Henry (1997). Atlas of human anatomy. East Hanover, N.J., Novartis.
- Nichols, Sarah C., Amanda D. Ramsey, Sue V. G. Cobb, Helen Neale, Mirabelle D'Cruz and John Wilson, R. (2000). Incidence of Virtual Reality Induced Symptoms and Effects (VRISE) in desktop and projection display systems, Health and Safety Executive (HSE), U.K.
- Pang, T. K., K. Y. Lim and S.M. Quek (2000). Design and development of a Stereoscopic Display Rig for a Comparative Assessment of Remote of Freight Handling Performance. Proceedings of the Joint Conference of APCHI 2000 & ASEAN Ergonomics 2000, Singapore, Elsevier Science.
- Parker, Donald E. (2003). "Spatial perception changes associated with space flight: implications for adaptation to altered inertial environments." Journal of Vestibular Research **13**: 331-344.
- Pausch, Randy, Tom Crea and Matthew J. Conway (1993). "A Literature Survey for Virtual Environments: Military Flight Simulator Visual Systems and Simulator Sickness." Presence: Teleoperators and Virtual Environments **1**(3).
- Peterson, B., M. Wells, T. Furness and E. Hunt (1998). The Effects of the Interface on Navigation in Virtual Environments. Proceedings of Human Factors and Ergonomics Society 1998 Annual Meeting.

- Pollefeys, Marc., Reinhard Koch, Maarten Vergauwen and Luc J. Van Gool (1998). Metric 3D Surface Reconstruction from Uncalibrated Image Sequences. Proceedings of the European Workshop on 3D Structure from Multiple Images of Large-Scale Environments, Springer-Verlag.
- Pregibon (1981). "Logistic Regression Diagnostics." The Annals of Statistics 9(4): 705-724.
- Prothero, Jerrold D. (1998). The Role of rest Frames in Vection, Presence and Motion Sickness (Doctoral Dissertation), Department of Industrial Engineering, University of Washington
- Prothero, Jerrold D., Mark H. Draper, Thomas A. Furness, Donald E. Parker and M. J. Wells (1999). "The use of an independent visual background to reduce simulator side-effects." Aviation Space Environmental Medicine 70: 277-283.
- Prothero, Jerrold D., Hunter G. Hoffman, Thomas A. Furness, Donald E. Parker and Maxwell J. Wells (1995). Foreground/background manipulations Affect Presence. Proceedings of Human Factors and Ergonomics Society.
- Regan, David Martin, Lloyd Kaufman and Janet Lincoln (1986). Motion in Depth and Visual Acceleration (Chapter 19). Handbook of Perception and Human Performance. Boff, K. R., L. Kaufman and J. P. Thomas. New York, Wiley-Interscience. 1: 19.1- 19.46.
- Riecke, Bernhard E., Henricus A. H. C. van Veen and Heinrich H. Bulthoff (2002). "Visual Homing Is Possible Without Landmarks: A Path Integration Study in Virtual Reality." Presence: Teleoperators and Virtual Environments 11(5): 443-473.
- Robinett, Warren and R. Holloway (1992). Implementation of flying, scaling and grabbing in virtual environments. ACM Symposium on Interactive 3D Graphics, Cambridge, Mass., ACM.
- Rolfé, J. M. and K. J. Staples (1986). Flight Simulation. New York, Cambridge University Press.
- Schweigart, Georg (2002). "Neck proprioception compensates for age-related deterioration of vestibular self-motion perception."
- Slater, M. and A. Steed (2000). "A Virtual Presence Counter." Presence: Teleoperators and Virtual Environments 9(5): 413-434.
- Slater, M., M. Usoh and A. Steed (1995). "Taking steps: The influence of a walking metaphor on presence in virtual reality." ACM Transactions on Computer Human Interaction (TOCHI) 2(3): 201-219.
- Slater, Mel (1999). "Measuring presence: A response to the Witmer and Singer presence questionnaire." Presence: Teleoperators and Virtual Environments 8(5): 560-566.
- Slater, Mel, Anthony Steed, J. McCarthy and F. Margingelli (1998). "The Influence of Body Movement on Subjective Presence in Virtual Environments." Human Factors: The Journal of the Human Factors and Ergonomics Society 40(3): 469-477.
- Smith, Steven W. (1997). The Scientist and Engineer's Guide to Digital Signal Processing. San Diego, California Technical Publishing.
- Snodgrass, Joan Gay, Gail Levy-Berger and Martin Haydon (1985). Human Experimental Psychology. New York, Oxford University Press.
- Stark, Peter A. (1998). Low- and High-Pass Filters. Electronics 101 Textbook (in draft). Mt. Kisco, NY, Star-K Systems. 20.

- Stoakley, R., M. J. Conway and Randy Pausch (1995). Virtual Reality on a WIM: Interactive Worlds in Miniature. Proceedings of Human Factors and Computing Systems CHI'95.
- Stoffregen, T. A. and G. E. Riccio (1988). "An ecological theory of orientation and the vestibular system." Psychological Review **91**(1): 3-14.
- Stoffregen, T. A. and G. E. Riccio (1991). "An ecological critique of the sensory conflict theory of motion sickness." Ecological Psychology **3**(3): 15-194.
- Strachan, Ian W. (2001). Motion Cueing in the Real World and in Simulators (draft), Jane's Simulation and Training Systems.
- Taylor, M. M. and C. D. Creelman (1967). "PEST: Efficient estimates on probability functions." Journal of the Acoustical Society of America **41**: 782-787.
- Templeman, James N., Patricia S. Denbrook and Linda E. Sibert (1999). "Virtual Locomotion: Walking in Place through Virtual Environments." Presence: Teleoperators and Virtual Environments **8**(6): 598-617.
- Usoh, Martin, Kevin Arthur, Mary C. Whitton, Rui Bastos, Anthony Steed, Mel Slater and Frederick P. Brooks (1999). Walking > walking-in-place > flying in virtual environments. Proceedings of ACM SIGGRAPH 99, Los Angeles.
- "The VMS Motion Base." (2005), NASA Ames Research Center, Retrieved June 25, 2005, from <http://www.simlabs.arc.nasa.gov/vms/motionb.html>.
- Warren Jr., William H, Bruce A Kay, Wendy D Zosh, Andrew P. Duchon and Stephanie Sahuc (2001). "Optic flow is used to control human walking." nature neuroscience **4**(2): 213-216.
- Weber, K. D., W. A. Fletcher, C. R. Gordon, G Melvill Jones and E. W. Block (1998). "Motor learning in the "podokinetic" system and its role in the spatial orientation during locomotion." Experimental Brain Research **120**(3): 377-385.
- Welch, R. B. and D. H. Warren (1986). Intersensory interactions. Handbook of Perception and Human Performance. Boff, K. R., L. Kaufman and J. P. Thomas. New York, Wiley-Interscience.
- Whitton, Mary C., Joseph V. Cohn, Jeff Feasel, Paul Zimmons, Sharif Razzaque, Sarah J. Poulton, Brandi McLeod and Frederick P. Brooks (2005). Comparing VE Locomotion Interfaces. in Proceedings of Virtual Reality 2005, Bonn, Germany, IEEE Computer Society.
- Zimmons, Paul (2004). The Influence of Lighting Quality on Presence and Task Performance in Virtual Environments (Doctoral Dissertation), Department of Computer Science, University of North Carolina at Chapel Hill. Chapel Hill

